
The Python Library Reference

发布 3.6.8

**Guido van Rossum
and the Python development team**

六月 19, 2019

**Python Software Foundation
Email: docs@python.org**

1	概述	3
2	内置函数	5
3	内置常量	25
3.1	由 site 模块添加的常量	26
4	内置类型	27
4.1	逻辑值检测	27
4.2	Boolean Operations —and, or, not	28
4.3	比较	28
4.4	数字类型—int, float, complex	29
4.5	迭代器类型	34
4.6	序列类型—list, tuple, range	35
4.7	文本序列类型—str	40
4.8	二进制序列类型—bytes, bytearray, memoryview	49
4.9	集合类型—set, frozenset	69
4.10	映射类型—dict	71
4.11	上下文管理器类型	74
4.12	其他内置类型	75
4.13	特殊属性	77
5	内置异常	79
5.1	基类	80
5.2	具体异常	80
5.3	警告	85
5.4	异常层次结构	86
6	文本处理服务	89
6.1	string —常见的字符串操作	89
6.2	re —正则表达式操作	99
6.3	difflib —计算差异的辅助工具	117
6.4	textwrap —Text wrapping and filling	127
6.5	unicodedata —Unicode 数据库	130
6.6	stringprep —Internet String Preparation	132
6.7	readline —GNU readline interface	133
6.8	rlcompleter —GNU readline 的完成函数	138

7	二进制数据服务	139
7.1	struct — Interpret bytes as packed binary data	139
7.2	codecs — Codec registry and base classes	145
8	数据类型	163
8.1	datetime — 基础日期和时间数据类型	163
8.2	calendar — 日历相关函数	192
8.3	collections — 容器数据类型	195
8.4	collections.abc — 容器的抽象基类	210
8.5	heapq — 堆队列算法	214
8.6	bisect — Array bisection algorithm	218
8.7	array — Efficient arrays of numeric values	220
8.8	weakref — 弱引用	223
8.9	types — Dynamic type creation and names for built-in types	230
8.10	copy — 浅层 (shallow) 和深层 (deep) 复制操作	234
8.11	pprint — 数据美化输出	235
8.12	reprlib — Alternate repr() implementation	240
8.13	enum — Support for enumerations	242
9	数字和数学模块	261
9.1	numbers — 数字的抽象基类	261
9.2	math — 数学函数	264
9.3	cmath — 关于复数的数学函数	269
9.4	decimal — 十进制定点和浮点运算	272
9.5	fractions — 分数	298
9.6	random — 生成伪随机数	300
9.7	statistics — Mathematical statistics functions	306
10	函数式编程模块	313
10.1	itertools — 为高效循环而创建迭代器的函数	313
10.2	functools — 高阶函数和可调对象上的操作	327
10.3	operator — 标准运算符替代函数	333
11	文件和目录访问	341
11.1	pathlib — 面向对象的文件系统路径	341
11.2	os.path — 常见路径操作	357
11.3	fileinput — Iterate over lines from multiple input streams	361
11.4	stat — Interpreting stat() results	364
11.5	filecmp — File and Directory Comparisons	369
11.6	tempfile — Generate temporary files and directories	371
11.7	glob — Unix style pathname pattern expansion	375
11.8	fnmatch — Unix filename pattern matching	376
11.9	linecache — Random access to text lines	377
11.10	shutil — High-level file operations	378
11.11	macpath — Mac OS 9 路径操作函数	385
12	数据持久化	387
12.1	pickle — Python 对象序列化	387
12.2	copyreg — Register pickle support functions	400
12.3	shelve — Python object persistence	401
12.4	marshal — Internal Python object serialization	403
12.5	dbm — Interfaces to Unix “databases”	404
12.6	sqlite3 — SQLite 数据库 DB-API 2.0 接口模块	408
13	数据压缩和存档	429

13.1	zlib —与 gzip 兼容的压缩	429
13.2	gzip —对 gzip 格式的支持	432
13.3	bz2 —对 bzip2 压缩算法的支持	435
13.4	lzma —用 LZMA 算法压缩	438
13.5	zipfile —使用 ZIP 存档	443
13.6	tarfile —读写 tar 归档文件	451
14	文件格式	461
14.1	csv —CSV 文件读写	461
14.2	configparser —Configuration file parser	468
14.3	netrc —netrc file processing	485
14.4	xdrllib —Encode and decode XDR data	486
14.5	plistlib —Generate and parse Mac OS X .plist files	489
15	加密服务	493
15.1	hashlib —安全哈希与消息摘要	493
15.2	hmac —基于密钥的消息验证	503
15.3	secrets —Generate secure random numbers for managing secrets	505
16	通用操作系统服务	509
16.1	os —操作系统接口模块	509
16.2	io —处理流的核心工具	554
16.3	time —时间的访问和转换	566
16.4	argparse —命令行选项、参数和子命令解析器	573
16.5	getopt —C-style parser for command line options	603
16.6	模块 logging —Python 的日志记录工具	605
16.7	logging.config —日志记录配置	621
16.8	logging.handlers —Logging handlers	631
16.9	getpass —便携式密码输入工具	643
16.10	curses —终端字符单元显示的处理	644
16.11	curses.textpad —Text input widget for curses programs	661
16.12	curses.ascii —Utilities for ASCII characters	662
16.13	curses.panel —A panel stack extension for curses	665
16.14	platform —获取底层平台的标识数据	666
16.15	errno —Standard errno system symbols	669
16.16	ctypes —Python 的外部函数库	675
17	并发执行	709
17.1	threading —基于线程的并行	709
17.2	multiprocessing —基于进程的并行	720
17.3	concurrent 包	761
17.4	concurrent.futures —启动并行任务	761
17.5	subprocess —子进程管理	767
17.6	sched —事件调度器	782
17.7	queue —一个同步的队列类	783
17.8	dummy_threading —可直接替代 threading 模块。	785
17.9	_thread —底层多线程 API	786
17.10	_dummy_thread —_thread 的替代模块	787
18	Interprocess Communication and Networking	789
18.1	socket —底层网络接口	789
18.2	ssl —TLS/SSL wrapper for socket objects	810
18.3	select —Waiting for I/O completion	841
18.4	selectors —高级 I/O 复用库	848
18.5	asyncio —Asynchronous I/O, event loop, coroutines and tasks	852

18.6	asyncore —异步 socket 处理器	912
18.7	asynchat —异步 socket 指令/响应处理器	917
18.8	signal —Set handlers for asynchronous events	919
18.9	mmap —内存映射文件支持	924
19	互联网数据处理	929
19.1	email —电子邮件与 MIME 处理包	929
19.2	json —JSON 编码和解码器	986
19.3	mailcap —Mailcap file handling	995
19.4	mailbox —Manipulate mailboxes in various formats	996
19.5	mimetypes —Map filenames to MIME types	1014
19.6	base64 —Base16, Base32, Base64, Base85 数据编码	1017
19.7	binhex —对 binhex4 文件进行编码和解码	1019
19.8	binascii —二进制和 ASCII 码互转	1020
19.9	quopri —Encode and decode MIME quoted-printable data	1022
19.10	uu —Encode and decode uuencode files	1022
20	结构化标记处理工具	1025
20.1	html —超文本标记语言支持	1025
20.2	html.parser —简单的 HTML 和 XHTML 解析器	1026
20.3	html.entities —HTML 一般实体的定义	1030
20.4	XML 处理模块	1030
20.5	xml.etree.ElementTree —The ElementTree XML API	1032
20.6	xml.dom —The Document Object Model API	1047
20.7	xml.dom.minidom —Minimal DOM implementation	1057
20.8	xml.dom.pulldom —Support for building partial DOM trees	1061
20.9	xml.sax —Support for SAX2 parsers	1063
20.10	xml.sax.handler —Base classes for SAX handlers	1065
20.11	xml.sax.saxutils —SAX Utilities	1070
20.12	xml.sax.xmlreader —Interface for XML parsers	1071
20.13	xml.parsers.expat —Fast XML parsing using Expat	1075
21	互联网协议和支持	1085
21.1	webbrowser —方便的 Web 浏览器控制器	1085
21.2	cgi —Common Gateway Interface support	1087
21.3	cgitb —Traceback manager for CGI scripts	1094
21.4	wsgiref —WSGI Utilities and Reference Implementation	1095
21.5	urllib —URL 处理模块	1104
21.6	urllib.request —用于打开 URL 的可扩展库	1104
21.7	urllib.response —Response classes used by urllib	1122
21.8	urllib.parse —Parse URLs into components	1122
21.9	urllib.error —Exception classes raised by urllib.request	1129
21.10	urllib.robotparser —Parser for robots.txt	1130
21.11	http —HTTP 模块	1131
21.12	http.client —HTTP 协议客户端	1133
21.13	ftplib —FTP protocol client	1139
21.14	poplib —POP3 protocol client	1144
21.15	imaplib —IMAP4 protocol client	1147
21.16	nntplib —NNTP protocol client	1153
21.17	smtplib —SMTP 协议客户端	1160
21.18	smtpd —SMTP Server	1167
21.19	telnetlib —Telnet client	1170
21.20	uuid —UUID objects according to RFC 4122	1173
21.21	socketserver —A framework for network servers	1176

21.22	http.server —HTTP 服务器	1184
21.23	http.cookies —HTTP state management	1189
21.24	http.cookiejar —Cookie handling for HTTP clients	1192
21.25	xmlrpc —XMLRPC 服务端与客户端模块	1201
21.26	xmlrpc.client —XML-RPC client access	1201
21.27	xmlrpc.server —Basic XML-RPC servers	1209
21.28	ipaddress —IPv4/IPv6 manipulation library	1214
22	多媒体服务	1227
22.1	audioop —Manipulate raw audio data	1227
22.2	aifc —Read and write AIFF and AIFC files	1230
22.3	sunau —读写 Sun AU 文件	1232
22.4	wave —读写 WAV 格式文件	1235
22.5	chunk —Read IFF chunked data	1238
22.6	colorsys —Conversions between color systems	1239
22.7	imghdr —推测图像类型	1240
22.8	sndhdr —推测声音文件的类型	1240
22.9	ossaudiodev —Access to OSS-compatible audio devices	1241
23	国际化	1247
23.1	gettext —多语种国际化服务	1247
23.2	locale —国际化服务	1255
24	程序框架	1263
24.1	turtle —海龟绘图	1263
24.2	cmd —支持面向行的命令解释器	1296
24.3	shlex —Simple lexical analysis	1301
25	Tk 图形用户界面 (GUI)	1307
25.1	tkinter —Tcl/Tk 的 Python 接口	1307
25.2	tkinter.ttk —Tk themed widgets	1318
25.3	tkinter.tix —Extension widgets for Tk	1335
25.4	tkinter.scrolledtext —滚动文字控件	1340
25.5	IDLE	1340
25.6	其他图形用户界面 (GUI) 包	1349
26	开发工具	1351
26.1	typing —类型标注支持	1351
26.2	pydoc —Documentation generator and online help system	1366
26.3	doctest —测试交互性的 Python 示例	1368
26.4	unittest —单元测试框架	1391
26.5	unittest.mock —mock object library	1418
26.6	unittest.mock 上手指南	1453
26.7	2to3 - 自动将 Python 2 代码转为 Python 3 代码	1472
26.8	test —Regression tests package for Python	1477
26.9	test.support —Utilities for the Python test suite	1480
27	调试和分析	1487
27.1	bdb —Debugger framework	1487
27.2	faulthandler —Dump the Python traceback	1491
27.3	pdb —Python 的调试器	1494
27.4	The Python Profilers	1499
27.5	timeit —测量小代码片段的执行时间	1507
27.6	trace —Trace or track Python statement execution	1512
27.7	tracemalloc —Trace memory allocations	1515

28 软件打包和分发	1525
28.1 distutils —构建和安装 Python 模块	1525
28.2 ensurepip —Bootstrapping the pip installer	1526
28.3 venv —创建虚拟环境	1527
28.4 zipapp —Manage executable python zip archives	1535
29 Python 运行时服务	1543
29.1 sys —系统相关的参数和函数	1543
29.2 sysconfig —Provide access to Python' s configuration information	1558
29.3 builtins —内建对象	1562
29.4 __main__ —顶层脚本环境	1563
29.5 warnings —Warning control	1563
29.6 contextlib —Utilities for with-statement contexts	1568
29.7 abc —抽象基类	1579
29.8 atexit —退出处理器	1584
29.9 traceback —打印或检索堆栈回溯	1585
29.10 __future__ —Future 语句定义	1591
29.11 gc —垃圾回收器接口	1592
29.12 inspect —检查对象	1595
29.13 site —Site-specific configuration hook	1610
29.14 fpectl —Floating point exception control	1613
30 自定义 Python 解释器	1615
30.1 code —Interpreter base classes	1615
30.2 codeop —Compile Python code	1617
31 导入模块	1619
31.1 zipimport —Import modules from Zip archives	1619
31.2 pkgutil —Package extension utility	1621
31.3 modulefinder —查找脚本使用的模块	1624
31.4 runpy —Locating and executing Python modules	1625
31.5 importlib —The implementation of import	1627
32 Python 语言服务	1645
32.1 parser —Access Python parse trees	1645
32.2 ast —抽象语法树	1649
32.3 symtable —Access to the compiler' s symbol tables	1654
32.4 symbol —与 Python 解析树一起使用的常量	1657
32.5 token —与 Python 解析树一起使用的常量	1657
32.6 keyword —检验 Python 关键字	1658
32.7 tokenize —Tokenizer for Python source	1659
32.8 tabnanny —模糊缩进检测	1662
32.9 pyclbr —Python class browser support	1663
32.10 py_compile —Compile Python source files	1664
32.11 compileall —Byte-compile Python libraries	1665
32.12 dis —Python 字节码反汇编器	1668
32.13 pickletools —Tools for pickle developers	1681
33 杂项服务	1683
33.1 formatter —Generic output formatting	1683
34 Windows 系统相关模块	1689
34.1 msilib —Read and write Microsoft Installer files	1689
34.2 msvcrt —Useful routines from the MS VC++ runtime	1694
34.3 winreg —Windows 注册表访问	1696

34.4	winsound —Sound-playing interface for Windows	1704
35	Unix 专有服务	1707
35.1	posix —The most common POSIX system calls	1707
35.2	pwd —用户密码数据库	1708
35.3	spwd —The shadow password database	1709
35.4	grp —The group database	1710
35.5	crypt —Function to check Unix passwords	1710
35.6	termios —POSIX style tty control	1712
35.7	tty —终端控制功能	1713
35.8	pty —Pseudo-terminal utilities	1714
35.9	fcntl —The fcntl and ioctl system calls	1715
35.10	pipes —Interface to shell pipelines	1717
35.11	resource —Resource usage information	1718
35.12	nis —Interface to Sun' s NIS (Yellow Pages)	1722
35.13	Unix syslog 库例程	1723
36	被取代的模块	1725
36.1	optparse —Parser for command line options	1725
36.2	imp —Access the import internals	1751
37	未创建文档的模块	1757
37.1	平台特定模块	1757
A	术语对照表	1759
B	文档说明	1771
B.1	Python 文档贡献者	1771
C	历史和许可证	1773
C.1	软件历史	1773
C.2	访问 Python 或以其他方式使用 Python 的条款和条件	1774
C.3	集成软件的许可和认可	1777
D	版权	1791
	Bibliography	1793
	Python 模块索引	1795
	索引	1799

`reference-index` 描述了 Python 语言的具体语法和语义，这份库参考则介绍了与 Python 一同发行的标准库。它还描述了通常包含在 Python 发行版中的一些可选组件。

Python 标准库非常庞大，所提供的组件涉及范围十分广泛，正如以下内容目录所显示的。这个库包含了多个内置模块 (以 C 编写)，Python 程序员必须依靠它们来实现系统级功能，例如文件 I/O，此外还有大量以 Python 编写的模块，提供了日常编程中许多问题的标准解决方案。其中有些模块经过专门设计，通过将特定平台功能抽象化为平台中立的 API 来鼓励和加强 Python 程序的可移植性。

Windows 版本的 Python 安装程序通常包含整个标准库，往往还包含许多额外组件。对于类 Unix 操作系统，Python 通常会分成一系列的软件包，因此可能需要使用操作系统所提供的包管理工具来获取部分或全部可选组件。

在这个标准库以外还存在成千上万并且不断增加的其他组件 (从单独的程序、模块、软件包直到完整的应用开发框架)，访问 [Python 包索引](#) 即可获取这些第三方包。

概述

“Python 库”中包含了几种不同的组件。

它包含通常被视为语言“核心”中的一部分的数据类型，例如数字和列表。对于这些类型，Python 语言核心定义了文字的形式，并对它们的语义设置了一些约束，但没有完全定义语义。（另一方面，语言核心确实定义了语法属性，如操作符的拼写和优先级。）

这个库也包含了内置函数和异常—不需要 `import` 语句就可以在所有 Python 代码中使用的对象。有一些是由语言核心定义的，但是许多对于核心语义不是必需的，并且仅在这里描述。

不过这个库主要是由一系列的模块组成。这些模块集可以不同方式分类。有些模块是用 C 编写并内置于 Python 解释器中；另一些模块则是用 Python 编写并以源码形式导入。有些模块提供专用于 Python 的接口，例如打印栈追踪信息；有些模块提供专用于特定操作系统的接口，例如操作特定的硬件；另一些模块则提供针对特定应用领域的接口，例如万维网。有些模块在所有更新和移植版本的 Python 中可用；另一些模块仅在底层系统支持或要求时可用；还有些模块则仅当编译和安装 Python 时选择了特定配置选项时才可用。

本手册以“从内到外”的顺序组织：首先描述内置函数、数据类型和异常，最后是根据相关性进行分组的各种模块。

这意味着如果你从头开始阅读本手册，并在感到厌烦时跳到下一章，你仍能对 Python 库的可用模块和所支持的应用领域有个大致了解。当然，你并非必须如同读小说一样从头读到尾—你也可以先浏览内容目录（在手册开头），或在索引（在手册末尾）中查找某个特定函数、模块或条目。最后，如果你喜欢随意学习某个主题，你可以选择一个随机页码（参见 `random` 模块）并读上一两小节。无论你想以怎样的顺序阅读本手册，还是建议先从内置函数这一章开始，因为本手册的其余内容都需要你熟悉其中的基本概念。

让我们开始吧！

内置函数

Python 解释器内置了很多函数和类型，您可以在任何时候使用它们。以下按字母表顺序列出它们。

		内置函数		
<i>abs()</i>	<i>dict()</i>	<i>help()</i>	<i>min()</i>	<i>setattr()</i>
<i>all()</i>	<i>dir()</i>	<i>hex()</i>	<i>next()</i>	<i>slice()</i>
<i>any()</i>	<i>divmod()</i>	<i>id()</i>	<i>object()</i>	<i>sorted()</i>
<i>ascii()</i>	<i>enumerate()</i>	<i>input()</i>	<i>oct()</i>	<i>staticmethod()</i>
<i>bin()</i>	<i>eval()</i>	<i>int()</i>	<i>open()</i>	<i>str()</i>
<i>bool()</i>	<i>exec()</i>	<i>isinstance()</i>	<i>ord()</i>	<i>sum()</i>
<i>bytearray()</i>	<i>filter()</i>	<i>issubclass()</i>	<i>pow()</i>	<i>super()</i>
<i>bytes()</i>	<i>float()</i>	<i>iter()</i>	<i>print()</i>	<i>tuple()</i>
<i>callable()</i>	<i>format()</i>	<i>len()</i>	<i>property()</i>	<i>type()</i>
<i>chr()</i>	<i>frozenset()</i>	<i>list()</i>	<i>range()</i>	<i>vars()</i>
<i>classmethod()</i>	<i>getattr()</i>	<i>locals()</i>	<i>repr()</i>	<i>zip()</i>
<i>compile()</i>	<i>globals()</i>	<i>map()</i>	<i>reversed()</i>	<i>__import__()</i>
<i>complex()</i>	<i>hasattr()</i>	<i>max()</i>	<i>round()</i>	
<i>delattr()</i>	<i>hash()</i>	<i>memoryview()</i>	<i>set()</i>	

abs(x)
返回一个数的绝对值。实参可以是整数或浮点数。如果实参是一个复数，返回它的模。

all(iterable)
如果 *iterable* 的所有元素为真（或迭代器为空），返回 True。等价于:

```
def all(iterable):
    for element in iterable:
        if not element:
            return False
    return True
```

any(iterable)
如果 *iterable* 的任一元素为真则返回 True。如果迭代器为空，返回 False。等价于:

```
def any(iterable):
    for element in iterable:
        if element:
            return True
    return False
```

ascii(object)

就像函数 `repr()`，返回一个对象可打印的字符串，但是 `repr()` 返回的字符串中非 ASCII 编码的字符，会使用 `\x`、`\u` 和 `\U` 来转义。生成的字符串和 Python 2 的 `repr()` 返回的结果相似。

bin(x)

将一个整数转变为一个前缀为 “0b” 的二进制字符串。结果是一个合法的 Python 表达式。如果 `x` 不是 Python 的 `int` 对象，那它需要定义 `__index__()` 方法返回一个整数。一些例子：

```
>>> bin(3)
'0b11'
>>> bin(-10)
'-0b1010'
```

如果不一定需要前缀 “0b”，还可以使用如下的方法。

```
>>> format(14, '#b'), format(14, 'b')
('0b1110', '1110')
>>> f'{14:#b}', f'{14:b}'
('0b1110', '1110')
```

另见 `format()` 获取更多信息。

class bool([x])

返回一个布尔值，True 或者 False。`x` 使用标准的真值测试过程来转换。如果 `x` 是假的或者被省略，返回 False；其他情况返回 True。`bool` 类是 `int` 的子类（参见数字类型—`int`、`float`、`complex`）。其他类不能继承自它。它只有 False 和 True 两个实例（参见布尔值）。

class bytearray([source[, encoding[, errors]]])

返回一个新的 bytes 数组。`bytearray` 类是一个可变序列，包含范围为 $0 \leq x < 256$ 的整数。它有可变序列大部分常见的方法，见可变序列类型的描述；同时有 `bytes` 类型的大部分方法，参见 `bytes` 和 `bytearray` 操作。

可选形参 `source` 可以用不同的方式来初始化数组：

- 如果是一个 `string`，您必须提供 `encoding` 参数（`errors` 参数仍是可选的）；`bytearray()` 会使用 `str.encode()` 方法来将 `string` 转变成 `bytes`。
- 如果是一个 `integer`，会初始化大小为该数字的数组，并使用 `null` 字节填充。
- 如果是一个符合 `buffer` 接口的对象，该对象的只读 `buffer` 会用来初始化字节数组。
- 如果是一个 `iterable` 可迭代对象，它的元素的范围必须是 $0 \leq x < 256$ 的整数，它会被用作数组的初始内容。

如果没有实参，则创建大小为 0 的数组。

另见二进制序列类型—`bytes`、`bytearray`、`memoryview` 和 `bytearray` 对象。

class bytes([source[, encoding[, errors]]])

返回一个新的 “bytes” 对象，是一个不可变序列，包含范围为 $0 \leq x < 256$ 的整数。`bytes` 是 `bytearray` 的不可变版本 - 它有其中不改变序列的方法和相同的索引、切片操作。

因此，构造函数的实参和 `bytearray()` 相同。

字节对象还可以用字面值创建，参见 `strings`。

另见二进制序列类型—*bytes*, *bytearray*, *memoryview*, *bytes* 对象和*bytes*和*bytearray*操作。

callable (*object*)

如果实参 *object* 是可调用的, 返回 *True*, 否则返回 *False*。如果返回真, 调用仍可能会失败; 但如果返回假, 则调用 *object* 肯定会失败。注意类是可调用的 (调用类会返回一个新的实例)。如果实例的类有 `__call__()` 方法, 则它是可调用。

3.2 新版功能: 这个函数一开始在 Python 3.0 被移除了, 但在 Python 3.2 被重新加入。

chr (*i*)

返回 Unicode 码位为整数 *i* 的字符的字符串格式。例如, `chr(97)` 返回字符串 'a', `chr(8364)` 返回字符串 '€'。这是 `ord()` 的逆函数。

实参的合法范围是 0 到 1,114,111 (16 进制表示是 0x10FFFF)。如果 *i* 超过这个范围, 会触发 *ValueError* 异常。

@classmethod

把一个方法封装成类方法。

一个类方法把类自己作为第一个实参, 就像一个实例方法把实例自己作为第一个实参。请用以下习惯来声明类方法:

```
class C:
    @classmethod
    def f(cls, arg1, arg2, ...): ...
```

The `@classmethod` form is a function *decorator* – see the description of function definitions in function for details.

It can be called either on the class (such as `C.f()`) or on an instance (such as `C().f()`). The instance is ignored except for its class. If a class method is called for a derived class, the derived class object is passed as the implied first argument.

Class methods are different than C++ or Java static methods. If you want those, see `staticmethod()` in this section.

For more information on class methods, consult the documentation on the standard type hierarchy in types.

compile (*source*, *filename*, *mode*, *flags=0*, *dont_inherit=False*, *optimize=-1*)

将 *source* 编译成代码或 AST 对象。代码对象可以被 `exec()` 或 `eval()` 执行。*source* 可以是常规的字符串、字节字符串, 或者 AST 对象。参见 *ast* 模块的文档了解如何使用 AST 对象。

filename 实参需要是代码读取的文件名; 如果代码不需要从文件中读取, 可以传入一些可辨识的值 (经常会使用 '`<string>`').

mode 实参指定了编译代码必须用的模式。如果 *source* 是语句序列, 可以是 'exec'; 如果是单一表达式, 可以是 'eval'; 如果是单个交互式语句, 可以是 'single'。(在最后一种情况下, 如果表达式执行结果不是 *None* 将会被打印出来。)

可选参数 *flags* 和 *dont_inherit* 控制在编译 *source* 时要用到哪个 *future* 语句。如果两者都未提供 (或都为零) 则会使用调用 `compile()` 的代码中有效的 *future* 语句来编译代码。如果给出了 *flags* 参数但没有 *dont_inherit* (或是为零) 则 *flags* 参数所指定的以及那些无论如何都有效的 *future* 语句会被使用。如果 *dont_inherit* 为一个非零整数, 则只使用 *flags* 参数 – 在调用外围有效的 *future* 语句将被忽略。

Future 语句使用比特位来指定, 多个语句可以通过按位或来指定。具体特性的比特位可以通过 `__future__` 模块中的 `_Feature` 类的实例的 `compiler_flag` 属性来获得。

optimize 实参指定编译器的优化级别; 默认值 -1 选择与解释器的 -O 选项相同的优化级别。显式级别为 0 (没有优化; `__debug__` 为真)、1 (断言被删除, `__debug__` 为假) 或 2 (文档字符串也被删除)。

如果编译的源码不合法, 此函数会触发 *SyntaxError* 异常; 如果源码包含 *null* 字节, 则会触发 *ValueError* 异常。

如果您想分析 Python 代码的 AST 表示, 请参阅[ast.parse\(\)](#)。

注解: 在 'single' 或 'eval' 模式编译多行代码字符串时, 输入必须以至少一个换行符结尾。这使 `code` 模块更容易检测语句的完整性。

警告: 在将足够大或者足够复杂的字符串编译成 AST 对象时, Python 解释器有可能会因为 Python AST 编译器的栈深度限制而崩溃。

在 3.2 版更改: 允许使用 Windows 和 Mac 的换行符。在 'exec' 模式不再需要以换行符结尾。增加了 *optimize* 形参。

在 3.5 版更改: 之前 *source* 中包含 null 字节的话会触发 *TypeError* 异常。

class complex(*real*[, *imag*])

返回值为 *real* + *imag**1j 的复数, 或将字符串或数字转换为复数。如果第一个形参是字符串, 则它被解释为一个复数, 并且函数调用时必须没有第二个形参。第二个形参不能是字符串。每个实参都可以是任意的数值类型 (包括复数)。如果省略了 *imag*, 则默认值为零, 构造函数会像 *int* 和 *float* 一样进行数值转换。如果两个实参都省略, 则返回 0j。

注解: 当从字符串转换时, 字符串在 + 或 - 的周围必须不能有空格。例如 `complex('1+2j')` 是合法的, 但 `complex('1 + 2j')` 会触发 *ValueError* 异常。

数字类型—*int*, *float*, *complex* 描述了复数类型。

在 3.6 版更改: 您可以使用下划线将代码文字中的数字进行分组。

delattr(*object*, *name*)

[setattr\(\)](#) 相关的函数。实参是一个对象和一个字符串。该字符串必须是对象的某个属性。如果对象允许, 该函数将删除指定的属性。例如 `delattr(x, 'foobar')` 等价于 `del x.foobar`。

class dict(***kwarg*)

class dict(*mapping*, ***kwarg*)

class dict(*iterable*, ***kwarg*)

创建一个新的字典。*dict* 对象是一个字典类。参见 *dict* 和映射类型—*dict* 了解这个类。

其他容器类型, 请参见内置的 *list*、*set* 和 *tuple* 类, 以及 *collections* 模块。

dir(*[object]*)

如果没有实参, 则返回当前本地作用域中的名称列表。如果有实参, 它会尝试返回该对象的有效属性列表。

如果对象有一个名为 `__dir__()` 的方法, 那么该方法将被调用, 并且必须返回一个属性列表。这允许实现自定义 `__getattr__()` 或 `__getattribute__()` 函数的对象能够自定义 *dir()* 来报告它们的属性。

如果对象不提供 `__dir__()`, 这个函数会尝试从对象已定义的 `__dict__` 属性和类型对象收集信息。结果列表并不总是完整的, 如果对象有自定义 `__getattr__()`, 那结果可能不准确。

默认的 *dir()* 机制对不同类型的对象行为不同, 它会试图返回最相关而不是最全的信息:

- 如果对象是模块对象, 则列表包含模块的属性名称。
- 如果对象是类型或类对象, 则列表包含它们的属性名称, 并且递归查找所有基类的属性。
- 否则, 列表包含对象的属性名称, 它的类属性名称, 并且递归查找它的类的所有基类的属性。

返回的列表按字母表排序。例如:

```

>>> import struct
>>> dir()      # show the names in the module namespace
['__builtins__', '__name__', 'struct']
>>> dir(struct) # show the names in the struct module # doctest: +SKIP
['Struct', '__all__', '__builtins__', '__cached__', '__doc__', '__file__',
 '__initializing__', '__loader__', '__name__', '__package__',
 '__clearcache', 'calcsiz', 'error', 'pack', 'pack_into',
 'unpack', 'unpack_from']
>>> class Shape:
...     def __dir__(self):
...         return ['area', 'perimeter', 'location']
>>> s = Shape()
>>> dir(s)
['area', 'location', 'perimeter']

```

注解：因为`dir()`主要是为了便于在交互式时使用，所以它会试图返回人们感兴趣的名字集合，而不是试图保证结果的严格性或一致性，它具体的行为也可能在不同版本之间改变。例如，当实参是一个类时，`metaclass`的属性不包含在结果列表中。

`divmod(a, b)`

它将两个（非复数）数字作为实参，并在执行整数除法时返回一对商和余数。对于混合操作数类型，适用双目算术运算符的规则。对于整数，结果和 $(a // b, a \% b)$ 一致。对于浮点数，结果是 $(q, a \% b)$ ， q 通常是`math.floor(a / b)`但可能会比1小。在任何情况下， $q * b + a \% b$ 和 a 基本相等；如果 $a \% b$ 非零，它的符号和 b 一样，并且 $0 \leq \text{abs}(a \% b) < \text{abs}(b)$ 。

`enumerate(iterable, start=0)`

返回一个枚举对象。`iterable`必须是一个序列，或`iterator`，或其他支持迭代的对象。`enumerate()`返回的迭代器的`__next__()`方法返回一个元组，里面包含一个计数值（从`start`开始，默认为0）和通过迭代`iterable`获得的值。

```

>>> seasons = ['Spring', 'Summer', 'Fall', 'Winter']
>>> list(enumerate(seasons))
[(0, 'Spring'), (1, 'Summer'), (2, 'Fall'), (3, 'Winter')]
>>> list(enumerate(seasons, start=1))
[(1, 'Spring'), (2, 'Summer'), (3, 'Fall'), (4, 'Winter')]

```

等价于：

```

def enumerate(sequence, start=0):
    n = start
    for elem in sequence:
        yield n, elem
        n += 1

```

`eval(expression, globals=None, locals=None)`

实参是一个字符串，以及可选的`globals`和`locals`。`globals`实参必须是一个字典。`locals`可以是任何映射对象。

`expression`参数会作为一个Python表达式（从技术上说是一个条件列表）被解析并求值，使用`globals`和`locals`字典作为全局和局部命名空间。如果`globals`字典存在且不包含以`__builtins__`为键的值，则会在解析`expression`之前插入以此为键的对内置模块`builtins`的字典的引用。这意味着`expression`通常具有对标准`builtins`模块的完全访问权限且受限的环境会被传播。如果省略`locals`字典则其默认值为`globals`字典。如果两个字典同时省略，表达式会在`eval()`被调用的环境中执行。返回值为表达式求值的结果。语法错误将作为异常被报告。例如：

```
>>> x = 1
>>> eval('x+1')
2
```

这个函数也可以用来执行任何代码对象（如`compile()`创建的）。这种情况下，参数是代码对象，而不是字符串。如果编译该对象时的 *mode* 实参是 'exec' 那么`eval()` 返回值为 `None`。

提示：`exec()` 函数支持动态执行语句。`globals()` 和 `locals()` 函数各自返回当前的全局和本地字典，因此您可以将它们传递给`eval()` 或`exec()` 来使用。

另外可以参阅`ast.literal_eval()`，该函数可以安全执行仅包含文字的表达式字符串。

exec(*object* [, *globals* [, *locals*]])

这个函数支持动态执行 Python 代码。*object* 必须是字符串或者代码对象。如果是字符串，那么该字符串将被解析为一系列 Python 语句并执行（除非发生语法错误）。¹ 如果是代码对象，它将被直接执行。在任何情况下，被执行的代码都需要和文件输入一样是有效的（见参考手册中关于文件输入的章节）。请注意即使在传递给`exec()` 函数的代码的上下文中，`return` 和 `yield` 语句也不能在函数定义之外使用。该函数返回值为 `None`。

无论哪种情况，如果省略了可选参数，代码将在当前范围内执行。如果提供了 *globals* 参数，就必须是字典类型，而且会被用作全局和本地变量。如果同时提供了 *globals* 和 *locals* 参数，它们分别被用作全局和本地变量。如果提供了 *locals* 参数，则它可以是任何映射型的对象。请记住在模块层级，全局和本地变量是相同的字典。如果 `exec` 有两个不同的 *globals* 和 *locals* 对象，代码就像嵌入在类定义中一样执行。

如果 *globals* 字典不包含 `__builtins__` 键值，则将为该键插入对内置 `builtins` 模块字典的引用。因此，在将执行的代码传递给`exec()` 之前，可以通过将自己的 `__builtins__` 字典插入到 *globals* 中来控制可以使用哪些内置代码。

注解： 内置`globals()` 和`locals()` 函数各自返回当前的全局和本地字典，因此可以将它们传递给`exec()` 的第二个和第三个实参。

注解： 默认情况下，*locals* 的行为如下面`locals()` 函数描述的一样：不要试图改变默认的 *locals* 字典。如果您想在`exec()` 函数返回时知道代码对 *locals* 的变动，请明确地传递 *locals* 字典。

filter(*function*, *iterable*)

用 *iterable* 中函数 *function* 返回真的那些元素，构建一个新的迭代器。*iterable* 可以是一个序列，一个支持迭代的容器，或一个迭代器。如果 *function* 是 `None`，则会假设它是一个身份函数，即 *iterable* 中所有返回假的元素会被移除。

请注意，`filter(function, iterable)` 相当于一个生成器表达式，当 *function* 不是 `None` 的时候为 `(item for item in iterable if function(item))`；*function* 是 `None` 的时候为 `(item for item in iterable if item)`。

请参阅`itertools.filterfalse()` 了解，只有 *function* 返回 `false` 时才选取 *iterable* 中元素的补充函数。

class float(*x*)

返回从数字或字符串 *x* 生成的浮点数。

如果实参是字符串，则它必须是包含十进制数字的字符串，字符串前面可以有符号，之前也可以有空格。可选的符号有 '+' 和 '-'；'+' 对创建的值没有影响。实参也可以是 NaN（非数字）、正负无穷大的字符串。确切地说，除去首尾的空格后，输入必须遵循以下语法：

¹ 解析器只接受 Unix 风格的行结束符。如果您从文件中读取代码，请确保用换行符转换模式转换 Windows 或 Mac 风格的换行符。

```

sign          ::=  "+" | "-"
infinity      ::=  "Infinity" | "inf"
nan           ::=  "nan"
numeric_value ::=  floatnumber | infinity | nan
numeric_string ::= [sign] numeric_value

```

这里，`floatnumber` 是 Python 浮点数的字符串形式，详见 `floating`。字母大小写都可以，例如，“`inf`”、“`Inf`”、“`INFINITY`”、“`iNfINity`”都可以表示正无穷大。

另一方面，如果实参是整数或浮点数，则返回具有相同值（在 Python 浮点精度范围内）的浮点数。如果实参在 Python 浮点精度范围外，则会触发 `OverflowError`。

对于一般的 Python 对象 `x`，`float(x)` 指派给 `x.__float__()`。

如果没有实参，则返回 `0.0`。

例如：

```

>>> float('+1.23')
1.23
>>> float('  -12345\n')
-12345.0
>>> float('1e-003')
0.001
>>> float('+1E6')
1000000.0
>>> float('-Infinity')
-inf

```

数字类型—`int`、`float`、`complex` 描述了浮点类型。

在 3.6 版更改：您可以使用下划线将代码文字中的数字进行分组。

format (*value* [, *format_spec*])

将 *value* 转换为 *format_spec* 控制的“格式化”表示。*format_spec* 的解释取决于 *value* 实参的类型，但是大多数内置类型使用标准格式化语法：格式规格迷你语言。

默认的 *format_spec* 是一个空字符串，它通常和调用 `str(value)` 的结果相同。

调用 `format(value, format_spec)` 会转换成 `type(value).__format__(value, format_spec)`，所以实例字典中的 `__format__()` 方法将不会调用。如果搜索到 *object* 有这个方法但 *format_spec* 不为空，*format_spec* 或返回值不是字符串，会触发 `TypeError` 异常。

在 3.4 版更改：当 *format_spec* 不是空字符串时，`object().__format__(format_spec)` 会触发 `TypeError`。

class frozenset ([*iterable*])

返回一个新的 *frozenset* 对象，它包含可选参数 *iterable* 中的元素。*frozenset* 是一个内置的类。有关此类的文档，请参阅 *frozenset* 和集合类型—*set*、*frozenset*。

请参阅内建的 *set*、*list*、*tuple* 和 *dict* 类，以及 *collections* 模块来了解其它的容器。

getattr (*object*, *name* [, *default*])

返回对象命名属性的值。*name* 必须是字符串。如果该字符串是对象的属性之一，则返回该属性的值。例如，`getattr(x, 'foobar')` 等同于 `x.foobar`。如果指定的属性不存在，且提供了 *default* 值，则返回它，否则触发 `AttributeError`。

globals ()

返回表示当前全局符号表的字典。这总是当前模块的字典（在函数或方法中，不是调用它的模块，而是定义它的模块）。

hasattr (*object*, *name*)

该实参是一个对象和一个字符串。如果字符串是对象的属性之一的名称，则返回 `True`，否则返回 `False`。（此功能是通过调用 `getattr(object, name)` 看是否有 `AttributeError` 异常来实现的。）

hash (*object*)

返回该对象的哈希值（如果它有的话）。哈希值是整数。它们在字典查找元素时用来快速比较字典的键。相同大小的数字变量有相同的哈希值（即使它们类型不同，如 `1` 和 `1.0`）。

注解： 如果对象实现了自己的 `__hash__()` 方法，请注意，`hash()` 根据机器的字长来截断返回值。另请参阅 `__hash__()`。

help ([*object*])

启动内置的帮助系统（此函数主要在交互式中使用）。如果没有实参，解释器控制台里会启动交互式帮助系统。如果实参是一个字符串，则在模块、函数、类、方法、关键字或文档主题中搜索该字符串，并在控制台上打印帮助信息。如果实参是其他任意对象，则会生成该对象的帮助页。

该函数通过 `site` 模块加入到内置命名空间。

在 3.4 版更改：`pydoc` 和 `inspect` 的变更使得可调用对象的签名信息更加全面和一致。

hex (*x*)

将整数转换为以 “0x” 为前缀的小写十六进制字符串。如果 *x* 不是 Python `int` 对象，则必须定义返回整数的 `__index__()` 方法。一些例子：

```
>>> hex(255)
'0xff'
>>> hex(-42)
'-0x2a'
```

如果要将整数转换为大写或小写的十六进制字符串，并可选择有无 “0x” 前缀，则可以使用如下方法：

```
>>> '%#x' % 255, '%x' % 255, '%X' % 255
('0xff', 'ff', 'FF')
>>> format(255, '#x'), format(255, 'x'), format(255, 'X')
('0xff', 'ff', 'FF')
>>> f'{255:#x}', f'{255:x}', f'{255:X}'
('0xff', 'ff', 'FF')
```

另见 `format()` 获取更多信息。

另请参阅 `int()` 将十六进制字符串转换为以 16 为基数的整数。

注解： 如果要获取浮点数的十六进制字符串形式，请使用 `float.hex()` 方法。

id (*object*)

返回对象的 “标识值”。该值是一个整数，在此对象的生命周期中保证是唯一且恒定的。两个生命期不重叠的对象可能具有相同的 `id()` 值。

CPython implementation detail: This is the address of the object in memory.

input ([*prompt*])

如果存在 *prompt* 实参，则将其写入标准输出，末尾不带换行符。接下来，该函数从输入中读取一行，将其转换为字符串（除了末尾的换行符）并返回。当读取到 EOF 时，则触发 `EOFError`。例如：

```
>>> s = input('--> ')
--> Monty Python's Flying Circus
>>> s
"Monty Python's Flying Circus"
```

如果加载了 `readline` 模块, `input()` 将使用它来提供复杂的行编辑和历史记录功能。

class int (x=0)

class int (x, base=10)

返回一个使用数字或字符串 `x` 生成的整数对象, 或者没有实参的时候返回 0。如果 `x` 定义了 `__int__()`, `int(x)` 返回 `x.__int__()`。如果 `x` 定义了 `__trunc__()`, 它返回 `x.__trunc__()`。对于浮点数, 它向零舍入。

如果 `x` 不是数字, 或者有 `base` 参数, `x` 必须是字符串、`bytes`、表示进制为 `base` 的整数字面值的 `bytearray` 实例。该文字前可以有 + 或 - (中间不能有空格), 前后可以有空格。一个进制为 `n` 的数字包含 0 到 `n-1` 的数, 其中 a 到 z (或 A 到 Z) 表示 10 到 35。默认的 `base` 为 10, 允许的进制有 0、2-36。2、8、16 进制的数字可以在代码中用 `0b/0B`、`0o/0O`、`0x/0X` 前缀来表示。进制为 0 将按照代码的字面量来精确解释, 最后的结果会是 2、8、10、16 进制中的一个。所以 `int('010', 0)` 是非法的, 但 `int('010')` 和 `int('010', 8)` 是合法的。

整数类型定义请参阅 [数字类型—int, float, complex](#)。

在 3.4 版更改: 如果 `base` 不是 `int` 的实例, 但 `base` 对象有 `base.__index__` 方法, 则会调用该方法来获取进制数。以前的版本使用 `base.__int__` 而不是 `base.__index__`。

在 3.6 版更改: 您可以使用下划线将代码文字中的数字进行分组。

isinstance (object, classinfo)

如果 `object` 实参是 `classinfo` 实参的实例, 或者是 (直接、间接或虚拟) 子类的实例, 则返回 `true`。如果 `object` 不是给定类型的对象, 函数始终返回 `false`。如果 `classinfo` 是对象类型 (或多个递归元组) 的元组, 如果 `object` 是其中的任何一个的实例则返回 `true`。如果 `classinfo` 既不是类型, 也不是类型元组或类型的递归元组, 那么会触发 `TypeError` 异常。

issubclass (class, classinfo)

如果 `class` 是 `classinfo` 的子类 (直接、间接或虚拟的), 则返回 `true`。`classinfo` 可以是类对象的元组, 此时 `classinfo` 中的每个元素都会被检查。其他情况, 会触发 `TypeError` 异常。

iter (object[, sentinel])

返回一个 `iterator` 对象。根据是否存在第二个实参, 第一个实参的解释是非常不同的。如果没有第二个实参, `object` 必须是支持迭代协议 (有 `__iter__()` 方法) 的集合对象, 或必须支持序列协议 (有 `__getitem__()` 方法, 且数字参数从 0 开始)。如果它不支持这些协议, 会触发 `TypeError`。如果有第二个实参 `sentinel`, 那么 `object` 必须是可调用的对象。这种情况下生成的迭代器, 每次迭代调用它的 `__next__()` 方法时都会不带实参地调用 `object`; 如果返回的结果是 `sentinel` 则触发 `StopIteration`, 否则返回调用结果。

另请参阅 [迭代器类型](#)。

One useful application of the second form of `iter()` is to read lines of a file until a certain line is reached. The following example reads a file until the `readline()` method returns an empty string:

```
with open('mydata.txt') as fp:
    for line in iter(fp.readline, ''):
        process_line(line)
```

len (s)

返回对象的长度 (元素个数)。实参可以是序列 (如 `string`、`bytes`、`tuple`、`list` 或 `range` 等) 或集合 (如 `dictionary`、`set` 或 `frozen set` 等)。

class list ([iterable])

除了是函数, `list` 也是可变序列类型, 详情请参阅 [列表和序列类型—list, tuple, range](#)。

locals()

Update and return a dictionary representing the current local symbol table. Free variables are returned by `locals()` when it is called in function blocks, but not in class blocks.

注解： 不要更改此字典的内容；更改不会影响解释器使用的局部变量或自由变量的值。

map(function, iterable, ...)

产生一个将 *function* 应用于迭代器中所有元素并返回结果的迭代器。如果传递了额外的 *iterable* 实参，*function* 必须接受相同个数的实参，并使用所有迭代器中并行获取的元素。当有多个迭代器时，最短的迭代器耗尽则整个迭代结束。如果函数的输入已经是元组实参，请参阅 `itertools.starmap()`。

max(iterable, *, key, default)**max(arg1, arg2, *args, key)**

返回可迭代对象中最大的元素，或者返回两个及以上实参中最大的。

如果只提供了一个位置参数，它必须是非空 *iterable*，返回可迭代对象中最大的元素；如果提供了两个及以上的位置参数，则返回最大的位置参数。

有两个可选只能用关键字的实参。*key* 实参指定排序函数用的参数，如传给 `list.sort()` 的。*default* 实参是当可迭代对象为空时返回的值。如果可迭代对象为空，并且没有给 *default*，则会触发 `ValueError`。

如果有多个最大元素，则此函数将返回第一个找到的。这和其他稳定排序工具如 `sorted(iterable, key=keyfunc, reverse=True)[0]` 和 `heapq.nlargest(1, iterable, key=keyfunc)` 保持一致。

3.4 新版功能: keyword-only 实参 *default*。

memoryview(obj)

返回由给定实参创建的“内存视图”对象。有关详细信息，请参阅[内存视图](#)。

min(iterable, *, key, default)**min(arg1, arg2, *args, key)**

返回可迭代对象中最小的元素，或者返回两个及以上实参中最小的。

如果只提供了一个位置参数，它必须是 *iterable*，返回可迭代对象中最小的元素；如果提供了两个及以上的位置参数，则返回最小的位置参数。

有两个可选只能用关键字的实参。*key* 实参指定排序函数用的参数，如传给 `list.sort()` 的。*default* 实参是当可迭代对象为空时返回的值。如果可迭代对象为空，并且没有给 *default*，则会触发 `ValueError`。

如果有多个最小元素，则此函数将返回第一个找到的。这和其他稳定排序工具如 `sorted(iterable, key=keyfunc)[0]` 和 `heapq.nsmallest(1, iterable, key=keyfunc)` 保持一致。

3.4 新版功能: keyword-only 实参 *default*。

next(iterator[, default])

通过调用 *iterator* 的 `__next__()` 方法获取下一个元素。如果迭代器耗尽，则返回给定的 *default*，如果没有默认值则触发 `StopIteration`。

class object

返回一个没有特征的新对象。*object* 是所有类的基类。它具有所有 Python 类实例的通用方法。这个函数不接受任何实参。

注解： 由于 *object* 没有 `__dict__`，因此无法将任意属性赋给 *object* 的实例。

oct(x)

将一个整数转变为一个前缀为“0o”的八进制字符串。结果是一个合法的 Python 表达式。如果 *x* 不是 Python 的 *int* 对象，那它需要定义 `__index__()` 方法返回一个整数。一些例子：

```
>>> oct(8)
'0o10'
>>> oct(-56)
'-0o70'
```

如果要整数转换为八进制字符串，并可选择有无“0o”前缀，则可以使用如下方法：

```
>>> '%#o' % 10, '%o' % 10
('0o12', '12')
>>> format(10, '#o'), format(10, 'o')
('0o12', '12')
>>> f'{10:#o}', f'{10:o}'
('0o12', '12')
```

另见 `format()` 获取更多信息。

open (*file*, *mode*='r', *buffering*=-1, *encoding*=None, *errors*=None, *newline*=None, *closefd*=True, *opener*=None)

打开 *file* 并返回对应的 *file object*。如果该文件不能打开，则触发 `OSError`。

file 是一个 *path-like object*，表示将要打开的文件的路径（绝对路径或者当前工作目录的相对路径），也可以是要被封装的整数类型文件描述符。（如果是文件描述符，它会随着返回的 I/O 对象关闭而关闭，除非 *closefd* 被设为 `False`。）

mode 是一个可选字符串，用于指定打开文件的模式。默认值是 'r'，这意味着它以文本模式打开并读取。其他常见模式有：写入 'w'（截断已经存在的文件）；排它性创建 'x'；追加写 'a'（在一些 Unix 系统上，无论当前的文件指针在什么位置，所有写入都会追加到文件末尾）。在文本模式，如果 *encoding* 没有指定，则根据平台来决定使用的编码：使用 `locale.getpreferredencoding(False)` 来获取本地编码。（要读取和写入原始字节，请使用二进制模式并不要指定 *encoding*。）可用的模式有：

字符	意义
'r'	读取（默认）
'w'	写入，并先截断文件
'x'	排它性创建，如果文件已存在则失败
'a'	写入，如果文件存在则在末尾追加
'b'	二进制模式
't'	文本模式（默认）
'+'	更新磁盘文件（读取并写入）
'U'	<i>universal newlines mode</i> (deprecated)

默认的模式是 'r'（打开并读取文本，同 'rt'）。对于二进制写入，'w+b' 模式打开并把文件截断成 0 字节；'r+b' 则不会截断。

正如在概述中提到的，Python 区分二进制和文本 I/O。以二进制模式打开的文件（包括 *mode* 参数中的 'b'）返回的内容为 `bytes` 对象，不进行任何解码。在文本模式下（默认情况下，或者在 **mode** 参数中包含 't'）时，文件内容返回为 *str*，首先使用指定的 *encoding*（如果给定）或者使用平台默认的字节编码解码。

注解：Python 不依赖于底层操作系统的文本文件概念；所有处理都由 Python 本身完成，因此与平台无关。

buffering 是一个可选的整数，用于设置缓冲策略。传递 0 以切换缓冲关闭（仅允许在二进制模式下），1 选择行缓冲（仅在文本模式下可用），并且 >1 的整数以指示固定大小的块缓冲区的大小（以字节为单位）。如果没有给出 *buffering* 参数，则默认缓冲策略的工作方式如下：

- 二进制文件以固定大小的块进行缓冲；使用启发式方法选择缓冲区的大小，尝试确定底层设备的“块大小”或使用 `io.DEFAULT_BUFFER_SIZE`。在许多系统上，缓冲区的长度通常为 4096 或 8192 字节。
- “交互式”文本文件（`isatty()` 返回 True 的文件）使用行缓冲。其他文本文件使用上述策略用于二进制文件。

`encoding` 是用于解码或编码文件的编码的名称。这应该只在文本模式下使用。默认编码是依赖于平台的（不管 `locale.getpreferredencoding()` 返回何值），但可以使用任何 Python 支持的 *text encoding*。有关支持的编码列表，请参阅 `codecs` 模块。

`errors` 是一个可选的字符串参数，用于指定如何处理编码和解码错误 - 这不能在二进制模式下使用。可以使用各种标准错误处理程序（列在 *Error Handlers*），但是使用 `codecs.register_error()` 注册的任何错误处理名称也是有效的。标准名称包括：

- 如果存在编码错误，'strict' 会引发 `ValueError` 异常。默认值 `None` 具有相同的效果。
- 'ignore' 忽略错误。请注意，忽略编码错误可能会导致数据丢失。
- 'replace' 会将替换标记（例如 '?'）插入有错误数据的地方。
- 'surrogateescape' 将表示任何不正确的字节作为 Unicode 专用区中的代码点，范围从 U+DC80 到 U+DCFF。当在写入数据时使用 `surrogateescape` 错误处理程序时，这些私有代码点将被转回到相同的字节中。这对于处理未知编码的文件很有用。
- 只有在写入文件时才支持 'xmlcharrefreplace'。编码不支持的字符将替换为相应的 XML 字符引用 `&#nnn;`。
- 'backslashreplace' 用 Python 的反向转义序列替换格式错误的数据。
- 'namereplace'（也只在编写时支持）用 `\N{...}` 转义序列替换不支持的字符。

`newline` 控制 *universal newlines* 模式如何生效（它仅适用于文本模式）。它可以是 `None`，`''`，`'\n'`，`'\r'` 和 `'\r\n'`。它的工作原理：

- 从流中读取输入时，如果 `newline` 为 `None`，则启用通用换行模式。输入中的行可以以 `'\n'`，`'\r'` 或 `'\r\n'` 结尾，这些行被翻译成 `'\n'` 在返回调用者之前。如果它是 `''`，则启用通用换行模式，但行结尾将返回给调用者未翻译。如果它具有任何其他合法值，则输入行仅由给定字符串终止，并且行结尾将返回给未调用的调用者。
- 将输出写入流时，如果 `newline` 为 `None`，则写入的任何 `'\n'` 字符都将转换为系统默认行分隔符 `os.linesep`。如果 `newline` 是 `''` 或 `'\n'`，则不进行翻译。如果 `newline` 是任何其他合法值，则写入的任何 `'\n'` 字符将被转换为给定的字符串。

如果 `closefd` 是 `False` 并且给出了文件描述符而不是文件名，那么当文件关闭时，底层文件描述符将保持打开状态。如果给出文件名则 `closefd` 必须为 `True`（默认值），否则将引发错误。

可以通过传递可调用的 `opener` 来使用自定义开启器。然后通过使用参数（`file`，`flags`）调用 `opener` 获得文件对象的基础文件描述符。`opener` 必须返回一个打开的文件描述符（使用 `os.open as opener` 时与传递 `None` 的效果相同）。

新创建的文件是 *不可继承的*。

下面的示例使用 `os.open()` 函数的 `dir_fd` 的形参，从给定的目录中用相对路径打开文件：

```
>>> import os
>>> dir_fd = os.open('somedir', os.O_RDONLY)
>>> def opener(path, flags):
...     return os.open(path, flags, dir_fd=dir_fd)
...
>>> with open('spamspam.txt', 'w', opener=opener) as f:
...     print('This will be written to somedir/spamspam.txt', file=f)
```

(下页继续)

(续上页)

```
...
>>> os.close(dir_fd)  # don't leak a file descriptor
```

`open()` 函数所返回的 *file object* 类型取决于所用模式。当使用 `open()` 以文本模式 ('w', 'r', 'wt', 'rt' 等) 打开文件时, 它将返回 `io.TextIOBase` (特别是 `io.TextIOWrapper`) 的一个子类。当使用缓冲以二进制模式打开文件时, 返回的类是 `io.BufferedIOBase` 的一个子类。具体的类会有多种: 在只读的二进制模式下, 它将返回 `io.BufferedReader`; 在写入二进制和追加二进制模式下, 它将返回 `io.BufferedWriter`, 而在读/写模式下, 它将返回 `io.BufferedRandom`。当禁用缓冲时, 则会返回原始流, 即 `io.RawIOBase` 的一个子类 `io.FileIO`。

另请参阅文件操作模块, 例如 `fileinput`、`io` (声明了 `open()`)、`os`、`os.path`、`tempfile` 和 `shutil`。

在 3.3 版更改:

- 增加了 `opener` 形参。
- 增加了 'x' 模式。
- 过去触发的 `IOError`, 现在是 `OSError` 的别名。
- 如果文件已存在但使用了排它性创建模式 ('x'), 现在会触发 `FileExistsError`。

在 3.4 版更改:

- 文件现在禁止继承。

Deprecated since version 3.4, will be removed in version 4.0: 'U' 模式。

在 3.5 版更改:

- 如果系统调用被中断, 但信号处理程序没有触发异常, 此函数现在会重试系统调用, 而不是触发 `InterruptedError` 异常 (原因详见 [PEP 475](#))。
- 增加了 'namereplace' 错误处理接口。

在 3.6 版更改:

- 增加对实现了 `os.PathLike` 对象的支持。
- 在 Windows 上, 打开一个控制台缓冲区将返回 `io.RawIOBase` 的子类, 而不是 `io.FileIO`。

`ord(c)`

对表示单个 Unicode 字符的字符串, 返回代表它 Unicode 码点的整数。例如 `ord('a')` 返回整数 97, `ord('€')` (欧元符号) 返回 8364。这是 `chr()` 的逆函数。

`pow(x, y[, z])`

返回 x 的 y 次幂; 如果 z 存在, 则对 z 取余 (比直接 `pow(x, y) % z` 计算更高效)。两个参数形式的 `pow(x, y)` 等价于幂运算符: `x**y`。

参数必须为数值类型。对于混用的操作数类型, 则适用二元算术运算符的类型强制转换规则。对于 `int` 操作数, 结果具有与操作数相同的类型 (转换后), 除非第二个参数为负值; 在这种情况下, 所有参数将被转换为浮点数并输出浮点数结果。例如, `10**2` 返回 100, 但 `10**-2` 返回 0.01。如果第二个参数为负值, 则第三个参数必须省略。如果存在 z , 则 x 和 y 必须为整数类型, 且 y 必须为非负数。

`print(*objects, sep=' ', end='\n', file=sys.stdout, flush=False)`

将 *objects* 打印到 *file* 指定的文本流, 以 *sep* 分隔并在末尾加上 *end*。 *sep*, *end*, *file* 和 *flush* 如果存在, 它们必须以关键字参数的形式给出。

所有非关键字参数都会被转换为字符串，就像是执行了 `str()` 一样，并会被写入到流，以 `sep` 且在末尾加上 `end`。`sep` 和 `end` 都必须为字符串；它们也可以为 `None`，这意味着使用默认值。如果没有给出 `objects`，则 `print()` 将只写入 `end`。

`file` 参数必须是一个具有 `write(string)` 方法的对象；如果参数不存在或为 `None`，则将使用 `sys.stdout`。由于要打印的参数会被转换为文本字符串，因此 `print()` 不能用于二进制模式的文件对象。对于这些对象，应改用 `file.write(...)`。

输出是否被缓存通常决定于 `file`，但如果 `flush` 关键字参数为真值，流会被强制刷新。

在 3.3 版更改：增加了 `flush` 关键字参数。

class property (`fget=None`, `fset=None`, `fdel=None`, `doc=None`)

返回 `property` 属性。

`fget` 是获取属性值的函数。`fset` 是用于设置属性值的函数。`fdel` 是用于删除属性值的函数。并且 `doc` 为属性对象创建文档字符串。

一个典型的用法是定义一个托管属性 `x`：

```
class C:
    def __init__(self):
        self._x = None

    def getx(self):
        return self._x

    def setx(self, value):
        self._x = value

    def delx(self):
        del self._x

    x = property(getx, setx, delx, "I'm the 'x' property.")
```

如果 `c` 是 `C` 的实例，`c.x` 将调用 `getter`，`c.x = value` 将调用 `setter`，`del c.x` 将调用 `deleter`。

如果给出，`doc` 将成为该 `property` 属性的文档字符串。否则该 `property` 将拷贝 `fget` 的文档字符串（如果存在）。这令使用 `property()` 作为 *decorator* 来创建只读的特征属性可以很容易地实现：

```
class Parrot:
    def __init__(self):
        self._voltage = 100000

    @property
    def voltage(self):
        """Get the current voltage."""
        return self._voltage
```

以上 `@property` 装饰器会将 `voltage()` 方法转化为一个具有相同名称的只读属性的“getter”，并将 `voltage` 的文档字符串设置为“Get the current voltage.”

特征属性对象具有 `getter`、`setter` 以及 `deleter` 方法，它们可用作装饰器来创建该特征属性的副本，并将相应的访问函数设为所装饰的函数。这最好是用一个例子来解释：

```
class C:
    def __init__(self):
        self._x = None

    @property
```

(下页继续)

(续上页)

```

def x(self):
    """I'm the 'x' property."""
    return self._x

@x.setter
def x(self, value):
    self._x = value

@x.deleter
def x(self):
    del self._x

```

上述代码与第一个例子完全等价。注意一定要给附加函数与原始的特征属性相同的名称 (在本例中为 `x`。)

返回的特征属性对象同样具有与构造器参数相对应的属性 `fget`, `fset` 和 `fdel`。

在 3.5 版更改: 特性属性对象的文档字符串现在是可写的。

range (*stop*)

range (*start*, *stop* [, *step*])

虽然被称为函数, 但 *range* 实际上是一个不可变的序列类型, 参见在 *range* 对象 与 序列类型—*list*, *tuple*, *range* 中的文档说明。

repr (*object*)

返回包含一个对象的可打印表示形式的字符串。对于许多类型来说, 该函数会尝试返回的字符串将会与该对象被传递给 `eval()` 时所生成的对象具有相同的值, 在其他情况下表示形式会是一个括在尖括号中的字符串, 其中包含对象类型的名称与通常包括对象名称和地址的附加信息。类可以通过定义 `__repr__()` 方法来控制此函数为它的实例所返回的内容。

reversed (*seq*)

返回一个反向的 *iterator*。 *seq* 必须是一个具有 `__reversed__()` 方法的对象或者是支持该序列协议 (具有从 “0” 开始的整数类型参数的 `__len__()` 方法和 `__getitem__()` 方法)。

round (*number* [, *ndigits*])

返回 *number* 舍入到小数点后 *ndigits* 位精度的值。如果 *ndigits* 被省略或为 `None`, 则返回最接近输入值的整数。

对于支持 `round()` 的内置类型, 值会被舍入到最接近的 10 的负 *ndigits* 次幂的倍数; 如果与两个倍数的距离相等, 则选择偶数 (因此, `round(0.5)` 和 `round(-0.5)` 均为 0 而 `round(1.5)` 为 2)。任何整数值都可作为有效的 *ndigits* (正数、零或负数)。如果 *ndigits* 被省略或为 `None` 则返回值将为整数。否则返回值与 *number* 的类型相同。

对于一般的 Python 对象 *number*, `round` 将委托给 `number.__round__`。

注解: 对浮点数执行 `round()` 的行为可能会令人惊讶: 例如, `round(2.675, 2)` 将给出 2.67 而不是期望的 2.68。这不算是程序错误: 这一结果是由于大多数十进制小数实际上都不能以浮点数精确地表示。请参阅 `tut-fp-issues` 了解更多信息。

class set ([*iterable*])

返回一个新的 *set* 对象, 可以选择带有从 *iterable* 获取的元素。*set* 是一个内置类型。请查看 *set* 和 集合类型—*set*, *frozenset* 获取关于这个类的文档。

有关其他容器请参看内置的 *frozenset*, *list*, *tuple* 和 *dict* 类, 以及 *collections* 模块。

setattr (*object*, *name*, *value*)

此函数与 `getattr()` 两相对应。其参数为一个对象、一个字符串和一个任意值。字符串指定一个现有属

性或者新增属性。函数会将值赋给该属性，只要对象允许这种操作。例如，`setattr(x, 'foobar', 123)` 等价于 `x.foobar = 123`。

class slice (*stop*)

class slice (*start, stop* [, *step*])

返回一个表示由 `range(start, stop, step)` 所指定索引集的 *slice* 对象。其中 *start* 和 *step* 参数默认为 `None`。切片对象具有仅会返回对应参数值（或其默认值）的只读数据属性 *start*, *stop* 和 *step*。它们没有其他的显式功能；不过它们会被 NumPy 以及其他第三方扩展所使用。切片对象也会在使用扩展索引语法时被生成。例如：`a[start:stop:step]` 或 `a[start:stop, i]`。请参阅 `itertools.islice()` 了解返回迭代器的一种替代版本。

sorted (*iterable*, *, *key=None*, *reverse=False*)

根据 *iterable* 中的项返回一个新的已排序列表。

具有两个可选参数，它们都必须指定为关键字参数。

key 指定带有单个参数的函数，用于从 *iterable* 的每个元素中提取用于比较的键（例如 `key=str.lower`）。默认值为 `None`（直接比较元素）。

reverse 为一个布尔值。如果设为 `True`，则每个列表元素将按反向顺序比较进行排序。

使用 `functools.cmp_to_key()` 可将老式的 *cmp* 函数转换为 *key* 函数。

内置的 `sorted()` 确保是稳定的。如果一个排序确保不会改变比较结果相等的元素的相对顺序就称其为稳定——这有利于进行多重排序（例如先按部门、再按薪级排序）。

有关排序示例和简要排序教程，请参阅 `sortinghowto`。

@staticmethod

将方法转换为静态方法。

静态方法不会接收隐式的第一个参数。要声明一个静态方法，请使用此语法

```
class C:
    @staticmethod
    def f(arg1, arg2, ...): ...
```

The `@staticmethod` form is a function *decorator* – see the description of function definitions in function for details.

It can be called either on the class (such as `C.f()`) or on an instance (such as `C().f()`). The instance is ignored except for its class.

Python 中的静态方法与 Java 或 C++ 中的静态方法类似。另请参阅 `classmethod()`，用于创建备用类构造函数的变体。

像所有装饰器一样，也可以像常规函数一样调用 `staticmethod`，并对其结果执行某些操作。比如某些情况下需要从类主体引用函数并且您希望避免自动转换为实例方法。对于这些情况，请使用此语法：

```
class C:
    builtin_open = staticmethod(open)
```

For more information on static methods, consult the documentation on the standard type hierarchy in types.

class str (*object=""*)

class str (*object=b", encoding='utf-8', errors='strict'*)

返回一个 *str* 版本的 *object*。有关详细信息，请参阅 `str()`。

str 是内置字符串 *class*。更多关于字符串的信息查看 [文本序列类型—str](#)。

sum (*iterable* [, *start*])

从 *start* 开始自左向右对 *iterable* 中的项求和并返回总计值。*start* 默认为 0。*iterable* 的项通常为数字，开始值则不允许为字符串。

对某些用例来说，存在`sum()`的更好替代。拼接字符串序列的更好更快方式是调用`''.join(sequence)`。要以扩展精度对浮点值求和，请参阅`math.fsum()`。要拼接一系列可迭代对象，请考虑使用`itertools.chain()`。

super (`[type[, object-or-type]]`)

返回一个代理对象，它会将方法调用委托给 `type` 指定的父类或兄弟类。这对于访问已在类中被重载的继承方法很有用。搜索顺序与`getattr()`所使用的相同，只是 `type` 指定的类型本身会被跳过。

`type` 的`__mro__` 属性列出了`getattr()` 和`super()` 所使用的方法解析顺序。该属性是动态的，可以在继承层级结构更新的时候任意改变。

如果省略第二个参数，则返回的超类对象是未绑定的。如果第二个参数为一个对象，则`isinstance(obj, type)` 必须为真值。如果第二个参数为一个类型，则`issubclass(type2, type)` 必须为真值（这适用于类方法）。

`super` 有两个典型用例。在具有单继承的类层级结构中，`super` 可用来引用父类而不必显式地指定它们的名称，从而令代码更易维护。这种用法与其他编程语言中 `super` 的用法非常相似。

第二个用例是在动态执行环境中支持协作多重继承。此用例为 Python 所独有，在静态编译语言或仅支持单继承的语言中是不存在的。这使得实现“菱形图”成为可能，在这时会有多个基类实现相同的方法。好的设计强制要求这种方法在每个情况下具有相同的调用签名（因为调用顺序是在运行时确定的，也因为该顺序要适应类层级结构的更改，还因为该顺序可能包含在运行时之前未知的兄弟类）。

对于以上两个用例，典型的超类调用看起来是这样的：

```
class C(B):
    def method(self, arg):
        super().method(arg)      # This does the same thing as:
                                # super(C, self).method(arg)
```

请注意`super()` 是作为显式加点属性查找的绑定过程的一部分来实现的，例如`super().__getitem__(name)`。它做到这一点是通过实现自己的`__getattribute__()` 方法，这样就能以可预测的顺序搜索类，并且支持协作多重继承。对应地，`super()` 在像`super()[name]` 这样使用语句或操作符进行隐式查找时则未被定义。

还要注意的，除了零个参数的形式以外，`super()` 并不限于在方法内部调用。两个参数的形式明确指定参数并进行相应的引用。零个参数的形式仅适用于类定义内部，因为编译器需要填入必要的细节以正确地检索到被定义的类，还需要为普通访问当前实例。

对于有关如何使用`super()` 来如何设计协作类的实用建议，请参阅 使用 `super()` 的指南。

tuple (`[iterable]`)

虽然被称为函数，但`tuple` 实际上是一个不可变的序列类型，参见在元组 与 序列类型—`list`, `tuple`, `range` 中的文档说明。

class type (`object`)

class type (`name, bases, dict`)

传入一个参数时，返回 `object` 的类型。返回值是一个 `type` 对象，通常与`object.__class__` 所返回的对象相同。

推荐使用`isinstance()` 内置函数来检测对象的类型，因为它会考虑子类的情况。

传入三个参数时，返回一个新的 `type` 对象。这在本质上是 `class` 语句的一种动态形式。`name` 字符串即类名并且会成为`__name__` 属性；`bases` 元组列出基类并且会成为`__bases__` 属性；而 `dict` 字典为包含类主体定义的命名空间并且会被复制到一个标准字典成为`__dict__` 属性。例如，下面两条语句会创建相同的`type` 对象：

```
>>> class X:
...     a = 1
...
>>> X = type('X', (object,), dict(a=1))
```

另请参阅类型对象。

在 3.6 版更改: `type` 的子类如果未重载 `type.__new__`, 将不再能使用一个参数的形式来获取对象的类型。

vars (`[object]`)

返回模块、类、实例或任何其它具有 `__dict__` 属性的对象的 `__dict__` 属性。

模块和实例这样的对象具有可更新的 `__dict__` 属性; 但是, 其它对象的 `__dict__` 属性可能会设为限制写入 (例如, 类会使用 `types.MappingProxyType` 来防止直接更新字典)。

不带参数时, `vars()` 的行为类似 `locals()`。请注意, `locals` 字典仅对于读取起作用, 因为对 `locals` 字典的更新会被忽略。

zip (`*iterables`)

创建一个聚合了来自每个可迭代对象中的元素的迭代器。

返回一个元组的迭代器, 其中的第 *i* 个元组包含来自每个参数序列或可迭代对象的第 *i* 个元素。当所输入可迭代对象中最短的一个被耗尽时, 迭代器将停止迭代。当只有一个可迭代对象参数时, 它将返回一个单元组的迭代器。不带参数时, 它将返回一个空迭代器。相当于:

```
def zip(*iterables):
    # zip('ABCD', 'xy') --> Ax By
    sentinel = object()
    iterators = [iter(it) for it in iterables]
    while iterators:
        result = []
        for it in iterators:
            elem = next(it, sentinel)
            if elem is sentinel:
                return
            result.append(elem)
        yield tuple(result)
```

函数会保证可迭代对象按从左至右的顺序被求值。使得可以通过 `zip(*[iter(s)]*n)` 这样的惯用形式将一系列数据聚类为长度为 *n* 的分组。这将重复 同样的迭代器 *n* 次, 以便每个输出的元组具有第 *n* 次调用该迭代器的结果。它的作用效果就是将输入拆分为长度为 *n* 的数据块。

当你不用关心较长可迭代对象末尾不匹配的值时, 则 `zip()` 只须使用长度不相等的输入即可。如果那些值很重要, 则应改用 `itertools.zip_longest()`。

`zip()` 与 `*` 运算符相结合可以用来拆解一个列表:

```
>>> x = [1, 2, 3]
>>> y = [4, 5, 6]
>>> zipped = zip(x, y)
>>> list(zipped)
[(1, 4), (2, 5), (3, 6)]
>>> x2, y2 = zip(*zip(x, y))
>>> x == list(x2) and y == list(y2)
True
```

__import__ (`name, globals=None, locals=None, fromlist=(), level=0`)

注解: 与 `importlib.import_module()` 不同, 这是一个日常 Python 编程中不需要用到的高级函数。

This function is invoked by the `import` statement. It can be replaced (by importing the `builtins` module and assigning to `builtins.__import__`) in order to change semantics of the `import` statement, but doing so is **strongly** discouraged as it is usually simpler to use import hooks (see [PEP 302](#)) to attain the same goals and does not cause issues with code which assumes the default import implementation is in use. Direct use of `__import__()` is also discouraged in favor of `importlib.import_module()`.

该函数会导入 *name* 模块，有可能使用给定的 *globals* 和 *locals* 来确定如何在包的上下文中解读名称。*fromlist* 给出了应该从由 *name* 指定的模块导入对象或子模块的名称。标准实现完全不使用其 *locals* 参数，而仅使用 *globals* 参数来确定 `import` 语句的包上下文。

level 指定是使用绝对还是相对导入。0 (默认值) 意味着仅执行绝对导入。*level* 为正数值表示相对于模块调用 `__import__()` 的目录，将要搜索的父目录层数 (详情参见 [PEP 328](#))。

当 *name* 变量的形式为 `package.module` 时，通常将会返回最高层级的包 (第一个点号之前的名称)，而不是以 *name* 命名的模块。但是，当给出了非空的 *fromlist* 参数时，则将返回以 *name* 命名的模块。

例如，语句 `import spam` 的结果将为与以下代码作用相同的字节码：

```
spam = __import__('spam', globals(), locals(), [], 0)
```

语句 `import spam.ham` 的结果将为以下调用：

```
spam = __import__('spam.ham', globals(), locals(), [], 0)
```

请注意在这里 `__import__()` 是如何返回顶层模块的，因为这是通过 `import` 语句被绑定到特定名称的对象。

另一方面，语句 `from spam.ham import eggs, sausage as saus` 的结果将为

```
_temp = __import__('spam.ham', globals(), locals(), ['eggs', 'sausage'], 0)
eggs = _temp.eggs
saus = _temp.sausage
```

在这里，`spam.ham` 模块会由 `__import__()` 返回。要导入的对象将从此对象中提取并赋值给它们对应的名称。

如果您只想按名称导入模块 (可能在包中)，请使用 `importlib.import_module()`

在 3.3 版更改: Negative values for *level* are no longer supported (which also changes the default value to 0).

内置常量

有少数的常量存在于内置命名空间中。它们是：

False

bool 类型的假值。给 False 赋值是非法的并会引发 *SyntaxError*。

True

bool 类型的真值。给 True 赋值是非法的并会引发 *SyntaxError*。

None

NoneType 类型的唯一值。None 经常用于表示缺少值，当因为默认参数未传递给函数时。给 None 赋值是非法的并会引发 *SyntaxError*。

NotImplemented

二进制特殊方法应返回的特殊值（例如，`__eq__()`、`__lt__()`、`__add__()`、`__rsub__()` 等）表示操作没有针对其他类型实现；为了相同的目的，可以通过就地二进制特殊方法（例如，`__imul__()`、`__righnd__()` 等）返回。它的逻辑值为真。

注解：当二进制（或就地）方法返回“NotImplemented”时，解释器将尝试对另一种类型（或其他一些回滚操作，取决于运算符）的反射操作。如果所有尝试都返回“NotImplemented”，则解释器将引发适当的异常。错误返回的“NotImplemented”将导致误导性错误消息或返回到 Python 代码中的“NotImplemented”值。

参见 *Implementing the arithmetic operations* 为例。

注解：NotImplementedError 和 NotImplemented 不可互换，即使它们有相似的名称和用途。有关何时使用它的详细信息，请参阅 *NotImplementedError*。

Ellipsis

与省略号文字字面“...”相同。特殊值主要与用户定义的容器数据类型的扩展切片语法结合使用。

__debug__

如果 Python 没有以 -O 选项启动，则此常量为真值。另请参见 assert 语句。

注解： 变量名 `None`, `False`, `True` 和 `__debug__` 无法重新赋值（赋值给它们，即使是属性名，将引发 `SyntaxError`），所以它们可以被认为是“真正的”常数。

3.1 由 `site` 模块添加的常量

`site` 模块（在启动期间自动导入，除非给出 `-s` 命令行选项）将几个常量添加到内置命名空间。它们对交互式解释器 `shell` 很有用，并且不应在程序中使用。

quit (`code=None`)

exit (`code=None`)

当打印此对象时，会打印出一条消息，例如 “Use quit() or Ctrl-D (i.e. EOF) to exit”，当调用此对象时，将使用指定的退出代码来引发 `SystemExit`。

copyright

credits

打印或调用的对象分别打印版权或作者的文本。

license

当打印此对象时，会打印出一条消息 “Type license() to see the full license text”，当调用此对象时，将以分页形式显示完整的许可证文本（每次显示一屏）。

以下部分描述了解释器中内置的标准类型。

主要内置类型有数字、序列、映射、类、实例和异常。

有些多项集类是可变的。它们用于添加、移除或重排其成员的方法将原地执行，并不返回特定的项，绝对不会返回多项集实例自身而是返回 `None`。

有些操作受多种对象类型的支持；特别地，实际上所有对象都可以被比较、检测逻辑值，以及转换为字符串（使用 `repr()` 函数或略有差异的 `str()` 函数）。后一个函数是在对象由 `print()` 函数输出时被隐式地调用的。

4.1 逻辑值检测

任何对象都可以进行逻辑值的检测，以便在 `if` 或 `while` 作为条件或是作为下文所述布尔运算的操作数来使用。

一个对象在默认情况下均被视为真值，除非当该对象被调用时其所属类定义了 `__bool__()` 方法且返回 `False` 或是定义了 `__len__()` 方法且返回零。¹ 下面基本完整地列出了会被视为假值的内置对象：

- 被定义为假值的常量: `None` 和 `False`。
- 任何数值类型的零: `0`, `0.0`, `0j`, `Decimal(0)`, `Fraction(0, 1)`
- 空的序列和多项集: `''`, `()`, `[]`, `{}`, `set()`, `range(0)`

产生布尔值结果的运算和内置函数总是返回 `0` 或 `False` 作为假值，`1` 或 `True` 作为真值，除非另行说明。（重要例外：布尔运算 `or` 和 `and` 总是返回其中一个操作数。）

¹ 有关这些特殊方法的额外信息可参看 Python 参考指南 (customization)。

4.2 Boolean Operations —and, or, not

这些属于布尔运算，按优先级升序排列：

运算	结果	注释
<code>x or y</code>	if <i>x</i> is false, then <i>y</i> , else <i>x</i>	(1)
<code>x and y</code>	if <i>x</i> is false, then <i>x</i> , else <i>y</i>	(2)
<code>not x</code>	if <i>x</i> is false, then True, else False	(3)

注释：

- (1) 这是个短路运算符，因此只有在第一个参数为假值时才会对第二个参数求值。
- (2) 这是个短路运算符，因此只有在第一个参数为真值时才会对第二个参数求值。
- (3) `not` 的优先级比非布尔运算符低，因此 `not a == b` 会被解读为 `not (a == b)` 而 `a == not b` 会引发语法错误。

4.3 比较

在 Python 中有八种比较运算符。它们的优先级相同（比布尔运算的优先级高）。比较运算可以任意串连；例如，`x < y <= z` 等价于 `x < y and y <= z`，前者的不同之处在于 *y* 只被求值一次（但在两种情况下当 `x < y` 结果为假值时 *z* 都不会被求值）。

此表格汇总了比较运算：

运算	含义
<code><</code>	严格小于
<code><=</code>	小于或等于
<code>></code>	严格大于
<code>>=</code>	大于或等于
<code>==</code>	等于
<code>!=</code>	不等于
<code>is</code>	对象标识
<code>is not</code>	否定的对象标识

除了不同数字类型以外，不同类型的对象比较时绝对不会相等。而且，某些类型（例如函数对象）仅支持简化比较形式，即任何两个该种类型的对象必定不相等。`<`、`<=`、`>` 和 `>=` 运算符在以下情况中将引发 `TypeError` 异常：当比较复数与另一个内置数字类型时，当两个对象具有无法被比较的不同类型时，或在未定义次序的其他情况时。

具有不同标识的类的实例比较结果通常为不相等，除非类定义了 `__eq__()` 方法。

一个类实例不能与相同类或的其他实例或其他类型的对象进行排序，除非该类定义了足够多的方法，包括 `__lt__()`、`__le__()`、`__gt__()` 以及 `__ge__()`（而如果你想实现常规意义上的比较操作，通常只要有 `__lt__()` 和 `__eq__()` 就可以了）。

`is` 和 `is not` 运算符无法自定义；并且它们可以被应用于任意两个对象而不会引发异常。

还有两种具有相同语法优先级的运算 `in` 和 `not in`，它们被 *iterable* 或实现了 `__contains__()` 方法的类型所支持。

4.4 数字类型—int, float, complex

存在三种不同的数字类型：整数，浮点数和复数。此外，布尔值属于整数的子类型。整数具有无限的精度。浮点数通常使用 C 中的 `double` 来实现；有关你的程序运行所在机器上浮点数的精度和内部表示法可在 `sys.float_info` 中查看。复数包含实部和虚部，分别以一个浮点数表示。要从一个复数 `z` 中提取这两个部分，可使用 `z.real` 和 `z.imag`。（标准库包含附加的数字类型，如表示有理数的 `fractions` 以及以用户定制精度表示浮点数的 `decimal`。）

数字是由数字面值或内置函数与运算符的结果来创建的。不带修饰的整数字面值（包括十六进制、八进制和二进制数）会生成整数。包含小数点或幂运算符的数字面值会生成浮点数。在数字面值末尾加上 `'j'` 或 `'J'` 会生成虚数（实部为零的复数），你可以将其与整数或浮点数相加来得到具有实部和虚部的复数。

Python 完全支持混合算术：当一个二元运算符用于不同数字类型的操作数时，具有“较窄”类型的操作数会被扩展为另一个操作数的类型，整数比浮点数更窄，浮点数又比复数更窄。混合类型数字之间的比较也使用相同的规则。² 构造器 `int()`、`float()` 和 `complex()` 可被用于生成特定类型的数字。

所有数字类型（复数除外）都支持下列运算，按优先级升序排序（所有数字运算的优先级都高于比较运算）：

运算	结果	注释	完整文档
<code>x + y</code>	<code>x</code> 和 <code>y</code> 的和		
<code>x - y</code>	<code>x</code> 和 <code>y</code> 的差		
<code>x * y</code>	<code>x</code> 和 <code>y</code> 的乘积		
<code>x / y</code>	<code>x</code> 和 <code>y</code> 的商		
<code>x // y</code>	<code>x</code> 和 <code>y</code> 的商数	(1)	
<code>x % y</code>	remainder of <code>x / y</code>	(2)	
<code>-x</code>	<code>x</code> 取反		
<code>+x</code>	<code>x</code> 不变		
<code>abs(x)</code>	<code>x</code> 的绝对值或大小		<code>abs()</code>
<code>int(x)</code>	将 <code>x</code> 转换为整数	(3)(6)	<code>int()</code>
<code>float(x)</code>	将 <code>x</code> 转换为浮点数	(4)(6)	<code>float()</code>
<code>complex(re, im)</code>	一个带有实部 <code>re</code> 和虚部 <code>im</code> 的复数。 <code>im</code> 默认为 0。	(6)	<code>complex()</code>
<code>c.conjugate()</code>	复数 <code>c</code> 的共轭		
<code>divmod(x, y)</code>	<code>(x // y, x % y)</code>	(2)	<code>divmod()</code>
<code>pow(x, y)</code>	<code>x</code> 的 <code>y</code> 次幂	(5)	<code>pow()</code>
<code>x ** y</code>	<code>x</code> 的 <code>y</code> 次幂	(5)	

注释：

- (1) 也称为整数除法。结果值是一个整数，但结果的类型不一定是 `int`。运算结果总是向负无穷的方向舍入：`1//2` 为 0，`(-1)//2` 为 -1，`1//(-2)` 为 -1 而 `(-1)//(-2)` 为 0。
- (2) 不可用于复数。而应在适当条件下使用 `abs()` 转换为浮点数。
- (3) 从浮点数转换为整数会被舍入或是像在 C 语言中一样被截断；请参阅 `math.floor()` 和 `math.ceil()` 函数查看转换的完整定义。
- (4) `float` 也接受字符串“nan”和附带可选前缀“+”或“-”的“inf”分别表示非数字 (NaN) 以及正或负无穷。
- (5) Python 将 `pow(0, 0)` 和 `0 ** 0` 定义为 1，这是编程语言的普遍做法。
- (6) 接受的数字面值包括数码 0 到 9 或任何等效的 Unicode 字符（具有 `Nd` 特征属性的代码点）。

See <http://www.unicode.org/Public/9.0.0/ucd/extracted/DerivedNumericType.txt> for a complete list of code points with the `Nd` property.

² 作为结果，列表 `[1, 2]` 与 `[1.0, 2.0]` 是相等的，元组的情况也类似。

所有 `numbers.Real` 类型 (`int` 和 `float`) 还包括下列运算:

运算	结果
<code>math.trunc(x)</code>	x 截断为 <i>Integral</i>
<code>round(x[, n])</code>	x 舍入到 n 位小数, 半数值会舍入到偶数。如果省略 n , 则默认为 0。
<code>math.floor(x)</code>	$\leq x$ 的最大 <i>Integral</i>
<code>math.ceil(x)</code>	$\geq x$ 的最小 <i>Integral</i>

有关更多的数字运算请参阅 `math` 和 `cmath` 模块。

4.4.1 整数类型的按位运算

按位运算只对整数有意义。计算按位运算的结果, 就相当于使用无穷多个二进制符号位对二的补码执行操作。

二进制按位运算的优先级全都低于数字运算, 但又高于比较运算; 一元运算 `~` 具有与其他一元算术运算 (`+` 和 `-`) 相同的优先级。

此表格是以优先级升序排序的按位运算列表:

运算	结果	注释
<code>x y</code>	x 和 y 按位 或	(4)
<code>x ^ y</code>	x 和 y 按位 异或	(4)
<code>x & y</code>	x 和 y 按位 与	(4)
<code>x << n</code>	x 左移 n 位	(1)(2)
<code>x >> n</code>	x 右移 n 位	(1)(3)
<code>~x</code>	x 逐位取反	

注释:

- (1) 负的移位数是非法的, 会导致引发 `ValueError`。
- (2) 左移 n 位等价于不带溢出检测地乘以 `pow(2, n)`。
- (3) 右移 n 位等价于不带溢出检测地除以 `pow(2, n)`。
- (4) 使用带有至少一个额外符号扩展位的有限个二进制补码表示 (有效位宽度为 `1 + max(x.bit_length(), y.bit_length())` 或以上) 执行这些计算就足以获得相当于有无数个符号位时的同样结果。

4.4.2 整数类型的附加方法

`int` 类型实现了 `numbers.Integral abstract base class`。此外, 它还提供了其他几个方法:

`int.bit_length()`

返回以二进制表示一个整数所需要的位数, 不包括符号位和前面的零:

```
>>> n = -37
>>> bin(n)
'-0b100101'
>>> n.bit_length()
6
```

更准确地说, 如果 x 非零, 则 `x.bit_length()` 是使得 $2^{k-1} \leq \text{abs}(x) < 2^k$ 的唯一正整数 k 。同样地, 当 $\text{abs}(x)$ 小到足以具有正确的舍入对数时, 则 $k = 1 + \text{int}(\log(\text{abs}(x), 2))$ 。如果 x 为零, 则 `x.bit_length()` 返回 0。

等价于:

```
def bit_length(self):
    s = bin(self)          # binary representation: bin(-37) --> '-0b100101'
    s = s.lstrip('-0b')    # remove leading zeros and minus sign
    return len(s)          # len('100101') --> 6
```

3.1 新版功能.

`int.to_bytes(length, byteorder, *, signed=False)`

返回表示一个整数的字节数组。

```
>>> (1024).to_bytes(2, byteorder='big')
b'\x04\x00'
>>> (1024).to_bytes(10, byteorder='big')
b'\x00\x00\x00\x00\x00\x00\x00\x00\x04\x00'
>>> (-1024).to_bytes(10, byteorder='big', signed=True)
b'\xff\xff\xff\xff\xff\xff\xff\xff\xfc\x00'
>>> x = 1000
>>> x.to_bytes((x.bit_length() + 7) // 8, byteorder='little')
b'\xe8\x03'
```

整数会使用 `length` 个字节来表示。如果整数不能用给定的字节数来表示则会引发 `OverflowError`。

`byteorder` 参数确定用于表示整数的字节顺序。如果 `byteorder` 为 "big", 则最高位字节放在字节数组的开头。如果 `byteorder` 为 "little", 则最高位字节放在字节数组的末尾。要请求主机系统上的原生字节顺序, 请使用 `sys.byteorder` 作为字节顺序值。

`signed` 参数确定是否使用二的补码来表示整数。如果 `signed` 为 `False` 并且给出的是负整数, 则会引发 `OverflowError`。`signed` 的默认值为 `False`。

3.2 新版功能.

`classmethod int.from_bytes(bytes, byteorder, *, signed=False)`

返回由给定字节数组所表示的整数。

```
>>> int.from_bytes(b'\x00\x10', byteorder='big')
16
>>> int.from_bytes(b'\x00\x10', byteorder='little')
4096
>>> int.from_bytes(b'\xfc\x00', byteorder='big', signed=True)
-1024
>>> int.from_bytes(b'\xfc\x00', byteorder='big', signed=False)
64512
>>> int.from_bytes([255, 0, 0], byteorder='big')
16711680
```

`bytes` 参数必须为一个 *bytes-like object* 或是生成字节的可迭代对象。

`byteorder` 参数确定用于表示整数的字节顺序。如果 `byteorder` 为 "big", 则最高位字节放在字节数组的开头。如果 `byteorder` 为 "little", 则最高位字节放在字节数组的末尾。要请求主机系统上的原生字节顺序, 请使用 `sys.byteorder` 作为字节顺序值。

`signed` 参数指明是否使用二的补码来表示整数。

3.2 新版功能.

4.4.3 浮点类型的附加方法

`float` 类型实现了 *numbers.Real abstract base class*。`float` 还具有以下附加方法。

`float.as_integer_ratio()`

返回一对整数，其比率正好等于原浮点数并且分母为正数。无穷大会引发 *OverflowError* 而 NaN 则会引发 *ValueError*。

`float.is_integer()`

如果 `float` 实例可用有限位整数表示则返回 `True`，否则返回 `False`：

```
>>> (-2.0).is_integer()
True
>>> (3.2).is_integer()
False
```

两个方法均支持与十六进制数字字符串之间的转换。由于 Python 浮点数在内部存储为二进制数，因此浮点数与十进制数字字符串之间的转换往往会导致微小的舍入错误。而十六进制数字字符串却允许精确地表示和描述浮点数。这在进行调试和数值工作时非常有用。

`float.hex()`

以十六进制字符串的形式返回一个浮点数表示。对于有限浮点数，这种表示法将总是包含前导的 `0x` 和尾随的 `p` 加指数。

classmethod `float.fromhex(s)`

返回以十六进制字符串 `s` 表示的浮点数的类方法。字符串 `s` 可以带有前导和尾随的空格。

请注意 `float.hex()` 是实例方法，而 `float.fromhex()` 是类方法。

十六进制字符串采用的形式为：

```
[sign] ['0x'] integer ['.' fraction] ['p' exponent]
```

可选的 `sign` 可以是 `+` 或 `-`，`integer` 和 `fraction` 是十六进制数码组成的字符串，`exponent` 是带有可选前导符的十进制整数。大小写没有影响，在 `integer` 或 `fraction` 中必须至少有一个十六进制数码。此语法类似于 C99 标准的 6.4.4.2 小节中所描述的语法，也是 Java 1.5 以上所使用的语法。特别地，`float.hex()` 的输出可以用作 C 或 Java 代码中的十六进制浮点数字面值，而由 C 的 `%a` 格式字符或 Java 的 `Double.toHexString` 所生成的十六进制数字字符串由 `float.fromhex()` 所接受。

请注意 `exponent` 是十进制数而非十六进制数，它给出要与系数相乘的 2 的幂次。例如，十六进制数字字符串 `0x3.a7p10` 表示浮点数 $(3 + 10./16 + 7./16**2) * 2.0**10$ 即 `3740.0`：

```
>>> float.fromhex('0x3.a7p10')
3740.0
```

对 `3740.0` 应用反向转换会得到另一个代表相同数值的十六进制数字字符串：

```
>>> float.hex(3740.0)
'0x1.d380000000000p+11'
```

4.4.4 数字类型的哈希运算

对于可能为不同类型的数字 `x` 和 `y`，要求 `x == y` 时必定 `hash(x) == hash(y)`（详情参见 `__hash__()` 方法的文档）。为了便于在各种数字类型（包括 `int`、`float`、`decimal.Decimal` 和 `fractions.Fraction`）上实现并保证效率，Python 对数字类型的哈希运算是基于为任意有理数定义统一的数学函数，因此该运算对 `int` 和 `fractions.Fraction` 的全部实例，以及 `float` 和 `decimal.Decimal` 的全部有限实例均可用。

从本质上说，此函数是通过以一个固定质数 P 进行 P 降模给出的。 P 的值在 Python 中可以 `sys.hash_info` 的 `modulus` 属性的形式被访问。

CPython implementation detail: 目前所用的质数设定，在 C long 为 32 位的机器上 $P = 2^{31} - 1$ 而在 C long 为 64 位的机器上 $P = 2^{61} - 1$ 。

详细规则如下所述：

- 如果 $x = m / n$ 是一个非负的有理数且 n 不可被 P 整除，则定义 `hash(x)` 为 $m * \text{invmod}(n, P) \% P$ ，其中 `invmod(n, P)` 是对 n 模 P 取反。
- 如果 $x = m / n$ 是一个非负的有理数且 n 可被 P 整除（但 m 不能）则 n 不能对 P 降模，以上规则不适用；在此情况下则定义 `hash(x)` 为常数值 `sys.hash_info.inf`。
- 如果 $x = m / n$ 是一个负的有理数则定义 `hash(x)` 为 `-hash(-x)`。如果结果哈希值为 `-1` 则将其替换为 `-2`。
- 特定值 `sys.hash_info.inf`, `-sys.hash_info.inf` 和 `sys.hash_info.nan` 被用作正无穷、负无穷和空值（所分别对应的）哈希值。（所有可哈希的空值都具有相同的哈希值。）
- 对于一个 `complex` 值 z ，会通过计算 `hash(z.real) + sys.hash_info.imag * hash(z.imag)` 将实部和虚部的哈希值结合起来，并进行降模 $2^{**} \text{sys.hash_info.width}$ 以使其处于 `range(-2^{**}(\text{sys.hash_info.width} - 1), 2^{**}(\text{sys.hash_info.width} - 1))` 范围之内。同样地，如果结果为 `-1` 则将其替换为 `-2`。

为了阐明上述规则，这里有一些等价于内置哈希算法的 Python 代码示例，可用于计算有理数、`float` 或 `complex` 的哈希值：

```
import sys, math

def hash_fraction(m, n):
    """Compute the hash of a rational number m / n.

    Assumes m and n are integers, with n positive.
    Equivalent to hash(fractions.Fraction(m, n)).

    """
    P = sys.hash_info.modulus
    # Remove common factors of P. (Unnecessary if m and n already coprime.)
    while m % P == n % P == 0:
        m, n = m // P, n // P

    if n % P == 0:
        hash_value = sys.hash_info.inf
    else:
        # Fermat's Little Theorem: pow(n, P-1, P) is 1, so
        # pow(n, P-2, P) gives the inverse of n modulo P.
        hash_value = (abs(m) % P) * pow(n, P - 2, P) % P
    if m < 0:
        hash_value = -hash_value
    if hash_value == -1:
        hash_value = -2
    return hash_value

def hash_float(x):
    """Compute the hash of a float x."""

    if math.isnan(x):
        return sys.hash_info.nan
    elif math.isinf(x):
```

(下页继续)

(续上页)

```

        return sys.hash_info.inf if x > 0 else -sys.hash_info.inf
    else:
        return hash_fraction(*x.as_integer_ratio())

def hash_complex(z):
    """Compute the hash of a complex number z."""

    hash_value = hash_float(z.real) + sys.hash_info.imag * hash_float(z.imag)
    # do a signed reduction modulo 2**sys.hash_info.width
    M = 2**(sys.hash_info.width - 1)
    hash_value = (hash_value & (M - 1)) - (hash_value & M)
    if hash_value == -1:
        hash_value = -2
    return hash_value

```

4.5 迭代器类型

Python 支持在容器中进行迭代的概念。这是通过使用两个单独方法来实现的；它们被用于允许用户自定义类对迭代的支持。将在下文中详细描述序列总是支持迭代方法。

容器对象要提供迭代支持，必须定义一个方法：

`container.__iter__()`

返回一个迭代器对象。该对象需要支持下文所述的迭代器协议。如果容器支持不同的迭代类型，则可以提供额外的方法来专门地请求不同迭代类型的迭代器。（支持多种迭代形式的对象的例子有同时支持广度优先和深度优先遍历的树结构。）此方法对应于 Python/C API 中 Python 对象类型结构体的 `tp_iter` 槽位。

迭代器对象自身需要支持以下两个方法，它们共同组成了迭代器协议：

`iterator.__iter__()`

返回迭代器对象本身。这是同时允许容器和迭代器配合 `for` 和 `in` 语句使用所必须的。此方法对应于 Python/C API 中 Python 对象类型结构体的 `tp_iter` 槽位。

`iterator.__next__()`

从容器中返回下一项。如果已经没有项可返回，则会引发 `StopIteration` 异常。此方法对应于 Python/C API 中 Python 对象类型结构体的 `tp_iternext` 槽位。

Python 定义了几种迭代器对象以支持对一般和特定序列类型、字典和其他更特别的形式进行迭代。除了迭代器协议的实现，特定类型的其他性质对迭代操作来说都不重要。

一旦迭代器的 `__next__()` 方法引发了 `StopIteration`，它必须一直对后续调用引发同样的异常。不遵循此行为特性的实现将无法正常使用。

4.5.1 生成器类型

Python 的 `generator` 提供了一种实现迭代器协议的便捷方式。如果容器对象 `__iter__()` 方法被实现为一个生成器，它将自动返回一个迭代器对象（从技术上说是一个生成器对象），该对象提供 `__iter__()` 和 `__next__()` 方法。有关生成器的更多信息可以参阅 `yield` 表达式的文档。

4.6 序列类型—list, tuple, range

有三种基本序列类型：list, tuple 和 range 对象。为处理二进制数据和文本字符串而特别定制的附加序列类型会在专门的小节中描述。

4.6.1 通用序列操作

大多数序列类型，包括可变类型和不可变类型都支持下表中的操作。`collections.abc.Sequence` ABC 被提供用来更容易地在自定义序列类型上正确地实现这些操作。

此表按优先级升序列出了序列操作。在表格中，*s* 和 *t* 是具有相同类型的序列，*n*, *i*, *j* 和 *k* 是整数而 *x* 是任何满足 *s* 所规定的类型和值限制的任意对象。

`in` 和 `not in` 操作具有与比较操作相同的优先级。`+` (拼接) 和 `*` (重复) 操作具有与对应数值运算相同的优先级。³

运算	结果	注释
<code>x in s</code>	如果 <i>s</i> 中的某项等于 <i>x</i> 则结果为 <code>True</code> ，否则为 <code>False</code>	(1)
<code>x not in s</code>	如果 <i>s</i> 中的某项等于 <i>x</i> 则结果为 <code>False</code> ，否则为 <code>True</code>	(1)
<code>s + t</code>	<i>s</i> 与 <i>t</i> 相拼接	(6)(7)
<code>s * n</code> 或 <code>n * s</code>	相当于 <i>s</i> 与自身进行 <i>n</i> 次拼接	(2)(7)
<code>s[i]</code>	<i>s</i> 的第 <i>i</i> 项，起始为 0	(3)
<code>s[i:j]</code>	<i>s</i> 从 <i>i</i> 到 <i>j</i> 的切片	(3)(4)
<code>s[i:j:k]</code>	<i>s</i> 从 <i>i</i> 到 <i>j</i> 步长为 <i>k</i> 的切片	(3)(5)
<code>len(s)</code>	<i>s</i> 的长度	
<code>min(s)</code>	<i>s</i> 的最小项	
<code>max(s)</code>	<i>s</i> 的最大项	
<code>s.index(x[, i[, j]])</code>	<i>x</i> 在 <i>s</i> 中首次出现项的索引号（索引号在 <i>i</i> 或其后且在 <i>j</i> 之前）	(8)
<code>s.count(x)</code>	<i>x</i> 在 <i>s</i> 中出现的总次数	

相同类型的序列也支持比较。特别地，tuple 和 list 的比较是通过比较对应元素的字典顺序。这意味着想要比较结果相等，则每个元素比较结果都必须相等，并且两个序列长度必须相同。（完整细节请参阅语言参考的 `comparisons` 部分。）

注释:

- (1) 虽然 `in` 和 `not in` 操作在通常情况下仅被用于简单的成员检测，某些专门化序列 (例如 `str`, `bytes` 和 `bytearray`) 也使用它们进行子序列检测:

```
>>> "gg" in "eggs"
True
```

- (2) 小于 0 的 *n* 值会被当作 0 来处理 (生成一个与 *s* 同类型的空序列)。请注意序列 *s* 中的项并不会被拷贝；它们会被多次引用。这一点经常会令 Python 编程新手感到困扰；例如:

```
>>> lists = [[]] * 3
>>> lists
[[], [], []]
>>> lists[0].append(3)
>>> lists
[[3], [3], [3]]
```

³ 它们必须如此，因为解析器无法区分这些操作数的类型。

具体的原因在于 `[]` 是一个包含了一个空列表的单元列表，所以 `[] * 3` 结果中的三个元素都是对这个空列表的引用。修改 `lists` 中的任何一个元素实际上都是对这个空列表的修改。你可以用以下方式创建以不同列表为元素的列表：

```
>>> lists = [[] for i in range(3)]
>>> lists[0].append(3)
>>> lists[1].append(5)
>>> lists[2].append(7)
>>> lists
[[3], [5], [7]]
```

进一步的解释可以在 FAQ 条目 [faq-multidimensional-list](#) 中查看。

- (3) 如果 *i* 或 *j* 为负值，则索引顺序是相对于序列 *s* 的末尾：索引号会被替换为 `len(s) + i` 或 `len(s) + j`。但要注意 `-0` 仍然为 `0`。
- (4) *s* 从 *i* 到 *j* 的切片被定义为所有满足 `i ≤ k < j` 的索引号 *k* 的项组成的序列。如果 *i* 或 *j* 大于 `len(s)`，则使用 `len(s)`。如果 *i* 被省略或为 `None`，则使用 `0`。如果 *j* 被省略或为 `None`，则使用 `len(s)`。如果 *i* 大于等于 *j*，则切片为空。
- (5) *s* 从 *i* 到 *j* 步长为 *k* 的切片被定义为所有满足 `0 ≤ n < (j-i)/k` 的索引号 `x = i + n*k` 的项组成的序列。换句话说，索引号为 `i, i+k, i+2*k, i+3*k`，以此类推，当达到 *j* 时停止（但一定不包括 *j*）。当 *k* 为正值时，*i* 和 *j* 会被减至不大于 `len(s)`。当 *k* 为负值时，*i* 和 *j* 会被减至不大于 `len(s) - 1`。如果 *i* 或 *j* 被省略或为 `None`，它们会成为“终止”值（是哪一端的终止值则取决于 *k* 的符号）。请注意，*k* 不可为零。如果 *k* 为 `None`，则当作 `1` 处理。
- (6) 拼接不可变序列总是会生成新的对象。这意味着通过重复拼接来构建序列的运行时开销将会基于序列总长度的乘方。想要获得线性的运行时开销，你必须改用下列替代方案之一：
 - 如果拼接 `str` 对象，你可以构建一个列表并在最后使用 `str.join()` 或是写入一个 `io.StringIO` 实例并在结束时获取它的值
 - 如果拼接 `bytes` 对象，你可以类似地使用 `bytes.join()` 或 `io.BytesIO`，或者你也可以使用 `bytearray` 对象进行原地拼接。`bytearray` 对象是可变的，并且具有高效的重分配机制
 - 如果拼接 `tuple` 对象，请改为扩展 `list` 类
 - 对于其它类型，请查看相应的文档
- (7) 某些序列类型（例如 `range`）仅支持遵循特定模式的项序列，因此并不支持序列拼接或重复。
- (8) 当 *x* 在 *s* 中找不到时 `index` 会引发 `ValueError`。不是所有实现都支持传入额外参数 *i* 和 *j*。这两个参数允许高效地搜索序列的子序列。传入这两个额外参数大致相当于使用 `s[i:j].index(x)`，但是不会复制任何数据，并且返回的索引是相对于序列的开头而非切片的开头。

4.6.2 不可变序列类型

不可变序列类型普遍实现而可变序列类型未实现的唯一操作就是对 `hash()` 内置函数的支持。

这种支持允许不可变类型，例如 `tuple` 实例被用作 `dict` 键，以及存储在 `set` 和 `frozenset` 实例中。

尝试对包含有不可哈希值的不可变序列进行哈希运算将会导致 `TypeError`。

4.6.3 可变序列类型

以下表格中的操作是在可变序列类型上定义的。`collections.abc.MutableSequence` ABC 被提供用来更容易地在自定义序列类型上正确实现这些操作。

表格中的 *s* 是可变序列类型的实例, *t* 是任意可迭代对象, 而 *x* 是符合对 *s* 所规定类型与值限制的任何对象 (例如, `bytearray` 仅接受满足 $0 \leq x \leq 255$ 值限制的整数)。

运算	结果	注释
<code>s[i] = x</code>	将 <i>s</i> 的第 <i>i</i> 项替换为 <i>x</i>	
<code>s[i:j] = t</code>	将 <i>s</i> 从 <i>i</i> 到 <i>j</i> 的切片替换为可迭代对象 <i>t</i> 的内容	
<code>del s[i:j]</code>	等同于 <code>s[i:j] = []</code>	
<code>s[i:j:k] = t</code>	将 <code>s[i:j:k]</code> 的元素替换为 <i>t</i> 的元素	(1)
<code>del s[i:j:k]</code>	从列表中移除 <code>s[i:j:k]</code> 的元素	
<code>s.append(x)</code>	将 <i>x</i> 添加到序列的末尾 (等同于 <code>s[len(s):len(s)] = [x]</code>)	
<code>s.clear()</code>	从 <i>s</i> 中移除所有项 (等同于 <code>del s[:]</code>)	(5)
<code>s.copy()</code>	创建 <i>s</i> 的浅拷贝 (等同于 <code>s[:]</code>)	(5)
<code>s.extend(t)</code> 或 <code>s += t</code>	用 <i>t</i> 的内容扩展 <i>s</i> (基本上等同于 <code>s[len(s):len(s)] = t</code>)	
<code>s *= n</code>	使用 <i>s</i> 的内容重复 <i>n</i> 次来对其进行更新	(6)
<code>s.insert(i, x)</code>	在由 <i>i</i> 给出的索引位置将 <i>x</i> 插入 <i>s</i> (等同于 <code>s[i:i] = [x]</code>)	
<code>s.pop([i])</code>	提取在 <i>i</i> 位置上的项, 并将其从 <i>s</i> 中移除	(2)
<code>s.remove(x)</code>	remove the first item from <i>s</i> where <code>s[i] == x</code>	(3)
<code>s.reverse()</code>	就地将列表中的元素逆序。	(4)

注释:

- (1) *t* 必须与它所替换的切片具有相同的长度。
- (2) 可选参数 *i* 默认为 -1, 因此在默认情况下会移除并返回最后一项。
- (3) 当在 *s* 中找不到 *x* 时 `remove` 操作会引发 `ValueError`。
- (4) 当反转大尺寸序列时 `reverse()` 方法会原地修改该序列以保证空间经济性。为提醒用户此操作是通过间接影响进行的, 它并不会返回反转后的序列。
- (5) 包括 `clear()` 和 `copy()` 是为了与不支持切片操作的可变容器 (例如 `dict` 和 `set`) 的接口保持一致
3.3 新版功能: `clear()` 和 `copy()` 方法。
- (6) *n* 值为一个整数, 或是一个实现了 `__index__()` 的对象。*n* 值为零或负数将清空序列。序列中的项不会被拷贝; 它们会被多次引用, 正如通用序列操作中有关 `s * n` 的说明。

4.6.4 列表

列表是可变序列, 通常用于存放同类项目的集合 (其中精确的相似程度将根据应用而变化)。

class list([iterable])

可以用多种方式构建列表:

- 使用一对方括号来表示空列表: `[]`
- 使用方括号, 其中的项以逗号分隔: `[a], [a, b, c]`
- 使用列表推导式: `[x for x in iterable]`
- 使用类型的构造器: `list()` 或 `list(iterable)`

构造器将构造一个列表, 其中的项与 *iterable* 中的项具有相同的值与顺序。*iterable* 可以是序列、支持迭代的容器或其它可迭代对象。如果 *iterable* 已经是一个列表, 将创建并返回其副本, 类似于 `iterable[:]`。例如, `list('abc')` 返回 `['a', 'b', 'c']` 而 `list((1, 2, 3))` 返回 `[1, 2, 3]`。如果没有给出参数, 构造器将创建一个空列表 `[]`。

其它许多操作也会产生列表, 包括 `sorted()` 内置函数。

列表实现了所有一般和可变序列的操作。列表还额外提供了以下方法:

sort (*, key=None, reverse=False)

此方法会对列表进行原地排序，只使用 < 来进行各项间比较。异常不会被屏蔽——如果有任何比较操作失败，整个排序操作将失败（而列表可能会处于被部分修改的状态）。

`sort()` 接受两个仅限以关键字形式传入的参数（仅限关键字参数）：

`key` 指定带有一个参数的函数，用于从每个列表元素中提取比较键（例如 `key=str.lower`）。对应于列表中每一项的键会被计算一次，然后在整个排序过程中使用。默认值 `None` 表示直接对列表项排序而不计算一个单独的键值。

可以使用 `functools.cmp_to_key()` 将 2.x 风格的 `cmp` 函数转换为 `key` 函数。

`reverse` 为一个布尔值。如果设为 `True`，则每个列表元素将按反向顺序比较进行排序。

当顺序大尺寸序列时此方法会原地修改该序列以保证空间经济性。为提醒用户此操作是通过间接影响进行的，它并不会返回排序后的序列（请使用 `sorted()` 显示地请求一个新的已排序列表实例）。

`sort()` 方法确保是稳定的。如果一个排序确保不会改变比较结果相等的元素的相对顺序就称其为稳定——这有利于进行多重排序（例如先按部门、再接薪级排序）。

CPython implementation detail: 在一个列表被排序期间，尝试改变甚至进行检测也会造成未定义的影响。Python 的 C 实现会在排序期间将列表显示为空，如果发现列表在排序期间被改变将会引发 `ValueError`。

4.6.5 元组

元组是不可变序列，通常用于储存异构数据的多项集（例如由 `enumerate()` 内置函数所产生的二元组）。元组也被用于需要同构数据的不可变序列的情况（例如允许存储到 `set` 或 `dict` 的实例）。

class tuple ([*iterable*])

可以用多种方式构建元组：

- 使用一对圆括号来表示空元组：()
- 使用一个后缀的逗号来表示单元组：`a,` 或 `(a,)`
- 使用以逗号分隔的多个项：`a, b, c` 或 `(a, b, c)`
- 使用内置的 `tuple()`：`tuple()` 或 `tuple(iterable)`

构造器将构造一个元组，其中的项与 *iterable* 中的项具有相同的值与顺序。*iterable* 可以是序列、支持迭代的容器或其他可迭代对象。如果 *iterable* 已经是一个元组，会不加改变地将其返回。例如，`tuple('abc')` 返回 ('a', 'b', 'c') 而 `tuple([1, 2, 3])` 返回 (1, 2, 3)。如果没有给出参数，构造器将创建一个空元组 ()。

请注意决定生成元组的其实是逗号而不是圆括号。圆括号只是可选的，生成空元组或需要避免语法歧义的情况除外。例如，`f(a, b, c)` 是在调用函数时附带三个参数，而 `f((a, b, c))` 则是在调用函数时附带一个三元组。

元组实现了所有一般序列的操作。

对于通过名称访问相比通过索引访问更清晰的异构数据多项集，`collections.namedtuple()` 可能是比简单元组对象更为合适的选择。

4.6.6 range 对象

`range` 类型表示不可变的数字序列，通常用于在 `for` 循环中循环指定的次数。

class range (stop)

class range (*start*, *stop* [, *step*])

`range` 构造器的参数必须为整数（可以是内置的 `int` 或任何实现了 `__index__` 特殊方法的对象）。如果省略 `step` 参数，其默认值为 1。如果省略 `start` 参数，其默认值为 0，如果 `step` 为零则会引发 `ValueError`。

如果 `step` 为正值，确定 `range` `r` 内容的公式为 $r[i] = start + step * i$ 其中 $i \geq 0$ 且 $r[i] < stop$ 。

如果 `step` 为负值，确定 `range` 内容的公式仍然为 $r[i] = start + step * i$ ，但限制条件改为 $i \geq 0$ 且 $r[i] > stop$ 。

如果 `r[0]` 不符合值的限制条件，则该 `range` 对象为空。`range` 对象确实支持负索引，但是会将其解读为从正索引所确定的序列的末尾开始索引。

元素绝对值大于 `sys.maxsize` 的 `range` 对象是被允许的，但某些特性（例如 `len()`）可能引发 `OverflowError`。

一些 `range` 对象的例子：

```
>>> list(range(10))
[0, 1, 2, 3, 4, 5, 6, 7, 8, 9]
>>> list(range(1, 11))
[1, 2, 3, 4, 5, 6, 7, 8, 9, 10]
>>> list(range(0, 30, 5))
[0, 5, 10, 15, 20, 25]
>>> list(range(0, 10, 3))
[0, 3, 6, 9]
>>> list(range(0, -10, -1))
[0, -1, -2, -3, -4, -5, -6, -7, -8, -9]
>>> list(range(0))
[]
>>> list(range(1, 0))
[]
```

`range` 对象实现了一般序列的所有操作，但拼接和重复除外（这是由于 `range` 对象只能表示符合严格模式的序列，而重复和拼接通常都会违反这样的模式）。

start

`start` 形参的值（如果该形参未提供则为 0）

stop

`stop` 形参的值

step

`step` 形参的值（如果该形参未提供则为 1）

`range` 类型相比常规 `list` 或 `tuple` 的优势在于一个 `range` 对象总是占用固定数量的（较小）内存，不论其所表示的范围有多大（因为它只保存了 `start`, `stop` 和 `step` 值，并会根据需要计算具体单项或子范围的值）。

`range` 对象实现了 `collections.abc.Sequence` ABC，提供如包含检测、元素索引查找、切片等特性，并支持负索引（参见序列类型—`list`, `tuple`, `range`）：

```
>>> r = range(0, 20, 2)
>>> r
range(0, 20, 2)
>>> 11 in r
False
>>> 10 in r
True
>>> r.index(10)
```

（下页继续）

(续上页)

```

5
>>> r[5]
10
>>> r[:5]
range(0, 10, 2)
>>> r[-1]
18

```

使用 `==` 和 `!=` 检测 `range` 对象是否相等是将其作为序列来比较。也就是说，如果两个 `range` 对象表示相同的值序列就认为它们是相等的。（请注意比较结果相等的两个 `range` 对象可能会具有不同的 `start`、`stop` 和 `step` 属性，例如 `range(0) == range(2, 1, 3)` 而 `range(0, 3, 2) == range(0, 4, 2)`。）

在 3.2 版更改：实现 Sequence ABC。支持切片和负数索引。使用 `int` 对象在固定时间内进行成员检测，而不是逐一迭代所有项。

在 3.3 版更改：定义 `'=='` 和 `'!='` 以根据 `range` 对象所定义的值序列来进行比较（而不是根据对象的标识）。

3.3 新版功能：`start`、`stop` 和 `step` 属性。

参见：

- `linspace recipe` 演示了如何实现一个延迟求值版本的适合浮点数应用的 `range` 对象。

4.7 文本序列类型—`str`

在 Python 中处理文本数据是使用 `str` 对象，也称为 字符串。字符串是由 Unicode 码位构成的不可变序列。字符串字面值有多种不同的写法：

- 单引号：`' '` 允许包含有 `" 双" 引号'`
- 双引号：`" "` 允许包含有 `' 单' 引号"`。
- 三重引号：`''' 三重单引号'''`，`""" 三重双引号"""`

使用三重引号的字符串可以跨越多行——其中所有的空白字符都将包含在该字符串字面值中。

作为单一表达式组成部分，之间只由空格分隔的多个字符串字面值会被隐式地转换为单个字符串字面值。也就是说，`("spam " "eggs") == "spam eggs"`。

请参阅 `strings` 有解有关不同字符串字面值的更多信息，包括所支持的转义序列，以及使用 `r`（“raw”）前缀来禁用大多数转义序列的处理。

字符串也可以通过使用 `str` 构造器从其他对象创建。

由于不存在单独的“字符”类型，对字符串做索引操作将产生一个长度为 1 的字符串。也就是说，对于一个非空字符串 `s`，`s[0] == s[0:1]`。

不存在可变的字符串类型，但是 `str.join()` 或 `io.StringIO` 可以被用来根据多个片段高效率地构建字符串。

在 3.3 版更改：为了与 Python 2 系列的向下兼容，再次允许字符串字面值使用 `u` 前缀。它对字符串字面值的含义没有影响，并且不能与 `r` 前缀同时出现。

```
class str (object=")
```

```
class str (object=b", encoding='utf-8', errors='strict')
```

返回 *object* 的字符串版本。如果未提供 *object* 则返回空字符串。在其他情况下 `str()` 的行为取决于 *encoding* 或 *errors* 是否有给出，具体见下。

如果 *encoding* 或 *errors* 均未给出, `str(object)` 返回 `object.__str__()`, 这是 *object* 的“非正式”或格式良好的字符串表示。对于字符串对象, 这是该字符串本身。如果 *object* 没有 `__str__()` 方法, 则 `str()` 将回退为返回 `repr(object)`。

如果 *encoding* 或 *errors* 至少给出其中之一, 则 *object* 应该是一个 *bytes-like object* (例如 *bytes* 或 *bytearray*)。在此情况下, 如果 *object* 是一个 *bytes* (或 *bytearray*) 对象, 则 `str(bytes, encoding, errors)` 等价于 `bytes.decode(encoding, errors)`。否则的话, 会在调用 `bytes.decode()` 之前获取缓冲区对象下层的 *bytes* 对象。请参阅二进制序列类型—*bytes*, *bytearray*, *memoryview* 与 *bufferobjects* 了解有关缓冲区对象的信息。

将一个 *bytes* 对象传入 `str()` 而不给出 *encoding* 或 *errors* 参数的操作属于第一种情况, 将返回非正式的字符串表示 (另请参阅 Python 的 `-b` 命令行选项)。例如:

```
>>> str(b'Zoot!')
"b'Zoot!'"
```

有关 `str` 类及其方法的更多信息, 请参阅下面的文本序列类型—*str* 和字符串的方法 小节。要输出格式化字符串, 请参阅 *f-strings* 和格式字符串语法 小节。此外还可以参阅文本处理服务 小节。

4.7.1 字符串的方法

字符串实现了所有一般序列的操作, 还额外提供了以下列出的一些附加方法。

字符串还支持两种字符串格式化样式, 一种提供了很大程度的灵活性和可定制性 (参阅 `str.format()`, 格式字符串语法 和自定义字符串格式化) 而另一种是基于 `C printf` 样式的格式化, 它可处理的类型范围较窄, 并且更难以正确使用, 但对于它可处理的情况往往会更为快速 (*printf* 风格的字符串格式化)。

标准库的文本处理服务 部分涵盖了许多其他模块, 提供各种文本相关工具 (例如包含于 *re* 模块中的正则表达式支持)。

`str.capitalize()`

返回原字符串的副本, 其首个字符大写, 其余为小写。

`str.casefold()`

返回原字符串消除大小写的副本。消除大小写的字符串可用于忽略大小写的匹配。

消除大小写类似于转为小写, 但是更加彻底一些, 因为它会移除字符串中的所有大小写变化形式。例如, 德语小写字母 'ß' 相当于 "ss"。由于它已经是小写了, `lower()` 不会对 'ß' 做任何改变; 而 `casefold()` 则会将其转换为 "ss"。

消除大小写算法的描述请参见 Unicode 标准的 3.13 节。

3.3 新版功能.

`str.center(width[, fillchar])`

返回长度为 *width* 的字符串, 原字符串在其正中。使用指定的 *fillchar* 填充两边的空位 (默认使用 ASCII 空格符)。如果 *width* 小于等于 `len(s)` 则返回原字符串的副本。

`str.count(sub[, start[, end]])`

返回子字符串 *sub* 在 `[start, end]` 范围内非重叠出现的次数。可选参数 *start* 与 *end* 会被解读为切片表示法。

`str.encode(encoding="utf-8", errors="strict")`

返回原字符串编码为字节串对象的版本。默认编码为 'utf-8'。可以给出 *errors* 来设置不同的错误处理方案。*errors* 的默认值为 'strict', 表示编码错误会引发 *UnicodeError*。其他可用的值为 'ignore', 'replace', 'xmlcharrefreplace', 'backslashreplace' 以及任何其他通过 `codecs.register_error()` 注册的值, 请参阅 *Error Handlers* 小节。要查看可用的编码列表, 请参阅标准编码 小节。

在 3.1 版更改: 加入了对关键字参数的支持。

`str.endswith(suffix[, start[, end]])`

如果字符串以指定的 *suffix* 结束返回 `True`，否则返回 `False`。*suffix* 也可以为由多个供查找的后缀构成的元组。如果有可选项 *start*，将从所指定位置开始检查。如果有可选项 *end*，将在所指定位置停止比较。

`str.expandtabs(tabsize=8)`

返回字符串的副本，其中所有的制表符会由一个或多个空格替换，具体取决于当前列位置和给定的制表符宽度。每 *tabsize* 个字符设为一个制表位（默认值 8 时设定的制表位在列 0, 8, 16 依次类推）。要展开字符串，当前列将被设为零并逐一检查字符串中的每个字符。如果字符为制表符 (`\t`)，则会在结果中插入一个或多个空格符，直到当前列等于下一个制表位。（制表符本身不会被复制。）如果字符为换行符 (`\n`) 或回车符 (`\r`)，它会被复制并将当前列重设为零。任何其他字符会被不加修改地复制并将当前列加一，不论该字符在被打印时会如何显示。

```
>>> '01\t012\t0123\t01234'.expandtabs()
'01      012      0123      01234'
>>> '01\t012\t0123\t01234'.expandtabs(4)
'01  012 0123   01234'
```

`str.find(sub[, start[, end]])`

返回子字符串 *sub* 在 `s[start:end]` 切片内被找到的最小索引。可选参数 *start* 与 *end* 会被解读为切片表示法。如果 *sub* 未被找到则返回 `-1`。

注解： `find()` 方法应该只在你需要知道 *sub* 所在位置时使用。要检查 *sub* 是否为子字符串，请使用 `in` 操作符：

```
>>> 'Py' in 'Python'
True
```

`str.format(*args, **kwargs)`

执行字符串格式化操作。调用此方法的字符串可以包含字符串面值或者以花括号 `{}` 括起来的替换域。每个替换域可以包含一个位置参数的数字索引，或者一个关键字参数的名称。返回的字符串副本中每个替换域都会被替换为对应参数的字符串值。

```
>>> "The sum of 1 + 2 is {0}.".format(1+2)
'The sum of 1 + 2 is 3'
```

请参阅[格式字符串语法](#) 了解有关可以在格式字符串中指定的各种格式选项的说明。

注解： 当使用 *n* 类型 (例如: `'{:n}'.format(1234)`) 来格式化数字 (*int*, *float*, *complex*, *decimal.Decimal* 及其子类) 的时候，该函数会临时性地将 `LC_CTYPE` 区域设置为 `LC_NUMERIC` 区域以解码 `localeconv()` 的 `decimal_point` 和 `thousands_sep` 字段，如果它们是非 `ASCII` 字符或长度超过 1 字节的话，并且 `LC_NUMERIC` 区域会与 `LC_CTYPE` 区域不一致。这个临时更改会影响其他线程。

在 3.6.5 版更改: 当使用 *n* 类型格式化数字时，该函数在某些情况下会临时性地将 `LC_CTYPE` 区域设置为 `LC_NUMERIC` 区域。

`str.format_map(mapping)`

类似于 `str.format(**mapping)`，不同之处在于 *mapping* 会被直接使用而不是复制到一个 *dict*。适宜使用此方法的一个例子是当 *mapping* 为 *dict* 的子类的情况：

```
>>> class Default(dict):
...     def __missing__(self, key):
...         return key
```

(下页继续)

(续上页)

```
...
>>> '{name} was born in {country}'.format_map(Default(name='Guido'))
'Guido was born in country'
```

3.2 新版功能.

`str.index(sub[, start[, end]])`

类似于 `find()`，但在找不到子类时会引发 `ValueError`。

`str.isalnum()`

如果字符串中至少有一个字符且所有字符均为字母或数字则返回真值，否则返回假值。如果以下方法中的一个返回 `True` 则字符 `c` 为字母或数字: `c.isalpha()`, `c.isdecimal()`, `c.isdigit()`, or `c.isnumeric()`。

`str.isalpha()`

如果字符串中至少有一个字符且所有字符均为字母则返回真值，否则返回假值。字母类字符是在 Unicode 字符数据库中被定义为 “Letter” 的字符，即一般分类特征属性为 “Lm”，“Lt”，“Lu”，“LI” 或 “Lo” 其中之一。请注意这不同于 Unicode 标准所定义的 “Alphabetic” 特征属性。

`str.isdecimal()`

如果字符串中至少有一个字符且所有字符均为十进制数字字符则返回真值，否则返回假值。十进制数字字符是以 10 为基数的计数制会用来组成数值的字符，例如 U+0660, ARABIC-INDIC DIGIT ZERO。正式的定义为：十进制数字字符就是 Unicode 一般分类 “Nd” 中的字符。

`str.isdigit()`

如果字符串中至少有一个字符且所有字符均为数字字符则返回真值，否则返回假值。数字字符包括十进制数字字符和需要特别处理的数字，例如兼容性上标数字。这也涵盖了不能被用来组成以 10 为基数的数值的数字，例如 Kharosthi 数字。正式的定义为：数字字符就是特征属性值 `Numeric_Type=Digit` 或 `Numeric_Type=Decimal` 的字符。

`str.isidentifier()`

如果字符串根据语言定义属于有效的标识符则返回真值，参见 `identifiers`。

请使用 `keyword.iskeyword()` 来检测保留标识符，例如 `def` 和 `class`。

`str.islower()`

如果字符串中至少有一个区分大小写的字符⁴ 且此类字符均为小写则返回真值，否则返回假值。

`str.isnumeric()`

如果字符串中至少有一个字符且所有字符均为数值字符则返回真值，否则返回假值。数值字符包括数字字符，以及所有在 Unicode 中设置了数值特性属性的字符，例如 U+2155, VULGAR FRACTION ONE FIFTH。正式的定义为：数值字符就是具有特征属性值 `Numeric_Type=Digit`, `Numeric_Type=Decimal` 或 `Numeric_Type=Numeric` 的字符。

`str.isprintable()`

如果字符串中所有字符均为可打印字符或字符串为空则返回真值，否则返回假值。不可打印字符是在 Unicode 字符数据库中被定义为 “Other” 或 “Separator” 的字符，例外情况是 ASCII 空格字符 (0x20) 被视作可打印字符。（请注意在此语境下可打印字符是指当对一个字符串发起调用 `repr()` 时不必被转义的字符。它们与字符串写入 `sys.stdout` 或 `sys.stderr` 时所需的处理无关。）

`str.isspace()`

如果字符串中至少有一个字符且所有字符均为空白字符则返回真值，否则返回假值。空白字符是在 Unicode 字符数据库中被定义为 “Other” 或 “Separator” 并且其双向特征属性为 “WS”，“B” 或 “S” 之一的字符。

`str.istitle()`

如果字符串中至少有一个字符且为标题字符串则返回真值，例如大写字符之后只能带非大写字符而小

⁴ 区分大小写的字符是指所属一般类别属性为 “Lu” (Letter, uppercase), “Ll” (Letter, lowercase) 或 “Lt” (Letter, titlecase) 之一的字符。

写字符必须有大写字符打头。否则返回假值。

`str.isupper()`

如果字符串中至少有一个区分大小写的字符⁴ 且此类字符均为大写则返回真值，否则返回假值。

`str.join(iterable)`

返回一个由 *iterable* 中的字符串拼接而成的字符串。如果 *iterable* 中存在任何非字符串值包括 *bytes* 对象则会引发 *TypeError*。调用该方法的字符串将作为元素之间的分隔。

`str.ljust(width[, fillchar])`

返回长度为 *width* 的字符串，原字符串在其中靠左对齐。使用指定的 *fillchar* 填充空位 (默认使用 ASCII 空格符)。如果 *width* 小于等于 `len(s)` 则返回原字符串的副本。

`str.lower()`

返回原字符串的副本，其所有区分大小写的字符⁴ 均转换为小写。

所用转换小写算法的描述请参见 Unicode 标准的 3.13 节。

`str.lstrip([chars])`

返回原字符串的副本，移除其中的前导字符。*chars* 参数为指定要移除字符的字符串。如果省略或为 *None*，则 *chars* 参数默认移除空格符。实际上 *chars* 参数并非指定单个前缀；而是会移除参数值的所有组合：

```
>>> '   spacious   '.lstrip()
'spacious'
>>> 'www.example.com'.lstrip('cmowz.')
'example.com'
```

`static str.maketrans(x[, y[, z]])`

此静态方法返回一个可供 *str.translate()* 使用的转换对照表。

如果只有一个参数，则它必须是一个将 Unicode 码位序号（整数）或字符（长度为 1 的字符串）映射到 Unicode 码位序号、（任意长度的）字符串或 *None* 的字典。字符键将会被转换为码位序号。

如果有两个参数，则它们必须是两个长度相等的字符串，并且在结果字典中，*x* 中每个字符将被映射到 *y* 中相同位置的字符。如果有第三个参数，它必须是一个字符串，其中的字符将在结果中被映射到 *None*。

`str.partition(sep)`

在 *sep* 首次出现的位置拆分字符串，返回一个 3 元组，其中包含分隔符之前的部分、分隔符本身，以及分隔符之后的部分。如果分隔符未找到，则返回的 3 元组中包含字符本身以及两个空字符串。

`str.replace(old, new[, count])`

返回字符串的副本，其中出现的所有子字符串 *old* 都将被替换为 *new*。如果给出了可选参数 *count*，则只替换前 *count* 次出现。

`str.rfind(sub[, start[, end]])`

返回子字符串 *sub* 在字符串内被找到的最大（最右）索引，这样 *sub* 将包含在 `s[start:end]` 当中。可选参数 *start* 与 *end* 会被解读为切片表示法。如果未找到则返回 -1。

`str.rindex(sub[, start[, end]])`

类似于 *rfind()*，但在子字符串 *sub* 未找到时会引发 *ValueError*。

`str.rjust(width[, fillchar])`

返回长度为 *width* 的字符串，原字符串在其中靠右对齐。使用指定的 *fillchar* 填充空位 (默认使用 ASCII 空格符)。如果 *width* 小于等于 `len(s)` 则返回原字符串的副本。

`str.rpartition(sep)`

在 *sep* 最后一次出现的位置拆分字符串，返回一个 3 元组，其中包含分隔符之前的部分、分隔符本身，以及分隔符之后的部分。如果分隔符未找到，则返回的 3 元组中包含两个空字符串以及字符串本身。

`str.rsplitleft (sep=None, maxsplit=-1)`

返回一个由字符串内单词组成的列表，使用 *sep* 作为分隔字符串。如果给出了 *maxsplit*，则最多进行 *maxsplit* 次拆分，从最右边开始。如果 *sep* 未指定或为 `None`，任何空白字符串都会被作为分隔符。除了从右边开始拆分，*rsplitleft()* 的其他行为都类似于下文所述的 *split()*。

`str.rstrip ([chars])`

返回原字符串的副本，移除其中的末尾字符。*chars* 参数为指定要移除字符的字符串。如果省略或为 `None`，则 *chars* 参数默认移除空格符。实际上 *chars* 参数并非指定单个后缀；而是会移除参数值的所有组合：

```
>>> '   spacious   '.rstrip()
'   spacious'
>>> 'mississippi'.rstrip('ipz')
'mississ'
```

`str.split (sep=None, maxsplit=-1)`

返回一个由字符串内单词组成的列表，使用 *sep* 作为分隔字符串。如果给出了 *maxsplit*，则最多进行 *maxsplit* 次拆分（因此，列表最多会有 *maxsplit*+1 个元素）。如果 *maxsplit* 未指定或为 -1，则不限制拆分次数（进行所有可能的拆分）。

如果给出了 *sep*，则连续的分隔符不会被组合在一起而是被视为分隔空字符串（例如 `'1,,2'.split(',')` 将返回 `['1', '', '2']`）。*sep* 参数可能由多个字符组成（例如 `'1<>2<>3'.split('<>')` 将返回 `['1', '2', '3']`）。使用指定的分隔符拆分空字符串将返回 `['']`。

例如：

```
>>> '1,2,3'.split(',')
['1', '2', '3']
>>> '1,2,3'.split(',', maxsplit=1)
['1', '2,3']
>>> '1,2,,3'.split(',')
['1', '2', '', '3', '']
```

如果 *sep* 未指定或为 `None`，则会应用另一种拆分算法：连续的空格会被视为单个分隔符，其结果将不包含开头或末尾的空字符串，如果字符串包含前缀或后缀空格的话。因此，使用 `None` 拆分空字符串或仅包含空格的字符串将返回 `[]`。

例如：

```
>>> '1 2 3'.split()
['1', '2', '3']
>>> '1 2 3'.split(maxsplit=1)
['1', '2 3']
>>> '   1   2   3   '.split()
['1', '2', '3']
```

`str.splitlines ([keepends])`

返回由原字符串中各行组成的列表，在行边界的位置拆分。结果列表中不包含行边界，除非给出了 *keepends* 且为真值。

此方法会以下列行边界进行拆分。特别地，行边界是 *universal newlines* 的一个超集。

表示符	描述
\n	换行
\r	回车
\r\n	回车 + 换行
\v 或 \x0b	行制表符
\f 或 \x0c	换表单
\x1c	文件分隔符
\x1d	组分隔符
\x1e	记录分隔符
\x85	下一行 (C1 控制码)
\u2028	行分隔符
\u2029	段分隔符

在 3.2 版更改: \v 和 \f 被添加到行边界列表

例如:

```
>>> 'ab c\n\nde fg\rkl\r\n'.splitlines()
['ab c', '', 'de fg', 'kl']
>>> 'ab c\n\nde fg\rkl\r\n'.splitlines(keepends=True)
['ab c\n', '\n', 'de fg\r', 'kl\r\n']
```

不同于 `split()`, 当给出了分隔字符串 `sep` 时, 对于空字符串此方法将返回一个空列表, 而末尾的换行不会令结果中增加额外的行:

```
>>> "".splitlines()
[]
>>> "One line\n".splitlines()
['One line']
```

作为比较, `split('\n')` 的结果为:

```
>>> ''.split('\n')
['']
>>> 'Two lines\n'.split('\n')
['Two lines', '']
```

`str.startswith(prefix[, start[, end]])`

如果字符串以指定的 `prefix` 开始则返回 `True`, 否则返回 `False`。 `prefix` 也可以为由多个供查找的前缀构成的元组。如果有可选项 `start`, 将从所指定位置开始检查。如果有可选项 `end`, 将在所指定位置停止比较。

`str.strip([chars])`

返回原字符串的副本, 移除其中的前导和末尾字符。 `chars` 参数为指定要移除字符的字符串。如果省略或为 `None`, 则 `chars` 参数默认移除空格符。实际上 `chars` 参数并非指定单个前缀或后缀; 而是会移除参数值的所有组合:

```
>>> '   spacious   '.strip()
'spacious'
>>> 'www.example.com'.strip('cmowz.')
'example'
```

最外侧的前导和末尾 `chars` 参数值将从字符串中移除。开头端的字符的移除将在遇到一个未包含于 `chars` 所指定字符集的字符时停止。类似的操作也将在结尾端发生。例如:

```
>>> comment_string = '#..... Section 3.2.1 Issue #32 .....'
>>> comment_string.strip('.#! ')
'Section 3.2.1 Issue #32'
```

str.swapcase()

返回原字符串的副本，其中大写字符转换为小写，反之亦然。请注意 `s.swapcase().swapcase() == s` 并不一定为真值。

str.title()

返回原字符串的标题版本，其中每个单词第一个字母为大写，其余字母为小写。

例如：

```
>>> 'Hello world'.title()
'Hello World'
```

该算法使用一种简单的与语言无关的定义，将连续的字母组合视为单词。该定义在多数情况下都很有效，但它也意味着代表缩写形式与所有格的撇号也会成为单词边界，这可能导致不希望的结果：

```
>>> "they're bill's friends from the UK".title()
'They'Re Bill'S Friends From The Uk'
```

可以使用正则表达式来构建针对撇号的特别处理：

```
>>> import re
>>> def titlecase(s):
...     return re.sub(r"[A-Za-z]+(' [A-Za-z]+)?",
...                     lambda mo: mo.group(0)[0].upper() +
...                                 mo.group(0)[1:].lower(),
...                     s)
...
>>> titlecase("they're bill's friends.")
'They're Bill's Friends.'
```

str.translate(table)

返回原字符串的副本，其中每个字符按给定的转换表进行映射。转换表必须是一个使用 `__getitem__()` 来实现索引操作的对象，通常为 *mapping* 或 *sequence*。当以 Unicode 码位序号（整数）为索引时，转换表对象可以做以下任何一种操作：返回 Unicode 序号或字符串，将字符映射为一个或多个字符；返回 None，将字符从结果字符串中删除；或引发 `LookupError` 异常，将字符映射为其自身。

你可以使用 `str.maketrans()` 基于不同格式的字符到字符映射来创建一个转换映射表。

另请参阅 `codecs` 模块以了解定制字符映射的更灵活方式。

str.upper()

返回原字符串的副本，其中所有区分大小写的字符⁴ 均转换为大写。请注意如果 `s` 包含不区分大小写的字符或者如果结果字符的 Unicode 类别不是 “Lu” (Letter, uppercase) 而是 “Lt” (Letter, titlecase) 则 `s.upper().isupper()` 有可能为 False。

所用转换大写算法的描述请参见 Unicode 标准的 3.13 节。

str.zfill(width)

返回原字符串的副本，在左边填充 ASCII '0' 数码使其长度变为 `width`。正负值前缀 ('+'/'-') 的处理方式是在正负符号之后填充而非在之前。如果 `width` 小于等于 `len(s)` 则返回原字符串的副本。

例如：

```
>>> "42".zfill(5)
'00042'
>>> "-42".zfill(5)
'-0042'
```

4.7.2 printf 风格的字符串格式化

注解： The formatting operations described here exhibit a variety of quirks that lead to a number of common errors (such as failing to display tuples and dictionaries correctly). Using the newer formatted string literals or the `str.format()` interface helps avoid these errors. These alternatives also provide more powerful, flexible and extensible approaches to formatting text.

字符串具有一种特殊的内置操作：使用 `%` (取模) 运算符。这也被称为字符串的 格式化或 插值运算符。对于 `format % values` (其中 *format* 为一个字符串)，在 *format* 中的 `%` 转换标记符将被替换为零个或多个 *values* 条目。其效果类似于在 C 语言中使用 `sprintf()`。

如果 *format* 要求一个单独参数，则 *values* 可以为一个非元组对象。⁵ 否则的话，*values* 必须或者是一个包含项数与格式字符串中指定的转换符项数相同的元组，或者是一个单独映射对象（例如字典）。

转换标记符包含两个或更多字符并具有以下组成，且必须遵循此处规定的顺序：

1. `'%'` 字符，用于标记转换符的起始。
2. 映射键（可选），由加圆括号的字符序列组成（例如 `(somename)`）。
3. 转换旗标（可选），用于影响某些转换类型的结果。
4. 最小字段宽度（可选）。如果指定为 `'*'` (星号)，则实际宽度会从 *values* 元组的下一元素中读取，要转换的对象则为最小字段宽度和可选的精度之后的元素。
5. 精度（可选），以在 `'.'` (点号) 之后加精度值的形式给出。如果指定为 `'*'` (星号)，则实际精度会从 *values* 元组的下一元素中读取，要转换的对象则为精度之后的元素。
6. 长度修饰符（可选）。
7. 转换类型。

当右边的参数为一个字典（或其他映射类型）时，字符串中的格式 必须包含加圆括号的映射键，对应 `'%'` 字符之后字典中的每一项。映射键将从映射中选取要格式化的值。例如：

```
>>> print('%(language)s has %(number)03d quote types.' %
...       {'language': "Python", "number": 2})
Python has 002 quote types.
```

在此情况下格式中不能出现 `*` 标记符（因其需要一个序列类的参数列表）。

转换旗标为：

标志	含义
<code>'#'</code>	值的转换将使用“替代形式”（具体定义见下文）。
<code>'0'</code>	转换将为数字值填充零字符。
<code>'-'</code>	转换值将靠左对齐（如果同时给出 <code>'0'</code> 转换，则会覆盖后者）。
<code>' '</code>	(空格) 符号位转换产生的正数（或空字符串）前将留出一个空格。
<code>'+'</code>	符号字符 (<code>'+'</code> 或 <code>'-'</code>) 将显示于转换结果的开头（会覆盖“空格”旗标）。

⁵ 要格式化单独一个元组，那么你应当提供一个单例元组，其唯一的元素就是要被格式化的元组。

可以给出长度修饰符 (h, l 或 L)，但会被忽略，因为对 Python 来说没有必要 – 所以 `%ld` 等价于 `%d`。

转换类型为：

转换符	含义	注释
'd'	有符号十进制整数。	
'i'	有符号十进制整数。	
'o'	有符号八进制数。	(1)
'u'	过时类型 – 等价于 'd'。	(6)
'x'	有符号十六进制数 (小写)。	(2)
'X'	有符号十六进制数 (大写)。	(2)
'e'	浮点指数格式 (小写)。	(3)
'E'	浮点指数格式 (大写)。	(3)
'f'	浮点十进制格式。	(3)
'F'	浮点十进制格式。	(3)
'g'	浮点格式。如果指数小于 -4 或不小于精度则使用小写指数格式，否则使用十进制格式。	(4)
'G'	浮点格式。如果指数小于 -4 或不小于精度则使用大写指数格式，否则使用十进制格式。	(4)
'c'	单个字符 (接受整数或单个字符的字符串)。	
'r'	字符串 (使用 <code>repr()</code> 转换任何 Python 对象)。	(5)
's'	字符串 (使用 <code>str()</code> 转换任何 Python 对象)。	(5)
'a'	字符串 (使用 <code>ascii()</code> 转换任何 Python 对象)。	(5)
'%'	不转换参数，在结果中输出一个 '%' 字符。	

注释：

- (1) 此替代形式会在第一个数码之前插入标示八进制数的前缀 ('0o')。
- (2) 此替代形式会在第一个数码之前插入 '0x' 或 '0X' 前缀 (取决于是使用 'x' 还是 'X' 格式)。
- (3) 此替代形式总是会在结果中包含一个小数点，即使其后并没有数码。
小数点后的数码位数由精度决定，默认为 6。
- (4) 此替代形式总是会在结果中包含一个小数点，末尾各位的零不会如其他情况下那样被移除。
小数点前后的有效数码位数由精度决定，默认为 6。
- (5) 如果精度为 N，输出将截短为 N 个字符。
- (6) 参见 [PEP 237](#)。

由于 Python 字符串显式指明长度，`%s` 转换不会将 `'\0'` 视为字符串的结束。

在 3.1 版更改：绝对值超过 `1e50` 的 `%f` 转换不会再被替换为 `%g` 转换。

4.8 二进制序列类型—bytes, bytearray, memoryview

操作二进制数据的核心内置类型是 `bytes` 和 `bytearray`。它们由 `memoryview` 提供支持，该对象使用缓冲区协议来访问其他二进制对象所在内存，不需要创建对象的副本。

`array` 模块支持高效地存储基本数据类型，例如 32 位整数和 IEEE754 双精度浮点值。

4.8.1 bytes 对象

`bytes` 对象是由单个字节构成的不可变序列。由于许多主要二进制协议都基于 ASCII 文本编码，因此 `bytes` 对象提供了一些仅在处理 ASCII 兼容数据时可用，并且在许多特性上与字符串对象紧密相关的方法。

class bytes ([*source* [, *encoding* [, *errors*]]])

首先, 表示 `bytes` 字面值的语法与字符串字面值的大致相同, 只是添加了一个 `b` 前缀:

- 单引号: `b'` 同样允许嵌入 `"` 双 `"` 引号 `'`。
- 双引号: `b"` 同样允许嵌入 `'` 单 `'` 引号 `"`。
- 三重引号: `b'''` 三重单引号 `'''`, `b"""` 三重双引号 `"""`

`bytes` 字面值中只允许 ASCII 字符 (无论源代码声明的编码为何)。任何超出 127 的二进制值必须使用相应的转义序列形式加入 `bytes` 字面值。

像字符串字面值一样, `bytes` 字面值也可以使用 `r` 前缀来禁用转义序列处理。请参阅 `strings` 了解有关各种 `bytes` 字面值形式的详情, 包括所支持的转义序列。

虽然 `bytes` 字面值和表示法是基于 ASCII 文本的, 但 `bytes` 对象的行为实际上更像是不可变的整数序列, 序列中的每个值的大小被限制为 $0 \leq x < 256$ (如果违反此限制将引发 `ValueError`)。这种限制是有意设计用以强调以下事实, 虽然许多二进制格式都包含基于 ASCII 的元素, 可以通过某些面向文本的算法进行有用的操作, 但情况对于任意二进制数据来说通常却并非如此 (盲目地将文本处理算法应用于不兼容 ASCII 的二进制数据格式往往将导致数据损坏)。

除了字面值形式, `bytes` 对象还可以通过其他几种方式来创建:

- 指定长度的以零值填充的 `bytes` 对象: `bytes(10)`
- 通过由整数组成的可迭代对象: `bytes(range(20))`
- 通过缓冲区协议复制现有的二进制数据: `bytes(obj)`

另请参阅 `bytes` 内置类型。

由于两个十六进制数码精确对应一个字节, 因此十六进制数是描述二进制数据的常用格式。相应地, `bytes` 类型具有从此种格式读取数据的附加类方法:

classmethod fromhex (*string*)

此 `bytes` 类方法返回一个解码给定字符串的 `bytes` 对象。字符串必须由表示每个字节的两个十六进制数码构成, 其中的 ASCII 空白符会被忽略。

```
>>> bytes.fromhex('2Ef0 F1f2 ')
b'\xf0\xf1\xf2'
```

存在一个反向转换函数, 可以将 `bytes` 对象转换为对应的十六进制表示。

hex ()

返回一个字符串对象, 该对象包含实例中每个字节的两个十六进制数字。

```
>>> b'\xf0\xf1\xf2'.hex()
'f0f1f2'
```

3.5 新版功能.

由于 `bytes` 对象是由整数构成的序列 (类似于元组), 因此对于一个 `bytes` 对象 `b`, `b[0]` 将为一个整数, 而 `b[0:1]` 将为一个长度为 1 的 `bytes` 对象。(这与文本字符串不同, 索引和切片所产生的将都是一个长度为 1 的字符串)。

`bytes` 对象的表示使用字面值格式 (`b'...'`), 因为它通常都要比像 `bytes([46, 46, 46])` 这样的格式更好用。你总是可以使用 `list(b)` 将 `bytes` 对象转换为一个由整数构成的列表。

注解: 针对 Python 2.x 用户的说明: 在 Python 2.x 系列中, 允许 8 位字符串 (2.x 所提供的最接近内置二进制数据类型的对象) 与 Unicode 字符串进行各种隐式转换。这是为了实现向下兼容的变通做法, 以适应 Python

最初只支持 8 位文本而 Unicode 文本是后来才被加入这一事实。在 Python 3.x 中，这些隐式转换已被取消——8 位二进制数据与 Unicode 文本间的转换必须显式地进行，bytes 与字符串对象的比较结果将总是不相等。

4.8.2 bytearray 对象

`bytearray` 对象是 `bytes` 对象的可变对应物。

class bytearray ([*source* [, *encoding* [, *errors*]]])

`bytearray` 对象没有专属的字面值语法，它们总是通过调用构造器来创建：

- 创建一个空实例: `bytearray()`
- 创建一个指定长度的以零值填充的实例: `bytearray(10)`
- 通过由整数组成的可迭代对象: `bytearray(range(20))`
- 通过缓冲区协议复制现有的二进制数据: `bytearray(b'Hi!')`

由于 `bytearray` 对象是可变的，该对象除了 `bytes` 和 `bytearray` 操作中所描述的 `bytes` 和 `bytearray` 共有操作之外，还支持可变序列操作。

另请参见 `bytearray` 内置类型。

由于两个十六进制数码精确对应一个字节，因此十六进制数是描述二进制数据的常用格式。相应地，`bytearray` 类型具有从此种格式读取数据的附加类方法：

classmethod fromhex (*string*)

`bytearray` 类方法返回一个解码给定字符串的 `bytearray` 对象。字符串必须由表示每个字节的两个十六进制数码构成，其中的 ASCII 空白符会被忽略。

```
>>> bytearray.fromhex('2Ef0 F1f2 ')
bytearray(b'\xf0\xf1\xf2')
```

存在一个反向转换函数，可以将 `bytearray` 对象转换为对应的十六进制表示。

hex ()

返回一个字符串对象，该对象包含实例中每个字节的两个十六进制数字。

```
>>> bytearray(b'\xf0\xf1\xf2').hex()
'f0f1f2'
```

3.5 新版功能.

由于 `bytearray` 对象是由整数构成的序列（类似于列表），因此对于一个 `bytearray` 对象 `b`，`b[0]` 将为一个整数，而 `b[0:1]` 将为一个长度为 1 的 `bytearray` 对象。（这与文本字符串不同，索引和切片所产生的将都是一个长度为 1 的字符串）。

`bytearray` 对象的表示使用 `bytes` 对象字面值格式 (`bytearray(b'...')`)，因为它通常都要比 `bytearray([46, 46, 46])` 这样的格式更好用。你总是可以使用 `list(b)` 将 `bytearray` 对象转换为一个由整数构成的列表。

4.8.3 bytes 和 bytearray 操作

`bytes` 和 `bytearray` 对象都支持通用序列操作。它们不仅能与相同类型的操作数，也能与任何 *bytes-like object* 进行互操作。由于这样的灵活性，它们可以在操作中自由地混合而不会导致错误。但是，操作结果的返回值类型可能取决于操作数的顺序。

注解：bytes 和 bytearray 对象的方法不接受字符串作为其参数，就像字符串的方法不接受 bytes 对象作为其参数一样。例如，你必须使用以下写法：

```
a = "abc"
b = a.replace("a", "f")
```

和：

```
a = b"abc"
b = a.replace(b"a", b"f")
```

某些 bytes 和 bytearray 操作假定使用兼容 ASCII 的二进制格式，因此在处理任意二进制数据时应当避免使用。这些限制会在下文中说明。

注解：使用这些基于 ASCII 的操作来处理未以基于 ASCII 的格式存储的二进制数据可能会导致数据损坏。

bytes 和 bytearray 对象的下列方法可以用于任意二进制数据。

bytes.count(sub[, start[, end]])

bytearray.count(sub[, start[, end]])

返回子序列 sub 在 [start, end] 范围内非重叠出现的次数。可选参数 start 与 end 会被解读为切片表示法。

要搜索的子序列可以是任意 bytes-like object 或是 0 至 255 范围内的整数。

在 3.3 版更改：也接受 0 至 255 范围内的整数作为子序列。

bytes.decode(encoding="utf-8", errors="strict")

bytearray.decode(encoding="utf-8", errors="strict")

返回从给定 bytes 解码出来的字符串。默认编码为 'utf-8'。可以给出 errors 来设置不同的错误处理方案。errors 的默认值为 'strict'，表示编码错误会引发 UnicodeError。其他可用的值为 'ignore'，'replace' 以及任何其他通过 codecs.register_error() 注册的名称，请参阅 Error Handlers 小节。要查看可用的编码列表，请参阅标准编码 小节。

注解：将 encoding 参数传给 str 允许直接解码任何 bytes-like object，无须创建临时的 bytes 或 bytearray 对象。

在 3.1 版更改：加入了对关键字参数的支持。

bytes.endswith(suffix[, start[, end]])

bytearray.endswith(suffix[, start[, end]])

如果二进制数据以指定的 suffix 结束则返回 True，否则返回 False。suffix 也可以为由多个供查找的后缀构成的元组。如果有可选项 start，将从所指定位置开始检查。如果有可选项 end，将在所指定位置停止比较。

要搜索的后缀可以是任意 bytes-like object。

bytes.find(sub[, start[, end]])

bytearray.find(sub[, start[, end]])

返回子序列 sub 在数据中被找到的最小索引，sub 包含于切片 s[start:end] 之内。可选参数 start 与 end 会被解读为切片表示法。如果 sub 未被找到则返回 -1。

要搜索的子序列可以是任意 bytes-like object 或是 0 至 255 范围内的整数。

注解: `find()` 方法应该只在你需要知道 `sub` 所在位置时使用。要检查 `sub` 是否为子串, 请使用 `in` 操作符:

```
>>> b'Py' in b'Python'
True
```

在 3.3 版更改: 也接受 0 至 255 范围内的整数作为子序列。

`bytes.index(sub[, start[, end]])`

`bytearray.index(sub[, start[, end]])`

类似于 `find()`, 但在找不到子序列时会引发 `ValueError`。

要搜索的子序列可以是任意 *bytes-like object* 或是 0 至 255 范围内的整数。

在 3.3 版更改: 也接受 0 至 255 范围内的整数作为子序列。

`bytes.join(iterable)`

`bytearray.join(iterable)`

返回一个由 `iterable` 中的二进制数据序列拼接而成的 `bytes` 或 `bytearray` 对象。如果 `iterable` 中存在任何非字节类对象 包括存在 `str` 对象值则会引发 `TypeError`。提供该方法的 `bytes` 或 `bytearray` 对象的内容将作为元素之间的分隔。

static `bytes.maketrans(from, to)`

static `bytearray.maketrans(from, to)`

此静态方法返回一个可用于 `bytes.translate()` 的转换对照表, 它将把 `from` 中的每个字符映射为 `to` 中相同位置上的字符; `from` 与 `to` 必须都是字节类对象 并且具有相同的长度。

3.1 新版功能。

`bytes.partition(sep)`

`bytearray.partition(sep)`

在 `sep` 首次出现的位置拆分序列, 返回一个 3 元组, 其中包含分隔符之前的部分、分隔符本身或其 `bytearray` 副本, 以及分隔符之后的部分。如果分隔符未找到, 则返回的 3 元组中包含原序列以及两个空的 `bytes` 或 `bytearray` 对象。

要搜索的分隔符可以是任意 *bytes-like object*。

`bytes.replace(old, new[, count])`

`bytearray.replace(old, new[, count])`

返回序列的副本, 其中出现的所有子序列 `old` 都将被替换为 `new`。如果给出了可选参数 `count`, 则只替换前 `count` 次出现。

要搜索的子序列及其替换序列可以是任意 *bytes-like object*。

注解: 此方法的 `bytearray` 版本 并非原地操作——它总是产生一个新对象, 即便没有做任何改变。

`bytes.rfind(sub[, start[, end]])`

`bytearray.rfind(sub[, start[, end]])`

返回子序列 `sub` 在序列内被找到的最大(最右)索引, 这样 `sub` 将包含在 `s[start:end]` 当中。可选参数 `start` 与 `end` 会被解读为切片表示法。如果未找到则返回 `-1`。

要搜索的子序列可以是任意 *bytes-like object* 或是 0 至 255 范围内的整数。

在 3.3 版更改: 也接受 0 至 255 范围内的整数作为子序列。

`bytes.rindex(sub[, start[, end]])`

`bytearray.rindex(sub[, start[, end]])`

类似于 `rfind()`, 但在子序列 `sub` 未找到时会引发 `ValueError`。

要搜索的子序列可以是任意 *bytes-like object* 或是 0 至 255 范围内的整数。

在 3.3 版更改: 也接受 0 至 255 范围内的整数作为子序列。

`bytes.rpartition(sep)`

`bytearray.rpartition(sep)`

Split the sequence at the last occurrence of *sep*, and return a 3-tuple containing the part before the separator, the separator itself or its bytearray copy, and the part after the separator. If the separator is not found, return a 3-tuple containing a copy of the original sequence, followed by two empty bytes or bytearray objects.

要搜索的分隔符可以是任意 *bytes-like object*。

`bytes.startswith(prefix[, start[, end]])`

`bytearray.startswith(prefix[, start[, end]])`

如果二进制数据以指定的 *prefix* 开头则返回 True, 否则返回 False。 *prefix* 也可以为由多个供查找的前缀构成的元组。如果有可选项 *start*, 将从所指定位置开始检查。如果有可选项 *end*, 将在所指定位置停止比较。

要搜索的前缀可以是任意 *bytes-like object*。

`bytes.translate(table, delete=b"")`

`bytearray.translate(table, delete=b"")`

返回原 bytes 或 bytearray 对象的副本, 移除其中所有在可选参数 *delete* 中出现的 bytes, 其余 bytes 将通过给定的转换表进行映射, 该转换表必须是长度为 256 的 bytes 对象。

你可以使用 `bytes.maketrans()` 方法来创建转换表。

对于仅需移除字符的转换, 请将 *table* 参数设为 None:

```
>>> b'read this short text'.translate(None, b'aeiou')
b'rd ths shrt txt'
```

在 3.6 版更改: 现在支持将 *delete* 作为关键字参数。

以下 bytes 和 bytearray 对象的方法的默认行为会假定使用兼容 ASCII 的二进制格式, 但通过传入适当的参数仍然可用于任意二进制数据。请注意本小节中所有的 bytearray 方法都不是原地执行操作, 而是会产生新的对象。

`bytes.center(width[, fillbyte])`

`bytearray.center(width[, fillbyte])`

返回原对象的副本, 在长度为 *width* 的序列内居中, 使用指定的 *fillbyte* 填充两边的空位 (默认使用 ASCII 空格符)。对于 *bytes* 对象, 如果 *width* 小于等于 `len(s)` 则返回原序列的副本。

注解: 此方法的 bytearray 版本并非原地操作——它总是产生一个新对象, 即便没有做任何改变。

`bytes.ljust(width[, fillbyte])`

`bytearray.ljust(width[, fillbyte])`

返回原对象的副本, 在长度为 *width* 的序列中靠左对齐。使用指定的 *fillbyte* 填充空位 (默认使用 ASCII 空格符)。对于 *bytes* 对象, 如果 *width* 小于等于 `len(s)` 则返回原序列的副本。

注解: 此方法的 bytearray 版本并非原地操作——它总是产生一个新对象, 即便没有做任何改变。

`bytes.lstrip([chars])`

`bytearray.lstrip([chars])`

返回原序列的副本, 移除指定的前导字节。 *chars* 参数为指定要移除字节值集合的二进制序列——这个名称表明此方法通常是用于 ASCII 字符。如果省略或为 None, 则 *chars* 参数默认移除 ASCII 空白符。 *chars* 参数并非指定单个前缀; 而是会移除参数值的所有组合:

```
>>> b'    spacious    '.lstrip()
b'spacious    '
>>> b'www.example.com'.lstrip(b'cmowz.')
b'example.com'
```

要移除的字节值二进制序列可以是任意 *bytes-like object*。

注解：此方法的 `bytearray` 版本 并非原地操作——它总是产生一个新对象，即便没有做任何改变。

`bytes.rjust(width[, fillbyte])`

`bytearray.rjust(width[, fillbyte])`

返回原对象的副本，在长度为 `width` 的序列中靠右对齐。使用指定的 `fillbyte` 填充空位（默认使用 ASCII 空格符）。对于 *bytes* 对象，如果 `width` 小于等于 `len(s)` 则返回原序列的副本。

注解：此方法的 `bytearray` 版本 并非原地操作——它总是产生一个新对象，即便没有做任何改变。

`bytes.rsplitlep=None, maxsplit=-1)`

`bytearray.rsplitlep=None, maxsplit=-1)`

将二进制序列拆分为相同类型的子序列，使用 `sep` 作为分隔符。如果给出了 `maxsplit`，则最多进行 `maxsplit` 次拆分，从最右边开始。如果 `sep` 未指定或为 `None`，任何只包含 ASCII 空白符的子序列都会被作为分隔符。除了从右边开始拆分，`rsplit()` 的其他行为都类似于下文所述的 `split()`。

`bytes.rstrip([chars])`

`bytearray.rstrip([chars])`

返回原序列的副本，移除指定的末尾字节。`chars` 参数为指定要移除字节值集合的二进制序列——这个名称表明此方法通常是用于 ASCII 字符。如果省略或为 `None`，则 `chars` 参数默认移除 ASCII 空白符。`chars` 参数并非指定单个后缀；而是会移除参数值的所有组合：

```
>>> b'    spacious    '.rstrip()
b'    spacious'
>>> b'mississippi'.rstrip(b'ipz')
b'mississ'
```

要移除的字节值二进制序列可以是任意 *bytes-like object*。

注解：此方法的 `bytearray` 版本 并非原地操作——它总是产生一个新对象，即便没有做任何改变。

`bytes.split(sep=None, maxsplit=-1)`

`bytearray.split(sep=None, maxsplit=-1)`

将二进制序列拆分为相同类型的子序列，使用 `sep` 作为分隔符。如果给出了 `maxsplit` 且非负值，则最多进行 `maxsplit` 次拆分（因此，列表最多会有 `maxsplit+1` 个元素）。如果 `maxsplit` 未指定或为 `-1`，则不限制拆分次数（进行所有可能的拆分）。

如果给出了 `sep`，则连续的分隔符不会被组合在一起而是被视为分隔空子序列（例如 `b'1,,2'.split(b',')` 将返回 `[b'1', b'', b'2']`）。`sep` 参数可能为一个多字节序列（例如 `b'1<>2<>3'.split(b'<>')` 将返回 `[b'1', b'2', b'3']`）。使用指定的分隔符拆分空序列将返回 `[b'']` 或 `[bytearray(b'')]`，具体取决于被拆分对象的类型。`sep` 参数可以是任意 *bytes-like object*。

例如：

```
>>> b'1,2,3'.split(b',')
[b'1', b'2', b'3']
```

(下页继续)

(续上页)

```
>>> b'1,2,3'.split(b',', maxsplit=1)
[b'1', b'2,3']
>>> b'1,2,,3,.'.split(b',')
[b'1', b'2', b'', b'3', b'']
```

如果 *sep* 未指定或为 `None`，则会应用另一种拆分算法：连续的 ASCII 空白符会被视为单个分隔符，其结果将不包含序列开头或末尾的空白符。因此，在不指定分隔符的情况下对空序列或仅包含 ASCII 空白符的序列进行拆分将返回 `[]`。

例如：

```
>>> b'1 2 3'.split()
[b'1', b'2', b'3']
>>> b'1 2 3'.split(maxsplit=1)
[b'1', b'2 3']
>>> b' 1 2 3 '.split()
[b'1', b'2', b'3']
```

`bytes.strip([chars])`

`bytearray.strip([chars])`

返回原序列的副本，移除指定的开头和末尾字节。*chars* 参数为指定要移除字节值集合的二进制序列——这个名称表明此方法通常是用于 ASCII 字符。如果省略或为 `None`，则 *chars* 参数默认移除 ASCII 空白符。*chars* 参数并非指定单个前缀或后缀；而是会移除参数值的所有组合：

```
>>> b'   spacious   '.strip()
b'spacious'
>>> b'www.example.com'.strip(b'cmowz.')
b'example'
```

要移除的字节值二进制序列可以是任意 *bytes-like object*。

注解：此方法的 `bytearray` 版本并非原地操作——它总是产生一个新对象，即便没有做任何改变。

以下 `bytes` 和 `bytearray` 对象的方法会假定使用兼容 ASCII 的二进制格式，不应当被应用于任意二进制数据。请注意本小节中所有的 `bytearray` 方法都不是原地执行操作，而是会产生新的对象。

`bytes.capitalize()`

`bytearray.capitalize()`

返回原序列的副本，其中每个字节都将被解读为一个 ASCII 字符，并且第一个字节的字符大写而其余的小写。非 ASCII 字节值将保持原样不变。

注解：此方法的 `bytearray` 版本并非原地操作——它总是产生一个新对象，即便没有做任何改变。

`bytes.expandtabs(tabsize=8)`

`bytearray.expandtabs(tabsize=8)`

返回序列的副本，其中所有的 ASCII 制表符会由一个或多个 ASCII 空格替换，具体取决于当前列位置和给定的制表符宽度。每 *tabsize* 个字节设为一个制表位（默认值 8 时设定的制表位在列 0, 8, 16 依次类推）。要展开序列，当前列位置将被设为零并逐一检查序列中的每个字节。如果字节为 ASCII 制表符 (`b'\t'`)，则并在结果中插入一个或多个空格符，直到当前列等于下一个制表位。（制表符本身不会被复制。）如果当前字节为 ASCII 换行符 (`b'\n'`) 或回车符 (`b'\r'`)，它会被复制并将当前列重设为零。任何其他字节会被不加修改地复制并将当前列加一，不论该字节值在被打印时会如何显示：

```
>>> b'01\t012\t0123\t01234'.expandtabs()
b'01      012      0123      01234'
>>> b'01\t012\t0123\t01234'.expandtabs(4)
b'01  012 0123  01234'
```

注解：此方法的 `bytearray` 版本并非原地操作——它总是产生一个新对象，即便没有做任何改变。

`bytes.isalnum()`

`bytearray.isalnum()`

如果序列中所有字节都是字母类 ASCII 字符或 ASCII 十进制数码并且序列非空则返回真值，否则返回假值。字母类 ASCII 字符就是字节值包含在序列 `b'abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ'` 中的字符。ASCII 十进制数码就是字节值包含在序列 `b'0123456789'` 中的字符。

例如：

```
>>> b'ABCabc1'.isalnum()
True
>>> b'ABC abc1'.isalnum()
False
```

`bytes.isalpha()`

`bytearray.isalpha()`

如果序列中所有字节都是字母类 ASCII 字符并且序列非空则返回真值，否则返回假值。字母类 ASCII 字符就是字节值包含在序列 `b'abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ'` 中的字符。

例如：

```
>>> b'ABCabc'.isalpha()
True
>>> b'ABCabc1'.isalpha()
False
```

`bytes.isdigit()`

`bytearray.isdigit()`

如果序列中所有字节都是 ASCII 十进制数码并且序列非空则返回真值，否则返回假值。ASCII 十进制数码就是字节值包含在序列 `b'0123456789'` 中的字符。

例如：

```
>>> b'1234'.isdigit()
True
>>> b'1.23'.isdigit()
False
```

`bytes.islower()`

`bytearray.islower()`

如果序列中至少有一个小写 ASCII 字符并且没有大写 ASCII 字符则返回真值，否则返回假值。

例如：

```
>>> b'hello world'.islower()
True
>>> b'Hello world'.islower()
False
```

小写 ASCII 字符就是字节值包含在序列 `b'abcdefghijklmnopqrstuvwxyz'` 中的字符。大写 ASCII 字符就是字节值包含在序列 `b'ABCDEFGHIJKLMNOPQRSTUVWXYZ'` 中的字符。

`bytes.isspace()`

`bytearray.isspace()`

如果序列中所有字节都是 ASCII 空白符并且序列非空则返回真值，否则返回假值。ASCII 空白符就是字节值包含在序列 `b' \t\n\r\x0b\f'` (空格, 制表, 换行, 回车, 垂直制表, 换页) 中的字符。

`bytes.istitle()`

`bytearray.istitle()`

如果序列为 ASCII 标题形式并且序列非空则返回真值，否则返回假值。请参阅 `bytes.title()` 了解有关“标题形式”的详细定义。

例如:

```
>>> b'Hello World'.istitle()
True
>>> b'Hello world'.istitle()
False
```

`bytes.isupper()`

`bytearray.isupper()`

如果序列中至少有一个大写字母 ASCII 字符并且没有小写 ASCII 字符则返回真值，否则返回假值。

例如:

```
>>> b'HELLO WORLD'.isupper()
True
>>> b'Hello world'.isupper()
False
```

小写 ASCII 字符就是字节值包含在序列 `b'abcdefghijklmnopqrstuvwxyz'` 中的字符。大写 ASCII 字符就是字节值包含在序列 `b'ABCDEFGHIJKLMNOPQRSTUVWXYZ'` 中的字符。

`bytes.lower()`

`bytearray.lower()`

返回原序列的副本，其所有大写 ASCII 字符均转换为对应的小写形式。

例如:

```
>>> b'Hello World'.lower()
b'hello world'
```

小写 ASCII 字符就是字节值包含在序列 `b'abcdefghijklmnopqrstuvwxyz'` 中的字符。大写 ASCII 字符就是字节值包含在序列 `b'ABCDEFGHIJKLMNOPQRSTUVWXYZ'` 中的字符。

注解: 此方法的 `bytearray` 版本 并非原地操作——它总是产生一个新对象，即便没有做任何改变。

`bytes.splitlines(keepends=False)`

`bytearray.splitlines(keepends=False)`

返回由原二进制序列中各行组成的列表，在 ASCII 行边界符的位置拆分。此方法使用 *universal newlines* 方式来分行。结果列表中不包含换行符，除非给出了 *keepends* 且为真值。

例如:

```
>>> b'ab c\n\nde fg\rkl\r\n'.splitlines()
[b'ab c', b'', b'de fg', b'kl']
```

(下页继续)

(续上页)

```
>>> b'ab c\n\nde fg\rkl\r\n'.splitlines(keepends=True)
[b'ab c\n', b'\n', b'de fg\r', b'kl\r\n']
```

不同于 `split()`，当给出了分隔符 `sep` 时，对于空字符串此方法将返回一个空列表，而末尾的换行不会令结果中增加额外的行：

```
>>> b"".split(b'\n'), b"Two lines\n".split(b'\n')
([b''], [b'Two lines', b''])
>>> b"".splitlines(), b"One line\n".splitlines()
([], [b'One line'])
```

`bytes.swapcase()`

`bytearray.swapcase()`

返回原序列的副本，其所有小写 ASCII 字符均转换为对应的大写形式，反之亦反。

例如：

```
>>> b'Hello World'.swapcase()
b'hELLO wORLD'
```

小写 ASCII 字符就是字节值包含在序列 `b'abcdefghijklmnopqrstuvwxyz'` 中的字符。大写 ASCII 字符就是字节值包含在序列 `b'ABCDEFGHIJKLMNOPQRSTUVWXYZ'` 中的字符。

不同于 `str.swapcase()`，在些二进制版本下 `bin.swapcase().swapcase() == bin` 总是成立。大小写转换在 ASCII 中是对称的，即使其对于任意 Unicode 码位来说并不总是成立。

注解：此方法的 `bytearray` 版本并非原地操作——它总是产生一个新对象，即便没有做任何改变。

`bytes.title()`

`bytearray.title()`

返回原二进制序列的标题版本，其中每个单词以一个大写 ASCII 字符为开头，其余字母为小写。不区别大小写的字节值将保持原样不变。

例如：

```
>>> b'Hello world'.title()
b'Hello World'
```

小写 ASCII 字符就是字节值包含在序列 `b'abcdefghijklmnopqrstuvwxyz'` 中的字符。大写 ASCII 字符就是字节值包含在序列 `b'ABCDEFGHIJKLMNOPQRSTUVWXYZ'` 中的字符。所有其他字节值都不区分大小写。

该算法使用一种简单的与语言无关的定义，将连续的字母组合视为单词。该定义在多数情况下都很有效，但它也意味着代表缩写形式与所有格的撇号也会成为单词边界，这可能导致不希望的结果：

```
>>> b"they're bill's friends from the UK".title()
b"They'Re Bill'S Friends From The Uk"
```

可以使用正则表达式来构建针对撇号的特别处理：

```
>>> import re
>>> def titlecase(s):
...     return re.sub(rb"[A-Za-z]+('[A-Za-z]+)?",
...                     lambda mo: mo.group(0)[0:1].upper() +
...                               mo.group(0)[1:].lower(),
```

(下页继续)

(续上页)

```

...         s)
...
>>> titlecase(b"they're bill's friends.")
b"They're Bill's Friends."

```

注解：此方法的 `bytearray` 版本 并非原地操作——它总是产生一个新对象，即便没有做任何改变。

`bytes.upper()`

`bytearray.upper()`

返回原序列的副本，其所有小写 ASCII 字符均转换为对应的大写形式。

例如：

```

>>> b'Hello World'.upper()
b'HELLO WORLD'

```

小写 ASCII 字符就是字节值包含在序列 `b'abcdefghijklmnopqrstuvwxyz'` 中的字符。大写 ASCII 字符就是字节值包含在序列 `b'ABCDEFGHIJKLMNOPQRSTUVWXYZ'` 中的字符。

注解：此方法的 `bytearray` 版本 并非原地操作——它总是产生一个新对象，即便没有做任何改变。

`bytes.zfill(width)`

`bytearray.zfill(width)`

返回原序列的副本，在左边填充 `b'0'` 数码使序列长度为 `width`。正负值前缀 (`b'+' / b'-'`) 的处理方式是在正负符号 之后填充而非在之前。对于 `bytes` 对象，如果 `width` 小于等于 `len(seq)` 则返回原序列。

例如：

```

>>> b"42".zfill(5)
b'00042'
>>> b"-42".zfill(5)
b'-0042'

```

注解：此方法的 `bytearray` 版本 并非原地操作——它总是产生一个新对象，即便没有做任何改变。

4.8.4 printf 风格的字节串格式化

注解：此处介绍的格式化操作具有多种怪异特性，可能导致许多常见错误（例如无法正确显示元组和字典）。如果要打印的值可能为元组或字典，请将其放入一个元组中。

字节串对象 (`bytes/bytearray`) 具有一种特殊的内置操作：使用 `%` (取模) 运算符。这也被称为字节串的 格式化或 插值运算符。对于 `format % values` (其中 `format` 为一个字节串对象)，在 `format` 中的 `%` 转换标记符将被替换为零个或多个 `values` 条目。其效果类似于在 C 语言中使用 `sprintf()`。

如果 `format` 要求一个单独参数，则 `values` 可以为一个非元组对象⁵。否则的话，`values` 必须或是是一个包含项数与格式字节串对象中指定的转换符项数相同的元组，或者是一个单独的映射对象（例如元组）。

转换标记符包含两个或更多字符并具有以下组成，且必须遵循此处规定的顺序：

1. '%' 字符，用于标记转换符的起始。
2. 映射键（可选），由加圆括号的字符序列组成（例如 (somename)）。
3. 转换旗标（可选），用于影响某些转换类型的结果。
4. 最小字段宽度（可选）。如果指定为 '*'（星号），则实际宽度会从 *values* 元组的下一元素中读取，要转换的对象则为最小字段宽度和可选的精度之后的元素。
5. 精度（可选），以在 '.'（点号）之后加精度值的形式给出。如果指定为 '*'（星号），则实际精度会从 *values* 元组的下一元素中读取，要转换的对象则为精度之后的元素。
6. 长度修饰符（可选）。
7. 转换类型。

当右边的参数为一个字典（或其他映射类型）时，字节串对象中的格式 必须包含加圆括号的映射键，对应 '%' 字符之后字典中的每一项。映射键将从映射中选取要格式化的值。例如：

```
>>> print(b'%(language)s has %(number)03d quote types.' %
...       {b'language': b"Python", b"number": 2})
b'Python has 002 quote types.'
```

在此情况下格式中不能出现 * 标记符（因其需要一个序列类的参数列表）。

转换旗标为：

标志	含义
'#'	值的转换将使用“替代形式”（具体定义见下文）。
'0'	转换将为数字值填充零字符。
'-'	转换值将靠左对齐（如果同时给出 '0' 转换，则会覆盖后者）。
' '	(空格) 符号位转换产生的正数（或空字符串）前将留出一个空格。
'+'	符号字符 ('+' 或 '-') 将显示于转换结果的开头（会覆盖“空格”旗标）。

可以给出长度修饰符 (h, l 或 L)，但会被忽略，因为对 Python 来说没有必要 – 所以 %ld 等价于 %d。

转换类型为：

转 换 符	含义	注 释
'd'	有符号十进制整数。	
'i'	有符号十进制整数。	
'o'	有符号八进制数。	(1)
'u'	过时类型 – 等价于 'd'。	(8)
'x'	有符号十六进制数 (小写)。	(2)
'X'	有符号十六进制数 (大写)。	(2)
'e'	浮点指数格式 (小写)。	(3)
'E'	浮点指数格式 (大写)。	(3)
'f'	浮点十进制格式。	(3)
'F'	浮点十进制格式。	(3)
'g'	浮点格式。如果指数小于 -4 或不小于精度则使用小写指数格式，否则使用十进制格式。	(4)
'G'	浮点格式。如果指数小于 -4 或不小于精度则使用大写指数格式，否则使用十进制格式。	(4)
'c'	单个字节 (接受整数或单个字节对象)。	
'b'	字节串 (任何遵循 缓冲区协议或是具有 <code>__bytes__()</code> 的对象)。	(5)
's'	's' 是 'b' 的一个别名，只应当在基于 Python2/3 的代码中使用。	(6)
'a'	字节串 (使用 <code>repr(obj).encode('ascii', 'backslashreplace)</code> 转换任何 Python 对象)。	(5)
'r'	'r' 是 'a' 的一个别名，只应当在基于 Python2/3 的代码中使用。	(7)
'%'	不转换参数，在结果中输出一个 '%' 字符。	

注释:

- (1) 此替代形式会在第一个数码之前插入标示八进制数的前缀 ('0o')。
- (2) 此替代形式会在第一个数码之前插入 '0x' 或 '0X' 前缀 (取决于是使用 'x' 还是 'X' 格式)。
- (3) 此替代形式总是会在结果中包含一个小数点，即使其后并没有数码。
小数点后的数码位数由精度决定，默认为 6。
- (4) 此替代形式总是会在结果中包含一个小数点，末尾各位的零不会如其他情况下那样被移除。
小数点前后的有效数码位数由精度决定，默认为 6。
- (5) 如果精度为 N，输出将截短为 N 个字符。
- (6) `b'%s'` 已弃用，但在 3.x 系列中将不会被移除。
- (7) `b'%r'` 已弃用，但在 3.x 系列中将不会被移除。
- (8) 参见 [PEP 237](#)。

注解: 此方法的 `bytearray` 版本 并非原地操作——它总是产生一个新对象，即便没有做任何改变。

参见:

[PEP 461](#) - 为 `bytes` 和 `bytearray` 添加 % 格式化

3.5 新版功能.

4.8.5 内存视图

`memoryview` 对象允许 Python 代码访问一个对象的内部数据，只要该对象支持 缓冲区协议而无需进行拷贝。

class memoryview(obj)

创建一个引用 *obj* 的 *memoryview*。 *obj* 必须支持缓冲区协议。支持缓冲区协议的内置对象包括 *bytes* 和 *bytearray*。

memoryview 具有 元素的概念，即由原始对象 *obj* 所处理的基本内存单元。对于许多简单类型例如 *bytes* 和 *bytearray* 来说，一个元素就是一个字节，但是其他的类型例如 *array.array* 可能有更大的元素。

`len(view)` 与 *tolist* 的长度相等。如果 `view.ndim = 0`，则其长度为 1。如果 `view.ndim = 1`，则其长度等于 *view* 中元素的数量。对于更高的维度，其长度等于表示 *view* 的嵌套列表的长度。*itemsize* 属性可向你给出单个元素所占的字节数。

memoryview 支持通过切片和索引访问其元素。一维切片的结果将是一个子视图：

```
>>> v = memoryview(b'abcefg')
>>> v[1]
98
>>> v[-1]
103
>>> v[1:4]
<memory at 0x7f3ddc9f4350>
>>> bytes(v[1:4])
b'bce'
```

如果 *format* 是一个来自于 *struct* 模块的原生格式说明符，则也支持使用整数或由整数构成的元组进行索引，并返回具有正确类型的单个 元素。一维内存视图可以使用一个整数或由一个整数构成的元组进行索引。多维内存视图可以使用由恰好 *ndim* 个整数构成的元素进行索引，*ndim* 即其维度。零维内存视图可以使用空元组进行索引。

这里是一个使用非字节格式的例子：

```
>>> import array
>>> a = array.array('l', [-11111111, 22222222, -33333333, 44444444])
>>> m = memoryview(a)
>>> m[0]
-11111111
>>> m[-1]
44444444
>>> m[::2].tolist()
[-11111111, -33333333]
```

如果下层对象是可写的，则内存视图支持一维切片赋值。改变大小则不被允许：

```
>>> data = bytearray(b'abcefg')
>>> v = memoryview(data)
>>> v.readonly
False
>>> v[0] = ord(b'z')
>>> data
bytearray(b'zbcefg')
>>> v[1:4] = b'123'
>>> data
bytearray(b'z123fg')
>>> v[2:3] = b'spam'
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
ValueError: memoryview assignment: lvalue and rvalue have different structures
>>> v[2:6] = b'spam'
```

(下页继续)

(续上页)

```
>>> data
bytearray(b'z!spam')
```

由带有格式符号 ‘B’，‘b’ 或 ‘c’ 的可哈希（只读）类型构成的一维内存视图同样是可哈希的。哈希定义为 `hash(m) == hash(m.tobytes())`：

```
>>> v = memoryview(b'abcefg')
>>> hash(v) == hash(b'abcefg')
True
>>> hash(v[2:4]) == hash(b'ce')
True
>>> hash(v[::-2]) == hash(b'abcefg'[::-2])
True
```

在 3.3 版更改：一维内存视图现在可以被切片。带有格式符号 ‘B’，‘b’ 或 ‘c’ 的一维内存视图现在是可哈希的。

在 3.4 版更改：内存视图现在会自动注册为 `collections.abc.Sequence`

在 3.5 版更改：内存视图现在可使用整数元组进行索引。

`memoryview` 具有以下一些方法：

`__eq__` (exporter)

`memoryview` 与 **PEP 3118** 中的导出器这两者如果形状相同，并且如果当使用 `struct` 语法解读操作数的相应格式代码时所有对应值都相同，则它们就是等价的。

对于 `tolist()` 当前所支持的 `struct` 格式字符串子集，如果 `v.tolist() == w.tolist()` 则 `v` 和 `w` 相等：

```
>>> import array
>>> a = array.array('I', [1, 2, 3, 4, 5])
>>> b = array.array('d', [1.0, 2.0, 3.0, 4.0, 5.0])
>>> c = array.array('b', [5, 3, 1])
>>> x = memoryview(a)
>>> y = memoryview(b)
>>> x == a == y == b
True
>>> x.tolist() == a.tolist() == y.tolist() == b.tolist()
True
>>> z = y[::-2]
>>> z == c
True
>>> z.tolist() == c.tolist()
True
```

如果两边的格式字符串都不被 `struct` 模块所支持，则两对象比较结果总是不相等（即使格式字符串和缓冲区内容相同）：

```
>>> from ctypes import BigEndianStructure, c_long
>>> class BEPoint(BigEndianStructure):
...     _fields_ = [("x", c_long), ("y", c_long)]
...
>>> point = BEPoint(100, 200)
>>> a = memoryview(point)
>>> b = memoryview(point)
>>> a == point
False
```

(下页继续)

(续上页)

```
>>> a == b
False
```

请注意，与浮点数的情况一样，对于内存视图对象来说，`v is w` 也并不意味着 `v == w`。
在 3.3 版更改：之前的版本比较原始内存时会忽略条目的格式与逻辑数组结构。

tobytes()

将缓冲区中的数据作为字节串返回。这相当于在内存视图上调用 `bytes` 构造器。

```
>>> m = memoryview(b"abc")
>>> m.tobytes()
b'abc'
>>> bytes(m)
b'abc'
```

对于非连续数组，结果等于平面化表示的列表，其中所有元素都转换为字节串。`tobytes()` 支持所有格式字符串，不符合 `struct` 模块语法的那些也包括在内。

hex()

返回一个字符串对象，其中分别以两个十六进制数码表示缓冲区里的每个字节。

```
>>> m = memoryview(b"abc")
>>> m.hex()
'616263'
```

3.5 新版功能.**tolist()**

将缓冲区内的数据以一个元素列表的形式返回。

```
>>> memoryview(b'abc').tolist()
[97, 98, 99]
>>> import array
>>> a = array.array('d', [1.1, 2.2, 3.3])
>>> m = memoryview(a)
>>> m.tolist()
[1.1, 2.2, 3.3]
```

在 3.3 版更改：`tolist()` 现在支持 `struct` 模块语法中的所有单字符原生格式以及多维表示形式。

release()

释放由内存视图对象所公开的底层缓冲区。许多对象在被视图所获取时都会采取特殊动作（例如，`bytearray` 将会暂时禁止调整大小）；因此，调用 `release()` 可以方便地尽早去除这些限制（并释放任何多余的资源）。

在此方法被调用后，任何对视图的进一步操作将引发 `ValueError` (`release()` 本身除外，它可以被多次调用)：

```
>>> m = memoryview(b'abc')
>>> m.release()
>>> m[0]
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
ValueError: operation forbidden on released memoryview object
```

使用 `with` 语句，可以通过上下文管理协议达到类似的效果：

```
>>> with memoryview(b'abc') as m:
...     m[0]
...
97
>>> m[0]
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
ValueError: operation forbidden on released memoryview object
```

3.2 新版功能.

cast (*format* [, *shape*])

将内存视图转化为新的格式或形状。*shape* 默认为 `[byte_length//new_itemsize]`，这意味着结果视图将是一维的。返回值是一个新的内存视图，但缓冲区本身不会被复制。支持的转化有 1D -> C-contiguous 和 C-contiguous -> 1D。

目标格式仅限于 *struct* 语法中的单一元素原生格式。其中一种格式必须为字节格式（‘B’，‘b’或‘c’）。结果的字节长度必须与原始长度相同。

将 1D/long 转换为 1D/unsigned bytes:

```
>>> import array
>>> a = array.array('l', [1,2,3])
>>> x = memoryview(a)
>>> x.format
'l'
>>> x.itemsize
8
>>> len(x)
3
>>> x.nbytes
24
>>> y = x.cast('B')
>>> y.format
'B'
>>> y.itemsize
1
>>> len(y)
24
>>> y.nbytes
24
```

将 1D/unsigned bytes 转换为 1D/char:

```
>>> b = bytearray(b'zyz')
>>> x = memoryview(b)
>>> x[0] = b'a'
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
ValueError: memoryview: invalid value for format "B"
>>> y = x.cast('c')
>>> y[0] = b'a'
>>> b
bytearray(b'ayz')
```

将 1D/bytes 转换为 3D/ints 再转换为 1D/signed char:

```

>>> import struct
>>> buf = struct.pack("i"*12, *list(range(12)))
>>> x = memoryview(buf)
>>> y = x.cast('i', shape=[2,2,3])
>>> y.tolist()
[[[0, 1, 2], [3, 4, 5]], [[6, 7, 8], [9, 10, 11]]]
>>> y.format
'i'
>>> y.itemsize
4
>>> len(y)
2
>>> y.nbytes
48
>>> z = y.cast('b')
>>> z.format
'b'
>>> z.itemsize
1
>>> len(z)
48
>>> z.nbytes
48

```

将 1D/unsigned char 转换为 2D/unsigned long:

```

>>> buf = struct.pack("L"*6, *list(range(6)))
>>> x = memoryview(buf)
>>> y = x.cast('L', shape=[2,3])
>>> len(y)
2
>>> y.nbytes
48
>>> y.tolist()
[[0, 1, 2], [3, 4, 5]]

```

3.3 新版功能.

在 3.5 版更改: 当转换为字节视图时, 源格式将不再受限。

还存在一些可用的只读属性:

obj

内存视图的下层对象:

```

>>> b = bytearray(b'xyz')
>>> m = memoryview(b)
>>> m.obj is b
True

```

3.3 新版功能.

nbytes

`nbytes == product(shape) * itemsize == len(m.tobytes())`。这是数组在连续表示时将会占用的空间总字节数。它不一定等于 `len(m)`:

```

>>> import array
>>> a = array.array('i', [1,2,3,4,5])

```

(下页继续)

(续上页)

```

>>> m = memoryview(a)
>>> len(m)
5
>>> m.nbytes
20
>>> y = m[::2]
>>> len(y)
3
>>> y.nbytes
12
>>> len(y.tobytes())
12

```

多维数组:

```

>>> import struct
>>> buf = struct.pack("d"*12, *[1.5*x for x in range(12)])
>>> x = memoryview(buf)
>>> y = x.cast('d', shape=[3,4])
>>> y.tolist()
[[0.0, 1.5, 3.0, 4.5], [6.0, 7.5, 9.0, 10.5], [12.0, 13.5, 15.0, 16.5]]
>>> len(y)
3
>>> y.nbytes
96

```

3.3 新版功能.

readonly

一个表明内存是否只读的布尔值。

format

一个字符串，包含视图中每个元素的格式（表示为`struct`模块样式）。内存视图可以从具有任意格式字符串的导出器创建，但某些方法（例如`tolist()`）仅限于原生的单元元素格式。

在 3.3 版更改: 格式 'B' 现在会按照 `struct` 模块语法来处理。这意味着 `memoryview(b'abc')[0] == b'abc'[0] == 97`。

itemsize

`memoryview` 中每个元素以字节表示的大小:

```

>>> import array, struct
>>> m = memoryview(array.array('H', [32000, 32001, 32002]))
>>> m.itemsize
2
>>> m[0]
32000
>>> struct.calcsize('H') == m.itemsize
True

```

ndim

一个整数，表示内存所代表的多维数组具有多少个维度。

shape

一个整数元组，通过 `ndim` 的长度值给出内存所代表的 N 维数组的形状。

在 3.3 版更改: 当 `ndim = 0` 时值为空元组而不再为 `None`。

strides

一个整数元组，通过`ndim` 的长度给出以字节表示的大小，以便访问数组中每个维度上的每个元素。

在 3.3 版更改：当 `ndim = 0` 时值为空元组而不再为 `None`。

suboffsets

供 PIL 风格的数组内部使用。该值仅作为参考信息。

c_contiguous

一个表明内存是否为 C-*contiguous* 的布尔值。

3.3 新版功能。

f_contiguous

一个表明内存是否为 Fortran *contiguous* 的布尔值。

3.3 新版功能。

contiguous

一个表明内存是否为 *contiguous* 的布尔值。

3.3 新版功能。

4.9 集合类型—`set`, `frozenset`

`set` 对象是由具有唯一性的`hashable` 对象所组成的无序多项集。常见的用途包括成员检测、从序列中去除重复项以及数学中的集合类计算，例如交集、并集、差集与对称差集等等。（关于其他容器对象请参看`dict`, `list` 与 `tuple` 等内置类，以及`collections` 模块。）

与其他多项集一样，集合也支持 `x in set`, `len(set)` 和 `for x in set`。作为一种无序的多项集，集合并不记录元素位置或插入顺序。相应地，集合不支持索引、切片或其他序列类的操作。

目前有两种内置集合类型，`set` 和 `frozenset`。`set` 类型是可变的—其内容可以使用 `add()` 和 `remove()` 这样的方法来改变。由于是可变类型，它没有哈希值，且不能被用作字典的键或其他集合的元素。`frozenset` 类型是不可变并且为`hashable`—其内容在被创建后不能再改变；因此它可以被用作字典的键或其他集合的元素。

除了可以使用`set` 构造器，非空的 `set` (不是 `frozenset`) 还可以通过将以逗号分隔的元素列表包含于花括号之内来创建，例如：`{'jack', 'sjoerd'}`。

两个类的构造器具有相同的作用方式：

```
class set([iterable])
class frozenset([iterable])
```

返回一个新的 `set` 或 `frozenset` 对象，其元素来自于 `iterable`。集合的元素必须为`hashable`。要表示由集合对象构成的集合，所有的内层集合必须为`frozenset` 对象。如果未指定 `iterable`，则将返回一个新的空集合。

`set` 和 `frozenset` 的实例提供以下操作：

len(s)

返回集合 `s` 中的元素数量（即 `s` 的基数）。

x in s

检测 `x` 是否为 `s` 中的成员。

x not in s

检测 `x` 是否非 `s` 中的成员。

isdisjoint (*other*)

如果集合中没有与 *other* 共有的元素则返回 True。当且仅当两个集合的交集为空集合时，两者为不相交集。

issubset (*other*)

set <= **other**

检测是否集合中的每个元素都在 *other* 之中。

set < **other**

检测集合是否为 *other* 的真子集，即 `set <= other and set != other`。

issuperset (*other*)

set >= **other**

检测是否 *other* 中的每个元素都在集合之中。

set > **other**

检测集合是否为 *other* 的真超集，即 `set >= other and set != other`。

union (**others*)

set | **other** | ...

返回一个新集合，其中包含来自原集合以及 *others* 指定的所有集合中的元素。

intersection (**others*)

set & **other** & ...

返回一个新集合，其中包含原集合以及 *others* 指定的所有集合中共有的元素。

difference (**others*)

set - **other** - ...

返回一个新集合，其中包含原集合中在 *others* 指定的其他集合中不存在的元素。

symmetric_difference (*other*)

set ^ **other**

返回一个新集合，其中的元素或属于原集合或属于 *other* 指定的其他集合，但不能同时属于两者。

copy ()

Return a new set with a shallow copy of *s*.

请注意，非运算符版本的 `union()`、`intersection()`、`difference()`，以及 `symmetric_difference()`、`issubset()` 和 `issuperset()` 方法会接受任意可迭代对象作为参数。相比之下，它们所对应的运算符版本则要求其参数为集合。这就排除了容易出错的构造形式例如 `set('abc') & 'cbs'`，而推荐可读性更强的 `set('abc').intersection('cbs')`。

`set` 和 `frozenset` 均支持集合与集合的比较。两个集合当且仅当每个集合中的每个元素均包含于另一个集合之内（即各为对方的子集）时则相等。一个集合当且仅当其为另一个集合的真子集（即为后者的子集但两者不相等）时则小于另一个集合。一个集合当且仅当其为另一个集合的真超集（即为后者的超集但两者不相等）时则大于另一个集合。

`set` 的实例与 `frozenset` 的实例之间基于它们的成员进行比较。例如 `set('abc') == frozenset('abc')` 返回 True，`set('abc') in set([frozenset('abc')])` 也一样。

子集与相等比较并不能推广为完全排序函数。例如，任意两个非空且不相交的集合不相等且互不为对方的子集，因此以下所有比较均返回 False: `a < b`, `a == b`, or `a > b`。

由于集合仅定义了部分排序（子集关系），因此由集合构成的列表 `list.sort()` 方法的输出并无定义。集合的元素，与字典的键类似，必须为 *hashable*。

混合了 `set` 实例与 `frozenset` 的二进制位运算将返回与第一个操作数相同的类型。例如：`frozenset('ab') | set('bc')` 将返回 `frozenset` 的实例。

下表列出了可用于 `set` 而不能用于不可变的 `frozenset` 实例的操作：

update (**others*)

```

set |= other | ...
    更新集合，添加来自 others 中的所有元素。

intersection_update(*others)
set &= other & ...
    更新集合，只保留其中在所有 others 中也存在的元素。

difference_update(*others)
set -= other | ...
    更新集合，移除其中也存在于 others 中的元素。

symmetric_difference_update(other)
set ^= other
    更新集合，只保留存在于集合的一方而非共同存在的元素。

add(elem)
    将元素 elem 添加到集合中。

remove(elem)
    从集合中移除元素 elem。如果 elem 不存在于集合中则会引发 KeyError。

discard(elem)
    如果元素 elem 存在于集合中则将其移除。

pop()
    从集合中移除并返回任意一个元素。如果集合为空则会引发 KeyError。

clear()
    从集合中移除所有元素。

```

请注意，非运算符版本的 `update()`、`intersection_update()`、`difference_update()` 和 `symmetric_difference_update()` 方法将接受任意可迭代对象作为参数。

请注意，`__contains__()`、`remove()` 和 `discard()` 方法的 `elem` 参数可能是一个 `set`。为支持对一个等价的 `frozenset` 进行搜索，会根据 `elem` 临时创建一个该类型对象。

4.10 映射类型—dict

mapping 对象会将 *hashable* 值映射到任意对象。映射属于可变对象。目前仅有一种标准映射类型 字典。（关于其他容器对象请参看 *list*、*set* 与 *tuple* 等内置类，以及 *collections* 模块。）

字典的键 几乎可以是任何值。非 *hashable* 的值，即包含列表、字典或其他可变类型的值（此类对象基于值而非对象标识进行比较）不可用作键。数字类型用作键时遵循数字比较的一般规则：如果两个数值相等（例如 1 和 1.0）则两者可以被用来索引同一字典条目。（但是请注意，由于计算机对于浮点数存储的只是近似值，因此将其用作字典键是不明智的。）

字典可以通过将以逗号分隔的 键：值对列表包含于花括号之内来创建，例如：{'jack': 4098, 'sjoerd': 4127} 或 {4098: 'jack', 4127: 'sjoerd'}，也可以通过 *dict* 构造器来创建。

```

class dict(**kwarg)
class dict(mapping, **kwarg)
class dict(iterable, **kwarg)

```

返回一个新的字典，基于可选的位置参数和可能为空的关键字参数集来初始化。

如果没有给出位置参数，将创建一个空字典。如果给出一个位置参数并且其属于映射对象，将创建一个具有与映射对象相同键值对的字典。否则的话，位置参数必须为一个 *iterable* 对象。该可迭代对象中的每一项本身必须为一个刚好包含两个元素的可迭代对象。每一项中的第一个对象将成为新字典的一个键，第二个对象将成为其对应的值。如果一个键出现一次以上，该键的最后一个值将成为其在新字典中对应的值。

如果给出了关键字参数，则关键字参数及其值会被加入到基于位置参数创建的字典。如果要加入的键已存在，来自关键字参数的值将替代来自位置参数的值。

作为演示，以下示例返回的字典均等于 `{"one": 1, "two": 2, "three": 3}`:

```
>>> a = dict(one=1, two=2, three=3)
>>> b = {'one': 1, 'two': 2, 'three': 3}
>>> c = dict(zip(['one', 'two', 'three'], [1, 2, 3]))
>>> d = dict([('two', 2), ('one', 1), ('three', 3)])
>>> e = dict({'three': 3, 'one': 1, 'two': 2})
>>> a == b == c == d == e
True
```

像第一个例子那样提供关键字参数的方式只能使用有效的 Python 标识符作为键。其他方式则可使用任何有效的键。

这些是字典所支持的操作（因而自定义的映射类型也应当支持）：

len(d)

返回字典 *d* 中的项数。

d[key]

返回 *d* 中以 *key* 为键的项。如果映射中不存在 *key* 则会引发 *KeyError*。

如果字典的子类定义了方法 `__missing__()` 并且 *key* 不存在，则 *d[key]* 操作将调用该方法并附带键 *key* 作为参数。*d[key]* 随后将返回或引发 `__missing__(key)` 调用所返回或引发的任何对象或异常。没有其他操作或方法会发起调用 `__missing__()`。如果未定义 `__missing__()`，则会引发 *KeyError*。`__missing__()` 必须是一个方法；它不能是一个实例变量：

```
>>> class Counter(dict):
...     def __missing__(self, key):
...         return 0
>>> c = Counter()
>>> c['red']
0
>>> c['red'] += 1
>>> c['red']
1
```

上面的例子显示了 `collections.Counter` 实现的部分代码。还有另一个不同的 `__missing__` 方法是由 `collections.defaultdict` 所使用的。

d[key] = value

将 *d[key]* 设为 *value*。

del d[key]

将 *d[key]* 从 *d* 中移除。如果映射中不存在 *key* 则会引发 *KeyError*。

key in d

如果 *d* 中存在键 *key* 则返回 `True`，否则返回 `False`。

key not in d

等价于 `not key in d`。

iter(d)

返回以字典的键为元素的迭代器。这是 `iter(d.keys())` 的快捷方式。

clear()

移除字典中的所有元素。

copy()

返回原字典的浅拷贝。

classmethod fromkeys (*seq*[, *value*])

Create a new dictionary with keys from *seq* and values set to *value*.

fromkeys() 属于类方法，会返回一个新字典。*value* 默认为 `None`。

get (*key*[, *default*])

如果 *key* 存在于字典中则返回 *key* 的值，否则返回 *default*。如果 *default* 未给出则默认为 `None`，因而此方法绝不会引发 `KeyError`。

items ()

返回由字典项 ((键, 值) 对) 组成的一个新视图。参见视图对象文档。

keys ()

返回由字典键组成的一个新视图。参见视图对象文档。

pop (*key*[, *default*])

如果 *key* 存在于字典中则将其移除并返回其值，否则返回 *default*。如果 *default* 未给出且 *key* 不存在于字典中，则会引发 `KeyError`。

popitem ()

Remove and return an arbitrary (*key*, *value*) pair from the dictionary.

popitem() 适用于对字典进行消耗性的迭代，这在集合算法中经常被使用。如果字典为空，调用 *popitem()* 将引发 `KeyError`。

setdefault (*key*[, *default*])

如果字典存在键 *key*，返回它的值。如果不存在，插入值为 *default* 的键 *key*，并返回 *default*。*default* 默认为 `None`。

update ([*other*])

使用来自 *other* 的键/值对更新字典，覆盖原有的键。返回 `None`。

update() 接受另一个字典对象，或者一个包含键/值对（以长度为二的元组或其他可迭代对象表示）的可迭代对象。如果给出了关键字参数，则会以其所指定的键/值对更新字典：`d.update(red=1, blue=2)`。

values ()

返回由字典值组成的一个新视图。参见视图对象文档。

两个字典的比较当且仅当具有相同的（键，值）对时才会相等。顺序比较（‘<’，‘<=’，‘>=’，‘>’）会引发 `TypeError`。

参见：

`types.MappingProxyType` 可被用来创建一个 *dict* 的只读视图。

4.10.1 字典视图对象

由 `dict.keys()`、`dict.values()` 和 `dict.items()` 所返回的对象是视图对象。该对象提供字典条目的一个动态视图，这意味着当字典改变时，视图也会相应改变。

字典视图可以被迭代以产生与其对应的数据，并支持成员检测：

len(dictview)

返回字典中的条目数。

iter(dictview)

返回字典中的键、值或项（以（键，值）为元素的元组表示）的迭代器。

Keys and values are iterated over in an arbitrary order which is non-random, varies across Python implementations, and depends on the dictionary’s history of insertions and deletions. If keys, values and items views are iterated over with no intervening modifications to the dictionary, the order of items will directly correspond. This allows the

creation of (value, key) pairs using `zip()`: `pairs = zip(d.values(), d.keys())`. Another way to create the same list is `pairs = [(v, k) for (k, v) in d.items()]`.

在添加或删除字典中的条目期间对视图进行迭代可能引发 `RuntimeError` 或者无法完全迭代所有条目。

x in dictview

如果 `x` 是对应字典中存在的键、值或项（在最后一种情况下 `x` 应为一个（键，值）元组）则返回 `True`。

键视图类似于集合，因为其条目不重复且可哈希。如果所有值都是可哈希的，即（键，值）对也是不重复且可哈希的，那么条目视图也会类似于集合。（值视图则不被视为类似于集合，因其条目通常都是有重复的。）对于类似于集合的视图，为抽象基类 `collections.abc.Set` 所定义的全部操作都是有效的（例如 `==`, `<` 或 `^`）。

一个使用字典视图的示例：

```
>>> dishes = {'eggs': 2, 'sausage': 1, 'bacon': 1, 'spam': 500}
>>> keys = dishes.keys()
>>> values = dishes.values()

>>> # iteration
>>> n = 0
>>> for val in values:
...     n += val
>>> print(n)
504

>>> # keys and values are iterated over in the same order
>>> list(keys)
['eggs', 'bacon', 'sausage', 'spam']
>>> list(values)
[2, 1, 1, 500]

>>> # view objects are dynamic and reflect dict changes
>>> del dishes['eggs']
>>> del dishes['sausage']
>>> list(keys)
['spam', 'bacon']

>>> # set operations
>>> keys & {'eggs', 'bacon', 'salad'}
{'bacon'}
>>> keys ^ {'sausage', 'juice'}
{'juice', 'sausage', 'bacon', 'spam'}
```

4.11 上下文管理器类型

Python 的 `with` 语句支持通过上下文管理器所定义的运行时上下文这一概念。此对象的实现使用了一对专门方法，允许用户自定义类来定义运行时上下文，在语句体被执行前进入该上下文，并在语句执行完毕时退出该上下文：

`contextmanager.__enter__()`

Enter the runtime context and return either this object or another object related to the runtime context. The value returned by this method is bound to the identifier in the `as` clause of `with` statements using this context manager.

一个返回其自身的上下文管理器的例子是 `file object`。文件对象会从 `__enter__()` 返回其自身，以允许 `open()` 被用作 `with` 语句中的上下文表达式。

An example of a context manager that returns a related object is the one returned by `decimal.localcontext()`. These managers set the active decimal context to a copy of the original decimal context and then return the copy. This allows changes to be made to the current decimal context in the body of the `with` statement without affecting code outside the `with` statement.

`contextmanager.__exit__(exc_type, exc_val, exc_tb)`

退出运行时上下文并返回一个布尔值旗标来表明所发生的任何异常是否应当被屏蔽。如果在执行 `with` 语句的语句体期间发生了异常，则参数会包含异常的类型、值以及回溯信息。在其他情况下三个参数均为 `None`。

Returning a true value from this method will cause the `with` statement to suppress the exception and continue execution with the statement immediately following the `with` statement. Otherwise the exception continues propagating after this method has finished executing. Exceptions that occur during execution of this method will replace any exception that occurred in the body of the `with` statement.

传入的异常绝对不应当被显式地重新引发——相反地，此方法应当返回一个假值以表明方法已成功完成并且不希望屏蔽被引发的异常。这允许上下文管理代码方便地检测 `__exit__()` 方法是否确实已失败。

Python 定义了一些上下文管理器来支持简易的线程同步、文件或其他对象的快速关闭，以及更方便地操作活动的十进制算术上下文。除了实现上下文管理协议以外，不同类型不会被特殊处理。请参阅 `contextlib` 模块查看相关的示例。

Python 的 `generator` 和 `contextlib.contextmanager` 装饰器提供了实现这些协议的便捷方式。如果使用 `contextlib.contextmanager` 装饰器来装饰一个生成器函数，它将返回一个实现了必要的 `__enter__()` 和 `__exit__()` 方法的上下文管理器，而不再是由未经装饰的生成器函数所产生的迭代器。

请注意，Python/C API 中 Python 对象的类型结构中并没有针对这些方法的专门槽位。想要定义这些方法的扩展类型必须将它们作为普通的 Python 可访问方法来提供。与设置运行时上下文的开销相比，单个类字典查找的开销可以忽略不计。

4.12 其他内置类型

解释器支持一些其他种类的对象。这些对象大都仅支持一两种操作。

4.12.1 模块

模块唯一的特殊操作是属性访问: `m.name`，这里 `m` 为一个模块而 `name` 访问定义在 `m` 的符号表中的一个名称。模块属性可以被赋值。（请注意 `import` 语句严格来说也是对模块对象的一种操作；`import foo` 不求存在一个名为 `foo` 的模块对象，而是要求存在一个对于名为 `foo` 的模块的（永久性）定义。）

每个模块都有一个特殊属性 `__dict__`。这是包含模块的符号表的字典。修改此字典将实际改变模块的符号表，但是无法直接对 `__dict__` 赋值（你可以写 `m.__dict__['a'] = 1`，这会将 `m.a` 定义为 1，但是你不能写 `m.__dict__ = {}`）。不建议直接修改 `__dict__`。

内置于解释器中的模块会写成这样: `<module 'sys' (built-in)>`。如果是从一个文件加载，则会写成 `<module 'os' from '/usr/local/lib/pythonX.Y/os.pyc'>`。

4.12.2 类与类实例

关于这些类型请参阅 `objects` 和 `class`。

4.12.3 函数

函数对象是通过函数定义创建的。对函数对象的唯一操作是调用它: `func(argument-list)`。

实际上存在两种不同的函数对象: 内置函数和用户自定义函数。两者支持同样的操作 (调用函数), 但实现方式不同, 因此对象类型也不同。

更多信息请参阅 `function`。

4.12.4 方法

方法是使用属性表示法来调用的函数。存在两种形式: 内置方法 (例如列表的 `append()` 方法) 和类实例方法。内置方法由支持它们的类型来描述。

如果你通过一个实例来访问方法 (即定义在类命名空间内的函数), 你会得到一个特殊对象: 绑定方法 (或称实例方法) 对象。当被调用时, 它会将 `self` 参数添加到参数列表。绑定方法具有两个特殊的只读属性: `m.__self__` 操作该方法的对象, 而 `m.__func__` 是实现该方法的函数。调用 `m(arg-1, arg-2, ..., arg-n)` 完全等价于调用 `m.__func__(m.__self__, arg-1, arg-2, ..., arg-n)`。

与函数对象类似, 绑定方法对象也支持获取任意属性。但是, 由于方法属性实际上保存于下层的函数对象中 (`meth.__func__`), 因此不允许设置绑定方法的方法属性。尝试设置方法的属性将会导致引发 `AttributeError`。想要设置方法属性, 你必须在下层的函数对象中显式地对其进行设置:

```
>>> class C:
...     def method(self):
...         pass
...
>>> c = C()
>>> c.method.whoami = 'my name is method' # can't set on the method
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
AttributeError: 'method' object has no attribute 'whoami'
>>> c.method.__func__.whoami = 'my name is method'
>>> c.method.whoami
'my name is method'
```

更多信息请参阅 `types`。

4.12.5 代码对象

代码对象被具体实现用来表示“伪编译”的可执行 Python 代码, 例如一个函数体。它们不同于函数对象, 因为它们不包含对其全局执行环境的引用。代码对象由内置的 `compile()` 函数返回, 并可通过从函数对象的 `__code__` 属性从中提取。另请参阅 `code` 模块。

可以通过将代码对象 (而非源码字符串) 传给 `exec()` 或 `eval()` 内置函数来执行或求值。

更多信息请参阅 `types`。

4.12.6 类型对象

类型对象表示各种对象类型。对象的类型可通过内置函数 `type()` 来获取。类型没有特殊的操作。标准库模块 `types` 定义了所有标准内置类型的名称。

类型以这样的写法来表示: `<class 'int'>`。

4.12.7 空对象

此对象会由不显式地返回值的函数所返回。它不支持任何特殊的操作。空对象只有一种值 `None` (这是个内置名称)。`type(None)()` 会生成同一个单例。

该对象的写法为 `None`。

4.12.8 省略符对象

此对象常被用于切片 (参见 `slicings`)。它不支持任何特殊的操作。省略符对象只有一种值 `Ellipsis` (这是个内置名称)。`type(Ellipsis)()` 会生成 `Ellipsis` 单例。

该对象的写法为 `Ellipsis` 或 `...`。

4.12.9 未实现对象

此对象会被作为比较和二元运算被应用于它们所不支持的类型时的返回值。请参阅 `comparisons` 了解更多信息。未实现对象只有一种值 `NotImplemented`。`type(NotImplemented)()` 会生成这个单例。

该对象的写法为 `NotImplemented`。

4.12.10 布尔值

布尔值是两个常量对象 `False` 和 `True`。它们被用来表示逻辑上的真假 (不过其他值也可被当作真值或假值)。在数字类的上下文中 (例如被用作算术运算符的参数时), 它们的行为分别类似于整数 0 和 1。内置函数 `bool()` 可被用来将任意值转换为布尔值, 只要该值可被解析为一个逻辑值 (参见之前的 [逻辑值检测](#) 部分)。

该对象的写法分别为 `False` 和 `True`。

4.12.11 内部对象

有关此对象的信息请参阅 `types`。其中描述了栈帧对象、回溯对象以及切片对象等等。

4.13 特殊属性

语言实现为部分对象类型添加了一些特殊的只读属性, 它们具有各自的作用。其中一些并不会被 `dir()` 内置函数所列出。

`object.__dict__`

一个字典或其他类型的映射对象, 用于存储对象的 (可写) 属性。

`instance.__class__`

类实例所属的类。

`class.__bases__`

由类对象的基类所组成的元组。

`definition.__name__`

类、函数、方法、描述器或生成器实例的名称。

definition. **__qualname__**

类、函数、方法、描述器或生成器实例的`qualified name`。

3.3 新版功能.

class. **__mro__**

此属性是由类组成的元组，在方法解析期间会基于它来查找基类。

class. **mro()**

此方法可被一个元类来重载，以为其实例定制方法解析顺序。它会在类实例化时被调用，其结果存储于 `__mro__` 之中。

class. **__subclasses__()**

每个类会保存由对其直接子类的弱引用组成的列表。此方法将返回一个由仍然存在的所有此类引用组成的列表。例如：

```
>>> int.__subclasses__()
[<class 'bool'>]
```

内置异常

在 Python 中，所有异常必须为一个派生自 `BaseException` 的类的实例。在带有提及一个特定类的 `except` 子句的 `try` 语句中，该子句也会处理任何派生自该类的异常类（但不处理它所派生出的异常类）。通过子类化创建的两个不相关异常类永远是不等效的，即使它们具有相同的名称。

下面列出的内置异常可通过解释器或内置函数来生成。除非另有说明，它们都会具有一个提示导致错误详细原因的“关联值”。这可以是一个字符串或由多个信息项（例如一个错误码和一个解释错误的字符串）组成的元组。关联值通常会作为参数被传递给异常类的构造器。

用户代码可以引发内置异常。这可被用于测试异常处理程序或报告错误条件，“就像”在解释器引发了相同异常的情况时一样；但是请注意，没有任何机制能防止用户代码引发不适当的错误。

内置异常类可以被子类化以定义新的异常；鼓励程序员从 `Exception` 类或它的某个子类而不是从 `BaseException` 来派生新的异常。关于定义异常的更多信息可以在 Python 教程的 `tut-userexceptions` 部分查看。

当在 `except` 或 `finally` 子句中引发（或重新引发）异常时，`__context__` 会被自动设为所捕获的最后一个异常；如果新的异常未被处理，则最终显示的回溯信息将包括原始的异常和最后的异常。

当引发一个新的异常（而不是简单地使用 `raise` 来重新引发当前在处理的异常）时，隐式的异常上下文可以通过使用带有 `raise` 的 `from` 来补充一个显式的原因：

```
raise new_exc from original_exc
```

跟在 `from` 之后的表达式必须为一个异常或 `None`。它将在所引发的异常上被设置为 `__cause__`。设置 `__cause__` 还会隐式地将 `__suppress_context__` 属性设为 `True`，这样使用 `raise new_exc from None` 可以有效地将旧异常替换为新异常来显示其目的（例如将 `KeyError` 转换为 `AttributeError`），同时让旧异常在 `__context__` 中保持可用状态以便在调试时进行内省。

除了异常本身的回溯以外，默认的回溯还会显示这些串连的异常。`__cause__` 中的显式串连异常如果存在将总是显示。`__context__` 中的隐式串连异常仅在 `__cause__` 为 `None` 并且 `__suppress_context__` 为假值时显示。

不论在哪种情况下，异常本身总会在任何串连异常之后显示，以便回溯的最后一行总是显示所引发的最后一个异常。

5.1 基类

下列异常主要被用作其他异常的基类。

exception `BaseException`

所有内置异常的基类。它不应该被用户自定义类直接继承 (这种情况请使用 `Exception`)。如果在此类的实例上调用 `str()`，则会返回实例的参数表示，或者当没有参数时返回空字符串。

`args`

传给异常构造器的参数元组。某些内置异常 (例如 `OSError`) 接受特定数量的参数并赋予此元组中的元素特殊的含义，而其他异常通常只接受一个给出错误信息的单独字符串。

`with_traceback(tb)`

此方法将 `tb` 设为异常的新回溯信息并返回该异常对象。它通常以如下的形式在异常处理程序中使用：

```
try:
    ...
except SomeException:
    tb = sys.exc_info()[2]
    raise OtherException(...).with_traceback(tb)
```

exception `Exception`

所有内置的非系统退出类异常都派生自此类。所有用户自定义异常也应当派生自此类。

exception `ArithmeticError`

此基类用于派生针对各种算术类错误而引发的内置异常：`OverflowError`、`ZeroDivisionError`、`FloatingPointError`。

exception `BufferError`

当与缓冲区相关的操作无法执行时将被引发。

exception `LookupError`

此基类用于派生当映射或序列所使用的键或索引无效时引发的异常：`IndexError`、`KeyError`。这可以通过 `codecs.lookup()` 来直接引发。

5.2 具体异常

以下异常属于经常被引发的异常。

exception `AssertionError`

当 `assert` 语句失败时将被引发。

exception `AttributeError`

当属性引用 (参见 `attribute-references`) 或赋值失败时将被引发。(当一个对象根本不支持属性引用或属性赋值时则将引发 `TypeError`。)

exception `EOFError`

当 `input()` 函数未读取任何数据即达到文件结束条件 (EOF) 时将被引发。(另外, `io.IOBase.read()` 和 `io.IOBase.readline()` 方法在遇到 EOF 则将返回一个空字符串。)

exception `FloatingPointError`

Raised when a floating point operation fails. This exception is always defined, but can only be raised when Python is configured with the `--with-fpectl` option, or the `WANT_SIGFPE_HANDLER` symbol is defined in the `pyconfig.h` file.

exception GeneratorExit

当一个 *generator* 或 *coroutine* 被关闭时将被引发；参见 `generator.close()` 和 `coroutine.close()`。它直接继承自 *BaseException* 而不是 *Exception*，因为从技术上来说它并不是一个错误。

exception ImportError

当 `import` 语句尝试加载模块遇到麻烦时将被引发。并且当 `from ... import` 中的“from list”存在无法找到的名称时也会被引发。

`name` 与 `path` 属性可通过对构造器使用仅关键字参数来设定。设定后它们将分别表示被尝试导入的模块名称与触发异常的任意文件所在路径。

在 3.3 版更改：添加了 `name` 与 `path` 属性。

exception ModuleNotFoundError

ImportError 的子类，当一个模块无法被定位时将由 `import` 引发。当在 `sys.modules` 中找到 `None` 时也会被引发。

3.6 新版功能。

exception IndexError

当序列抽取超出范围时将被引发。（切片索引会被静默截短到允许的范围；如果指定索引不是整数则 *TypeError* 会被引发。）

exception KeyError

当在现有键集中找不到指定的映射（字典）键时将被引发。

exception KeyboardInterrupt

当用户按下中断键（通常为 `Control-C` 或 `Delete`）时将被引发。在执行期间，会定期检测中断信号。该异常继承自 *BaseException* 以确保不会被处理 *Exception* 的代码意外捕获，这样可以避免退出解释器。

exception MemoryError

当一个操作耗尽内存但情况仍可（通过删除一些对象）进行挽救时将被引发。关联的值是一个字符串，指明是哪种（内部）操作耗尽了内存。请注意由于底层的内存管理架构（C 的 `malloc()` 函数），解释器也许并不总是能够从这种情况下完全恢复；但它毕竟可以引发一个异常，这样就能打印出栈回溯信息，以便找出导致问题的失控程序。

exception NameError

当某个局部或全局名称未找到时将被引发。此异常仅用于非限定名称。关联的值是一条错误信息，其中包含未找到的名称。

exception NotImplementedError

此异常派生自 *RuntimeError*。在用户自定义的基类中，抽象方法应当在其要求所派生类重载该方法，或是在其要求所开发的类提示具体实现尚待添加时引发此异常。

注解： 它不应当用来表示一个运算符或方法根本不能被支持 – 在此情况下应当让特定运算符 / 方法保持未定义，或者在子类中将其设为 *None*。

注解： *NotImplementedError* 和 *NotImplemented* 不可互换，即使它们有相似的名称和用途。请参阅 *NotImplemented* 了解有关何时使用它们的详细说明。

exception OSError ([arg])**exception OSError (errno, strerror[, filename[, winerror[, filename2]]])**

此异常在一个系统函数返回系统相关的错误时将被引发，此类错误包括 I/O 操作失败例如“文件未找到”或“磁盘已满”等（不包括非法参数类型或其他偶然性错误）。

构造器的第二种形式可设置如下所述的相应属性。如果未指定这些属性则默认为`None`。为了能向下兼容，如果传入了三个参数，则`args`属性将仅包含由前两个构造器参数组成的2元组。

构造器实际返回的往往是`OSError`的某个子类，如下文`OS exceptions`中所描述的。具体的子类取决于最终的`errno`值。此行为仅在直接或通过别名来构造`OSError`时发生，并且在子类化时不会被继承。

errno

来自于C变量`errno`的数字错误码。

winerror

在Windows下，此参数将给出原生的Windows错误码。而`errno`属性将是该原生错误码在POSIX平台下的近似转换形式。

在Windows下，如果`winerror`构造器参数是一个整数，则`errno`属性会根据Windows错误码来确定，而`errno`参数会被忽略。在其他平台上，`winerror`参数会被忽略，并且`winerror`属性将不存在。

strerror

操作系统所提供的相应错误信息。它在POSIX平台中由C函数`perror()`来格式化，在Windows中则是由`FormatMessage()`。

filename

filename2

对于与文件系统路径有关(例如`open()`或`os.unlink()`)的异常，`filename`是传给函数的文件名。对于涉及两个文件系统路径的函数(例如`os.rename()`)，`filename2`将是传给函数的第二个文件名。

在3.3版更改：`EnvironmentError`、`IOError`、`WindowsError`、`socket.error`、`select.error`与`mmap.error`已被合并到`OSError`，构造器可能返回其中一个子类。

在3.4版更改：`filename`属性现在将是传给函数的原始文件名，而不是经过编码或基于文件系统编码进行解码之后的名称。此外还添加了`filename2`构造器参数和属性。

exception OverflowError

当算术运算的结果大到无法表示时将被引发。这对整数来说不可能发生(宁可引发`MemoryError`也不会放弃尝试)。但是出于历史原因，有时也会在整数超出要求范围的情况下引发`OverflowError`。因为在C中缺少对浮点异常处理的标准化，大多数浮点运算都不会做检查。

exception RecursionError

此异常派生自`RuntimeError`。它会在解释器检测发现超过最大递归深度(参见`sys.getrecursionlimit()`)时被引发。

3.5新版功能: 在此之前将只引发`RuntimeError`。

exception ReferenceError

此异常将在使用`weakref.proxy()`函数所创建的弱引用来访问该引用的某个已被作为垃圾回收的属性时被引发。有关弱引用的更多信息请参阅`weakref`模块。

exception RuntimeError

当检测到一个不归属于任何其他类别的错误时将被引发。关联的值是一个指明究竟发生了什么问题的字符串。

exception StopIteration

由内置函数`next()`和`iterator`的`__next__()`方法所引发，用来表示该迭代器不能产生下一项。

该异常对象只有一个属性`value`，它在构造该异常时作为参数给出，默认值为`None`。

当一个`generator`或`coroutine`函数返回时，将引发一个新的`StopIteration`实例，函数返回的值将被用作异常构造器的`value`形参。

If a generator function defined in the presence of a `from __future__ import generator_stop` directive raises `StopIteration`, it will be converted into a `RuntimeError` (retaining the `StopIteration` as

the new exception' s cause).

在 3.3 版更改: 添加了 `value` 属性及其被生成器函数用作返回值的功能。

在 3.5 版更改: Introduced the `RuntimeError` transformation.

exception `StopAsyncIteration`

必须由一个 *asynchronous iterator* 对象的 `__anext__()` 方法来引发以停止迭代操作。

3.5 新版功能。

exception `SyntaxError`

当解析器遇到语法错误时将被引发。这可以发生在 `import` 语句, 对内置函数 `exec()` 或 `eval()` 的调用, 或者读取原始脚本或标准输入 (也包括交互模式) 的时候。

该类的实例包含有属性 `filename`, `lineno`, `offset` 和 `text` 用于方便地访问相应的详细信息。异常实例的 `str()` 仅返回消息文本。

exception `IndentationError`

与不正确的缩进相关的语法错误的基类。这是 `SyntaxError` 的一个子类。

exception `TabError`

当缩进包含对制表符和空格符不一致的使用时将被引发。这是 `IndentationError` 的一个子类。

exception `SystemError`

当解释器发现内部错误, 但情况看起来尚未严重到要放弃所有希望时将被引发。关联的值是一个指明发生了什么的字符串 (表示为低层级的符号)。

你应当将此问题报告给你所用 Python 解释器的作者或维护人员。请确认报告 Python 解释器的版本号 (`sys.version`; 它也会在交互式 Python 会话开始时被打印出来), 具体的错误消息 (异常所关联的值) 以及可能触发该错误的程序源码。

exception `SystemExit`

此异常由 `sys.exit()` 函数引发。它继承自 `BaseException` 而不是 `Exception` 以确保不会被处理 `Exception` 的代码意外捕获。这允许此异常正确地向上传播并导致解释器退出。如果它未被处理, 则 Python 解释器就将退出; 不会打印任何栈回溯信息。构造器接受的可选参数与传递给 `sys.exit()` 的相同。如果该值为一个整数, 则它指明系统退出状态码 (会传递给 C 的 `exit()` 函数); 如果该值为 `None`, 则退出状态码为零; 如果该值为其他类型 (例如字符串), 则会打印对象的值并将退出状态码设为一。

对 `sys.exit()` 的调用会被转换为一个异常以便能执行清理处理程序 (`try` 语句的 `finally` 子句), 并且使得调试器可以执行一段脚本而不必冒失去控制的风险。如果绝对确实地需要立即退出 (例如在调用 `os.fork()` 之后的子进程中) 则可使用 `os._exit()`。

code

传给构造器的退出状态码或错误信息 (默认为 `None`。)

exception `TypeError`

当一个操作或函数被应用于类型不适当的对象时将被引发。关联的值是一个字符串, 给出有关类型不匹配的详情。

此异常可以由用户代码引发, 以表明尝试对某个对象进行的操作不受支持也不应当受支持。如果某个对象应当支持给定的操作但尚未提供相应的实现, 所要引发的适当异常应为 `NotImplementedError`。

传入参数的类型错误 (例如在要求 `int` 时却传入了 `list`) 应当导致 `TypeError`, 但传入参数的值错误 (例如传入要求范围之外的数值) 则应当导致 `ValueError`。

exception `UnboundLocalError`

当在函数或方法中对某个局部变量进行引用, 但该变量并未绑定任何值时将被引发。此异常是 `NameError` 的一个子类。

exception `UnicodeError`

当发生与 Unicode 相关的编码或解码错误时将被引发。此异常是 `ValueError` 的一个子类。

UnicodeError 具有一些描述编码或解码错误的属性。例如 `err.object[err.start:err.end]` 会给出导致编解码器失败的特定无效输入。

encoding

引发错误的编码名称。

reason

描述特定编解码器错误的字符串。

object

编解码器试图要编码或解码的对象。

start

object 中无效数据的开始位置索引。

end

object 中无效数据的末尾位置索引（不含）。

exception UnicodeEncodeError

当在编码过程中发生与 Unicode 相关的错误时将被引发。此异常是 *UnicodeError* 的一个子类。

exception UnicodeDecodeError

当在解码过程中发生与 Unicode 相关的错误时将被引发。此异常是 *UnicodeError* 的一个子类。

exception UnicodeTranslateError

在转写过程中发生与 Unicode 相关的错误时将被引发。此异常是 *UnicodeError* 的一个子类。

exception ValueError

当操作或函数接收到具有正确类型但值不适合的参数，并且情况不能用更精确的异常例如 *IndexError* 来描述时将被引发。

exception ZeroDivisionError

当除法或取余运算的第二个参数为零时将被引发。关联的值是一个字符串，指明操作数和运算的类型。

下列异常被保留以与之前的版本相兼容；从 Python 3.3 开始，它们都是 *OSError* 的别名。

exception EnvironmentError**exception IOError****exception WindowsError**

限在 Windows 中可用。

5.2.1 OS 异常

下列异常均为 *OSError* 的子类，它们将根据系统错误代码被引发。

exception BlockingIOError

当一个操作会被某个设置为非阻塞操作的对象（例如套接字）所阻塞时将被引发。对应于 `errno` `EAGAIN`, `EALREADY`, `EWOULDBLOCK` 和 `EINPROGRESS`。

除了 *OSError* 已有的属性，*BlockingIOError* 还有一个额外属性：

characters_written

一个整数，表示在被阻塞前已写入到流的字符数。当使用来自 *io* 模块的带缓冲 I/O 类时此属性可用。

exception ChildProcessError

当一个子进程上的操作失败时将被引发。对应于 `errno` `ECHILD`。

exception ConnectionError

与连接相关问题的基类。

其子类有 *BrokenPipeError*, *ConnectionAbortedError*, *ConnectionRefusedError* 和 *ConnectionResetError*。

exception BrokenPipeError

ConnectionError 的子类, 当试图写入另一端已被关闭的管道, 或是试图写入已关闭写入的套接字时将被引发。对应于 `errno EPIPE` 和 `ESHUTDOWN`。

exception ConnectionAbortedError

ConnectionError 的子类, 当连接尝试被对端中止时将被引发。对应于 `errno ECONNABORTED`。

exception ConnectionRefusedError

ConnectionError 的子类, 当连接尝试被对端拒绝时将被引发。对应于 `errno ECONNREFUSED`。

exception ConnectionResetError

ConnectionError 的子类, 当连接被对端重置时将被引发。对应于 `errno ECONNRESET`。

exception FileExistsError

当试图创建一个已存在的文件或目录时将被引发。对应于 `errno EEXIST`。

exception FileNotFoundError

当所请求的文件或目录不存在时将被引发。对应于 `errno ENOENT`。

exception InterruptedError

当系统调用被输入信号中断时将被引发。对应于 `errno EINTR`。

在 3.5 版更改: 当系统调用被某个信号中断时, Python 现在会重试系统调用, 除非该信号的处理程序引发了其它异常 (原理参见 [PEP 475](#)) 而不是引发 *InterruptedError*。

exception IsADirectoryError

当请求对一个目录执行文件操作 (例如 `os.remove()`) 将被引发。对应于 `errno EISDIR`。

exception NotADirectoryError

当请求对一个非目录对象执行目录操作 (例如 `os.listdir()`) 时将被引发。对应于 `errno ENOTDIR`。

exception PermissionError

当在具有足够操作权限的情况下试图执行某个操作时将被引发——例如缺少文件系统权限。对应于 `errno EACCES` 和 `EPERM`。

exception ProcessLookupError

当给定的进程不存在时将被引发。对应于 `errno ESRCH`。

exception TimeoutError

当一个系统函数发生系统级超时的情况下将被引发。对应于 `errno ETIMEDOUT`。

3.3 新版功能: 添加了以上所有 *OSError* 的子类。

参见:

[PEP 3151](#) - 重写 OS 和 IO 异常的层次结构

5.3 警告

The following exceptions are used as warning categories; see the *warnings* module for more information.

exception Warning

警告类别的基类。

exception UserWarning

用户代码所产生警告的基类。

exception DeprecationWarning

Base class for warnings about deprecated features.

exception PendingDeprecationWarning

Base class for warnings about features which will be deprecated in the future.

exception SyntaxWarning

与模糊的语法相关的警告的基类。

exception RuntimeWarning

与模糊的运行时行为相关的警告的基类。

exception FutureWarning

Base class for warnings about constructs that will change semantically in the future.

exception ImportError

与在模块导入中可能的错误相关的警告的基类。

exception UnicodeWarning

与 Unicode 相关的警告的基类。

exception BytesWarning

与 `bytes` 和 `bytearray` 相关的警告的基类。

exception ResourceWarning

Base class for warnings related to resource usage.

3.2 新版功能.

5.4 异常层次结构

内置异常的分类层级结构如下：

```
BaseException
+-- SystemExit
+-- KeyboardInterrupt
+-- GeneratorExit
+-- Exception
    +-- StopIteration
    +-- StopAsyncIteration
    +-- ArithmeticError
    |   +-- FloatingPointError
    |   +-- OverflowError
    |   +-- ZeroDivisionError
    +-- AssertionError
    +-- AttributeError
    +-- BufferError
    +-- EOFError
    +-- ImportError
    |   +-- ModuleNotFoundError
    +-- LookupError
    |   +-- IndexError
    |   +-- KeyError
    +-- MemoryError
    +-- NameError
    |   +-- UnboundLocalError
    +-- OSError
    |   +-- BlockingIOError
    |   +-- ChildProcessError
    |   +-- ConnectionError
    |   |   +-- BrokenPipeError
```

(下页继续)

(续上页)

```
| | +-- ConnectionAbortedError
| | +-- ConnectionRefusedError
| | +-- ConnectionResetError
| +-- FileExistsError
| +-- FileNotFoundError
| +-- InterruptedError
| +-- IsADirectoryError
| +-- NotADirectoryError
| +-- PermissionError
| +-- ProcessLookupError
| +-- TimeoutError
+-- ReferenceError
+-- RuntimeError
| +-- NotImplementedError
| +-- RecursionError
+-- SyntaxError
| +-- IndentationError
| +-- TabError
+-- SystemError
+-- TypeError
+-- ValueError
| +-- UnicodeError
| +-- UnicodeDecodeError
| +-- UnicodeEncodeError
| +-- UnicodeTranslateError
+-- Warning
| +-- DeprecationWarning
| +-- PendingDeprecationWarning
| +-- RuntimeWarning
| +-- SyntaxWarning
| +-- UserWarning
| +-- FutureWarning
| +-- ImportWarning
| +-- UnicodeWarning
| +-- BytesWarning
| +-- ResourceWarning
```


本章介绍的模块提供了广泛的字符串操作和其他文本处理服务。

在二进制数据服务之下描述的 *codecs* 模块也与文本处理高度相关。此外也请参阅 Python 内置字符串类型的文档文本序列类型—*str*。

6.1 string 一常见的字符串操作

源代码： `Lib/string.py`

参见：

文本序列类型—*str*

字符串的方法

6.1.1 字符串常量

此模块中定义的常量为：

`string.ascii_letters`

下文所述 *ascii_lowercase* 和 *ascii_uppercase* 常量的拼连。该值不依赖于语言区域。

`string.ascii_lowercase`

小写字母 `'abcdefghijklmnopqrstuvwxyz'`。该值不依赖于语言区域，不会发生改变。

`string.ascii_uppercase`

大写字母 `'ABCDEFGHIJKLMNOPQRSTUVWXYZ'`。该值不依赖于语言区域，不会发生改变。

`string.digits`

字符串 `'0123456789'`。

`string.hexdigits`

字符串 '0123456789abcdefABCDEF'。

`string.octdigits`

字符串 '01234567'。

`string.punctuation`

由在 C 语言区域中被视为标点符号的 ASCII 字符组成的字符串。

`string.printable`

由被视为可打印符号的 ASCII 字符组成的字符串。这是 *digits*, *ascii_letters*, *punctuation* 和 *whitespace* 的总和。

`string.whitespace`

由被视为空白符号的 ASCII 字符组成的字符串。其中包括空格、制表、换行、回车、进纸和纵向制表符。

6.1.2 自定义字符串格式化

内置的字符串类提供了通过使用 **PEP 3101** 所描述的 *format()* 方法进行复杂变量替换和值格式化的能力。*string* 模块中的 *Formatter* 类允许你使用与内置 *format()* 方法相同的实现来创建并定制你自己的字符串格式化行为。

class `string.Formatter`

Formatter 类包含下列公有方法：

format (*format_string*, **args*, ***kwargs*)

首要的 API 方法。它接受一个格式字符串和任意一组位置和关键字参数。它只是一个调用 *vformat()* 的包装器。

3.5 版后已移除: Passing a format string as keyword argument *format_string* has been deprecated.

vformat (*format_string*, *args*, *kwargs*)

此函数执行实际的格式化操作。它被公开为一个单独的函数，用于需要传入一个预定义字母作为参数，而不是使用 **args* 和 ***kwargs* 语法将字典解包为多个单独参数并重打包的情况。*vformat()* 完成将格式字符串分解为字符数据和替换字段的工作。它会调用下文所述的几种不同方法。

此外，*Formatter* 还定义了一些旨在被子类替换的方法：

parse (*format_string*)

循环遍历 *format_string* 并返回一个由可迭代对象组成的元组 (*literal_text*, *field_name*, *format_spec*, *conversion*)。它会被 *vformat()* 用来将字符串分解为文本字面值或替换字段。

元组中的值在概念上表示一段字面文本加上一个替换字段。如果没有字面文本（如果连续出现两个替换字段就会发生这种情况），则 *literal_text* 将是一个长度为零的字符串。如果没有替换字段，则 *field_name*, *format_spec* 和 *conversion* 的值将为 *None*。

get_field (*field_name*, *args*, *kwargs*)

给定 *field_name* 作为 *parse()* (见上文) 的返回值，将其转换为要格式化的对象。返回一个元组 (*obj*, *used_key*)。默认版本接受在 **PEP 3101** 所定义形式的字符串，例如 “0[name]” 或 “label.title”。*args* 和 *kwargs* 与传给 *vformat()* 的一样。返回值 *used_key* 与 *get_value()* 的 *key* 形参具有相同的含义。

get_value (*key*, *args*, *kwargs*)

提取给定的字段值。*key* 参数将为整数或字符串。如果是整数，它表示 *args* 中位置参数的索引；如果是字符串，它表示 *kwargs* 中的关键字参数名。

args 形参会被设为 *vformat()* 的位置参数列表，而 *kwargs* 形参会被设为由关键字参数组成的字典。

对于复合字段名称，仅会为字段名称的第一个组件调用这些函数；后续组件会通过普通属性和索引操作来进行处理。

因此举例来说，字段表达式 `'0.name'` 将导致调用 `get_value()` 时附带 `key` 参数值 `0`。在 `get_value()` 通过调用内置的 `getattr()` 函数返回后将会查找 `name` 属性。

如果索引或关键字引用了一个不存在的项，则将引发 `IndexError` 或 `KeyError`。

check_unused_args (*used_args, args, kwargs*)

在必要时实现对未使用参数进行检测。此函数的参数是是格式字符串中实际引用的所有参数键的集合（整数表示位置参数，字符串表示名称参数），以及被传给 `vformat` 的 `args` 和 `kwargs` 的引用。未使用参数的集合可以根据这些形参计算出来。如果检测失败则 `check_unused_args()` 应会引发一个异常。

format_field (*value, format_spec*)

`format_field()` 会简单地调用内置全局函数 `format()`。提供该方法是为了让子类能够重载它。

convert_field (*value, conversion*)

使用给定的转换类型（来自 `parse()` 方法所返回的元组）来转换（由 `get_field()` 所返回的）值。默认版本支持 `'s'` (`str`)，`'r'` (`repr`) 和 `'a'` (`ascii`) 等转换类型。

6.1.3 格式字符串语法

`str.format()` 方法和 `Formatter` 类共享相同的格式字符串语法（虽然对于 `Formatter` 来说，其子类可以定义它们自己的格式字符串语法）。具体语法与 格式化字符串字面值相似，但也存在区别。

格式字符串包含有以花括号 `{}` 括起来的“替换字段”。不在花括号之内的内容被视为字面文本，会不加修改地复制到输出中。如果你需要在字面文本中包含花括号字符，可以通过重复来转义：`{{ and }}`。

替换字段的语法如下：

```
replacement_field ::= "{" [field_name] ["!" conversion] [":" format_spec] "}"
field_name         ::= arg_name ("." attribute_name | "[" element_index "]")*
arg_name           ::= [identifier | digit+]
attribute_name     ::= identifier
element_index      ::= digit+ | index_string
index_string       ::= <any source character except "]"> +
conversion         ::= "r" | "s" | "a"
format_spec        ::= <described in the next section>
```

用不太正式的术语来描述，替换字段开头可以用一个 *field_name* 指定要对值进行格式化并取代替换字符被插入到输出结果的对象。*field_name* 之后有可选的 *conversion* 字段，它是一个感叹号 `!` 加一个 *format_spec*，并以一个冒号 `:` 打头。这些指明了替换值的非默认格式。

另请参阅[格式规格迷你语言](#)一节。

field_name 本身以一个数字或关键字 *arg_name* 打头。如果为数字，则它指向一个位置参数，而如果为关键字，则它指向一个命名关键字参数。如果格式字符串中的数字 *arg_names* 为 `0, 1, 2, …` 的序列，它们可以全部省略（而非部分省略），数字 `0, 1, 2, …` 将会按顺序自动插入。由于 *arg_name* 不使用引号分隔，因此无法在格式字符串中指定任意的字典键（例如字符串 `'10'` 或 `'[:-]'`）。*arg_name* 之后可以带上任意数量的索引或属性表达式。`'.name'` 形式的表达式会使用 `getattr()` 选择命名属性，而 `'[index]'` 形式的表达式会使用 `__getitem__()` 执行索引查找。

在 3.1 版更改：位置参数说明符对于 `str.format()` 可以省略，因此 `'{} {}'.format(a, b)` 等价于 `'{0} {1}'.format(a, b)`。

在 3.4 版更改: 位置参数说明符对于 *Formatter* 可以省略。

一些简单的格式字符串示例

```
"First, thou shalt count to {0}" # References first positional argument
"Bring me a {}"                 # Implicitly references the first positional_
↪ argument
"From {} to {}"                 # Same as "From {0} to {1}"
"My quest is {name}"            # References keyword argument 'name'
"Weight in tons {0.weight}"     # 'weight' attribute of first positional arg
"Units destroyed: {players[0]}" # First element of keyword argument 'players'.
```

使用 *conversion* 字段在格式化之前进行类型强制转换。通常, 格式化值的工作由值本身的 `__format__()` 方法来完成。但是, 在某些情况下最好强制将类型格式化为一个字符串, 覆盖其本身的格式化定义。通过在调用 `__format__()` 之前将值转换为字符串, 可以绕过正常的格式化逻辑。

目前支持的转换旗标有三种: `'!s'` 会对值调用 `str()`, `'!r'` 调用 `repr()` 而 `'!a'` 则调用 `ascii()`。

几个例子:

```
"Harold's a clever {0!s}"       # Calls str() on the argument first
"Bring out the holy {name!r}"   # Calls repr() on the argument first
"More {!a}"                    # Calls ascii() on the argument first
```

format_spec 字段包含值应如何呈现的规格描述, 例如字段宽度、对齐、填充、小数精度等细节信息。每种值类型可以定义自己的“格式化迷你语言”或对 *format_spec* 的解读方式。

大多数内置类型都支持同样的格式化迷你语言, 具体描述见下一节。

format_spec 字段还可以在其内部包含嵌套的替换字段。这些嵌套的替换字段可能包括字段名称、转换旗标和格式规格描述, 但是不再允许更深层的嵌套。*format_spec* 内部的替换字段会在解读 *format_spec* 字符串之前先被解读。这将允许动态地指定特定值的格式。

请参阅格式示例一节查看相关示例。

格式规格迷你语言

“格式规格”在格式字符串所包含的替换字段内部使用, 用于定义单个值应如何呈现 (参见格式字符串语法和 f-strings)。它们也可以被直接传给内置的 *format()* 函数。每种可格式化的类型都可以自行定义如何对格式规格进行解读。

大多数内置类型都为格式规格实现了下列选项, 不过某些格式化选项只被数值类型所支持。

一般约定空格式字符串 ("") 将产生与对值调用 *str()* 相同的结果。非空格式字符串通常会修改这一结果。

标准格式说明符的一般形式如下:

```
format_spec ::= [[fill]align][sign][#][0][width][grouping_option][.precision][type]
fill        ::= <any character>
align       ::= "<" | ">" | "=" | "^"
sign        ::= "+" | "-" | " "
width       ::= digit+
grouping_option ::= "_" | ","
precision  ::= digit+
type        ::= "b" | "c" | "d" | "e" | "E" | "f" | "F" | "g" | "G" | "n" | "o" | "s"
```

如果指定了一个有效的 *align* 值, 则可以在该值前面加一个 *fill* 字符, 它可以为任意字符, 如果省略则默认为空格符。在 格式化字符串字面值或使用 *str.format()* 方法时是无法使用花括号字面值 (“{” or “}”)

作为 *fill* 字符的。但是，通过嵌套替换字段插入花括号则是可以的。这个限制不会影响 `format()` 函数。

各种对齐选项的含义如下：

选项	意义
'<'	强制字段在可用空间内左对齐（这是大多数对象的默认值）。
'>'	强制字段在可用空间内右对齐（这是数字的默认值）。
'= '	强制将填充放置在符号（如果有）之后但在数字之前。这用于以 “+000000120” 形式打印字段。此对齐选项仅对数字类型有效。当 ‘0’ 紧接在字段宽度之前时，它成为默认值。
'^ '	强制字段在可用空间内居中。

请注意，除非定义了最小字段宽度，否则字段宽度将始终与填充它的数据大小相同，因此在这种情况下，对齐选项没有意义。

sign 选项仅对数字类型有效，可以是以下之一：

选项	意义
'+'	表示标志应该用于正数和负数。
'- '	表示标志应仅用于负数（这是默认行为）。
space	表示应在正数上使用前导空格，在负数上使用减号。

'#' 选项可以让“替代形式”被用于转换。替代形式可针对不同类型分别定义。此选项仅对整数、浮点、复数和 `Decimal` 类型有效。对于整数类型，当使用二进制、八进制或十六进制输出时，此选项会为输出值添加相应的 '0b', '0o' 或 '0x' 前缀。对于浮点数、复数和 `Decimal` 类型，替代形式会使得转换结果总是包含小数点符号，即使其不带小数。通常只有在带有小数的情况下，此类转换的结果中才会出现小数点符号。此外，对于 'g' 和 'G' 转换，末尾的零不会从结果中被移除。

',' 选项表示使用逗号作为千位分隔符。对于感应区域设置的分隔符，请改用 'n' 整数表示类型。

在 3.1 版更改：添加了 ',' 选项（另请参阅 [PEP 378](#)）。

'_' 选项表示对浮点表示类型和整数表示类型 'd' 使用下划线作为千位分隔符。对于整数表示类型 'b', 'o', 'x' 和 'X'，将为每 4 个数位插入一个下划线。对于其他表示类型指定此选项则将导致错误。

在 3.6 版更改：添加了 '_' 选项（另请参阅 [PEP 515](#)）。

width 是一个定义最小字段宽度的十进制整数。如果未指定，则字段宽度将由内容确定。

当未显式给出对齐方式时，在 *width* 字段前加一个零 ('0') 字段将为数字类型启用感知正负号的零填充。这相当于设置 *fill* 字符为 '0' 且 *alignment* 类型为 '='。

precision 是一个十进制数字，表示对于以 'f' and 'F' 格式化的浮点数值要在小数点后显示多少个数位，或者对于以 'g' 或 'G' 格式化的浮点数值要在小数点前后共显示多少个数位。对于非数字类型，该字段表示最大字段大小——换句话说就是要使用多少个来自字段内容的字符。对于整数值则不允许使用 *precision*。

最后，*type* 确定了数据应如何呈现。

可用的字符串表示类型是：

类型	意义
's'	字符串格式。这是字符串的默认类型，可以省略。
None	和 's' 一样。

可用的整数表示类型是：

类型	意义
'b'	二进制格式。输出以 2 为基数的数字。
'c'	字符。在打印之前将整数转换为相应的 <code>unicode</code> 字符。
'd'	十进制整数。输出以 10 为基数的数字。
'o'	八进制格式。输出以 8 为基数的数字。
'x'	十六进制格式。输出以 16 为基数的数字，使用小写字母表示 9 以上的数码。
'X'	十六进制格式。输出以 16 为基数的数字，使用大写字母表示 9 以上的数码。
'n'	数字。这与 'd' 相似，不同之处在于它会使用当前区域设置来插入适当的数字分隔字符。
None	和 'd' 相同。

在上述的表示类型之外，整数还可以通过下列的浮点表示类型来格式化（除了 'n' 和 None）。当这样做时，会在格式化之前使用 `float()` 将整数转换为浮点数。

浮点数和小数值可用的表示类型有：

类型	意义
'e'	指数表示。以使用字母 'e' 来标示指数的科学计数法打印数字。默认的精度为 6。
'E'	指数表示。与 'e' 相似，不同之处在于它使用大写字母 'E' 作为分隔字符。
'f'	定点表示。将数字显示为一个定点数。默认的精确度为 6。
'F'	定点表示。与 'f' 相似，但会将 <code>nan</code> 转为 <code>NAN</code> 并将 <code>inf</code> 转为 <code>INF</code> 。
'g'	常规格式。对于给定的精度 <code>p >= 1</code> ，这会将数值舍入到 <code>p</code> 位有效数字，再将结果以定点格式或科学计数法进行格式化，具体取决于其值的大小。 准确的规则如下：假设使用表示类型 'e' 和精度 <code>p-1</code> 进行格式化的结果具有指数值 <code>exp</code> 。则如果 $-4 \leq \text{exp} < p$ ，该数字将使用表示类型 'f' 和精度 <code>p-1-exp</code> 进行格式化。否则的话，该数字将使用表示类型 'e' 和精度 <code>p-1</code> 进行格式化。在两种情况下，都会从有效数字中移除无意义的末尾零，如果小数点之后没有数字则小数点也会被移除。 正负无穷，正负零和 <code>nan</code> 会分别被格式化为 <code>inf</code> , <code>-inf</code> , <code>0</code> , <code>-0</code> 和 <code>nan</code> ，无论精度如何设定。 精度 0 会被视为等同于精度 1。默认精度为 6。
'G'	常规格式。类似于 'g'，不同之处在于当数值非常大时会切换为 'E'。无穷与 <code>NaN</code> 也会表示为大写形式。
'n'	数字。这与 'g' 相似，不同之处在于它会使用当前区域设置来插入适当的数字分隔字符。
'%'	百分比。将数字乘以 100 并显示为定点 ('f') 格式，后面带一个百分号。
None	类似于 'g'，不同之处在于当使用定点表示法时，小数点后将至少显示一位。默认精度与表示给定值所需的精度一样。整体效果为与其他格式修饰符所调整的 <code>str()</code> 输出保持一致。

格式示例

本节包含 `str.format()` 语法的示例以及与旧式 `%` 格式化的比较。

该语法在大多数情况下与旧式的 `%` 格式化类似，只是增加了 `{}` 和 `:` 来取代 `%`。例如，`'%03.2f'` 可以被改写为 `'{:03.2f}'`。

新的格式语法还支持新增的不同选项，将在以下示例中说明。

按位置访问参数：

```
>>> '{0}, {1}, {2}'.format('a', 'b', 'c')
'a, b, c'
>>> '{}', {}, {}'.format('a', 'b', 'c')  # 3.1+ only
'a, b, c'
>>> '{2}, {1}, {0}'.format('a', 'b', 'c')
'c, b, a'
>>> '{2}, {1}, {0}'.format(*'abc')        # unpacking argument sequence
'c, b, a'
>>> '{0}{1}{0}'.format('abra', 'cad')    # arguments' indices can be repeated
'abracadabra'
```

按名称访问参数:

```
>>> 'Coordinates: {latitude}, {longitude}'.format(latitude='37.24N', longitude='-115.81W')
'Coordinates: 37.24N, -115.81W'
>>> coord = {'latitude': '37.24N', 'longitude': '-115.81W'}
>>> 'Coordinates: {latitude}, {longitude}'.format(**coord)
'Coordinates: 37.24N, -115.81W'
```

访问参数的属性:

```
>>> c = 3-5j
>>> ('The complex number {0} is formed from the real part {0.real} '
... 'and the imaginary part {0.imag}.').format(c)
'The complex number (3-5j) is formed from the real part 3.0 and the imaginary part -5.0.'
>>> class Point:
...     def __init__(self, x, y):
...         self.x, self.y = x, y
...     def __str__(self):
...         return 'Point({self.x}, {self.y})'.format(self=self)
...
>>> str(Point(4, 2))
'Point(4, 2)'
```

访问参数的项:

```
>>> coord = (3, 5)
>>> 'X: {0[0]}; Y: {0[1]}'.format(coord)
'X: 3; Y: 5'
```

替代 %s 和 %r:

```
>>> "repr() shows quotes: {!r}; str() doesn't: {!s}".format('test1', 'test2')
'repr() shows quotes: \'test1\'; str() doesn\'t: test2'
```

对齐文本以及指定宽度:

```
>>> '{:<30}'.format('left aligned')
'left aligned'
>>> '{:>30}'.format('right aligned')
'right aligned'
>>> '{:^30}'.format('centered')
'centered'
>>> '{:*^30}'.format('centered')  # use '*' as a fill char
'*****centered*****'
```

替代 `%+f`, `%-f` 和 `% f` 以及指定正负号:

```
>>> '{:+f}; {:+f}'.format(3.14, -3.14) # show it always
'+3.140000; -3.140000'
>>> '{: f}; {: f}'.format(3.14, -3.14) # show a space for positive numbers
' 3.140000; -3.140000'
>>> '{:-f}; {:-f}'.format(3.14, -3.14) # show only the minus -- same as '{:f}; {:f}'
'3.140000; -3.140000'
```

替代 `%x` 和 `%o` 以及转换基于不同进位制的值:

```
>>> # format also supports binary numbers
>>> "int: {0:d}; hex: {0:x}; oct: {0:o}; bin: {0:b}".format(42)
'int: 42; hex: 2a; oct: 52; bin: 101010'
>>> # with 0x, 0o, or 0b as prefix:
>>> "int: {0:d}; hex: {0:#x}; oct: {0:#o}; bin: {0:#b}".format(42)
'int: 42; hex: 0x2a; oct: 0o52; bin: 0b101010'
```

使用逗号作为千位分隔符:

```
>>> '{:,}'.format(1234567890)
'1,234,567,890'
```

表示为百分数:

```
>>> points = 19
>>> total = 22
>>> 'Correct answers: {:.2%}'.format(points/total)
'Correct answers: 86.36%'
```

使用特定类型的专属格式化:

```
>>> import datetime
>>> d = datetime.datetime(2010, 7, 4, 12, 15, 58)
>>> '{:%Y-%m-%d %H:%M:%S}'.format(d)
'2010-07-04 12:15:58'
```

嵌套参数以及更复杂的示例:

```
>>> for align, text in zip('<^>', ['left', 'center', 'right']):
...     '{0:{fill}{align}16}'.format(text, fill=align, align=align)
...
'left<<<<<<<<<<'
'^^^^^center^^^^^'
'>>>>>>>>>>right'
>>>
>>> octets = [192, 168, 0, 1]
>>> '{:02X}{:02X}{:02X}{:02X}'.format(*octets)
'C0A80001'
>>> int(_, 16)
3232235521
>>>
>>> width = 5
>>> for num in range(5,12):
...     for base in 'dXob':
...         print('{0:{width}{base}}'.format(num, base=base, width=width), end=' ')
...     print()
... 
```

(下页继续)

(续上页)

5	5	5	101
6	6	6	110
7	7	7	111
8	8	10	1000
9	9	11	1001
10	A	12	1010
11	B	13	1011

6.1.4 模板字符串

模板字符串提供了由 [PEP 292](#) 所描述的更简便的字符串替换方式。模板字符串的一个主要用例是文本国际化 (i18n)，因为在此场景下，更简单的语法和功能使得文本翻译过程比使用 Python 的其他内置字符串格式化工具更为方便。作为基于模板字符串构建以实现 i18n 的库的一个示例，请参看 [flufl.i18n](#) 包。

模板字符串支持基于 `$` 的替换，使用以下规则：

- `$$` 为转义符号；它会被替换为单个的 `$`。
- `$identifier` 为替换占位符，它会匹配一个名为 "identifier" 的映射键。在默认情况下，"identifier" 限制为任意 ASCII 字母数字（包括下划线）组成的字符串，不区分大小写，以下划线或 ASCII 字母开头。在 `$` 字符之后的第一个非标识符字符将表明占位符的终结。
- `${identifier}` 等价于 `$identifier`。当占位符之后紧跟着有效的但又不是占位符一部分的标识符字符时需要使用，例如 `"${noun}ification"`。

在字符串的其他位置出现 `$` 将导致引发 `ValueError`。

`string` 模块提供了实现这些规则的 `Template` 类。`Template` 有下列方法：

class `string.Template(template)`

该构造器接受一个参数作为模板字符串。

substitute(mapping, **kws)

执行模板替换，返回一个新字符串。`mapping` 为任意字典类对象，其中的键将匹配模板中的占位符。或者你也可以提供一组关键字参数，其中的关键字即对应占位符。当同时给出 `mapping` 和 `kws` 并且存在重复时，则以 `kws` 中的占位符为优先。

safe_substitute(mapping, **kws)

类似于 `substitute()`，不同之处是如果有占位符未在 `mapping` 和 `kws` 中找到，不是引发 `KeyError` 异常，而是将原始占位符不加修改地显示在结果字符串中。另一个与 `substitute()` 的差异是任何在其他情况下出现的 `$` 将简单地返回 `$` 而不是引发 `ValueError`。

此方法被认为“安全”，因为虽然仍有可能发生其他异常，但它总是尝试返回可用的字符串而不是引发一个异常。从另一方面来说，`safe_substitute()` 也可能根本算不上安全，因为它将静默地忽略错误格式的模板，例如包含多余的分隔符、不成对的花括号或不是合法 Python 标识符的占位符等等。

`Template` 的实例还提供一个公有数据属性：

template

这是作为构造器的 `template` 参数被传入的对象。一般来说，你不应该修改它，但并不强制要求只读访问。

以下是一个如何使用模版的示例：

```
>>> from string import Template
>>> s = Template('$who likes $what')
>>> s.substitute(who='tim', what='kung pao')
```

(下页继续)

(续上页)

```
'tim likes kung pao'
>>> d = dict(who='tim')
>>> Template('Give $who $100').substitute(d)
Traceback (most recent call last):
...
ValueError: Invalid placeholder in string: line 1, col 11
>>> Template('$who likes $what').substitute(d)
Traceback (most recent call last):
...
KeyError: 'what'
>>> Template('$who likes $what').safe_substitute(d)
'tim likes $what'
```

进阶用法：你可以派生 *Template* 的子类来自定义占位符语法、分隔符，或用于解析模板字符串的整个正则表达式。为此目的，你可以重载这些类属性：

- *delimiter* – This is the literal string describing a placeholder introducing delimiter. The default value is `$`. Note that this should *not* be a regular expression, as the implementation will call `re.escape()` on this string as needed.
- *idpattern* – This is the regular expression describing the pattern for non-braced placeholders (the braces will be added automatically as appropriate). The default value is the regular expression `(? -i: [_a-zA-Z] [_a-zA-Z0-9]*)`.

注解： Since default *flags* is `re.IGNORECASE`, pattern `[a-z]` can match with some non-ASCII characters. That's why we use local `-i` flag here.

While *flags* is kept to `re.IGNORECASE` for backward compatibility, you can override it to `0` or `re.IGNORECASE | re.ASCII` when subclassing.

- *flags* – 将在编译用于识别替换内容的正则表达式被应用的正则表达式旗标。默认值为 `re.IGNORECASE`。请注意 `re.VERBOSE` 总是会被加为旗标，因此自定义的 *idpattern* 必须遵循详细正则表达式的约定。

3.2 新版功能.

作为另一种选项，你可以通过重载类属性 *pattern* 来提供整个正则表达式模式。如果你这样做，该值必须为一个具有四个命名捕获组的正则表达式对象。这些捕获组对应于上面已经给出的规则，以及无效占位符的规则：

- *escaped* – 这个组匹配转义序列，在默认模式中即 `$$`。
- *named* – 这个组匹配不带花括号的占位符名称；它不应当包含捕获组中的分隔符。
- *braced* – 这个组匹配带有花括号的占位符名称；它不应当包含捕获组中的分隔符或者花括号。
- *invalid* – 这个组匹配任何其他分隔符模式（通常为单个分隔符），并且它应当出现在正则表达式的末尾。

6.1.5 辅助函数

`string.capwords(s, sep=None)`

使用 `str.split()` 将参数拆分为单词，使用 `str.capitalize()` 将单词转为大写形式，使用 `str.join()` 将大写的单词进行拼接。如果可选的第二个参数 *sep* 被省略或为 `None`，则连续的空白字符会被替换为单个空格符并且开头和末尾的空白字符会被移除，否则 *sep* 会被用来拆分和拼接单词。

6.2 re — 正则表达式操作

源代码: [Lib/re.py](#)

这个模块提供了与 Perl 语言类似的正则表达式匹配操作。

模式和被搜索的字符串既可以是 Unicode 字符串 (*str*)，也可以是 8 位字节串 (*bytes*)。但是，Unicode 字符串与 8 位字节串不能混用：也就是说，你不能用一个字节串模式去匹配 Unicode 字符串，反之亦然；类似地，当进行替换操作时，替换字符串的类型也必须与所用的模式和搜索字符串的类型一致。

正则表达式使用反斜杠 (`'\'`) 来表示特殊形式，或者把特殊字符转义成普通字符。而反斜杠在普通的 Python 字符串里也有相同的作用，所以就产生了冲突。比如说，要匹配一个字面上的反斜杠，正则表达式模式不得不写成 `'\\\\'`，因为正则表达式里匹配一个反斜杠必须是 `\\`，而每个反斜杠在普通的 Python 字符串里都要写成 `\\`。

解决办法是对于正则表达式样式使用 Python 的原始字符串表示法；在带有 `'r'` 前缀的字符串字面值中，反斜杠不必做任何特殊处理。因此 `r"\n"` 表示包含 `'\'` 和 `'n'` 两个字符的字符串，而 `"\n"` 则表示只包含一个换行符的字符串。样式在 Python 代码中通常都会使用这种原始字符串表示法来表示。

绝大部分正则表达式操作都提供为模块函数和方法，在[编译正则表达式](#)。这些函数是一个捷径，不需要先编译一个正则对象，但是损失了一些优化参数。

参见：

第三方模块 [regex](#)，提供了与标准库 `re` 模块兼容的 API 接口，同时还提供了额外的功能和更全面的 Unicode 支持。

6.2.1 正则表达式语法

一个正则表达式（或 RE）指定了一集与之匹配的字符串；模块内的函数可以让你检查某个字符串是否跟给定的正则表达式匹配（或者一个正则表达式是否匹配到一个字符串，这两种说法含义相同）。

正则表达式可以拼接；如果 *A* 和 *B* 都是正则表达式，那么 *AB* 也是正则表达式。通常，如果字符串 *p* 匹配 *A* 并且另一个字符串 *q* 匹配 *B*，那么 *pq* 可以匹配 *AB*。除非 *A* 或者 *B* 包含低优先级操作，*A* 和 *B* 存在边界条件；或者命名组引用。所以，复杂表达式可以很容易的从这里描述的简单源语表达式构建。了解更多正则表达式理论和实现，参考 the Friedl book [Frie09]，或者其他编译器构建的书籍。

以下是正则表达式格式的简要说明。更详细的信息和演示，参考 [regex-howto](#)。

正则表达式可以包含普通或者特殊字符。绝大部分普通字符，比如 `'A'`，`'a'`，或者 `'0'`，都是最简单的正则表达式。它们就匹配自身。你可以拼接普通字符，所以 `last` 匹配字符串 `'last'`。（在这一节的其他部分，我们将用 `this special style` 这种方式表示正则表达式，通常不带引号，要匹配的字符串用 `'in single quotes'`，单引号形式。）

有些字符，比如 `'|'` 或者 `'('`，属于特殊字符。特殊字符既可以表示它的普通含义，也可以影响它旁边的正则表达式的解释。

重复修饰符 (`*`，`+`，`?`，`{m,n}`，等) 不能直接嵌套。这样避免了非贪婪后缀？修饰符，和其他实现中的修饰符产生的多义性。要应用一个内层重复嵌套，可以使用括号。比如，表达式 `(?:a{6})*` 匹配 6 个 `'a'` 字符重复任意次数。

特殊字符是：

- (点) 在默认模式，匹配除了换行的任意字符。如果指定了标签 `DOTALL`，它将匹配包括换行符的任意字符。
- ^ (插入符号) 匹配字符串的开头，并且在 `MULTILINE` 模式也匹配换行后的首个符号。

\$ 匹配字符串尾或者换行符的前一个字符，在 *MULTILINE* 模式匹配换行符的前一个字符。`foo` 匹配 `'foo'` 和 `'foobar'`，但正则 `foo$` 只匹配 `'foo'`。更有趣的是，在 `'foo1\nfoo2\n'` 搜索 `foo.$`，通常匹配 `'foo2'`，但在 *MULTILINE* 模式，可以匹配到 `'foo1'`；在 `'foo\n'` 搜索 `$` 会找到两个空串：一个在换行前，一个在字符串最后。

***** 对它前面的正则式匹配 0 到任意次重复，尽量多的匹配字符串。`ab*` 会匹配 `'a'`，`'ab'`，或者 `'a'`` 后面跟随任意个 ``'b'`。

+ 对它前面的正则式匹配 1 到任意次重复。`ab+` 会匹配 `'a'` 后面跟随 1 个以上到任意个 `'b'`，它不会匹配 `'a'`。

? 对它前面的正则式匹配 0 到 1 次重复。`ab?` 会匹配 `'a'` 或者 `'ab'`。

***, +, ?, ??** `'*'`，`'+'`，和 `'?'` 修饰符都是贪婪的；它们在字符串进行尽可能多的匹配。有时候并不需要这种行为。如果正则式 `<.*>` 希望找到 `'<a> b <c>'`，它将会匹配整个字符串，而不仅是 `'<a>'`。在修饰符之后添加 `?` 将使样式以非贪婪方式或者 *dfn*:‘最小方式进行匹配；尽量少的字符将会被匹配。使用正则式 `<.*?>` 将会仅仅匹配 `'<a>'`。

{m}” 对其之前的正则式指定匹配 *m* 个重复；少于 *m* 的话就会导致匹配失败。比如，`a{6}` 将匹配 6 个 `'a'`，但是不能是 5 个。

{m, n}” 对正则式进行 *m* 到 *n* 次匹配，在 *m* 和 *n* 之间取尽量多。比如，`a{3, 5}` 将匹配 3 到 5 个 `'a'`。忽略 *m* 意为指定下界为 0，忽略 *n* 指定上界为无限次。比如 `a{4, }b` 将匹配 `'aaaab'` 或者 1000 个 `'a'` 尾随一个 `'b'`，但不能匹配 `'aaab'`。逗号不能省略，否则无法辨别修饰符应该忽略哪个边界。

{m, n}? 前一个修饰符的非贪婪模式，只匹配尽量少的字符次数。比如，对于 `'aaaaaa'`，`a{3, 5}` 匹配 5 个 `'a'`，而 `a{3, 5}?` 只匹配 3 个 `'a'`。

**** 转义特殊字符（允许你匹配 `'*'`，`'?'`，或者此类其他），或者表示一个特殊序列；特殊序列之后进行讨论。

如果你没有使用原始字符串 (`r'raw'`) 来表达样式，要牢记 Python 也使用反斜杠作为转义序列；如果转义序列不被 Python 的分析器识别，反斜杠和字符才能出现在字符串中。如果 Python 可以识别这个序列，那么反斜杠就应该重复两次。这将导致理解障碍，所以高度推荐，就算是最简单的表达式，也要使用原始字符串。

[] 用于表示一个字符集合。在一个集合中：

- 字符可以单独列出，比如 `[amk]` 匹配 `'a'`，`'m'`，或者 `'k'`。
- 可以表示字符范围，通过用 `'-'` 将两个字符连起来。比如 `[a-z]` 将匹配任何小写 ASCII 字符，`[0-5][0-9]` 将匹配从 00 到 59 的两位数字，`[0-9A-Fa-f]` 将匹配任何十六进制数位。如果 - 进行了转义（比如 `[a\-z]`）或者它的位置在首位或者末尾（如 `[-a]` 或 `[a-]`），它就只表示普通字符 `'-'`。
- 特殊字符在集合中，失去它的特殊含义。比如 `[(+*)]` 只会匹配这几个文法字符 `'('`，`'+'`，`'*'`，或 `')'`。
- 字符类如 `\w` 或者 `\s` (如下定义) 在集合内可以接受，它们可以匹配的字符由 *ASCII* 或者 *LOCALE* 模式决定。
- 不在集合范围内的字符可以通过取反来进行匹配。如果集合首字符是 `'^'`，所有不在集合内的字符将会被匹配，比如 `[^5]` 将匹配所有字符，除了 `'5'`，`[^~]` 将匹配所有字符，除了 `'^'`。^ 如果不在集合首位，就没有特殊含义。
- 在集合内要匹配一个字符 `']'`，有两种方法，要么就在它之前加上反斜杠，要么就把它放到集合首位。比如，`[(\)]{\}` 和 `[()]{}` 都可以匹配括号。

| A|B, *A* 和 *B* 可以是任意正则表达式，创建一个正则表达式，匹配 *A* 或者 *B*。任意个正则表达式可以用 `'|'` 连接。它也可以在组合（见下列）内使用。扫描目标字符串时，`'|'` 分隔开的正则样式从左到右进行匹配。当一个样式完全匹配时，这个分支就被接受。意思就是，一旦 *A* 匹配成功，*B* 就不再进行匹配，即便它能产生一个更好的匹配。或者说，`'|'` 操作符绝不贪婪。如果要匹配 `'|'` 字符，使用 `\|`，或者把它包含在字符集里，比如 `[|]`。

(...) (组合), 匹配括号内的任意正则表达式, 并标识出组合的开始和结尾。匹配完成后, 组合的内容可以被获取, 并可以在之后用 `\number` 转义序列进行再次匹配, 之后进行详细说明。要匹配字符 '(' 或者 ')', 用 `\(` 或 `\)`, 或者把它们包含在字符集合里: `[(), []]`。

(?...) 这是个扩展标记法 (一个 '?' 跟随 '(' 并无含义)。`'?'` 后面的第一个字符决定了这个构建采用什么样的语法。这种扩展通常并不创建新的组合; (`?P<name>...`) 是唯一的例外。以下是目前支持的扩展。

(**?aiLmsux**) ('a', 'i', 'L', 'm', 's', 'u', 'x' 中的一个或多个) 这个组合匹配一个空字符串; 这些字符对正则表达式设置以下标记 `re.A` (只匹配 ASCII 字符), `re.I` (忽略大小写), `re.L` (语言依赖), `re.M` (多行模式), `re.S` (点 `dot` 匹配全部字符), `re.U` (Unicode 匹配), and `re.X` (冗长模式)。(这些标记在 [模块内容](#) 中描述) 如果你想将这些标记包含在正则表达式中, 这个方法就很有用, 免去了在 `re.compile()` 中传递 `flag` 参数。标记应该在表达式字符串首位表示。

(?:...) 正则括号的非捕获版本。匹配在括号内的任何正则表达式, 但该分组所匹配的子字符串 不能在执行匹配后被获取或是之后在模式中被引用。

(**?imsx-imsx:...**) (Zero or more letters from the set 'i', 'm', 's', 'x', optionally followed by '-' followed by one or more letters from the same set.) The letters set or removes the corresponding flags: `re.I` (ignore case), `re.M` (multi-line), `re.S` (dot matches all), and `re.X` (verbose), for the part of the expression. (The flags are described in [模块内容](#).)

3.6 新版功能.

(**?P<name>...**) (命名组合) 类似正则组合, 但是匹配到的子串组在外部是通过定义的 `name` 来获取的。组合名必须是有效的 Python 标识符, 并且每个组合名只能用一个正则表达式定义, 只能定义一次。一个符号组合同样是一个数字组合, 就像这个组合没有被命名一样。

命名组合可以在三种上下文中引用。如果样式是 (`?P<quote>['"]`).`*?(?P=quote)` (也就是说, 匹配单引号或者双引号括起来的字符串):

引用组合 “quote” 的上下文	引用方法
在正则式自身内	<ul style="list-style-type: none"> (<code>?P=quote</code>) (如示) <code>\1</code>
处理匹配对象 <code>m</code>	<ul style="list-style-type: none"> <code>m.group('quote')</code> <code>m.end('quote')</code> (等)
传递到 <code>re.sub()</code> 里的 <code>repl</code> 参数中	<ul style="list-style-type: none"> <code>\g<quote></code> <code>\g<1></code> <code>\1</code>

(**?P=name**) 反向引用一个命名组合; 它匹配前面那个叫 `name` 的命名组中匹配到的串同样的字串。

(**?#...**) 注释; 里面的内容会被忽略。

(**?=...**) 匹配 ... 的内容, 但是并不消费样式的内容。这个叫做 *lookahead assertion*。比如, `Isaac (?=Asimov)` 匹配 'Isaac ' 只有在后面是 'Asimov' 的时候。

(**?!...**) 匹配 ... 不符合的情况。这个叫 *negative lookahead assertion* (前视取反)。比如说, `Isaac (?!Asimov)` 只有后面 不是 'Asimov' 的时候才匹配 'Isaac '。

(**?<=...**) 匹配字符串的当前位置, 它的前面匹配 ... 的内容到当前位置。这叫 *dfn:positive lookbehind assertion* (正向后视断定)。(`?<=abc`) `def` 会在 'abcdef' 中找到一个匹配, 因为后视会往后看 3 个字符并检查是否包含匹配的样式。包含的匹配样式必须是定长的, 意思就是 `abc` 或 `a|b` 是允许的, 但是 `a*` 和

`a{3,4}` 不可以。注意以 `positive lookbehind assertions` 开始的样式，如 `(?<=abc)def`，并不是从 `a` 开始搜索，而是从 `d` 往回看的。你可能更加愿意使用 `search()` 函数，而不是 `match()` 函数：

```
>>> import re
>>> m = re.search('(?!<=abc)def', 'abcdef')
>>> m.group(0)
'def'
```

这个例子搜索一个跟随在连字符后的单词：

```
>>> m = re.search(r'(?<=)\w+', 'spam-egg')
>>> m.group(0)
'egg'
```

在 3.5 版更改：添加定长组合引用的支持。

(?!...) 匹配当前位置之前不是 ... 的样式。这个叫 *dfn:negative lookbehind assertion*（后视断定取非）。类似正向后视断定，包含的样式匹配必须是定长的。由 *negative lookbehind assertion* 开始的样式可以从字符串搜索开始的位置进行匹配。

(?(id/name)yes-pattern|no-pattern) 如果给定的 *id* 或 *name* 存在，将会尝试匹配 *yes-pattern*，否则就尝试匹配 *no-pattern*，*no-pattern* 可选，也可以被忽略。比如，`(<)?(\w+@\w+(?:\.\w+)+)(?(1)>|$)` 是一个 email 样式匹配，将匹配 `<user@host.com>` 或 `'user@host.com'`，但不会匹配 `'<user@host.com'`，也不会匹配 `'user@host.com>'`。

由 `'\'` 和一个字符组成的特殊序列在以下列出。如果普通字符不是 ASCII 数位或者 ASCII 字母，那么正则样式将匹配第二个字符。比如，`\$` 匹配字符 `'$'`。

\number 匹配数字代表的组合。每个括号是一个组合，组合从 1 开始编号。比如 `(.+) \1` 匹配 `'the the'` 或者 `'55 55'`，但不会匹配 `'thethe'`（注意组合后面的空格）。这个特殊序列只能用于匹配前面 99 个组合。如果 *number* 的第一个数位是 0，或者 *number* 是三个八进制数，它将不会被看作是一个组合，而是八进制的数字值。在 `'['` 和 `']'` 字符集合内，任何数字转义都被看作是字符。

\A 只匹配字符串开始。

\b 匹配空字符串，但只在单词开始或结尾的位置。一个单词被定义为一个单词字符的序列。注意，通常 `\b` 定义为 `\w` 和 `\W` 字符之间，或者 `\w` 和字符串开始/结尾的边界，意思就是 `r'\bfoo\b'` 匹配 `'foo'`，`'foo.'`，`'(foo)'`，`'bar foo baz'` 但不匹配 `'foobar'` 或者 `'foo3'`。

默认情况下，Unicode 字母和数字是在 Unicode 样式中使用的，但是可以用 *ASCII* 标记来更改。如果 *LOCALE* 标记被设置的话，词的边界是由当前语言区域设置决定的，`\b` 表示退格字符，以便与 Python 字符串文本兼容。

\B 匹配空字符串，但不能在词的开头或者结尾。意思就是 `r'py\B'` 匹配 `'python'`，`'py3'`，`'py2'`，但不匹配 `'py'`，`'py.'`，或者 `'py!'`。`\B` 是 `\b` 的取非，所以 Unicode 样式的词语是由 Unicode 字母，数字或下划线构成的，虽然可以用 *ASCII* 标志来改变。如果使用了 *LOCALE* 标志，则词的边界由当前语言区域设置。

\d

对于 Unicode (str) 样式： Matches any Unicode decimal digit (that is, any character in Unicode character category `[Nd]`). This includes `[0-9]`, and also many other digit characters. If the *ASCII* flag is used only `[0-9]` is matched (but the flag affects the entire regular expression, so in such cases using an explicit `[0-9]` may be a better choice).

对于 8 位 (bytes) 样式： 匹配任何十进制数，就是 `[0-9]`。

\D Matches any character which is not a decimal digit. This is the opposite of `\d`. If the *ASCII* flag is used this becomes the equivalent of `[^0-9]` (but the flag affects the entire regular expression, so in such cases using an explicit `[^0-9]` may be a better choice).

\s

对于 Unicode (str) 样式： Matches Unicode whitespace characters (which includes [\t\n\r\f\v], and also many other characters, for example the non-breaking spaces mandated by typography rules in many languages). If the *ASCII* flag is used, only [\t\n\r\f\v] is matched (but the flag affects the entire regular expression, so in such cases using an explicit [\t\n\r\f\v] may be a better choice).

对于 8 位 (bytes) 样式： 匹配 ASCII 中的空白字符，就是 [\t\n\r\f\v] 。

\S Matches any character which is not a whitespace character. This is the opposite of \s. If the *ASCII* flag is used this becomes the equivalent of [^\t\n\r\f\v] (but the flag affects the entire regular expression, so in such cases using an explicit [^\t\n\r\f\v] may be a better choice).

\w

对于 Unicode (str) 样式： Matches Unicode word characters; this includes most characters that can be part of a word in any language, as well as numbers and the underscore. If the *ASCII* flag is used, only [a-zA-Z0-9_] is matched (but the flag affects the entire regular expression, so in such cases using an explicit [a-zA-Z0-9_] may be a better choice).

对于 8 位 (bytes) 样式： 匹配 ASCII 字符中的数字和字母和下划线，就是 [a-zA-Z0-9_]。如果设置了 *LOCALE* 标记，就匹配当前语言区域的数字和字母和下划线。

\W Matches any character which is not a word character. This is the opposite of \w. If the *ASCII* flag is used this becomes the equivalent of [^\a-zA-Z0-9_] (but the flag affects the entire regular expression, so in such cases using an explicit [^\a-zA-Z0-9_] may be a better choice). If the *LOCALE* flag is used, matches characters considered alphanumeric in the current locale and the underscore.

\Z 只匹配字符串尾。

绝大部分 Python 的标准转义字符也被正则表达式分析器支持。：

\a	\b	\f	\n
\r	\t	\u	\U
\v	\x	\\	

(注意 \b 被用于表示词语的边界，它只在字符集合内表示退格，比如 [\b]。)

'\u' and '\U' escape sequences are only recognized in Unicode patterns. In bytes patterns they are errors.

八进制转义包含为一个有限形式。如果首位数字是 0，或者有三个八进制数位，那么就认为它是八进制转义。其他的情况，就看作是组引用。对于字符串文本，八进制转义最多有三个数位长。

在 3.3 版更改：增加了 '\u' 和 '\U' 转义序列。

在 3.6 版更改：由 '\\' 和一个 ASCII 字符组成的未知转义会被看成错误。

6.2.2 模块内容

模块定义了几个函数，常量，和一个例外。有些函数是编译后的正则表达式方法的简化版本（少了一些特性）。绝大部分重要的应用，总是会先将正则表达式编译，之后在进行操作。

在 3.6 版更改：标志常量现在是 `RegexFlag` 类的实例，这个类是 `enum.IntFlag` 的子类。

`re.compile(pattern, flags=0)`

Compile a regular expression pattern into a *regular expression object*, which can be used for matching using its `match()`, `search()` and other methods, described below.

这个表达式的行为可以通过指定 标记的值来改变。值可以是以下任意变量，可以通过位的 OR 操作来结合 (| 操作符)。

序列

```
prog = re.compile(pattern)
result = prog.match(string)
```

等价于

```
result = re.match(pattern, string)
```

如果需要多次使用这个正则表达式的话，使用 `re.compile()` 和保存这个正则对象以便复用，可以让程序更加高效。

注解：通过 `re.compile()` 编译后的样式，和模块级的函数会被缓存，所以少数的正则表达式使用无需考虑编译的问题。

re.A

re.ASCII

让 `\w`, `\W`, `\b`, `\B`, `\d`, `\D`, `\s` 和 `\S` 只匹配 ASCII，而不是 Unicode。这只对 Unicode 样式有效，会被 byte 样式忽略。相当于前面语法中的内联标志 (`?a`)。

注意，为了保持向后兼容，`re.U` 标记依然存在（还有他的同义 `re.UNICODE` 和嵌入形式 (`?u`)），但是这些在 Python 3 是冗余的，因为默认字符串已经是 Unicode 了（并且 Unicode 匹配不允许 byte 出现）。

re.DEBUG

显示编译时的 debug 信息，没有内联标志。

re.I

re.IGNORECASE

进行忽略大小写匹配；表达式如 `[A-Z]` 也会匹配小写字符。Unicode 匹配（比如 `ü` 匹配 `ü`）同样有用，除非设置了 `re.ASCII` 标记来禁用非 ASCII 匹配。当前语言区域不会改变这个标记，除非设置了 `re.LOCALE` 标记。这个相当于内联标志 (`?i`)。

Note that when the Unicode patterns `[a-z]` or `[A-Z]` are used in combination with the `IGNORECASE` flag, they will match the 52 ASCII letters and 4 additional non-ASCII letters: ‘`İ`’ (U+0130, Latin capital letter I with dot above), ‘`ı`’ (U+0131, Latin small letter dotless i), ‘`ſ`’ (U+017F, Latin small letter long s) and ‘`℔`’ (U+212A, Kelvin sign). If the `ASCII` flag is used, only letters ‘`a`’ to ‘`z`’ and ‘`A`’ to ‘`Z`’ are matched (but the flag affects the entire regular expression, so in such cases using an explicit `(?-i: [a-zA-Z])` may be a better choice).

re.L

re.LOCALE

由当前语言区域决定 `\w`, `\W`, `\b`, `\B` 和大小写敏感匹配。这个标记只能对 byte 样式有效。这个标记不推荐使用，因为语言区域机制很不可靠，它一次只能处理一个“习惯”，而且只对 8 位字节有效。Unicode 匹配在 Python 3 里默认启用，并可以处理不同语言。这个对应内联标志 (`?L`)。

在 3.6 版更改：`re.LOCALE` 只能用于 byte 样式，而且不能和 `re.ASCII` 一起用。

re.M

re.MULTILINE

设置以后，样式字符 ‘`^`’ 匹配字符串的开始，和每一行的开始（换行符后面紧跟的符号）；样式字符 ‘`$`’ 匹配字符串尾，和每一行的结尾（换行符前面那个符号）。默认情况下，‘`^`’ 匹配字符串头，‘`$`’ 匹配字符串尾。对应内联标志 (`?m`)。

re.S

re.DOTALL

让 ‘`.`’ 特殊字符匹配任何字符，包括换行符；如果没有这个标记，‘`.`’ 就匹配除了换行符的其他任意字符。对应内联标志 (`?s`)。

re.X

re.VERBOSE

这个标记允许你编写更具可读性更友好的正则表达式。通过分段和添加注释。空白符号会被忽略，除非在一个字符集合当中或者由反斜杠转义，或者在 `*?`, `(?: or (?P<...>` 分组之内。当一个行内有 `#` 不在字符集和转义序列，那么它之后的所有字符都是注释。

意思就是下面两个正则表达式等价地匹配一个十进制数字：

```
a = re.compile(r"""\d + # the integral part
                  \.    # the decimal point
                  \d *  # some fractional digits""", re.X)
b = re.compile(r"\d+\.\d*")
```

对应内联标记 `(?x)`。

re.search(pattern, string, flags=0)

扫描整个字符串找到匹配样式的第一个位置，并返回一个相应的匹配对象。如果没有匹配，就返回一个 `None`；注意这和找到一个零长度匹配是不同的。

re.match(pattern, string, flags=0)

如果 `string` 开始的 0 或者多个字符匹配到了正则表达式样式，就返回一个相应的匹配对象。如果没有匹配，就返回 `None`；注意它跟零长度匹配是不同的。

注意即便是 `MULTILINE` 多行模式，`re.match()` 也只匹配字符串的开始位置，而不匹配每行开始。

如果你想定位 `string` 的任何位置，使用 `search()` 来替代（也可参考 `search()` vs. `match()`）

re.fullmatch(pattern, string, flags=0)

如果整个 `string` 匹配到正则表达式样式，就返回一个相应的匹配对象。否则就返回一个 `None`；注意这跟零长度匹配是不同的。

3.4 新版功能.**re.split(pattern, string, maxsplit=0, flags=0)**

用 `pattern` 分开 `string`。如果在 `pattern` 中捕获到括号，那么所有的组里的文字也会包含在列表里。如果 `maxsplit` 非零，最多进行 `maxsplit` 次分隔，剩下的字符全部返回到列表的最后一个元素。

```
>>> re.split(r'\W+', 'Words, words, words.')
['Words', 'words', 'words', '']
>>> re.split(r'(\W+)', 'Words, words, words.')
['Words', '', ' ', 'words', '', ' ', 'words', '.', '']
>>> re.split(r'\W+', 'Words, words, words.', 1)
['Words', 'words, words.']
>>> re.split('[a-f]+', '0a3B9', flags=re.IGNORECASE)
['0', '3', '9']
```

如果分隔符里有捕获组合，并且匹配到字符串的开始，那么结果将会以一个空字符串开始。对于结尾也是一样

```
>>> re.split(r'(\W+)', '...words, words...')
['', '...', 'words', '', ' ', 'words', '...', '']
```

这样的话，分隔组将会出现在结果列表中同样的位置。

注解： `split()` doesn't currently split a string on an empty pattern match. For example:

```
>>> re.split('x*', 'axbc')
['a', 'bc']
```

Even though `'x*'` also matches 0 `'x'` before `'a'`, between `'b'` and `'c'`, and after `'c'`, currently these matches are ignored. The correct behavior (i.e. splitting on empty matches too and returning `['', 'a', 'b',`

'c', ' ']) will be implemented in future versions of Python, but since this is a backward incompatible change, a *FutureWarning* will be raised in the meanwhile.

Patterns that can only match empty strings currently never split the string. Since this doesn't match the expected behavior, a *ValueError* will be raised starting from Python 3.5:

```
>>> re.split("^$", "foo\n\nbar\n", flags=re.M)
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
    ...
ValueError: split() requires a non-empty pattern match.
```

在 3.1 版更改: 增加了可选标记参数。

在 3.5 版更改: Splitting on a pattern that could match an empty string now raises a warning. Patterns that can only match empty strings are now rejected.

re.findall (*pattern*, *string*, *flags*=0)

对 *string* 返回一个不重复的 *pattern* 的匹配列表, *string* 从左到右进行扫描, 匹配按找到的顺序返回。如果样式里存在一到多个组, 就返回一个组合列表; 就是一个元组的列表 (如果样式里有超过一个组合的话)。空匹配也会包含在结果里。

注解: Due to the limitation of the current implementation the character following an empty match is not included in a next match, so `findall(r'^|\w+', 'two words')` returns `['', 'wo', 'words']` (note missed "t"). This is changed in Python 3.7.

re.finditer (*pattern*, *string*, *flags*=0)

Return an *iterator* yielding *match objects* over all non-overlapping matches for the RE *pattern* in *string*. The *string* is scanned left-to-right, and matches are returned in the order found. Empty matches are included in the result. See also the note about *findall* ().

re.sub (*pattern*, *repl*, *string*, *count*=0, *flags*=0)

Return the string obtained by replacing the leftmost non-overlapping occurrences of *pattern* in *string* by the replacement *repl*. If the pattern isn't found, *string* is returned unchanged. *repl* can be a string or a function; if it is a string, any backslash escapes in it are processed. That is, `\n` is converted to a single newline character, `\r` is converted to a carriage return, and so forth. Unknown escapes such as `\&` are left alone. Backreferences, such as `\6`, are replaced with the substring matched by group 6 in the pattern. For example:

```
>>> re.sub(r'def\s+([a-zA-Z_][a-zA-Z_0-9]*)\s*(\s*):',
...        r'static PyObject*\np_{1}(void)\n{',
...        'def myfunc():')
'static PyObject*\np_myfunc(void)\n{'
```

如果 *repl* 是一个函数, 那它会对每个非重复的 *pattern* 的情况调用。这个函数只能有一个 *匹配对象* 参数, 并返回一个替换后的字符串。比如

```
>>> def dashrepl(matchobj):
...     if matchobj.group(0) == '-': return ' '
...     else: return '-'
>>> re.sub('-{1,2}', dashrepl, 'pro---gram-files')
'pro--gram files'
>>> re.sub(r'\sAND\s', ' & ', 'Baked Beans And Spam', flags=re.IGNORECASE)
'Baked Beans & Spam'
```

样式可以是一个字符串或者一个样式对象。

The optional argument *count* is the maximum number of pattern occurrences to be replaced; *count* must be a non-negative integer. If omitted or zero, all occurrences will be replaced. Empty matches for the pattern are replaced only when not adjacent to a previous match, so `sub('x*', '-', 'abc')` returns `'-a-b-c-'`.

在字符串类型的 *repl* 参数里，如上所述的转义和向后引用中，`\g<name>` 会使用命名组合 *name*，（在 `(?P<name>...)` 语法中定义）`\g<number>` 会使用数字组；`\g<2>` 就是 `\2`，但它避免了二义性，如 `\g<2>0`。`\20` 就会被解释为组 20，而不是组 2 后面跟随一个字符 '0'。向后引用 `\g<0>` 把 *pattern* 作为一个整个组进行引用。

在 3.1 版更改：增加了可选标记参数。

在 3.5 版更改：不匹配的组合替换为空字符串。

在 3.6 版更改：*pattern* 中的未知转义（由 `'\'` 和一个 ASCII 字符组成）被视为错误。

Deprecated since version 3.5, will be removed in version 3.7: Unknown escapes in *repl* consisting of `'\'` and an ASCII letter now raise a deprecation warning and will be forbidden in Python 3.7.

re.subn (*pattern*, *repl*, *string*, *count=0*, *flags=0*)

行为与 `sub()` 相同，但是返回一个元组（字符串，替换次数）。

在 3.1 版更改：增加了可选标记参数。

在 3.5 版更改：不匹配的组合替换为空字符串。

re.escape (*pattern*)

Escape all the characters in *pattern* except ASCII letters, numbers and `'_'`. This is useful if you want to match an arbitrary literal string that may have regular expression metacharacters in it. For example:

```
>>> print(re.escape('python.exe'))
python\.exe

>>> legal_chars = string.ascii_lowercase + string.digits + "!#$%&'*+-.^_`|~:"
>>> print('[%s]+' % re.escape(legal_chars))
[abcdefghijklmnopqrstuvwxyz0123456789\!#\$\%&'\*\+\-\.\^_\`|\~\:] +

>>> operators = ['+', '-', '*', '/', '**']
>>> print(''.join(map(re.escape, sorted(operators, reverse=True))))
\|\/\-\|\/\+\|\/\*\|\/\*
```

这个函数不能用在 `sub()` 和 `subn()` 的替换字符串里，只有反斜杠应该被转义，比如说

```
>>> digits_re = r'\d+'
>>> sample = '/usr/sbin/sendmail - 0 errors, 12 warnings'
>>> print(re.sub(digits_re, digits_re.replace('\d', r'\\'), sample))
/usr/sbin/sendmail - \d+ errors, \d+ warnings
```

在 3.3 版更改：`'_'` 不再被转义。

re.purge ()

清除正则表达式缓存。

exception re.error (*msg*, *pattern=None*, *pos=None*)

raise 一个例外。当传递到函数的字符串不是一个有效正则表达式的时候（比如，包含一个不匹配的括号）或者其他错误在编译时或匹配时产生。如果字符串不包含样式匹配，是不会被视为错误的。错误实例有以下附加属性：

msg

未格式化的错误消息。

pattern

正则表达式样式。

pos

编译失败的 *pattern* 的位置索引 (可以是 None)。

lineno

对应 *pos* (可以是 None) 的行号。

colno

对应 *pos* (可以是 None) 的列号。

在 3.5 版更改: 添加了附加属性。

6.2.3 正则表达式对象 (正则对象)

编译后的正则表达式对象支持一下方法和属性:

regex.search (*string* [, *pos* [, *endpos*]])

扫描整个 *string* 寻找第一个匹配的位置, 并返回一个相应的匹配对象。如果没有匹配, 就返回 None; 注意它和零长度匹配是不同的。

可选的第二个参数 *pos* 给出了字符串中开始搜索的位置索引; 默认为 0, 它不完全等价于字符串切片; `'^'` 样式字符匹配字符串真正的开头, 和换行符后面的第一个字符, 但不会匹配索引规定开始的位置。

可选参数 *endpos* 限定了字符串搜索的结束; 它假定字符串长度到 *endpos*, 所以只有从 *pos* 到 *endpos* - 1 的字符会被匹配。如果 *endpos* 小于 *pos*, 就不会有匹配产生; 另外, 如果 *rx* 是一个编译后的正则对象, `rx.search(string, 0, 50)` 等价于 `rx.search(string[:50], 0)`。

```
>>> pattern = re.compile("d")
>>> pattern.search("dog")          # Match at index 0
<_sre.SRE_Match object; span=(0, 1), match='d'>
>>> pattern.search("dog", 1)      # No match; search doesn't include the "d"
```

regex.match (*string* [, *pos* [, *endpos*]])

如果 *string* 的开始位置能够找到这个正则样式的任意个匹配, 就返回一个相应的匹配对象。如果不匹配, 就返回 None; 注意它与零长度匹配是不同的。

The optional *pos* and *endpos* parameters have the same meaning as for the *search()* method.

```
>>> pattern = re.compile("o")
>>> pattern.match("dog")           # No match as "o" is not at the start of "dog".
>>> pattern.match("dog", 1)       # Match as "o" is the 2nd character of "dog".
<_sre.SRE_Match object; span=(1, 2), match='o'>
```

If you want to locate a match anywhere in *string*, use *search()* instead (see also *search()* vs. *match()*).

regex.fullmatch (*string* [, *pos* [, *endpos*]])

如果整个 *string* 匹配这个正则表达式, 就返回一个相应的匹配对象。否则就返回 None; 注意跟零长度匹配是不同的。

The optional *pos* and *endpos* parameters have the same meaning as for the *search()* method.

```
>>> pattern = re.compile("o[gh]")
>>> pattern.fullmatch("dog")       # No match as "o" is not at the start of "dog".
>>> pattern.fullmatch("ogre")     # No match as not the full string matches.
>>> pattern.fullmatch("doggie", 1, 3) # Matches within given limits.
<_sre.SRE_Match object; span=(1, 3), match='og'>
```

3.4 新版功能.

regex.split (*string*, *maxsplit*=0)

等价于 *split()* 函数, 使用了编译后的样式。

`regex.findall(string[, pos[, endpos]])`

类似函数 `findall()`，使用了编译后样式，但也可以接收可选参数 `pos` 和 `endpos`，限制搜索范围，就像 `search()`。

`regex.finditer(string[, pos[, endpos]])`

类似函数 `finditer()`，使用了编译后样式，但也可以接收可选参数 `pos` 和 `endpos`，限制搜索范围，就像 `search()`。

`regex.sub(repl, string, count=0)`

等价于 `sub()` 函数，使用了编译后的样式。

`regex.subn(repl, string, count=0)`

等价于 `subn()` 函数，使用了编译后的样式。

`regex.flags`

正则匹配标记。这是可以传递给 `compile()` 的参数，任何 (??...) 内联标记，隐性标记比如 UNICODE 的结合。

`regex.groups`

捕获组合的数量。

`regex.groupindex`

映射由 (P<id>) 定义的命名符号组合和数字组合的字典。如果没有符号组，那字典就是空的。

`regex.pattern`

The pattern string from which the RE object was compiled.

6.2.4 匹配对象

Match objects always have a boolean value of True. Since `match()` and `search()` return None when there is no match, you can test whether there was a match with a simple if statement:

```
match = re.search(pattern, string)
if match:
    process(match)
```

匹配对象支持以下方法和属性：

`match.expand(template)`

Return the string obtained by doing backslash substitution on the template string `template`, as done by the `sub()` method. Escapes such as `\n` are converted to the appropriate characters, and numeric backreferences (`\1`, `\2`) and named backreferences (`\g<1>`, `\g<name>`) are replaced by the contents of the corresponding group.

在 3.5 版更改：不匹配的组合替换为空字符串。

`match.group([group1, ...])`

返回一个或者多个匹配的子组。如果只有一个参数，结果就是一个字符串，如果有多个参数，结果就是一个元组（每个参数对应一个项），如果没有参数，组 1 默认到 0（整个匹配都被返回）。如果一个组 N 参数值为 0，相应的返回值就是整个匹配字符串；如果它是一个范围 [1..99]，结果就是相应的括号组字符串。如果一个组号是负数，或者大于样式中定义的组数，一个 `IndexError` 索引错误就 raise。如果一个组包含在样式的一部分，并被匹配多次，就返回最后一个匹配。：

```
>>> m = re.match(r"(\w+) (\w+)", "Isaac Newton, physicist")
>>> m.group(0)           # The entire match
'Isaac Newton'
>>> m.group(1)           # The first parenthesized subgroup.
'Isaac'
>>> m.group(2)           # The second parenthesized subgroup.
```

(下页继续)

(续上页)

```
'Newton'
>>> m.group(1, 2)    # Multiple arguments give us a tuple.
('Isaac', 'Newton')
```

如果正则表达式使用了 `(?P<name>...)` 语法, `groupN` 参数就也可能是命名组合的名字。如果一个字符串参数在样式中未定义为组合名, 一个 `IndexError` 就 `raise`。

一个相对复杂的例子

```
>>> m = re.match(r"(?P<first_name>\w+) (?P<last_name>\w+)", "Malcolm Reynolds")
>>> m.group('first_name')
'Malcolm'
>>> m.group('last_name')
'Reynolds'
```

命名组合同样可以通过索引值引用

```
>>> m.group(1)
'Malcolm'
>>> m.group(2)
'Reynolds'
```

如果一个组匹配成功多次, 就只返回最后一个匹配

```
>>> m = re.match(r"(..)+", "a1b2c3")    # Matches 3 times.
>>> m.group(1)                          # Returns only the last match.
'c3'
```

`match.__getitem__(g)`

这个等价于 `m.group(g)`。这允许更方便的引用一个匹配

```
>>> m = re.match(r"(\w+) (\w+)", "Isaac Newton, physicist")
>>> m[0]    # The entire match
'Isaac Newton'
>>> m[1]    # The first parenthesized subgroup.
'Isaac'
>>> m[2]    # The second parenthesized subgroup.
'Newton'
```

3.6 新版功能.

`match.groups(default=None)`

返回一个元组, 包含所有匹配的子组, 在样式中出现的从 1 到任意多的组合。`default` 参数用于不参与匹配的情况, 默认为 `None`。

例如

```
>>> m = re.match(r"(\d+)\.(\d+)", "24.1632")
>>> m.groups()
('24', '1632')
```

如果我们使小数点可选, 那么不是所有的组都会参与到匹配当中。这些组合默认会返回一个 `None`, 除非指定了 `default` 参数。

```
>>> m = re.match(r"(\d+)\.?(\\d+)?", "24")
>>> m.groups()    # Second group defaults to None.
('24', None)
```

(下页继续)

(续上页)

```
>>> m.groups('0')    # Now, the second group defaults to '0'.
('24', '0')
```

match.groupdict (*default=None*)

返回一个字典，包含了所有的命名子组。**key** 就是组名。*default* 参数用于不参与匹配的组合；默认为 *None*。例如

```
>>> m = re.match(r"(?P<first_name>\w+) (?P<last_name>\w+)", "Malcolm Reynolds")
>>> m.groupdict()
{'first_name': 'Malcolm', 'last_name': 'Reynolds'}
```

match.start (*[group]*)

match.end (*[group]*)

返回 *group* 匹配到的字串的开始和结束标号。*group* 默认为 0（意思是整个匹配的子串）。如果 *group* 存在，但未产生匹配，就返回 -1。对于一个匹配对象 *m*，和一个未参与匹配的组 *g*，组 *g*（等价于 *m.group(g)*）产生的匹配是

```
m.string[m.start(g):m.end(g)]
```

注意 *m.start(group)* 将会等于 *m.end(group)*，如果 *group* 匹配一个空字符串的话。比如，在 *m = re.search('b(c?)', 'cba')* 之后，*m.start(0)* 为 1，*m.end(0)* 为 2，*m.start(1)* 和 *m.end(1)* 都是 2，*m.start(2)* raise 一个 *IndexError* 例外。

这个例子会从 email 地址中移除掉 *remove_this*

```
>>> email = "tony@tiremove_thisger.net"
>>> m = re.search("remove_this", email)
>>> email[:m.start()] + email[m.end():]
'tony@tiger.net'
```

match.span (*[group]*)

对于一个匹配 *m*，返回一个二元组 (*m.start(group)*, *m.end(group)*)。注意如果 *group* 没有在这个匹配中，就返回 (-1, -1)。*group* 默认为 0，就是整个匹配。

match.pos

The value of *pos* which was passed to the *search()* or *match()* method of a *regex object*. This is the index into the string at which the RE engine started looking for a match.

match.endpos

The value of *endpos* which was passed to the *search()* or *match()* method of a *regex object*. This is the index into the string beyond which the RE engine will not go.

match.lastindex

捕获组的最后一个匹配的整数索引值，或者 *None* 如果没有匹配产生的话。比如，对于字符串 'ab'，表达式 (a)b, ((a)(b)), 和 ((ab)) 将得到 *lastindex == 1*，而 (a)(b) 会得到 *lastindex == 2*。

match.lastgroup

最后一个匹配的命名组名字，或者 *None* 如果没有产生匹配的话。

match.re

The *regular expression object* whose *match()* or *search()* method produced this match instance.

match.string

The string passed to *match()* or *search()*.

6.2.5 正则表达式例子

检查对子

在这个例子里，我们使用以下辅助函数来更好的显示匹配对象：

```
def displaymatch(match):
    if match is None:
        return None
    return '<Match: %r, groups=%r>' % (match.group(), match.groups())
```

假设你在写一个扑克程序，一个玩家的一手牌为五个字符的串，每个字符表示一张牌，”a”就是A，“k”K，“q”Q，“j”J，“t”为10，“2”到“9”表示2到9。

要看给定的字符串是否有效，我们可以按照以下步骤

```
>>> valid = re.compile(r"^[a2-9tjqk]{5}$")
>>> displaymatch(valid.match("akt5q")) # Valid.
"<Match: 'akt5q', groups=()>"
>>> displaymatch(valid.match("akt5e")) # Invalid.
>>> displaymatch(valid.match("akt")) # Invalid.
>>> displaymatch(valid.match("727ak")) # Valid.
"<Match: '727ak', groups=()>"
```

最后一手牌，“727ak”，包含了一个对子，或者两张同样数值的牌。要用正则表达式匹配它，应该使用向后引用如下

```
>>> pair = re.compile(r".*(.)\1")
>>> displaymatch(pair.match("717ak")) # Pair of 7s.
"<Match: '717', groups=('7',)>"
>>> displaymatch(pair.match("718ak")) # No pairs.
>>> displaymatch(pair.match("354aa")) # Pair of aces.
"<Match: '354aa', groups=('a',)>"
```

To find out what card the pair consists of, one could use the `group()` method of the match object in the following manner:

```
>>> pair.match("717ak").group(1)
'7'

# Error because re.match() returns None, which doesn't have a group() method:
>>> pair.match("718ak").group(1)
Traceback (most recent call last):
  File "<pyshell#23>", line 1, in <module>
    re.match(r".*(.)\1", "718ak").group(1)
AttributeError: 'NoneType' object has no attribute 'group'

>>> pair.match("354aa").group(1)
'a'
```

模拟 scanf()

Python 目前没有一个类似 c 函数 `scanf()` 的替代品。正则表达式通常比 `scanf()` 格式字符串要更强大一些，但也带来更多复杂性。下面的表格提供了 `scanf()` 格式符和正则表达式大致相同的映射。

scanf() 格式符	正则表达式
%c	.
%5c	.{5}
%d	[-+] ? \d +
%e, %E, %f, %g	[-+] ? (\d + (\. \d *) ? \. \d +) ([eE] [-+] ? \d +) ?
%i	[-+] ? (0 [xX] [\dA-Fa-f] + 0 [0-7] * \d +)
%o	[-+] ? [0-7] +
%s	\S +
%u	\d +
%x, %X	[-+] ? (0 [xX]) ? [\dA-Fa-f] +

从文件名和数字提取字符串

```
/usr/sbin/sendmail - 0 errors, 4 warnings
```

你可以使用 `scanf()` 格式化

```
%s - %d errors, %d warnings
```

等价的正则表达式是：

```
(\S+) - (\d+) errors, (\d+) warnings
```

search() vs. match()

Python 提供了两种不同的操作：基于 `re.match()` 检查字符串开头，或者 `re.search()` 检查字符串的任意位置（默认 Perl 中的行为）。

例如

```
>>> re.match("c", "abcdef")      # No match
>>> re.search("c", "abcdef")     # Match
<_sre.SRE_Match object; span=(2, 3), match='c'>
```

在 `search()` 中，可以用 `'^'` 作为开始来限制匹配到字符串的首位

```
>>> re.match("c", "abcdef")      # No match
>>> re.search("^c", "abcdef")    # No match
>>> re.search("^a", "abcdef")    # Match
<_sre.SRE_Match object; span=(0, 1), match='a'>
```

注意 `MULTILINE` 多行模式中函数 `match()` 只匹配字符串的开始，但使用 `search()` 和以 `'^'` 开始的正则表达式会匹配每行的开始

```
>>> re.match('X', 'A\nB\nX', re.MULTILINE) # No match
>>> re.search('^X', 'A\nB\nX', re.MULTILINE) # Match
<_sre.SRE_Match object; span=(4, 5), match='X'>
```

建立一个电话本

`split()` 将字符串用参数传递的样式分隔开。这个方法对于转换文本数据到易读而且容易修改的数据结构，是很有用的，如下面的例子证明。

首先，这里是输入。通常是一个文件，这里我们用三引号字符串语法

```
>>> text = """Ross McFluff: 834.345.1254 155 Elm Street
...
... Ronald Heathmore: 892.345.3428 436 Finley Avenue
... Frank Burger: 925.541.7625 662 South Dogwood Way
...
...
... Heather Albrecht: 548.326.4584 919 Park Place"""
```

条目用一个或者多个换行符分开。现在我们将字符串转换为一个列表，每个非空行都有一个条目：

```
>>> entries = re.split("\n+", text)
>>> entries
['Ross McFluff: 834.345.1254 155 Elm Street',
 'Ronald Heathmore: 892.345.3428 436 Finley Avenue',
 'Frank Burger: 925.541.7625 662 South Dogwood Way',
 'Heather Albrecht: 548.326.4584 919 Park Place']
```

最终，将每个条目分割为一个由名字、姓氏、电话号码和地址组成的列表。我们为`split()`使用了`maxsplit`形参，因为地址中包含有被我们作为分割模式的空格符：

```
>>> [re.split("?: ", entry, 3) for entry in entries]
[['Ross', 'McFluff', '834.345.1254', '155 Elm Street'],
 ['Ronald', 'Heathmore', '892.345.3428', '436 Finley Avenue'],
 ['Frank', 'Burger', '925.541.7625', '662 South Dogwood Way'],
 ['Heather', 'Albrecht', '548.326.4584', '919 Park Place']]
```

`?:` 样式匹配姓后面的冒号，因此它不出现在结果列表中。如果`maxsplit` 设置为 4，我们还可以从地址中获取到房间号：

```
>>> [re.split("?: ", entry, 4) for entry in entries]
[['Ross', 'McFluff', '834.345.1254', '155', 'Elm Street'],
 ['Ronald', 'Heathmore', '892.345.3428', '436', 'Finley Avenue'],
 ['Frank', 'Burger', '925.541.7625', '662', 'South Dogwood Way'],
 ['Heather', 'Albrecht', '548.326.4584', '919', 'Park Place']]
```

文字整理

`sub()` 替换字符串中出现的样式的每一个实例。这个例子证明了使用`sub()` 来整理文字，或者随机化每个字符的位置，除了首位和末尾字符

```
>>> def repl(m):
...     inner_word = list(m.group(2))
...     random.shuffle(inner_word)
...     return m.group(1) + "".join(inner_word) + m.group(3)
>>> text = "Professor Abdolmalek, please report your absences promptly."
>>> re.sub(r"(\w)(\w+)(\w)", repl, text)
'Poefsrorsr Aealmlbdk, pslaee reorpt your abnseces plmrptoy.'
>>> re.sub(r"(\w)(\w+)(\w)", repl, text)
'Pofsroser Aodlambelk, plasee reorpt yuor asnebces potlmpy.'
```

找到所有副词

`findall()` 匹配样式 所有的出现，不仅是像`search()` 中的第一个匹配。比如，如果一个作者希望找到文字中的所有副词，他可能会按照以下方法用`findall()`

```
>>> text = "He was carefully disguised but captured quickly by police."
>>> re.findall(r"\w+ly", text)
['carefully', 'quickly']
```

找到所有副词和位置

如果需要匹配样式的更多信息，`finditer()` 可以起到作用，它提供了匹配对象 作为返回值，而不是字符串。继续上面的例子，如果一个作者希望找到所有副词和它的位置，可以按照下面方法使用 `finditer()`

```
>>> text = "He was carefully disguised but captured quickly by police."
>>> for m in re.finditer(r"\w+ly", text):
...     print('%02d-%02d: %s' % (m.start(), m.end(), m.group(0)))
07-16: carefully
40-47: quickly
```

原始字符记法

原始字符串记法 (`r"text"`) 保持正则表达式正常。否则，每个正则式里的反斜杠 (`'\'`) 都必须前缀一个反斜杠来转义。比如，下面两行代码功能就是完全一致的

```
>>> re.match(r"\W(.)\1\W", " ff ")
<_sre.SRE_Match object; span=(0, 4), match=' ff '>
>>> re.match("\\W(.)\\1\\W", " ff ")
<_sre.SRE_Match object; span=(0, 4), match=' ff '>
```

当需要匹配一个字符反斜杠，它必须在正则表达式中转义。在原始字符串记法，就是 `r"\"`。否则就必须用 `"\\\"`，来表示同样的意思

```
>>> re.match(r"\\", r"\\")
<_sre.SRE_Match object; span=(0, 1), match='\\'>
>>> re.match("\\\\", r"\\")
<_sre.SRE_Match object; span=(0, 1), match='\\'>
```

写一个词法分析器

一个 词法器或词法分析器 分析字符串，并分类成目录组。这是写一个编译器或解释器的第一步。

文字目录是由正则表达式指定的。这个技术是通过将这些样式合并为一个主正则式，并且循环匹配来实现的

```
import collections
import re

Token = collections.namedtuple('Token', ['type', 'value', 'line', 'column'])

def tokenize(code):
    keywords = {'IF', 'THEN', 'ENDIF', 'FOR', 'NEXT', 'GOSUB', 'RETURN'}
    token_specification = [
        ('NUMBER',    r'\d+(\.\d*)?'), # Integer or decimal number
        ('ASSIGN',    r':='),          # Assignment operator
        ('END',       r';'),            # Statement terminator
        ('ID',        r'[A-Za-z]+'),   # Identifiers
        ('OP',        r'[+ \-*/]'),    # Arithmetic operators
```

(下页继续)

(续上页)

```

        ('NEWLINE', r'\n'),          # Line endings
        ('SKIP', r'[ \t]+'),        # Skip over spaces and tabs
        ('MISMATCH', r'.'),         # Any other character
    ]
    tok_regex = '|'.join('(?P<%s>%s)' % pair for pair in token_specification)
    line_num = 1
    line_start = 0
    for mo in re.finditer(tok_regex, code):
        kind = mo.lastgroup
        value = mo.group()
        column = mo.start() - line_start
        if kind == 'NUMBER':
            value = float(value) if '.' in value else int(value)
        elif kind == 'ID' and value in keywords:
            kind = value
        elif kind == 'NEWLINE':
            line_start = mo.end()
            line_num += 1
            continue
        elif kind == 'SKIP':
            continue
        elif kind == 'MISMATCH':
            raise RuntimeError(f'{value!r} unexpected on line {line_num}')
        yield Token(kind, value, line_num, column)

statements = '''
    IF quantity THEN
        total := total + price * quantity;
        tax := price * 0.05;
    ENDIF;
'''

for token in tokenize(statements):
    print(token)

```

这个词法器产生以下输出

```

Token(type='IF', value='IF', line=2, column=4)
Token(type='ID', value='quantity', line=2, column=7)
Token(type='THEN', value='THEN', line=2, column=16)
Token(type='ID', value='total', line=3, column=8)
Token(type='ASSIGN', value=':=', line=3, column=14)
Token(type='ID', value='total', line=3, column=17)
Token(type='OP', value='+', line=3, column=23)
Token(type='ID', value='price', line=3, column=25)
Token(type='OP', value='*', line=3, column=31)
Token(type='ID', value='quantity', line=3, column=33)
Token(type='END', value=';', line=3, column=41)
Token(type='ID', value='tax', line=4, column=8)
Token(type='ASSIGN', value=':=', line=4, column=12)
Token(type='ID', value='price', line=4, column=15)
Token(type='OP', value='*', line=4, column=21)
Token(type='NUMBER', value=0.05, line=4, column=23)
Token(type='END', value=';', line=4, column=27)
Token(type='ENDIF', value='ENDIF', line=5, column=4)
Token(type='END', value=';', line=5, column=9)

```

6.3 difflib — 计算差异的辅助工具

源代码: [Lib/difflib.py](#)

此模块提供用于比较序列的类和函数。例如，它可以用于比较文件，并可以产生各种格式的不同信息，包括 HTML 和上下文以及统一格式的差异点。有关目录和文件的比较，请参见 [filecmp](#) 模块。

class difflib.SequenceMatcher

This is a flexible class for comparing pairs of sequences of any type, so long as the sequence elements are *hashable*. The basic algorithm predates, and is a little fancier than, an algorithm published in the late 1980' s by Ratcliff and Obershelp under the hyperbolic name “gestalt pattern matching.” The idea is to find the longest contiguous matching subsequence that contains no “junk” elements; these “junk” elements are ones that are uninteresting in some sense, such as blank lines or whitespace. (Handling junk is an extension to the Ratcliff and Obershelp algorithm.) The same idea is then applied recursively to the pieces of the sequences to the left and to the right of the matching subsequence. This does not yield minimal edit sequences, but does tend to yield matches that “look right” to people.

Timing: The basic Ratcliff-Obershelp algorithm is cubic time in the worst case and quadratic time in the expected case. *SequenceMatcher* is quadratic time for the worst case and has expected-case behavior dependent in a complicated way on how many elements the sequences have in common; best case time is linear.

Automatic junk heuristic: *SequenceMatcher* supports a heuristic that automatically treats certain sequence items as junk. The heuristic counts how many times each individual item appears in the sequence. If an item' s duplicates (after the first one) account for more than 1% of the sequence and the sequence is at least 200 items long, this item is marked as “popular” and is treated as junk for the purpose of sequence matching. This heuristic can be turned off by setting the *autojunk* argument to *False* when creating the *SequenceMatcher*.

3.2 新版功能: The *autojunk* parameter.

class difflib.Differ

This is a class for comparing sequences of lines of text, and producing human-readable differences or deltas. Differ uses *SequenceMatcher* both to compare sequences of lines, and to compare sequences of characters within similar (near-matching) lines.

Each line of a *Differ* delta begins with a two-letter code:

状态码	意义
'- '	line unique to sequence 1
'+' '	line unique to sequence 2
' ' '	line common to both sequences
'?' '	line not present in either input sequence

Lines beginning with ‘?’ attempt to guide the eye to intraline differences, and were not present in either input sequence. These lines can be confusing if the sequences contain tab characters.

class difflib.HtmlDiff

This class can be used to create an HTML table (or a complete HTML file containing the table) showing a side by side, line by line comparison of text with inter-line and intra-line change highlights. The table can be generated in either full or contextual difference mode.

这个类的构造函数:

__init__ (*tabsize=8, wrapcolumn=None, linejunk=None, charjunk=IS_CHARACTER_JUNK*)

初始化 *HtmlDiff* 的实例。

tabsize is an optional keyword argument to specify tab stop spacing and defaults to 8.

wrapcolumn is an optional keyword to specify column number where lines are broken and wrapped, defaults to `None` where lines are not wrapped.

linejunk and *charjunk* are optional keyword arguments passed into `ndiff()` (used by `HtmlDiff` to generate the side by side HTML differences). See `ndiff()` documentation for argument default values and descriptions.

下列是公开的方法

make_file (*fromlines*, *tolines*, *fromdesc*="", *todesc*="", *context*=False, *numlines*=5, *, *charset*='utf-8')

Compares *fromlines* and *tolines* (lists of strings) and returns a string which is a complete HTML file containing a table showing line by line differences with inter-line and intra-line changes highlighted.

fromdesc and *todesc* are optional keyword arguments to specify from/to file column header strings (both default to an empty string).

context and *numlines* are both optional keyword arguments. Set *context* to `True` when contextual differences are to be shown, else the default is `False` to show the full files. *numlines* defaults to 5. When *context* is `True` *numlines* controls the number of context lines which surround the difference highlights. When *context* is `False` *numlines* controls the number of lines which are shown before a difference highlight when using the “next” hyperlinks (setting to zero would cause the “next” hyperlinks to place the next difference highlight at the top of the browser without any leading context).

在 3.5 版更改: *charset* keyword-only argument was added. The default charset of HTML document changed from 'ISO-8859-1' to 'utf-8'.

make_table (*fromlines*, *tolines*, *fromdesc*="", *todesc*="", *context*=False, *numlines*=5)

Compares *fromlines* and *tolines* (lists of strings) and returns a string which is a complete HTML table showing line by line differences with inter-line and intra-line changes highlighted.

The arguments for this method are the same as those for the `make_file()` method.

`Tools/scripts/diff.py` is a command-line front-end to this class and contains a good example of its use.

`difflib.context_diff(a, b, fromfile="", tofile="", fromfiledate="", tofiledate="", n=3, lineterm='\n')`

Compare *a* and *b* (lists of strings); return a delta (a *generator* generating the delta lines) in context diff format.

Context diffs are a compact way of showing just the lines that have changed plus a few lines of context. The changes are shown in a before/after style. The number of context lines is set by *n* which defaults to three.

By default, the diff control lines (those with `***` or `---`) are created with a trailing newline. This is helpful so that inputs created from `io.IOBase.readlines()` result in diffs that are suitable for use with `io.IOBase.writelines()` since both the inputs and outputs have trailing newlines.

For inputs that do not have trailing newlines, set the *lineterm* argument to `" "` so that the output will be uniformly newline free.

The context diff format normally has a header for filenames and modification times. Any or all of these may be specified using strings for *fromfile*, *tofile*, *fromfiledate*, and *tofiledate*. The modification times are normally expressed in the ISO 8601 format. If not specified, the strings default to blanks.

```
>>> s1 = ['bacon\n', 'eggs\n', 'ham\n', 'guido\n']
>>> s2 = ['python\n', 'eggy\n', 'hamster\n', 'guido\n']
>>> sys.stdout.writelines(context_diff(s1, s2, fromfile='before.py', tofile=
↪ 'after.py'))
*** before.py
--- after.py
*****
*** 1,4 ****
! bacon
! eggs
```

(下页继续)

(续上页)

```
! ham
  guido
--- 1,4 ----
! python
! eggy
! hamster
  guido
```

See [A command-line interface to difflib](#) for a more detailed example.

`difflib.get_close_matches(word, possibilities, n=3, cutoff=0.6)`

Return a list of the best “good enough” matches. *word* is a sequence for which close matches are desired (typically a string), and *possibilities* is a list of sequences against which to match *word* (typically a list of strings).

Optional argument *n* (default 3) is the maximum number of close matches to return; *n* must be greater than 0.

Optional argument *cutoff* (default 0.6) is a float in the range [0, 1]. Possibilities that don’t score at least that similar to *word* are ignored.

The best (no more than *n*) matches among the possibilities are returned in a list, sorted by similarity score, most similar first.

```
>>> get_close_matches('appel', ['ape', 'apple', 'peach', 'puppy'])
['apple', 'ape']
>>> import keyword
>>> get_close_matches('wheel', keyword.kwlist)
['while']
>>> get_close_matches('pineapple', keyword.kwlist)
[]
>>> get_close_matches('accept', keyword.kwlist)
['except']
```

`difflib.ndiff(a, b, linejunk=None, charjunk=IS_CHARACTER_JUNK)`

Compare *a* and *b* (lists of strings); return a *Differ*-style delta (a *generator* generating the delta lines).

Optional keyword parameters *linejunk* and *charjunk* are filtering functions (or None):

linejunk: A function that accepts a single string argument, and returns true if the string is junk, or false if not. The default is None. There is also a module-level function `IS_LINE_JUNK()`, which filters out lines without visible characters, except for at most one pound character (`'#'`) – however the underlying *SequenceMatcher* class does a dynamic analysis of which lines are so frequent as to constitute noise, and this usually works better than using this function.

charjunk: A function that accepts a character (a string of length 1), and returns if the character is junk, or false if not. The default is module-level function `IS_CHARACTER_JUNK()`, which filters out whitespace characters (a blank or tab; it’s a bad idea to include newline in this!).

`Tools/scripts/ndiff.py` is a command-line front-end to this function.

```
>>> diff = ndiff('one\ntwo\nthree\n'.splitlines(keepends=True),
...              'ore\ntree\nemu\n'.splitlines(keepends=True))
>>> print(''.join(diff), end="")
- one
?  ^
+ ore
?  ^
- two
- three
```

(下页继续)

(续上页)

```
? -
+ tree
+ emu
```

`difflib.restore(sequence, which)`

Return one of the two sequences that generated a delta.

Given a *sequence* produced by `Differ.compare()` or `ndiff()`, extract lines originating from file 1 or 2 (parameter *which*), stripping off line prefixes.

示例:

```
>>> diff = ndiff('one\ntwo\nthree\n'.splitlines(keepends=True),
...              'ore\ntree\nemu\n'.splitlines(keepends=True))
>>> diff = list(diff) # materialize the generated delta into a list
>>> print(''.join(restore(diff, 1)), end="")
one
two
three
>>> print(''.join(restore(diff, 2)), end="")
ore
tree
emu
```

`difflib.unified_diff(a, b, fromfile="", tofile="", fromfiledate="", tofiledate="", n=3, lineterm='\n')`

Compare *a* and *b* (lists of strings); return a delta (a *generator* generating the delta lines) in unified diff format.

Unified diffs are a compact way of showing just the lines that have changed plus a few lines of context. The changes are shown in an inline style (instead of separate before/after blocks). The number of context lines is set by *n* which defaults to three.

By default, the diff control lines (those with `---`, `+++`, or `@@`) are created with a trailing newline. This is helpful so that inputs created from `io.IOBase.readlines()` result in diffs that are suitable for use with `io.IOBase.writelines()` since both the inputs and outputs have trailing newlines.

For inputs that do not have trailing newlines, set the *lineterm* argument to `" "` so that the output will be uniformly newline free.

The context diff format normally has a header for filenames and modification times. Any or all of these may be specified using strings for *fromfile*, *tofile*, *fromfiledate*, and *tofiledate*. The modification times are normally expressed in the ISO 8601 format. If not specified, the strings default to blanks.

```
>>> s1 = ['bacon\n', 'eggs\n', 'ham\n', 'guido\n']
>>> s2 = ['python\n', 'eggy\n', 'hamster\n', 'guido\n']
>>> sys.stdout.writelines(unified_diff(s1, s2, fromfile='before.py', tofile=
↪ 'after.py'))
--- before.py
+++ after.py
@@ -1,4 +1,4 @@
-bacon
-eggs
-ham
+python
+eggy
+hamster
 guido
```

See [A command-line interface to difflib](#) for a more detailed example.

`difflib.diff_bytes(dfunc, a, b, fromfile=b'', tofile=b'', fromfiledate=b'', tofiledate=b'', n=3, lineterm=b'\n')`

Compare *a* and *b* (lists of bytes objects) using *dfunc*; yield a sequence of delta lines (also bytes) in the format returned by *dfunc*. *dfunc* must be a callable, typically either `unified_diff()` or `context_diff()`.

Allows you to compare data with unknown or inconsistent encoding. All inputs except *n* must be bytes objects, not str. Works by losslessly converting all inputs (except *n*) to str, and calling `dfunc(a, b, fromfile, tofile, fromfiledate, tofiledate, n, lineterm)`. The output of *dfunc* is then converted back to bytes, so the delta lines that you receive have the same unknown/inconsistent encodings as *a* and *b*.

3.5 新版功能.

`difflib.IS_LINE_JUNK(line)`

Return true for ignorable lines. The line *line* is ignorable if *line* is blank or contains a single '#', otherwise it is not ignorable. Used as a default for parameter *linejunk* in `ndiff()` in older versions.

`difflib.IS_CHARACTER_JUNK(ch)`

Return true for ignorable characters. The character *ch* is ignorable if *ch* is a space or tab, otherwise it is not ignorable. Used as a default for parameter *charjunk* in `ndiff()`.

参见:

Pattern Matching: The Gestalt Approach Discussion of a similar algorithm by John W. Ratcliff and D. E. Metzener. This was published in *Dr. Dobbs' s Journal* in July, 1988.

6.3.1 SequenceMatcher Objects

The `SequenceMatcher` class has this constructor:

class `difflib.SequenceMatcher(isjunk=None, a="", b="", autojunk=True)`

Optional argument *isjunk* must be `None` (the default) or a one-argument function that takes a sequence element and returns true if and only if the element is “junk” and should be ignored. Passing `None` for *isjunk* is equivalent to passing `lambda x: 0`; in other words, no elements are ignored. For example, pass:

```
lambda x: x in " \t"
```

if you're comparing lines as sequences of characters, and don't want to synch up on blanks or hard tabs.

The optional arguments *a* and *b* are sequences to be compared; both default to empty strings. The elements of both sequences must be *hashable*.

The optional argument *autojunk* can be used to disable the automatic junk heuristic.

3.2 新版功能: The *autojunk* parameter.

`SequenceMatcher` objects get three data attributes: *bjunk* is the set of elements of *b* for which *isjunk* is `True`; *bpopular* is the set of non-junk elements considered popular by the heuristic (if it is not disabled); *b2j* is a dict mapping the remaining elements of *b* to a list of positions where they occur. All three are reset whenever *b* is reset with `set_seqs()` or `set_seq2()`.

3.2 新版功能: The *bjunk* and *bpopular* attributes.

`SequenceMatcher` objects have the following methods:

set_seqs(a, b)

Set the two sequences to be compared.

`SequenceMatcher` computes and caches detailed information about the second sequence, so if you want to compare one sequence against many sequences, use `set_seq2()` to set the commonly used sequence once and call `set_seq1()` repeatedly, once for each of the other sequences.

set_seq1 (*a*)

Set the first sequence to be compared. The second sequence to be compared is not changed.

set_seq2 (*b*)

Set the second sequence to be compared. The first sequence to be compared is not changed.

find_longest_match (*alo, ahi, blo, bhi*)

Find longest matching block in *a*[*alo*:*ahi*] and *b*[*blo*:*bhi*].

If *isjunk* was omitted or *None*, *find_longest_match()* returns (*i*, *j*, *k*) such that *a*[*i*:*i*+*k*] is equal to *b*[*j*:*j*+*k*], where *alo* ≤ *i* ≤ *i*+*k* ≤ *ahi* and *blo* ≤ *j* ≤ *j*+*k* ≤ *bhi*. For all (*i'*, *j'*, *k'*) meeting those conditions, the additional conditions *k* ≥ *k'*, *i* ≤ *i'*, and if *i* == *i'*, *j* ≤ *j'* are also met. In other words, of all maximal matching blocks, return one that starts earliest in *a*, and of all those maximal matching blocks that start earliest in *a*, return the one that starts earliest in *b*.

```
>>> s = SequenceMatcher(None, " abcd", "abcd abcd")
>>> s.find_longest_match(0, 5, 0, 9)
Match(a=0, b=4, size=5)
```

If *isjunk* was provided, first the longest matching block is determined as above, but with the additional restriction that no junk element appears in the block. Then that block is extended as far as possible by matching (only) junk elements on both sides. So the resulting block never matches on junk except as identical junk happens to be adjacent to an interesting match.

Here's the same example as before, but considering blanks to be junk. That prevents ' abcd' from matching the ' abcd' at the tail end of the second sequence directly. Instead only the 'abcd' can match, and matches the leftmost 'abcd' in the second sequence:

```
>>> s = SequenceMatcher(lambda x: x==" ", " abcd", "abcd abcd")
>>> s.find_longest_match(0, 5, 0, 9)
Match(a=1, b=0, size=4)
```

If no blocks match, this returns (*alo*, *blo*, 0).

This method returns a *named tuple* *Match*(*a*, *b*, *size*).

get_matching_blocks ()

Return list of triples describing non-overlapping matching subsequences. Each triple is of the form (*i*, *j*, *n*), and means that *a*[*i*:*i*+*n*] == *b*[*j*:*j*+*n*]. The triples are monotonically increasing in *i* and *j*.

The last triple is a dummy, and has the value (*len*(*a*), *len*(*b*), 0). It is the only triple with *n* == 0. If (*i*, *j*, *n*) and (*i'*, *j'*, *n'*) are adjacent triples in the list, and the second is not the last triple in the list, then *i*+*n* < *i'* or *j*+*n* < *j'*; in other words, adjacent triples always describe non-adjacent equal blocks.

```
>>> s = SequenceMatcher(None, "abxcd", "abcd")
>>> s.get_matching_blocks()
[Match(a=0, b=0, size=2), Match(a=3, b=2, size=2), Match(a=5, b=4, size=0)]
```

get_opcodes ()

Return list of 5-tuples describing how to turn *a* into *b*. Each tuple is of the form (*tag*, *i1*, *i2*, *j1*, *j2*). The first tuple has *i1* == *j1* == 0, and remaining tuples have *i1* equal to the *i2* from the preceding tuple, and, likewise, *j1* equal to the previous *j2*.

The *tag* values are strings, with these meanings:

值	意义
'replace'	<code>a[i1:i2]</code> should be replaced by <code>b[j1:j2]</code> .
'delete'	<code>a[i1:i2]</code> should be deleted. Note that <code>j1 == j2</code> in this case.
'insert'	<code>b[j1:j2]</code> should be inserted at <code>a[i1:i1]</code> . Note that <code>i1 == i2</code> in this case.
'equal'	<code>a[i1:i2] == b[j1:j2]</code> (the sub-sequences are equal).

例如:

```
>>> a = "qabxcd"
>>> b = "abycdf"
>>> s = SequenceMatcher(None, a, b)
>>> for tag, i1, i2, j1, j2 in s.get_opcodes():
...     print('{:7}   a[{}:{}] --> b[{}:{}] {!r:>8} --> {!r}'.format(
...         tag, i1, i2, j1, j2, a[i1:i2], b[j1:j2]))
delete    a[0:1] --> b[0:0]      'q' --> ''
equal     a[1:3] --> b[0:2]      'ab' --> 'ab'
replace   a[3:4] --> b[2:3]      'x' --> 'y'
equal     a[4:6] --> b[3:5]      'cd' --> 'cd'
insert    a[6:6] --> b[5:6]      '' --> 'f'
```

get_grouped_opcodes (n=3)

Return a *generator* of groups with up to *n* lines of context.

Starting with the groups returned by `get_opcodes()`, this method splits out smaller change clusters and eliminates intervening ranges which have no changes.

The groups are returned in the same format as `get_opcodes()`.

ratio ()

Return a measure of the sequences' similarity as a float in the range [0, 1].

Where T is the total number of elements in both sequences, and M is the number of matches, this is $2.0 * M / T$. Note that this is 1.0 if the sequences are identical, and 0.0 if they have nothing in common.

This is expensive to compute if `get_matching_blocks()` or `get_opcodes()` hasn't already been called, in which case you may want to try `quick_ratio()` or `real_quick_ratio()` first to get an upper bound.

quick_ratio ()

Return an upper bound on `ratio()` relatively quickly.

real_quick_ratio ()

Return an upper bound on `ratio()` very quickly.

The three methods that return the ratio of matching to total characters can give different results due to differing levels of approximation, although `quick_ratio()` and `real_quick_ratio()` are always at least as large as `ratio()`:

```
>>> s = SequenceMatcher(None, "abcd", "bcde")
>>> s.ratio()
0.75
>>> s.quick_ratio()
0.75
>>> s.real_quick_ratio()
1.0
```

6.3.2 序列匹配的示例

This example compares two strings, considering blanks to be “junk” :

```
>>> s = SequenceMatcher(lambda x: x == " ",
...                       "private Thread currentThread;",
...                       "private volatile Thread currentThread;")
```

`ratio()` returns a float in $[0, 1]$, measuring the similarity of the sequences. As a rule of thumb, a `ratio()` value over 0.6 means the sequences are close matches:

```
>>> print(round(s.ratio(), 3))
0.866
```

If you're only interested in where the sequences match, `get_matching_blocks()` is handy:

```
>>> for block in s.get_matching_blocks():
...     print("a[%d] and b[%d] match for %d elements" % block)
a[0] and b[0] match for 8 elements
a[8] and b[17] match for 21 elements
a[29] and b[38] match for 0 elements
```

Note that the last tuple returned by `get_matching_blocks()` is always a dummy, `(len(a), len(b), 0)`, and this is the only case in which the last tuple element (number of elements matched) is 0.

If you want to know how to change the first sequence into the second, use `get_opcodes()`:

```
>>> for opcode in s.get_opcodes():
...     print("%6s a[%d:%d] b[%d:%d]" % opcode)
equal a[0:8] b[0:8]
insert a[8:8] b[8:17]
equal a[8:29] b[17:38]
```

参见:

- The `get_close_matches()` function in this module which shows how simple code building on `SequenceMatcher` can be used to do useful work.
- [Simple version control recipe](#) for a small application built with `SequenceMatcher`.

6.3.3 Differ Objects

Note that `Differ`-generated deltas make no claim to be **minimal** diffs. To the contrary, minimal diffs are often counter-intuitive, because they synch up anywhere possible, sometimes accidental matches 100 pages apart. Restricting synch points to contiguous matches preserves some notion of locality, at the occasional cost of producing a longer diff.

The `Differ` class has this constructor:

```
class difflib.Differ (linejunk=None, charjunk=None)
```

Optional keyword parameters `linejunk` and `charjunk` are for filter functions (or `None`):

linejunk: A function that accepts a single string argument, and returns true if the string is junk. The default is `None`, meaning that no line is considered junk.

charjunk: A function that accepts a single character argument (a string of length 1), and returns true if the character is junk. The default is `None`, meaning that no character is considered junk.

These junk-filtering functions speed up matching to find differences and do not cause any differing lines or characters to be ignored. Read the description of the `find_longest_match()` method's `isjunk` parameter for an explanation.

`Differ` objects are used (deltas generated) via a single method:

compare (*a*, *b*)

Compare two sequences of lines, and generate the delta (a sequence of lines).

Each sequence must contain individual single-line strings ending with newlines. Such sequences can be obtained from the `readlines()` method of file-like objects. The delta generated also consists of newline-terminated strings, ready to be printed as-is via the `writelines()` method of a file-like object.

6.3.4 Differ Example

This example compares two texts. First we set up the texts, sequences of individual single-line strings ending with newlines (such sequences can also be obtained from the `readlines()` method of file-like objects):

```
>>> text1 = ''' 1. Beautiful is better than ugly.
... 2. Explicit is better than implicit.
... 3. Simple is better than complex.
... 4. Complex is better than complicated.
... '''.splitlines(keepends=True)
>>> len(text1)
4
>>> text1[0][-1]
'\n'
>>> text2 = ''' 1. Beautiful is better than ugly.
... 3. Simple is better than complex.
... 4. Complicated is better than complex.
... 5. Flat is better than nested.
... '''.splitlines(keepends=True)
```

Next we instantiate a `Differ` object:

```
>>> d = Differ()
```

Note that when instantiating a `Differ` object we may pass functions to filter out line and character “junk.” See the `Differ()` constructor for details.

Finally, we compare the two:

```
>>> result = list(d.compare(text1, text2))
```

`result` is a list of strings, so let's pretty-print it:

```
>>> from pprint import pprint
>>> pprint(result)
[' 1. Beautiful is better than ugly.\n',
'- 2. Explicit is better than implicit.\n',
'- 3. Simple is better than complex.\n',
'+ 3. Simple is better than complex.\n',
'? ++\n',
'- 4. Complex is better than complicated.\n',
'? ^ ---- ^\n',
'+ 4. Complicated is better than complex.\n',
'? ++++ ^ ^\n',
'+ 5. Flat is better than nested.\n']
```

As a single multi-line string it looks like this:

```
>>> import sys
>>> sys.stdout.writelines(result)
```

(下页继续)

(续上页)

```

1. Beautiful is better than ugly.
- 2. Explicit is better than implicit.
- 3. Simple is better than complex.
+ 3.   Simple is better than complex.
?    ++
- 4. Complex is better than complicated.
?      ^          ---- ^
+ 4. Complicated is better than complex.
?      ++++ ^          ^
+ 5. Flat is better than nested.

```

6.3.5 A command-line interface to difflib

This example shows how to use difflib to create a diff-like utility. It is also contained in the Python source distribution, as Tools/scripts/diff.py.

```

#!/usr/bin/env python3
""" Command line interface to difflib.py providing diffs in four formats:

* ndiff:    lists every line and highlights interline changes.
* context:  highlights clusters of changes in a before/after format.
* unified:  highlights clusters of changes in an inline format.
* html:     generates side by side comparison with change highlights.

"""

import sys, os, difflib, argparse
from datetime import datetime, timezone

def file_mtime(path):
    t = datetime.fromtimestamp(os.stat(path).st_mtime,
                              timezone.utc)
    return t.astimezone().isoformat()

def main():

    parser = argparse.ArgumentParser()
    parser.add_argument('-c', action='store_true', default=False,
                        help='Produce a context format diff (default)')
    parser.add_argument('-u', action='store_true', default=False,
                        help='Produce a unified format diff')
    parser.add_argument('-m', action='store_true', default=False,
                        help='Produce HTML side by side diff '
                              '(can use -c and -l in conjunction)')
    parser.add_argument('-n', action='store_true', default=False,
                        help='Produce a ndiff format diff')
    parser.add_argument('-l', '--lines', type=int, default=3,
                        help='Set number of context lines (default 3)')
    parser.add_argument('fromfile')
    parser.add_argument('tofile')
    options = parser.parse_args()

    n = options.lines
    fromfile = options.fromfile
    tofile = options.tofile

```

(下页继续)

(续上页)

```

fromdate = file_mtime(fromfile)
todate = file_mtime(tofile)
with open(fromfile) as ff:
    fromlines = ff.readlines()
with open(tofile) as tf:
    tolines = tf.readlines()

if options.u:
    diff = difflib.unified_diff(fromlines, tolines, fromfile, tofile, fromdate,
→todate, n=n)
elif options.n:
    diff = difflib.ndiff(fromlines, tolines)
elif options.m:
    diff = difflib.HtmlDiff().make_file(fromlines, tolines, fromfile, tofile,
→context=options.c, numlines=n)
else:
    diff = difflib.context_diff(fromlines, tolines, fromfile, tofile, fromdate,
→todate, n=n)

sys.stdout.writelines(diff)

if __name__ == '__main__':
    main()

```

6.4 textwrap —Text wrapping and filling

Source code: [Lib/textwrap.py](#)

The *textwrap* module provides some convenience functions, as well as *TextWrapper*, the class that does all the work. If you're just wrapping or filling one or two text strings, the convenience functions should be good enough; otherwise, you should use an instance of *TextWrapper* for efficiency.

`textwrap.wrap(text, width=70, **kwargs)`

Wraps the single paragraph in *text* (a string) so every line is at most *width* characters long. Returns a list of output lines, without final newlines.

Optional keyword arguments correspond to the instance attributes of *TextWrapper*, documented below. *width* defaults to 70.

See the *TextWrapper.wrap()* method for additional details on how *wrap()* behaves.

`textwrap.fill(text, width=70, **kwargs)`

Wraps the single paragraph in *text*, and returns a single string containing the wrapped paragraph. *fill()* is shorthand for

```
"\n".join(wrap(text, ...))
```

In particular, *fill()* accepts exactly the same keyword arguments as *wrap()*.

`textwrap.shorten(text, width, **kwargs)`

Collapse and truncate the given *text* to fit in the given *width*.

First the whitespace in *text* is collapsed (all whitespace is replaced by single spaces). If the result fits in the *width*, it is returned. Otherwise, enough words are dropped from the end so that the remaining words plus the *placeholder* fit within *width*:

```
>>> textwrap.shorten("Hello world!", width=12)
'Hello world!'
>>> textwrap.shorten("Hello world!", width=11)
'Hello [...]'
>>> textwrap.shorten("Hello world", width=10, placeholder="...")
'Hello...'
```

Optional keyword arguments correspond to the instance attributes of *TextWrapper*, documented below. Note that the whitespace is collapsed before the text is passed to the *TextWrapper.fill()* function, so changing the value of *tabsize*, *expand_tabs*, *drop_whitespace*, and *replace_whitespace* will have no effect.

3.4 新版功能.

`textwrap.dedent` (*text*)

Remove any common leading whitespace from every line in *text*.

This can be used to make triple-quoted strings line up with the left edge of the display, while still presenting them in the source code in indented form.

Note that tabs and spaces are both treated as whitespace, but they are not equal: the lines " hello" and "\thello" are considered to have no common leading whitespace.

例如:

```
def test():
    # end first line with \ to avoid the empty line!
    s = '''\
    hello
        world
    '''
    print(repr(s))          # prints '    hello\n        world\n    '
    print(repr(dedent(s)))  # prints 'hello\n    world\n'
```

`textwrap.indent` (*text*, *prefix*, *predicate=None*)

Add *prefix* to the beginning of selected lines in *text*.

Lines are separated by calling `text.splitlines(True)`.

By default, *prefix* is added to all lines that do not consist solely of whitespace (including any line endings).

例如:

```
>>> s = 'hello\n\n \nworld'
>>> indent(s, ' ')
'  hello\n\n \n world'
```

The optional *predicate* argument can be used to control which lines are indented. For example, it is easy to add *prefix* to even empty and whitespace-only lines:

```
>>> print(indent(s, '+ ', lambda line: True))
+ hello
+
+
+ world
```

3.3 新版功能.

`wrap()`, `fill()` and `shorten()` work by creating a `TextWrapper` instance and calling a single method on it. That instance is not reused, so for applications that process many text strings using `wrap()` and/or `fill()`, it may be more efficient to create your own `TextWrapper` object.

Text is preferably wrapped on whitespaces and right after the hyphens in hyphenated words; only then will long words be broken if necessary, unless `TextWrapper.break_long_words` is set to false.

class `textwrap.TextWrapper` (***kwargs*)

The `TextWrapper` constructor accepts a number of optional keyword arguments. Each keyword argument corresponds to an instance attribute, so for example

```
wrapper = TextWrapper(initial_indent="* ")
```

is the same as

```
wrapper = TextWrapper()
wrapper.initial_indent = "* "
```

You can re-use the same `TextWrapper` object many times, and you can change any of its options through direct assignment to instance attributes between uses.

The `TextWrapper` instance attributes (and keyword arguments to the constructor) are as follows:

width

(default: 70) The maximum length of wrapped lines. As long as there are no individual words in the input text longer than `width`, `TextWrapper` guarantees that no output line will be longer than `width` characters.

expand_tabs

(default: True) If true, then all tab characters in `text` will be expanded to spaces using the `expandtabs()` method of `text`.

tabsize

(default: 8) If `expand_tabs` is true, then all tab characters in `text` will be expanded to zero or more spaces, depending on the current column and the given tab size.

3.3 新版功能.

replace_whitespace

(default: True) If true, after tab expansion but before wrapping, the `wrap()` method will replace each whitespace character with a single space. The whitespace characters replaced are as follows: tab, newline, vertical tab, formfeed, and carriage return (`'\t\n\v\f\r'`).

注解: If `expand_tabs` is false and `replace_whitespace` is true, each tab character will be replaced by a single space, which is *not* the same as tab expansion.

注解: If `replace_whitespace` is false, newlines may appear in the middle of a line and cause strange output. For this reason, text should be split into paragraphs (using `str.splitlines()` or similar) which are wrapped separately.

drop_whitespace

(default: True) If true, whitespace at the beginning and ending of every line (after wrapping but before indenting) is dropped. Whitespace at the beginning of the paragraph, however, is not dropped if non-whitespace follows it. If whitespace being dropped takes up an entire line, the whole line is dropped.

initial_indent

(default: ' ') String that will be prepended to the first line of wrapped output. Counts towards the length of the first line. The empty string is not indented.

subsequent_indent

(default: ' ') String that will be prepended to all lines of wrapped output except the first. Counts towards the length of each line except the first.

fix_sentence_endings

(default: False) If true, *TextWrapper* attempts to detect sentence endings and ensure that sentences are always separated by exactly two spaces. This is generally desired for text in a monospaced font. However, the sentence detection algorithm is imperfect: it assumes that a sentence ending consists of a lowercase letter followed by one of '.', '!', or '?', possibly followed by one of '"' or "'", followed by a space. One problem with this algorithm is that it is unable to detect the difference between “Dr.” in

```
[...] Dr. Frankenstein's monster [...]
```

and “Spot.” in

```
[...] See Spot. See Spot run [...]
```

fix_sentence_endings is false by default.

Since the sentence detection algorithm relies on `string.lowercase` for the definition of “lowercase letter,” and a convention of using two spaces after a period to separate sentences on the same line, it is specific to English-language texts.

break_long_words

(default: True) If true, then words longer than *width* will be broken in order to ensure that no lines are longer than *width*. If it is false, long words will not be broken, and some lines may be longer than *width*. (Long words will be put on a line by themselves, in order to minimize the amount by which *width* is exceeded.)

break_on_hyphens

(default: True) If true, wrapping will occur preferably on whitespaces and right after hyphens in compound words, as it is customary in English. If false, only whitespaces will be considered as potentially good places for line breaks, but you need to set *break_long_words* to false if you want truly insecable words. Default behaviour in previous versions was to always allow breaking hyphenated words.

max_lines

(default: None) If not None, then the output will contain at most *max_lines* lines, with *placeholder* appearing at the end of the output.

3.4 新版功能.

placeholder

(default: ' [...] ') String that will appear at the end of the output text if it has been truncated.

3.4 新版功能.

TextWrapper also provides some public methods, analogous to the module-level convenience functions:

wrap (*text*)

Wraps the single paragraph in *text* (a string) so every line is at most *width* characters long. All wrapping options are taken from instance attributes of the *TextWrapper* instance. Returns a list of output lines, without final newlines. If the wrapped output has no content, the returned list is empty.

fill (*text*)

Wraps the single paragraph in *text*, and returns a single string containing the wrapped paragraph.

6.5 unicodedata —Unicode 数据库

This module provides access to the Unicode Character Database (UCD) which defines character properties for all Unicode characters. The data contained in this database is compiled from the [UCD version 9.0.0](#).

该模块使用与 Unicode 标准附件 #44，“Unicode 字符数据库”定义的相同名称和符号。它定义了以下功能：

`unicodedata.lookup(name)`

按名称查找字符。如果找到具有给定名称的字符，则返回相应的字符。如果没有找到，则 `KeyError` 被引发。

在 3.3 版更改：已添加对名称别名¹ 和命名序列² 的支持。

`unicodedata.name(chr[, default])`

返回分配给字符 `chr` 的名称作为字符串。如果没有定义名称，则返回 `default`，如果没有给出，则 `ValueError` 被引发。

`unicodedata.decimal(chr[, default])`

返回分配给字符 `chr` 的十进制值作为整数。如果没有定义这样的值，则返回 `default`，如果没有给出，则 `ValueError` 被引发。

`unicodedata.digit(chr[, default])`

返回分配给字符 `chr` 的数字值作为整数。如果没有定义这样的值，则返回 `default`，如果没有给出，则 `ValueError` 被引发。

`unicodedata.numeric(chr[, default])`

返回分配给字符 `chr` 的数值作为浮点数。如果没有定义这样的值，则返回 `default`，如果没有给出，则 `ValueError` 被引发。

`unicodedata.category(chr)`

返回分配给字符 `chr` 的常规类别为字符串。

`unicodedata.bidirectional(chr)`

返回分配给字符 `chr` 的双向类作为字符串。如果未定义此类值，则返回空字符串。

`unicodedata.combining(chr)`

返回分配给字符 `chr` 的规范组合类作为整数。如果没有定义组合类，则返回 0。

`unicodedata.east_asian_width(chr)`

返回分配给字符 `chr` 的东亚宽度作为字符串。

`unicodedata.mirrored(chr)`

返回分配给字符 `chr` 的镜像属性为整数。如果字符在双向文本中被识别为“镜像”字符，则返回 1，否则返回 0。

`unicodedata.decomposition(chr)`

返回分配给字符 `chr` 的字符分解映射作为字符串。如果未定义此类映射，则返回空字符串。

`unicodedata.normalize(form, unistr)`

返回 Unicode 字符串 `unistr` 的正常形式 `form`。`form` 的有效值为 ‘NFC’、‘NFKC’、‘NFD’ 和 ‘NFKD’。

Unicode 标准基于规范等价和兼容性等效的定义定义了 Unicode 字符串的各种规范化形式。在 Unicode 中，可以以各种方式表示多个字符。例如，字符 U+00C7（带有 CEDILLA 的 LATIN CAPITAL LETTER C）也可以表示为序列 U+0043（LATIN CAPITAL LETTER C）U+0327（COMBINING CEDILLA）。

对于每个字符，有两种正规形式：正规形式 C 和正规形式 D。正规形式 D（NFD）也称为规范分解，并将每个字符转换为其分解形式。正规形式 C（NFC）首先应用规范分解，然后再次组合预组合字符。

除了这两种形式之外，还有两种基于兼容性等效的其他常规形式。在 Unicode 中，支持某些字符，这些字符通常与其他字符统一。例如，U+2160（ROMAN NUMERAL ONE）与 U+0049（LATIN CAPITAL LETTER I）完全相同。但是，Unicode 支持它与现有字符集（例如 gb2312）的兼容性。

¹ <http://www.unicode.org/Public/9.0.0/ucd/NameAliases.txt>

² <http://www.unicode.org/Public/9.0.0/ucd/NamedSequences.txt>

正规形式 KD (NFKD) 将应用兼容性分解, 即用其等价项替换所有兼容性字符。正规形式 KC (NFKC) 首先应用兼容性分解, 然后是规范组合。

即使两个 `unicode` 字符串被规范化并且人类读者看起来相同, 如果一个具有组合字符而另一个没有, 则它们可能无法相等。

此外, 该模块暴露了以下常量:

`unicodedata.unidata_version`

此模块中使用的 Unicode 数据库的版本。

`unicodedata.ucd_3_2_0`

这是一个与整个模块具有相同方法的对象, 但对于需要此特定版本的 Unicode 数据库 (如 IDNA) 的应用程序, 则使用 Unicode 数据库版本 3.2。

示例:

```
>>> import unicodedata
>>> unicodedata.lookup('LEFT CURLY BRACKET')
'{'
>>> unicodedata.name('/')
'SOLIDUS'
>>> unicodedata.decimal('9')
9
>>> unicodedata.decimal('a')
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
ValueError: not a decimal
>>> unicodedata.category('A') # 'L'etter, 'u'ppercase
'Lu'
>>> unicodedata.bidirectional('\u0660') # 'A'rabic, 'N'umber
'AN'
```

6.6 stringprep — Internet String Preparation

Source code: [Lib/stringprep.py](#)

When identifying things (such as host names) in the internet, it is often necessary to compare such identifications for “equality”. Exactly how this comparison is executed may depend on the application domain, e.g. whether it should be case-insensitive or not. It may be also necessary to restrict the possible identifications, to allow only identifications consisting of “printable” characters.

RFC 3454 defines a procedure for “preparing” Unicode strings in internet protocols. Before passing strings onto the wire, they are processed with the preparation procedure, after which they have a certain normalized form. The RFC defines a set of tables, which can be combined into profiles. Each profile must define which tables it uses, and what other optional parts of the `stringprep` procedure are part of the profile. One example of a `stringprep` profile is `nameprep`, which is used for internationalized domain names.

The module `stringprep` only exposes the tables from **RFC 3454**. As these tables would be very large to represent them as dictionaries or lists, the module uses the Unicode character database internally. The module source code itself was generated using the `mkstringprep.py` utility.

As a result, these tables are exposed as functions, not as data structures. There are two kinds of tables in the RFC: sets and mappings. For a set, `stringprep` provides the “characteristic function”, i.e. a function that returns true if the parameter is part of the set. For mappings, it provides the mapping function: given the key, it returns the associated value. Below is a list of all functions available in the module.

`stringprep.in_table_a1 (code)`
 Determine whether *code* is in tableA.1 (Unassigned code points in Unicode 3.2).

`stringprep.in_table_b1 (code)`
 Determine whether *code* is in tableB.1 (Commonly mapped to nothing).

`stringprep.map_table_b2 (code)`
 Return the mapped value for *code* according to tableB.2 (Mapping for case-folding used with NFKC).

`stringprep.map_table_b3 (code)`
 Return the mapped value for *code* according to tableB.3 (Mapping for case-folding used with no normalization).

`stringprep.in_table_c11 (code)`
 Determine whether *code* is in tableC.1.1 (ASCII space characters).

`stringprep.in_table_c12 (code)`
 Determine whether *code* is in tableC.1.2 (Non-ASCII space characters).

`stringprep.in_table_c11_c12 (code)`
 Determine whether *code* is in tableC.1 (Space characters, union of C.1.1 and C.1.2).

`stringprep.in_table_c21 (code)`
 Determine whether *code* is in tableC.2.1 (ASCII control characters).

`stringprep.in_table_c22 (code)`
 Determine whether *code* is in tableC.2.2 (Non-ASCII control characters).

`stringprep.in_table_c21_c22 (code)`
 Determine whether *code* is in tableC.2 (Control characters, union of C.2.1 and C.2.2).

`stringprep.in_table_c3 (code)`
 Determine whether *code* is in tableC.3 (Private use).

`stringprep.in_table_c4 (code)`
 Determine whether *code* is in tableC.4 (Non-character code points).

`stringprep.in_table_c5 (code)`
 Determine whether *code* is in tableC.5 (Surrogate codes).

`stringprep.in_table_c6 (code)`
 Determine whether *code* is in tableC.6 (Inappropriate for plain text).

`stringprep.in_table_c7 (code)`
 Determine whether *code* is in tableC.7 (Inappropriate for canonical representation).

`stringprep.in_table_c8 (code)`
 Determine whether *code* is in tableC.8 (Change display properties or are deprecated).

`stringprep.in_table_c9 (code)`
 Determine whether *code* is in tableC.9 (Tagging characters).

`stringprep.in_table_d1 (code)`
 Determine whether *code* is in tableD.1 (Characters with bidirectional property “R” or “AL”).

`stringprep.in_table_d2 (code)`
 Determine whether *code* is in tableD.2 (Characters with bidirectional property “L”).

6.7 readline —GNU readline interface

The `readline` module defines a number of functions to facilitate completion and reading/writing of history files from the Python interpreter. This module can be used directly, or via the `rlcompleter` module, which supports completion of Python identifiers at the interactive prompt. Settings made using this module affect the behaviour of both the interpreter's interactive prompt and the prompts offered by the built-in `input()` function.

注解: The underlying Readline library API may be implemented by the `libedit` library instead of GNU readline. On MacOS X the `readline` module detects which library is being used at run time.

The configuration file for `libedit` is different from that of GNU readline. If you programmatically load configuration strings you can check for the text “libedit” in `readline.__doc__` to differentiate between GNU readline and libedit.

Readline keybindings may be configured via an initialization file, typically `.inputrc` in your home directory. See [Readline Init File](#) in the GNU Readline manual for information about the format and allowable constructs of that file, and the capabilities of the Readline library in general.

6.7.1 Init file

The following functions relate to the init file and user configuration:

`readline.parse_and_bind(string)`
Execute the init line provided in the *string* argument. This calls `rl_parse_and_bind()` in the underlying library.

`readline.read_init_file([filename])`
Execute a readline initialization file. The default filename is the last filename used. This calls `rl_read_init_file()` in the underlying library.

6.7.2 Line buffer

The following functions operate on the line buffer:

`readline.get_line_buffer()`
Return the current contents of the line buffer (`rl_line_buffer` in the underlying library).

`readline.insert_text(string)`
Insert text into the line buffer at the cursor position. This calls `rl_insert_text()` in the underlying library, but ignores the return value.

`readline.redisplay()`
Change what's displayed on the screen to reflect the current contents of the line buffer. This calls `rl_redisplay()` in the underlying library.

6.7.3 History file

The following functions operate on a history file:

`readline.read_history_file([filename])`
Load a readline history file, and append it to the history list. The default filename is `~/.history`. This calls `read_history()` in the underlying library.

`readline.write_history_file([filename])`
Save the history list to a readline history file, overwriting any existing file. The default filename is `~/.history`. This calls `write_history()` in the underlying library.

`readline.append_history_file(nelements[, filename])`

Append the last *nelements* items of history to a file. The default filename is `~/.history`. The file must already exist. This calls `append_history()` in the underlying library. This function only exists if Python was compiled for a version of the library that supports it.

3.5 新版功能.

`readline.get_history_length()`

`readline.set_history_length(length)`

Set or return the desired number of lines to save in the history file. The `write_history_file()` function uses this value to truncate the history file, by calling `history_truncate_file()` in the underlying library. Negative values imply unlimited history file size.

6.7.4 History list

The following functions operate on a global history list:

`readline.clear_history()`

Clear the current history. This calls `clear_history()` in the underlying library. The Python function only exists if Python was compiled for a version of the library that supports it.

`readline.get_current_history_length()`

Return the number of items currently in the history. (This is different from `get_history_length()`, which returns the maximum number of lines that will be written to a history file.)

`readline.get_history_item(index)`

Return the current contents of history item at *index*. The item index is one-based. This calls `history_get()` in the underlying library.

`readline.remove_history_item(pos)`

Remove history item specified by its position from the history. The position is zero-based. This calls `remove_history()` in the underlying library.

`readline.replace_history_item(pos, line)`

Replace history item specified by its position with *line*. The position is zero-based. This calls `replace_history_entry()` in the underlying library.

`readline.add_history(line)`

Append *line* to the history buffer, as if it was the last line typed. This calls `add_history()` in the underlying library.

`readline.set_auto_history(enabled)`

Enable or disable automatic calls to `add_history()` when reading input via `readline`. The *enabled* argument should be a Boolean value that when true, enables auto history, and that when false, disables auto history.

3.6 新版功能.

CPython implementation detail: Auto history is enabled by default, and changes to this do not persist across multiple sessions.

6.7.5 Startup hooks

`readline.set_startup_hook([function])`

Set or remove the function invoked by the `rl_startup_hook` callback of the underlying library. If *function* is specified, it will be used as the new hook function; if omitted or `None`, any function already installed is removed. The hook is called with no arguments just before `readline` prints the first prompt.

`readline.set_pre_input_hook([function])`

Set or remove the function invoked by the `rl_pre_input_hook` callback of the underlying library. If *function* is specified, it will be used as the new hook function; if omitted or `None`, any function already installed is removed. The hook is called with no arguments after the first prompt has been printed and just before `readline` starts reading input characters. This function only exists if Python was compiled for a version of the library that supports it.

6.7.6 Completion

The following functions relate to implementing a custom word completion function. This is typically operated by the Tab key, and can suggest and automatically complete a word being typed. By default, `Readline` is set up to be used by `rlcompleter` to complete Python identifiers for the interactive interpreter. If the `readline` module is to be used with a custom completer, a different set of word delimiters should be set.

`readline.set_completer([function])`

Set or remove the completer function. If *function* is specified, it will be used as the new completer function; if omitted or `None`, any completer function already installed is removed. The completer function is called as `function(text, state)`, for *state* in 0, 1, 2, ..., until it returns a non-string value. It should return the next possible completion starting with *text*.

The installed completer function is invoked by the `entry_func` callback passed to `rl_completion_matches()` in the underlying library. The *text* string comes from the first parameter to the `rl_attempted_completion_function` callback of the underlying library.

`readline.get_completer()`

Get the completer function, or `None` if no completer function has been set.

`readline.get_completion_type()`

Get the type of completion being attempted. This returns the `rl_completion_type` variable in the underlying library as an integer.

`readline.get_begidx()`

`readline.get_endidx()`

Get the beginning or ending index of the completion scope. These indexes are the *start* and *end* arguments passed to the `rl_attempted_completion_function` callback of the underlying library.

`readline.set_completer_delims(string)`

`readline.get_completer_delims()`

Set or get the word delimiters for completion. These determine the start of the word to be considered for completion (the completion scope). These functions access the `rl_completer_word_break_characters` variable in the underlying library.

`readline.set_completion_display_matches_hook([function])`

Set or remove the completion display function. If *function* is specified, it will be used as the new completion display function; if omitted or `None`, any completion display function already installed is removed. This sets or clears the `rl_completion_display_matches_hook` callback in the underlying library. The completion display function is called as `function(substitution, [matches], longest_match_length)` once each time matches need to be displayed.

6.7.7 示例

The following example demonstrates how to use the `readline` module's history reading and writing functions to automatically load and save a history file named `.python_history` from the user's home directory. The code below would normally be executed automatically during interactive sessions from the user's `PYTHONSTARTUP` file.

```

import atexit
import os
import readline

histfile = os.path.join(os.path.expanduser("~"), ".python_history")
try:
    readline.read_history_file(histfile)
    # default history len is -1 (infinite), which may grow unruly
    readline.set_history_length(1000)
except FileNotFoundError:
    pass

atexit.register(readline.write_history_file, histfile)

```

This code is actually automatically run when Python is run in interactive mode (see [Readline configuration](#)).

The following example achieves the same goal but supports concurrent interactive sessions, by only appending the new history.

```

import atexit
import os
import readline

histfile = os.path.join(os.path.expanduser("~"), ".python_history")

try:
    readline.read_history_file(histfile)
    h_len = readline.get_current_history_length()
except FileNotFoundError:
    open(histfile, 'wb').close()
    h_len = 0

def save(prev_h_len, histfile):
    new_h_len = readline.get_current_history_length()
    readline.set_history_length(1000)
    readline.append_history_file(new_h_len - prev_h_len, histfile)
atexit.register(save, h_len, histfile)

```

The following example extends the `code.InteractiveConsole` class to support history save/restore.

```

import atexit
import code
import os
import readline

class HistoryConsole(code.InteractiveConsole):
    def __init__(self, locals=None, filename="<console>",
                 histfile=os.path.expanduser("~/console-history")):
        code.InteractiveConsole.__init__(self, locals, filename)
        self.init_history(histfile)

    def init_history(self, histfile):
        readline.parse_and_bind("tab: complete")
        if hasattr(readline, "read_history_file"):
            try:
                readline.read_history_file(histfile)
            except FileNotFoundError:
                pass
        atexit.register(self.save_history, histfile)

```

(下页继续)

(续上页)

```
def save_history(self, histfile):
    readline.set_history_length(1000)
    readline.write_history_file(histfile)
```

6.8 rlcompleter — GNU readline 的完成函数

源代码: [Lib/rlcompleter.py](#)

`rlcompeleter` 通过补全有效的 Python 标识符和关键字为 `readline` 模块定义了一个补全功能套装。

当此模块在具有可用的 `readline` 模块的 Unix 平台被导入, 一个 `Completer` 实例将被自动创建并且它的 `complete()` 方法将设置为 `readline` 的补全器。

示例:

```
>>> import rlcompleter
>>> import readline
>>> readline.parse_and_bind("tab: complete")
>>> readline. <TAB PRESSED>
readline.__doc__          readline.get_line_buffer(  readline.read_init_file(
readline.__file__        readline.insert_text(      readline.set_completer(
readline.__name__        readline.parse_and_bind(
>>> readline.
```

`rlcompleter` 模块是为了使用 Python 的交互模式而设计的。除非 Python 是通过 `-S` 选项运行, 这个模块总是自动地被导入且配置 (参考 [Readline configuration](#))。

在没有 `readline` 的平台, 此模块定义的 `Completer` 类仍然可以用于自定义行为。

6.8.1 Completer 对象

`Completer` 对象具有以下方法:

`Completer.complete(text, state)`

为 `text` 返回 `stateth` 补全。

如果要求 `text` 不包含句点字符 ('.'), 它将从当前定义的名称中完成 `__main__`, `builtins` 和 `keywords` (就和通过 `keyword` 模块定义的一样)。

如果为加点的名称执行调用, 它将尝试尽量求值直到最后一部分为止而不产生附带影响 (函数不会被求值, 但它生成对 `__getattr__()` 的调用), 并通过 `dir()` 函数来匹配剩余部分。在对表达式求值期间引发的任何异常都会被捕获、静默处理并返回 `None`。

二进制数据服务

本章介绍的模块提供了一些操作二进制数据的基本服务操作。有关二进制数据的其他操作，特别是与文件格式和网络协议有关的操作，将在相关章节中介绍。

下面描述的一些库文本处理服务也可以使用 ASCII 兼容的二进制格式（例如 `re`）或所有二进制数据（例如 `difflib`）。

另外，请参阅 Python 的内置二进制数据类型的文档二进制序列类型—`bytes`, `bytearray`, `memoryview`。

7.1 struct — Interpret bytes as packed binary data

源代码： [Lib/struct.py](#)

This module performs conversions between Python values and C structs represented as Python `bytes` objects. This can be used in handling binary data stored in files or from network connections, among other sources. It uses *Format Strings* as compact descriptions of the layout of the C structs and the intended conversion to/from Python values.

注解： By default, the result of packing a given C struct includes pad bytes in order to maintain proper alignment for the C types involved; similarly, alignment is taken into account when unpacking. This behavior is chosen so that the bytes of a packed struct correspond exactly to the layout in memory of the corresponding C struct. To handle platform-independent data formats or omit implicit pad bytes, use `standard` size and alignment instead of `native` size and alignment: see 字节顺序, 大小和对齐方式 for details.

Several `struct` functions (and methods of `Struct`) take a *buffer* argument. This refers to objects that implement the bufferobjects and provide either a readable or read-writable buffer. The most common types used for that purpose are `bytes` and `bytearray`, but many other types that can be viewed as an array of bytes implement the buffer protocol, so that they can be read/filled without additional copying from a `bytes` object.

7.1.1 Functions and Exceptions

The module defines the following exception and functions:

exception `struct.error`

Exception raised on various occasions; argument is a string describing what is wrong.

`struct.pack(fmt, v1, v2, ...)`

Return a bytes object containing the values *v1*, *v2*, ... packed according to the format string *fmt*. The arguments must match the values required by the format exactly.

`struct.pack_into(fmt, buffer, offset, v1, v2, ...)`

Pack the values *v1*, *v2*, ... according to the format string *fmt* and write the packed bytes into the writable buffer *buffer* starting at position *offset*. Note that *offset* is a required argument.

`struct.unpack(fmt, buffer)`

Unpack from the buffer *buffer* (presumably packed by `pack(fmt, ...)`) according to the format string *fmt*. The result is a tuple even if it contains exactly one item. The buffer's size in bytes must match the size required by the format, as reflected by `calcsize()`.

`struct.unpack_from(fmt, buffer, offset=0)`

Unpack from *buffer* starting at position *offset*, according to the format string *fmt*. The result is a tuple even if it contains exactly one item. The buffer's size in bytes, minus *offset*, must be at least the size required by the format, as reflected by `calcsize()`.

`struct.iter_unpack(fmt, buffer)`

Iteratively unpack from the buffer *buffer* according to the format string *fmt*. This function returns an iterator which will read equally-sized chunks from the buffer until all its contents have been consumed. The buffer's size in bytes must be a multiple of the size required by the format, as reflected by `calcsize()`.

Each iteration yields a tuple as specified by the format string.

3.4 新版功能.

`struct.calcsize(fmt)`

Return the size of the struct (and hence of the bytes object produced by `pack(fmt, ...)`) corresponding to the format string *fmt*.

7.1.2 Format Strings

Format strings are the mechanism used to specify the expected layout when packing and unpacking data. They are built up from 格式字符, which specify the type of data being packed/unpacked. In addition, there are special characters for controlling the 字节顺序, 大小和对齐方式.

字节顺序, 大小和对齐方式

默认情况下, C 类型以机器的本机格式和字节顺序表示, 并在必要时通过跳过填充字节进行正确对齐 (根据 C 编译器使用的规则)。

或者, 根据下表, 格式字符串的第一个字符可用于指示打包数据的字节顺序, 大小和对齐方式:

字符	字节顺序	大小	对齐方式
@	按原字节	按原字节	按原字节
=	按原字节	标准	无
<	小端	标准	无
>	大端	标准	无
!	网络 (= 大端)	标准	无

如果第一个字符不是其中之一，则假定为 '@' 。

Native byte order is big-endian or little-endian, depending on the host system. For example, Intel x86 and AMD64 (x86-64) are little-endian; Motorola 68000 and PowerPC G5 are big-endian; ARM and Intel Itanium feature switchable endianness (bi-endian). Use `sys.byteorder` to check the endianness of your system.

Native size and alignment are determined using the C compiler's `sizeof` expression. This is always combined with native byte order.

Standard size depends only on the format character; see the table in the 格式字符 section.

Note the difference between '@' and '=': both use native byte order, but the size and alignment of the latter is standardized.

The form '!' is available for those poor souls who claim they can't remember whether network byte order is big-endian or little-endian.

There is no way to indicate non-native byte order (force byte-swapping); use the appropriate choice of '<' or '>'.

注释:

- (1) Padding is only automatically added between successive structure members. No padding is added at the beginning or the end of the encoded struct.
- (2) No padding is added when using non-native size and alignment, e.g. with '<', '>', '=', and '!'.
- (3) To align the end of a structure to the alignment requirement of a particular type, end the format with the code for that type with a repeat count of zero. See 示例.

格式字符

Format characters have the following meaning; the conversion between C and Python values should be obvious given their types. The 'Standard size' column refers to the size of the packed value in bytes when using standard size; that is, when the format string starts with one of '<', '>', '!' or '='. When using native size, the size of the packed value is platform-dependent.

格式	C 类型	Python 数据类型	标准大小	注释
x	pad byte	没有价值		
c	char	长度为 1 的字节	1	
b	signed char	整数	1	(1),(3)
B	unsigned char	整数	1	(3)
?	_Bool	布尔	1	(1)
h	short	整数	2	(3)
H	unsigned short	整数	2	(3)
i	int	整数	4	(3)
I	unsigned int	整数	4	(3)
l	long	整数	4	(3)
L	unsigned long	整数	4	(3)
q	long long	整数	8	(2), (3)
Q	unsigned long long	整数	8	(2), (3)
n	ssize_t	整数		(4)
N	size_t	整数		(4)
e	(7)	浮点数	2	(5)
f	float	浮点数	4	(5)
d	double	浮点数	8	(5)
s	char[]	bytes		
p	char[]	bytes		
P	void *	整数		(6)

在 3.3 版更改: 增加了对 'n' 和 'N' 格式的支持

在 3.6 版更改: 增加了对 'e' 格式的支持。

注释:

- (1) The '?' conversion code corresponds to the `_Bool` type defined by C99. If this type is not available, it is simulated using a `char`. In standard mode, it is always represented by one byte.
- (2) The 'q' and 'Q' conversion codes are available in native mode only if the platform C compiler supports C long long, or, on Windows, `__int64`. They are always available in standard modes.
- (3) When attempting to pack a non-integer using any of the integer conversion codes, if the non-integer has a `__index__()` method then that method is called to convert the argument to an integer before packing.
在 3.2 版更改: Use of the `__index__()` method for non-integers is new in 3.2.
- (4) The 'n' and 'N' conversion codes are only available for the native size (selected as the default or with the '@' byte order character). For the standard size, you can use whichever of the other integer formats fits your application.
- (5) For the 'f', 'd' and 'e' conversion codes, the packed representation uses the IEEE 754 binary32, binary64 or binary16 format (for 'f', 'd' or 'e' respectively), regardless of the floating-point format used by the platform.
- (6) The 'P' format character is only available for the native byte ordering (selected as the default or with the '@' byte order character). The byte order character '=' chooses to use little- or big-endian ordering based on the host system. The struct module does not interpret this as native ordering, so the 'P' format is not available.
- (7) The IEEE 754 binary16 “half precision” type was introduced in the 2008 revision of the [IEEE 754 standard](#). It has a sign bit, a 5-bit exponent and 11-bit precision (with 10 bits explicitly stored), and can represent numbers between approximately $6.1e-05$ and $6.5e+04$ at full precision. This type is not widely supported by C compilers: on a typical machine, an unsigned short can be used for storage, but not for math operations. See the Wikipedia page on the [half-precision floating-point format](#) for more information.

A format character may be preceded by an integral repeat count. For example, the format string '4h' means exactly the same as 'hhhh'.

Whitespace characters between formats are ignored; a count and its format must not contain whitespace though.

For the 's' format character, the count is interpreted as the length of the bytes, not a repeat count like for the other format characters; for example, '10s' means a single 10-byte string, while '10c' means 10 characters. If a count is not given, it defaults to 1. For packing, the string is truncated or padded with null bytes as appropriate to make it fit. For unpacking, the resulting bytes object always has exactly the specified number of bytes. As a special case, '0s' means a single, empty string (while '0c' means 0 characters).

When packing a value *x* using one of the integer formats ('b', 'B', 'h', 'H', 'i', 'I', 'l', 'L', 'q', 'Q'), if *x* is outside the valid range for that format then `struct.error` is raised.

在 3.1 版更改: In 3.0, some of the integer formats wrapped out-of-range values and raised `DeprecationWarning` instead of `struct.error`.

The 'p' format character encodes a “Pascal string”, meaning a short variable-length string stored in a *fixed number of bytes*, given by the count. The first byte stored is the length of the string, or 255, whichever is smaller. The bytes of the string follow. If the string passed in to `pack()` is too long (longer than the count minus 1), only the leading `count-1` bytes of the string are stored. If the string is shorter than `count-1`, it is padded with null bytes so that exactly `count` bytes in all are used. Note that for `unpack()`, the 'p' format character consumes `count` bytes, but that the string returned can never contain more than 255 bytes.

For the '?' format character, the return value is either `True` or `False`. When packing, the truth value of the argument object is used. Either 0 or 1 in the native or standard bool representation will be packed, and any non-zero value will be `True` when unpacking.

示例

注解: All examples assume a native byte order, size, and alignment with a big-endian machine.

A basic example of packing/unpacking three integers:

```
>>> from struct import *
>>> pack('hhl', 1, 2, 3)
b'\x00\x01\x00\x02\x00\x00\x00\x03'
>>> unpack('hhl', b'\x00\x01\x00\x02\x00\x00\x00\x03')
(1, 2, 3)
>>> calcsize('hhl')
8
```

Unpacked fields can be named by assigning them to variables or by wrapping the result in a named tuple:

```
>>> record = b'raymond \x32\x12\x08\x01\x08'
>>> name, serialnum, school, gradelevel = unpack('<10sHHb', record)

>>> from collections import namedtuple
>>> Student = namedtuple('Student', 'name serialnum school gradelevel')
>>> Student._make(unpack('<10sHHb', record))
Student(name=b'raymond ', serialnum=4658, school=264, gradelevel=8)
```

The ordering of format characters may have an impact on size since the padding needed to satisfy alignment requirements is different:

```
>>> pack('ci', b'*', 0x12131415)
b'*\x00\x00\x00\x12\x13\x14\x15'
>>> pack('ic', 0x12131415, b'*')
```

(下页继续)

(续上页)

```
b'\x12\x13\x14\x15*'
>>> calcsiz('ci')
8
>>> calcsiz('ic')
5
```

The following format 'llh0l' specifies two pad bytes at the end, assuming longs are aligned on 4-byte boundaries:

```
>>> pack('llh0l', 1, 2, 3)
b'\x00\x00\x00\x01\x00\x00\x02\x00\x03\x00\x00'
```

This only works when native size and alignment are in effect; standard size and alignment does not enforce any alignment.

参见:

Module `array` Packed binary storage of homogeneous data.

Module `xdrlib` Packing and unpacking of XDR data.

7.1.3 类

The `struct` module also defines the following type:

class `struct.Struct` (*format*)

Return a new Struct object which writes and reads binary data according to the format string *format*. Creating a Struct object once and calling its methods is more efficient than calling the `struct` functions with the same format since the format string only needs to be compiled once.

Compiled Struct objects support the following methods and attributes:

pack (*v1*, *v2*, ...)

Identical to the `pack()` function, using the compiled format. (`len(result)` will equal *size*.)

pack_into (*buffer*, *offset*, *v1*, *v2*, ...)

Identical to the `pack_into()` function, using the compiled format.

unpack (*buffer*)

Identical to the `unpack()` function, using the compiled format. The buffer's size in bytes must equal *size*.

unpack_from (*buffer*, *offset*=0)

Identical to the `unpack_from()` function, using the compiled format. The buffer's size in bytes, minus *offset*, must be at least *size*.

iter_unpack (*buffer*)

Identical to the `iter_unpack()` function, using the compiled format. The buffer's size in bytes must be a multiple of *size*.

3.4 新版功能.

format

The format string used to construct this Struct object.

size

The calculated size of the struct (and hence of the bytes object produced by the `pack()` method) corresponding to *format*.

7.2 codecs —Codec registry and base classes

Source code: [Lib/codecs.py](#)

This module defines base classes for standard Python codecs (encoders and decoders) and provides access to the internal Python codec registry, which manages the codec and error handling lookup process. Most standard codecs are *text encodings*, which encode text to bytes, but there are also codecs provided that encode text to text, and bytes to bytes. Custom codecs may encode and decode between arbitrary types, but some module features are restricted to use specifically with *text encodings*, or with codecs that encode to *bytes*.

The module defines the following functions for encoding and decoding with any codec:

`codecs.encode(obj, encoding='utf-8', errors='strict')`

Encodes *obj* using the codec registered for *encoding*.

Errors may be given to set the desired error handling scheme. The default error handler is 'strict' meaning that encoding errors raise *ValueError* (or a more codec specific subclass, such as *UnicodeEncodeError*). Refer to *Codec Base Classes* for more information on codec error handling.

`codecs.decode(obj, encoding='utf-8', errors='strict')`

Decodes *obj* using the codec registered for *encoding*.

Errors may be given to set the desired error handling scheme. The default error handler is 'strict' meaning that decoding errors raise *ValueError* (or a more codec specific subclass, such as *UnicodeDecodeError*). Refer to *Codec Base Classes* for more information on codec error handling.

The full details for each codec can also be looked up directly:

`codecs.lookup(encoding)`

Looks up the codec info in the Python codec registry and returns a *CodecInfo* object as defined below.

Encodings are first looked up in the registry's cache. If not found, the list of registered search functions is scanned. If no *CodecInfo* object is found, a *LookupError* is raised. Otherwise, the *CodecInfo* object is stored in the cache and returned to the caller.

class `codecs.CodecInfo` (*encode*, *decode*, *streamreader=None*, *streamwriter=None*, *incrementalencoder=None*, *incrementaldecoder=None*, *name=None*)

Codec details when looking up the codec registry. The constructor arguments are stored in attributes of the same name:

name

编码名称

encode

decode

The stateless encoding and decoding functions. These must be functions or methods which have the same interface as the *encode()* and *decode()* methods of Codec instances (see *Codec Interface*). The functions or methods are expected to work in a stateless mode.

incrementalencoder

incrementaldecoder

Incremental encoder and decoder classes or factory functions. These have to provide the interface defined by the base classes *IncrementalEncoder* and *IncrementalDecoder*, respectively. Incremental codecs can maintain state.

streamwriter

streamreader

Stream writer and reader classes or factory functions. These have to provide the interface defined by the base classes *StreamWriter* and *StreamReader*, respectively. Stream codecs can maintain state.

To simplify access to the various codec components, the module provides these additional functions which use `lookup()` for the codec lookup:

`codecs.getencoder(encoding)`

Look up the codec for the given encoding and return its encoder function.

Raises a `LookupError` in case the encoding cannot be found.

`codecs.getdecoder(encoding)`

Look up the codec for the given encoding and return its decoder function.

Raises a `LookupError` in case the encoding cannot be found.

`codecs.getincrementalencoder(encoding)`

Look up the codec for the given encoding and return its incremental encoder class or factory function.

Raises a `LookupError` in case the encoding cannot be found or the codec doesn't support an incremental encoder.

`codecs.getincrementaldecoder(encoding)`

Look up the codec for the given encoding and return its incremental decoder class or factory function.

Raises a `LookupError` in case the encoding cannot be found or the codec doesn't support an incremental decoder.

`codecs.getreader(encoding)`

Look up the codec for the given encoding and return its `StreamReader` class or factory function.

Raises a `LookupError` in case the encoding cannot be found.

`codecs.getwriter(encoding)`

Look up the codec for the given encoding and return its `StreamWriter` class or factory function.

Raises a `LookupError` in case the encoding cannot be found.

Custom codecs are made available by registering a suitable codec search function:

`codecs.register(search_function)`

Register a codec search function. Search functions are expected to take one argument, being the encoding name in all lower case letters, and return a `CodecInfo` object. In case a search function cannot find a given encoding, it should return `None`.

注解: Search function registration is not currently reversible, which may cause problems in some cases, such as unit testing or module reloading.

While the builtin `open()` and the associated `io` module are the recommended approach for working with encoded text files, this module provides additional utility functions and classes that allow the use of a wider range of codecs when working with binary files:

`codecs.open(filename, mode='r', encoding=None, errors='strict', buffering=1)`

Open an encoded file using the given `mode` and return an instance of `StreamReaderWriter`, providing transparent encoding/decoding. The default file mode is `'r'`, meaning to open the file in read mode.

注解: Underlying encoded files are always opened in binary mode. No automatic conversion of `'\n'` is done on reading and writing. The `mode` argument may be any binary mode acceptable to the built-in `open()` function; the `'b'` is automatically added.

`encoding` specifies the encoding which is to be used for the file. Any encoding that encodes to and decodes from bytes is allowed, and the data types supported by the file methods depend on the codec used.

errors may be given to define the error handling. It defaults to `'strict'` which causes a `ValueError` to be raised in case an encoding error occurs.

buffering has the same meaning as for the built-in `open()` function. It defaults to line buffered.

`codecs.EncodedFile` (*file*, *data_encoding*, *file_encoding=None*, *errors='strict'*)

Return a `StreamRecoder` instance, a wrapped version of *file* which provides transparent transcoding. The original file is closed when the wrapped version is closed.

Data written to the wrapped file is decoded according to the given *data_encoding* and then written to the original file as bytes using *file_encoding*. Bytes read from the original file are decoded according to *file_encoding*, and the result is encoded using *data_encoding*.

If *file_encoding* is not given, it defaults to *data_encoding*.

errors may be given to define the error handling. It defaults to `'strict'`, which causes `ValueError` to be raised in case an encoding error occurs.

`codecs.iterencode` (*iterator*, *encoding*, *errors='strict'*, ***kwargs*)

Uses an incremental encoder to iteratively encode the input provided by *iterator*. This function is a *generator*. The *errors* argument (as well as any other keyword argument) is passed through to the incremental encoder.

This function requires that the codec accept text `str` objects to encode. Therefore it does not support bytes-to-bytes encoders such as `base64_codec`.

`codecs.iterdecode` (*iterator*, *encoding*, *errors='strict'*, ***kwargs*)

Uses an incremental decoder to iteratively decode the input provided by *iterator*. This function is a *generator*. The *errors* argument (as well as any other keyword argument) is passed through to the incremental decoder.

This function requires that the codec accept `bytes` objects to decode. Therefore it does not support text-to-text encoders such as `rot_13`, although `rot_13` may be used equivalently with `iterencode()`.

The module also provides the following constants which are useful for reading and writing to platform dependent files:

```
codecs.BOM
codecs.BOM_BE
codecs.BOM_LE
codecs.BOM_UTF8
codecs.BOM_UTF16
codecs.BOM_UTF16_BE
codecs.BOM_UTF16_LE
codecs.BOM_UTF32
codecs.BOM_UTF32_BE
codecs.BOM_UTF32_LE
```

These constants define various byte sequences, being Unicode byte order marks (BOMs) for several encodings. They are used in UTF-16 and UTF-32 data streams to indicate the byte order used, and in UTF-8 as a Unicode signature. `BOM_UTF16` is either `BOM_UTF16_BE` or `BOM_UTF16_LE` depending on the platform's native byte order, `BOM` is an alias for `BOM_UTF16`, `BOM_LE` for `BOM_UTF16_LE` and `BOM_BE` for `BOM_UTF16_BE`. The others represent the BOM in UTF-8 and UTF-32 encodings.

7.2.1 Codec Base Classes

The `codecs` module defines a set of base classes which define the interfaces for working with codec objects, and can also be used as the basis for custom codec implementations.

Each codec has to define four interfaces to make it usable as codec in Python: stateless encoder, stateless decoder, stream reader and stream writer. The stream reader and writers typically reuse the stateless encoder/decoder to implement the file protocols. Codec authors also need to define how the codec will handle encoding and decoding errors.

Error Handlers

To simplify and standardize error handling, codecs may implement different error handling schemes by accepting the *errors* string argument. The following string values are defined and implemented by all standard Python codecs:

值	意义
'strict'	Raise <code>UnicodeError</code> (or a subclass); this is the default. Implemented in <code>strict_errors()</code> .
'ignore'	Ignore the malformed data and continue without further notice. Implemented in <code>ignore_errors()</code> .

The following error handlers are only applicable to *text encodings*:

值	意义
'replace'	Replace with a suitable replacement marker; Python will use the official U+FFFD REPLACEMENT CHARACTER for the built-in codecs on decoding, and '?' on encoding. Implemented in <code>replace_errors()</code> .
'xmlcharrefreplace'	Replace with the appropriate XML character reference (only for encoding). Implemented in <code>xmlcharrefreplace_errors()</code> .
'backslashreplace'	Replace with backslashed escape sequences. Implemented in <code>backslashreplace_errors()</code> .
'namereplace'	Replace with <code>\N{...}</code> escape sequences (only for encoding). Implemented in <code>namereplace_errors()</code> .
'surrogateescape'	On decoding, replace byte with individual surrogate code ranging from U+DC80 to U+DCFF. This code will then be turned back into the same byte when the 'surrogateescape' error handler is used when encoding the data. (See PEP 383 for more.)

In addition, the following error handler is specific to the given codecs:

值	Codecs	意义
'surrogatepass'	utf-8, 'utf-16, utf-32, utf-16-be, utf-16-le, utf-32-be, utf-32-le	Allow encoding and decoding of surrogate codes. These codecs normally treat the presence of surrogates as an error.

3.1 新版功能: The 'surrogateescape' and 'surrogatepass' error handlers.

在 3.4 版更改: The 'surrogatepass' error handlers now works with utf-16* and utf-32* codecs.

3.5 新版功能: The 'namereplace' error handler.

在 3.5 版更改: The 'backslashreplace' error handlers now works with decoding and translating.

The set of allowed values can be extended by registering a new named error handler:

`codecs.register_error` (*name*, *error_handler*)

Register the error handling function *error_handler* under the name *name*. The *error_handler* argument will be called during encoding and decoding in case of an error, when *name* is specified as the errors parameter.

For encoding, *error_handler* will be called with a `UnicodeEncodeError` instance, which contains information about the location of the error. The error handler must either raise this or a different exception, or return a tuple with a replacement for the unencodable part of the input and a position where encoding should continue. The replacement may be either *str* or *bytes*. If the replacement is bytes, the encoder will simply copy them into the output buffer. If the replacement is a string, the encoder will encode the replacement. Encoding continues on original input at the specified position. Negative position values will be treated as being relative to the end of the input string. If the resulting position is out of bound an `IndexError` will be raised.

Decoding and translating works similarly, except `UnicodeDecodeError` or `UnicodeTranslateError` will be passed to the handler and that the replacement from the error handler will be put into the output directly.

Previously registered error handlers (including the standard error handlers) can be looked up by name:

`codecs.lookup_error(name)`

Return the error handler previously registered under the name *name*.

Raises a `LookupError` in case the handler cannot be found.

The following standard error handlers are also made available as module level functions:

`codecs.strict_errors(exception)`

Implements the 'strict' error handling: each encoding or decoding error raises a `UnicodeError`.

`codecs.replace_errors(exception)`

Implements the 'replace' error handling (for *text encodings* only): substitutes '?' for encoding errors (to be encoded by the codec), and '\ufffd' (the Unicode replacement character) for decoding errors.

`codecs.ignore_errors(exception)`

Implements the 'ignore' error handling: malformed data is ignored and encoding or decoding is continued without further notice.

`codecs.xmlcharrefreplace_errors(exception)`

Implements the 'xmlcharrefreplace' error handling (for encoding with *text encodings* only): the unencodable character is replaced by an appropriate XML character reference.

`codecs.backslashreplace_errors(exception)`

Implements the 'backslashreplace' error handling (for *text encodings* only): malformed data is replaced by a backslashed escape sequence.

`codecs.namereplace_errors(exception)`

Implements the 'namereplace' error handling (for encoding with *text encodings* only): the unencodable character is replaced by a `\N{...}` escape sequence.

3.5 新版功能.

Stateless Encoding and Decoding

The base `Codec` class defines these methods which also define the function interfaces of the stateless encoder and decoder:

`Codec.encode(input[, errors])`

Encodes the object *input* and returns a tuple (output object, length consumed). For instance, *text encoding* converts a string object to a bytes object using a particular character set encoding (e.g., `cp1252` or `iso-8859-1`).

The *errors* argument defines the error handling to apply. It defaults to 'strict' handling.

The method may not store state in the `Codec` instance. Use `StreamWriter` for codecs which have to keep state in order to make encoding efficient.

The encoder must be able to handle zero length input and return an empty object of the output object type in this situation.

`Codec.decode(input[, errors])`

Decodes the object *input* and returns a tuple (output object, length consumed). For instance, for a *text encoding*, decoding converts a bytes object encoded using a particular character set encoding to a string object.

For text encodings and bytes-to-bytes codecs, *input* must be a bytes object or one which provides the read-only buffer interface – for example, buffer objects and memory mapped files.

The *errors* argument defines the error handling to apply. It defaults to 'strict' handling.

The method may not store state in the `Codec` instance. Use `StreamReader` for codecs which have to keep state in order to make decoding efficient.

The decoder must be able to handle zero length input and return an empty object of the output object type in this situation.

Incremental Encoding and Decoding

The *IncrementalEncoder* and *IncrementalDecoder* classes provide the basic interface for incremental encoding and decoding. Encoding/decoding the input isn't done with one call to the stateless encoder/decoder function, but with multiple calls to the *encode()*/*decode()* method of the incremental encoder/decoder. The incremental encoder/decoder keeps track of the encoding/decoding process during method calls.

The joined output of calls to the *encode()*/*decode()* method is the same as if all the single inputs were joined into one, and this input was encoded/decoded with the stateless encoder/decoder.

IncrementalEncoder Objects

The *IncrementalEncoder* class is used for encoding an input in multiple steps. It defines the following methods which every incremental encoder must define in order to be compatible with the Python codec registry.

class codecs.*IncrementalEncoder* (*errors*='strict')

Constructor for an *IncrementalEncoder* instance.

All incremental encoders must provide this constructor interface. They are free to add additional keyword arguments, but only the ones defined here are used by the Python codec registry.

The *IncrementalEncoder* may implement different error handling schemes by providing the *errors* keyword argument. See *Error Handlers* for possible values.

The *errors* argument will be assigned to an attribute of the same name. Assigning to this attribute makes it possible to switch between different error handling strategies during the lifetime of the *IncrementalEncoder* object.

encode (*object* [, *final*])

Encodes *object* (taking the current state of the encoder into account) and returns the resulting encoded object. If this is the last call to *encode()* *final* must be true (the default is false).

reset ()

Reset the encoder to the initial state. The output is discarded: call *.encode(object, final=True)*, passing an empty byte or text string if necessary, to reset the encoder and to get the output.

getstate ()

Return the current state of the encoder which must be an integer. The implementation should make sure that 0 is the most common state. (States that are more complicated than integers can be converted into an integer by marshaling/pickling the state and encoding the bytes of the resulting string into an integer).

setstate (*state*)

Set the state of the encoder to *state*. *state* must be an encoder state returned by *getstate()*.

IncrementalDecoder Objects

The *IncrementalDecoder* class is used for decoding an input in multiple steps. It defines the following methods which every incremental decoder must define in order to be compatible with the Python codec registry.

class codecs.*IncrementalDecoder* (*errors*='strict')

Constructor for an *IncrementalDecoder* instance.

All incremental decoders must provide this constructor interface. They are free to add additional keyword arguments, but only the ones defined here are used by the Python codec registry.

The *IncrementalDecoder* may implement different error handling schemes by providing the *errors* keyword argument. See *Error Handlers* for possible values.

The *errors* argument will be assigned to an attribute of the same name. Assigning to this attribute makes it possible to switch between different error handling strategies during the lifetime of the *IncrementalDecoder* object.

decode (*object* [, *final*])

Decodes *object* (taking the current state of the decoder into account) and returns the resulting decoded object. If this is the last call to *decode()* *final* must be true (the default is false). If *final* is true the decoder must decode the input completely and must flush all buffers. If this isn't possible (e.g. because of incomplete byte sequences at the end of the input) it must initiate error handling just like in the stateless case (which might raise an exception).

reset ()

Reset the decoder to the initial state.

getstate ()

Return the current state of the decoder. This must be a tuple with two items, the first must be the buffer containing the still undecoded input. The second must be an integer and can be additional state info. (The implementation should make sure that 0 is the most common additional state info.) If this additional state info is 0 it must be possible to set the decoder to the state which has no input buffered and 0 as the additional state info, so that feeding the previously buffered input to the decoder returns it to the previous state without producing any output. (Additional state info that is more complicated than integers can be converted into an integer by marshaling/pickling the info and encoding the bytes of the resulting string into an integer.)

setstate (*state*)

Set the state of the encoder to *state*. *state* must be a decoder state returned by *getstate()*.

Stream Encoding and Decoding

The *StreamWriter* and *StreamReader* classes provide generic working interfaces which can be used to implement new encoding submodules very easily. See `encodings.utf_8` for an example of how this is done.

StreamWriter Objects

The *StreamWriter* class is a subclass of *Codec* and defines the following methods which every stream writer must define in order to be compatible with the Python codec registry.

class `codecs.StreamWriter` (*stream*, *errors*='strict')

Constructor for a *StreamWriter* instance.

All stream writers must provide this constructor interface. They are free to add additional keyword arguments, but only the ones defined here are used by the Python codec registry.

The *stream* argument must be a file-like object open for writing text or binary data, as appropriate for the specific codec.

The *StreamWriter* may implement different error handling schemes by providing the *errors* keyword argument. See *Error Handlers* for the standard error handlers the underlying stream codec may support.

The *errors* argument will be assigned to an attribute of the same name. Assigning to this attribute makes it possible to switch between different error handling strategies during the lifetime of the *StreamWriter* object.

write (*object*)

Writes the object's contents encoded to the stream.

writelines (*list*)

Writes the concatenated list of strings to the stream (possibly by reusing the `write()` method). The standard bytes-to-bytes codecs do not support this method.

reset ()

Flushes and resets the codec buffers used for keeping state.

Calling this method should ensure that the data on the output is put into a clean state that allows appending of new fresh data without having to rescan the whole stream to recover state.

In addition to the above methods, the `StreamWriter` must also inherit all other methods and attributes from the underlying stream.

StreamReader Objects

The `StreamReader` class is a subclass of `Codec` and defines the following methods which every stream reader must define in order to be compatible with the Python codec registry.

class `codecs.StreamReader` (*stream*, *errors*='strict')

Constructor for a `StreamReader` instance.

All stream readers must provide this constructor interface. They are free to add additional keyword arguments, but only the ones defined here are used by the Python codec registry.

The *stream* argument must be a file-like object open for reading text or binary data, as appropriate for the specific codec.

The `StreamReader` may implement different error handling schemes by providing the *errors* keyword argument. See [Error Handlers](#) for the standard error handlers the underlying stream codec may support.

The *errors* argument will be assigned to an attribute of the same name. Assigning to this attribute makes it possible to switch between different error handling strategies during the lifetime of the `StreamReader` object.

The set of allowed values for the *errors* argument can be extended with `register_error()`.

read ([*size*[, *chars*[, *firstline*]]])

Decodes data from the stream and returns the resulting object.

The *chars* argument indicates the number of decoded code points or bytes to return. The `read()` method will never return more data than requested, but it might return less, if there is not enough available.

The *size* argument indicates the approximate maximum number of encoded bytes or code points to read for decoding. The decoder can modify this setting as appropriate. The default value -1 indicates to read and decode as much as possible. This parameter is intended to prevent having to decode huge files in one step.

The *firstline* flag indicates that it would be sufficient to only return the first line, if there are decoding errors on later lines.

The method should use a greedy read strategy meaning that it should read as much data as is allowed within the definition of the encoding and the given size, e.g. if optional encoding endings or state markers are available on the stream, these should be read too.

readline ([*size*[, *keepends*]])

Read one line from the input stream and return the decoded data.

size, if given, is passed as size argument to the stream's `read()` method.

If *keepends* is false line-endings will be stripped from the lines returned.

readlines ([*sizehint*[, *keepends*]])

Read all lines available on the input stream and return them as a list of lines.

Line-endings are implemented using the codec's decoder method and are included in the list entries if *keepends* is true.

sizehint, if given, is passed as the *size* argument to the stream's *read()* method.

reset()

Resets the codec buffers used for keeping state.

Note that no stream repositioning should take place. This method is primarily intended to be able to recover from decoding errors.

In addition to the above methods, the *StreamReader* must also inherit all other methods and attributes from the underlying stream.

StreamReaderWriter Objects

The *StreamReaderWriter* is a convenience class that allows wrapping streams which work in both read and write modes.

The design is such that one can use the factory functions returned by the *lookup()* function to construct the instance.

class `codecs.StreamReaderWriter` (*stream*, *Reader*, *Writer*, *errors*='strict')

Creates a *StreamReaderWriter* instance. *stream* must be a file-like object. *Reader* and *Writer* must be factory functions or classes providing the *StreamReader* and *StreamWriter* interface resp. Error handling is done in the same way as defined for the stream readers and writers.

StreamReaderWriter instances define the combined interfaces of *StreamReader* and *StreamWriter* classes. They inherit all other methods and attributes from the underlying stream.

StreamRecoder Objects

The *StreamRecoder* translates data from one encoding to another, which is sometimes useful when dealing with different encoding environments.

The design is such that one can use the factory functions returned by the *lookup()* function to construct the instance.

class `codecs.StreamRecoder` (*stream*, *encode*, *decode*, *Reader*, *Writer*, *errors*='strict')

Creates a *StreamRecoder* instance which implements a two-way conversion: *encode* and *decode* work on the frontend—the data visible to code calling *read()* and *write()*, while *Reader* and *Writer* work on the backend—the data in *stream*.

You can use these objects to do transparent transcodings from e.g. Latin-1 to UTF-8 and back.

The *stream* argument must be a file-like object.

The *encode* and *decode* arguments must adhere to the *Codec* interface. *Reader* and *Writer* must be factory functions or classes providing objects of the *StreamReader* and *StreamWriter* interface respectively.

Error handling is done in the same way as defined for the stream readers and writers.

StreamRecoder instances define the combined interfaces of *StreamReader* and *StreamWriter* classes. They inherit all other methods and attributes from the underlying stream.

7.2.2 Encodings and Unicode

Strings are stored internally as sequences of code points in range 0x0–0x10FFFF. (See [PEP 393](#) for more details about the implementation.) Once a string object is used outside of CPU and memory, endianness and how these arrays are stored as bytes become an issue. As with other codecs, serialising a string into a sequence of bytes is known as *encoding*,

and recreating the string from the sequence of bytes is known as *decoding*. There are a variety of different text serialisation codecs, which are collectively referred to as *text encodings*.

The simplest text encoding (called 'latin-1' or 'iso-8859-1') maps the code points 0–255 to the bytes 0x0–0xff, which means that a string object that contains code points above U+00FF can't be encoded with this codec. Doing so will raise a `UnicodeEncodeError` that looks like the following (although the details of the error message may differ): `UnicodeEncodeError: 'latin-1' codec can't encode character '\u1234' in position 3: ordinal not in range(256)`.

There's another group of encodings (the so called charmap encodings) that choose a different subset of all Unicode code points and how these code points are mapped to the bytes 0x0–0xff. To see how this is done simply open e.g. `encodings/cp1252.py` (which is an encoding that is used primarily on Windows). There's a string constant with 256 characters that shows you which character is mapped to which byte value.

All of these encodings can only encode 256 of the 1114112 code points defined in Unicode. A simple and straightforward way that can store each Unicode code point, is to store each code point as four consecutive bytes. There are two possibilities: store the bytes in big endian or in little endian order. These two encodings are called UTF-32-BE and UTF-32-LE respectively. Their disadvantage is that if e.g. you use UTF-32-BE on a little endian machine you will always have to swap bytes on encoding and decoding. UTF-32 avoids this problem: bytes will always be in natural endianness. When these bytes are read by a CPU with a different endianness, then bytes have to be swapped though. To be able to detect the endianness of a UTF-16 or UTF-32 byte sequence, there's the so called BOM ("Byte Order Mark"). This is the Unicode character U+FEFF. This character can be prepended to every UTF-16 or UTF-32 byte sequence. The byte swapped version of this character (0xFFFE) is an illegal character that may not appear in a Unicode text. So when the first character in an UTF-16 or UTF-32 byte sequence appears to be a U+FFFE the bytes have to be swapped on decoding. Unfortunately the character U+FEFF had a second purpose as a ZERO WIDTH NO-BREAK SPACE: a character that has no width and doesn't allow a word to be split. It can e.g. be used to give hints to a ligature algorithm. With Unicode 4.0 using U+FEFF as a ZERO WIDTH NO-BREAK SPACE has been deprecated (with U+2060 (WORD JOINER) assuming this role). Nevertheless Unicode software still must be able to handle U+FEFF in both roles: as a BOM it's a device to determine the storage layout of the encoded bytes, and vanishes once the byte sequence has been decoded into a string; as a ZERO WIDTH NO-BREAK SPACE it's a normal character that will be decoded like any other.

There's another encoding that is able to encoding the full range of Unicode characters: UTF-8. UTF-8 is an 8-bit encoding, which means there are no issues with byte order in UTF-8. Each byte in a UTF-8 byte sequence consists of two parts: marker bits (the most significant bits) and payload bits. The marker bits are a sequence of zero to four 1 bits followed by a 0 bit. Unicode characters are encoded like this (with x being payload bits, which when concatenated give the Unicode character):

范围	编码
U-00000000 ... U-0000007F	0xxxxxxx
U-00000080 ... U-000007FF	110xxxxx 10xxxxxx
U-00000800 ... U-0000FFFF	1110xxxx 10xxxxxx 10xxxxxx
U-00010000 ... U-0010FFFF	11110xxx 10xxxxxx 10xxxxxx 10xxxxxx

The least significant bit of the Unicode character is the rightmost x bit.

As UTF-8 is an 8-bit encoding no BOM is required and any U+FEFF character in the decoded string (even if it's the first character) is treated as a ZERO WIDTH NO-BREAK SPACE.

Without external information it's impossible to reliably determine which encoding was used for encoding a string. Each charmap encoding can decode any random byte sequence. However that's not possible with UTF-8, as UTF-8 byte sequences have a structure that doesn't allow arbitrary byte sequences. To increase the reliability with which a UTF-8 encoding can be detected, Microsoft invented a variant of UTF-8 (that Python 2.5 calls "utf-8-sig") for its Notepad program: Before any of the Unicode characters is written to the file, a UTF-8 encoded BOM (which looks like this as a byte sequence: 0xef, 0xbb, 0xbf) is written. As it's rather improbable that any charmap encoded file starts with these byte values (which would e.g. map to

LATIN SMALL LETTER I WITH DIAERESIS
RIGHT-POINTING DOUBLE ANGLE QUOTATION MARK
INVERTED QUESTION MARK

in iso-8859-1), this increases the probability that a `utf-8-sig` encoding can be correctly guessed from the byte sequence. So here the BOM is not used to be able to determine the byte order used for generating the byte sequence, but as a signature that helps in guessing the encoding. On encoding the `utf-8-sig` codec will write `0xef, 0xbb, 0xbf` as the first three bytes to the file. On decoding `utf-8-sig` will skip those three bytes if they appear as the first three bytes in the file. In UTF-8, the use of the BOM is discouraged and should generally be avoided.

7.2.3 标准编码

Python comes with a number of codecs built-in, either implemented as C functions or with dictionaries as mapping tables. The following table lists the codecs by name, together with a few common aliases, and the languages for which the encoding is likely used. Neither the list of aliases nor the list of languages is meant to be exhaustive. Notice that spelling alternatives that only differ in case or use a hyphen instead of an underscore are also valid aliases; therefore, e.g. `'utf-8'` is a valid alias for the `'utf_8'` codec.

CPython implementation detail: Some common encodings can bypass the codecs lookup machinery to improve performance. These optimization opportunities are only recognized by CPython for a limited set of (case insensitive) aliases: `utf-8`, `utf8`, `latin-1`, `latin1`, `iso-8859-1`, `iso8859-1`, `mbcs` (Windows only), `ascii`, `us-ascii`, `utf-16`, `utf16`, `utf-32`, `utf32`, and the same using underscores instead of dashes. Using alternative aliases for these encodings may result in slower execution.

在 3.6 版更改: Optimization opportunity recognized for `us-ascii`.

Many of the character sets support the same languages. They vary in individual characters (e.g. whether the EURO SIGN is supported or not), and in the assignment of characters to code positions. For the European languages in particular, the following variants typically exist:

- an ISO 8859 codeset
- a Microsoft Windows code page, which is typically derived from an 8859 codeset, but replaces control characters with additional graphic characters
- an IBM EBCDIC code page
- an IBM PC code page, which is ASCII compatible

编码	别名	语言
ascii	646, us-ascii	英语
big5	big5-tw, csbig5	繁体中文
big5hkscs	big5-hkscs, hkscs	繁体中文
cp037	IBM037, IBM039	英语
cp273	273, IBM273, csIBM273	德语 3.4 新版功能.
cp424	EBCDIC-CP-HE, IBM424	希伯来语
cp437	437, IBM437	英语
cp500	EBCDIC-CP-BE, EBCDIC-CP-CH, IBM500	西欧
cp720		阿拉伯语
cp737		希腊语
cp775	IBM775	波罗的海语言
cp850	850, IBM850	西欧
cp852	852, IBM852	中欧和东欧

下页继续

表 1 - 续上页

编码	别名	语言
cp855	855, IBM855	保加利亚语, 白俄罗斯语, 马其顿语, 俄语, 塞尔维亚语
cp856		希伯来语
cp857	857, IBM857	土耳其语
cp858	858, IBM858	西欧
cp860	860, IBM860	葡萄牙语
cp861	861, CP-IS, IBM861	冰岛语
cp862	862, IBM862	希伯来语
cp863	863, IBM863	加拿大语
cp864	IBM864	阿拉伯语
cp865	865, IBM865	丹麦语/挪威语
cp866	866, IBM866	俄语
cp869	869, CP-GR, IBM869	希腊语
cp874		泰语
cp875		希腊语
cp932	932, ms932, mskanji, ms-kanji	日语
cp949	949, ms949, uhc	韩语
cp950	950, ms950	繁体中文
cp1006		乌尔都语
cp1026	ibm1026	土耳其语
cp1125	1125, ibm1125, cp866u, ruscii	乌克兰语 3.4 新版功能.
cp1140	ibm1140	西欧
cp1250	windows-1250	中欧和东欧
cp1251	windows-1251	保加利亚语, 白俄罗斯语, 马其顿语, 俄语, 塞尔维亚语
cp1252	windows-1252	西欧
cp1253	windows-1253	希腊语
cp1254	windows-1254	土耳其语
cp1255	windows-1255	希伯来语
cp1256	windows-1256	阿拉伯语
cp1257	windows-1257	波罗的海语言
cp1258	windows-1258	越南语
cp65001		仅 Windows: Windows UTF-8 (CP_UTF8) 3.3 新版功能.
euc_jp	eucjp, ujis, u-jis	日语
euc_jis_2004	jisx0213, eucjis2004	日语
euc_jisx0213	eucjisx0213	日语
euc_kr	euckr, korean, ksc5601, ks_c-5601, ks_c-5601-1987, ksx1001, ks_x-1001	韩语
gb2312	chinese, csiso58gb231280, euc-cn, euccn, eucgb2312-cn, gb2312-1980, gb2312-80, iso-ir-58	简体中文
gbk	936, cp936, ms936	统一汉语
gb18030	gb18030-2000	统一汉语
hz	hzgb, hz-gb, hz-gb-2312	简体中文
iso2022_jp	csiso2022jp, iso2022jp, iso-2022-jp	日语

下页继续

表 1 - 续上页

编码	别名	语言
iso2022_jp_1	iso2022jp-1, iso-2022-jp-1	日语
iso2022_jp_2	iso2022jp-2, iso-2022-jp-2	日语, 韩语, 简体中文, 西欧, 希腊语
iso2022_jp_2004	iso2022jp-2004, iso-2022-jp-2004	日语
iso2022_jp_3	iso2022jp-3, iso-2022-jp-3	日语
iso2022_jp_ext	iso2022jp-ext, iso-2022-jp-ext	日语
iso2022_kr	csiso2022kr, iso2022kr, iso-2022-kr	韩语
latin_1	iso-8859-1, iso8859-1, 8859, cp819, latin, latin1, L1	西欧
iso8859_2	iso-8859-2, latin2, L2	中欧和东欧
iso8859_3	iso-8859-3, latin3, L3	世界语, 马耳他语
iso8859_4	iso-8859-4, latin4, L4	波罗的海语言
iso8859_5	iso-8859-5, cyrillic	保加利亚语, 白俄罗斯语, 马其顿语, 俄语, 塞尔维亚语
iso8859_6	iso-8859-6, arabic	阿拉伯语
iso8859_7	iso-8859-7, greek, greek8	希腊语
iso8859_8	iso-8859-8, hebrew	希伯来语
iso8859_9	iso-8859-9, latin5, L5	土耳其语
iso8859_10	iso-8859-10, latin6, L6	北欧语言
iso8859_11	iso-8859-11, thai	泰语
iso8859_13	iso-8859-13, latin7, L7	波罗的海语言
iso8859_14	iso-8859-14, latin8, L8	凯尔特语
iso8859_15	iso-8859-15, latin9, L9	西欧
iso8859_16	iso-8859-16, latin10, L10	东南欧
johab	cp1361, ms1361	韩语
koi8_r		俄语
koi8_t		塔吉克 3.5 新版功能.
koi8_u		乌克兰语
kz1048	kz_1048, strk1048_2002, rk1048	哈萨克语 3.5 新版功能.
mac_cyrillic	maccyrillic	保加利亚语, 白俄罗斯语, 马其顿语, 俄语, 塞尔维亚语
mac_greek	macgreek	希腊语
mac_iceland	maciceland	冰岛语
mac_latin2	maclatin2, maccentraleurope	中欧和东欧
mac_roman	macroman, macintosh	西欧
mac_turkish	macturkish	土耳其语
ptcp154	csptcp154, pt154, cp154, cyrillic-asian	哈萨克语
shift_jis	csshiftjis, shiftjis, sjis, s_jis	日语
shift_jis_2004	shiftjis2004, sjis_2004, sjis2004	日语
shift_jisx0213	shiftjisx0213, sjisx0213, s_jisx0213	日语
utf_32	U32, utf32	所有语言
utf_32_be	UTF-32BE	所有语言
utf_32_le	UTF-32LE	所有语言
utf_16	U16, utf16	所有语言
utf_16_be	UTF-16BE	所有语言
utf_16_le	UTF-16LE	所有语言

下页继续

表 1 – 续上页

编码	别名	语言
utf_7	U7, unicode-1-1-utf-7	所有语言
utf_8	U8, UTF, utf8	所有语言
utf_8_sig		所有语言

在 3.4 版更改: The utf-16* and utf-32* encoders no longer allow surrogate code points (U+D800–U+DFFF) to be encoded. The utf-32* decoders no longer decode byte sequences that correspond to surrogate code points.

7.2.4 Python Specific Encodings

A number of predefined codecs are specific to Python, so their codec names have no meaning outside Python. These are listed in the tables below based on the expected input and output types (note that while text encodings are the most common use case for codecs, the underlying codec infrastructure supports arbitrary data transforms rather than just text encodings). For asymmetric codecs, the stated purpose describes the encoding direction.

文字编码

The following codecs provide *str* to *bytes* encoding and *bytes-like object* to *str* decoding, similar to the Unicode text encodings.

编码	别名	目的
idna		Implements RFC 3490 , see also encodings.idna . Only <code>errors='strict'</code> is supported.
mbcs	ansi, dbcs	Windows only: Encode operand according to the ANSI codepage (CP_ACP)
oem		Windows only: Encode operand according to the OEM codepage (CP_OEMCP) 3.6 新版功能.
palmos		Encoding of PalmOS 3.5
punycode		Implements RFC 3492 . Stateful codecs are not supported.
raw_unicode_escape		Latin-1 encoding with <code>\uXXXX</code> and <code>\UXXXXXXXX</code> for other code points. Existing backslashes are not escaped in any way. It is used in the Python pickle protocol.
undefined		Raise an exception for all conversions, even empty strings. The error handler is ignored.
unicode_escape		Encoding suitable as the contents of a Unicode literal in ASCII-encoded Python source code, except that quotes are not escaped. Decodes from Latin-1 source code. Beware that Python source code actually uses UTF-8 by default.
unicode_internal		Return the internal representation of the operand. Stateful codecs are not supported. 3.3 版后已移除: This representation is obsoleted by PEP 393 .

二进制转换

The following codecs provide binary transforms: *bytes-like object* to *bytes* mappings. They are not supported by `bytes.decode()` (which only produces *str* output).

编码	别名	目的	编码器/解码器
base64_codec ¹	base64, base_64	Convert operand to multiline MIME base64 (the result always includes a trailing '\n') 在 3.4 版更改: accepts any <i>bytes-like object</i> as input for encoding and decoding	<code>base64.encodebytes()</code> / <code>base64.decodebytes()</code>
bz2_codec	bz2	使用 bz2 压缩操作数	<code>bz2.compress()</code> / <code>bz2.decompress()</code>
hex_codec	hex	将操作数转换为十六进制表示, 每个字节有两位数	<code>binascii.b2a_hex()</code> / <code>binascii.a2b_hex()</code>
quopri_codec	quopri, quotedprintable, quoted_printable	Convert operand to MIME quoted printable	<code>quopri.encode()</code> with <code>quotetabs=True</code> / <code>quopri.decode()</code>
uu_codec	uu	使用 uuencode 转换操作数	<code>uu.encode()</code> / <code>uu.decode()</code>
zlib_codec	zip, zlib	使用 gzip 压缩操作数	<code>zlib.compress()</code> / <code>zlib.decompress()</code>

3.2 新版功能: 恢复二进制转换。

在 3.4 版更改: 恢复二进制转换的别名。

文字转换

The following codec provides a text transform: a *str* to *str* mapping. It is not supported by `str.encode()` (which only produces *bytes* output).

编码	别名	目的
rot_13	rot13	Returns the Caesar-cypher encryption of the operand

3.2 新版功能: Restoration of the `rot_13` text transform.

在 3.4 版更改: Restoration of the `rot13` alias.

7.2.5 encodings.idna — 应用程序中的国际化域名

This module implements **RFC 3490** (Internationalized Domain Names in Applications) and **RFC 3492** (Nameprep: A Stringprep Profile for Internationalized Domain Names (IDN)). It builds upon the `punycode` encoding and `stringprep`.

These RFCs together define a protocol to support non-ASCII characters in domain names. A domain name containing non-ASCII characters (such as `www.Alliancefrançaise.nu`) is converted into an ASCII-compatible encoding (ACE, such as `www.xn--alliancefranaise-npb.nu`). The ACE form of the domain name is then used in all places where arbitrary characters are not allowed by the protocol, such as DNS queries, HTTP *Host* fields, and so on. This conversion is carried out in the application; if possible invisible to the user: The application should transparently convert Unicode domain labels to IDNA on the wire, and convert back ACE labels to Unicode before presenting them to the user.

¹ In addition to *bytes-like objects*, 'base64_codec' also accepts ASCII-only instances of *str* for decoding

Python supports this conversion in several ways: the `idna` codec performs conversion between Unicode and ACE, separating an input string into labels based on the separator characters defined in [section 3.1 of RFC 3490](#) and converting each label to ACE as required, and conversely separating an input byte string into labels based on the `.` separator and converting any ACE labels found into unicode. Furthermore, the `socket` module transparently converts Unicode host names to ACE, so that applications need not be concerned about converting host names themselves when they pass them to the `socket` module. On top of that, modules that have host names as function parameters, such as `http.client` and `ftplib`, accept Unicode host names (`http.client` then also transparently sends an IDNA hostname in the `Host` field if it sends that field at all).

When receiving host names from the wire (such as in reverse name lookup), no automatic conversion to Unicode is performed: Applications wishing to present such host names to the user should decode them to Unicode.

The module `encodings.idna` also implements the nameprep procedure, which performs certain normalizations on host names, to achieve case-insensitivity of international domain names, and to unify similar characters. The nameprep functions can be used directly if desired.

`encodings.idna.nameprep(label)`

Return the nameprepped version of *label*. The implementation currently assumes query strings, so `AllowUnassigned` is `true`.

`encodings.idna.ToASCII(label)`

Convert a label to ASCII, as specified in [RFC 3490](#). Use `STD3ASCIIRules` is assumed to be `false`.

`encodings.idna.ToUnicode(label)`

Convert a label to Unicode, as specified in [RFC 3490](#).

7.2.6 `encodings.mbc`s —Windows ANSI 代码页

根据 ANSI 代码页（CP_ACP）对操作数进行编码。

Availability: Windows only.

在 3.3 版更改: 支持任何错误处理

在 3.2 版更改: Before 3.2, the *errors* argument was ignored; 'replace' was always used to encode, and 'ignore' to decode.

7.2.7 `encodings.utf_8_sig` —带 BOM 签名的 UTF-8 编解码器

This module implements a variant of the UTF-8 codec: On encoding a UTF-8 encoded BOM will be prepended to the UTF-8 encoded bytes. For the stateful encoder this is only done once (on the first write to the byte stream). For decoding an optional UTF-8 encoded BOM at the start of the data will be skipped.

数据类型

本章节描述的模块提供了一系列专门的数据类型例如日期与实践，固定类型的数组，堆队列，同步队列与集合。

Python 同样提供一些内置的数据类型，特别的，`dict`，`list`，`set` 与 `fromzenset` 以及 `tuple`。`str` 类通常指 Unicode 字符串，并且 `bytes` 通常指二进制数据。

本章记录以下模块：

8.1 `datetime` — 基础日期和时间数据类型

源代码： [Lib/datetime.py](#)

`datetime` 模块提供了可以通过多种方式操作日期和时间的类。在支持日期时间数学运算的同时，实现的关注点更着重于如何能够更有效地解析其属性用于格式化输出和数据操作。相关功能可以参阅 `time` 和 `calendar` 模块。

有两种日期和时间的对象：天真型和觉悟型。

觉悟型对象有着用足以支持一些应用层面算法和国家层面时间调整的信息，例如时区和夏令时，来让自己和其他的觉悟型对象区别开来。觉悟型对象是用来表达不对解释器开放的特定时间信息¹。

天真型对象没包含足够多的信息来明确定位与之相关的“date”对象和“time”对象。天真型对象代表的是世界标准时间 (UTC)，当地时间或者是其它时区的时间完全取决于程序就像一个数字是代表的是米、英里或者质量完全取决于程序。天真型对象以忽略了一些现实情况的为代价使得它容易理解和使用。

对于需要觉悟型对象的应用，`datetime` 对象和 `time` 对象有一个可选的时区信息属性 `tzinfo`，这个属性可以设置给 `tzinfo` 类的子类实例。这些 `tzinfo` 对象捕获关于相对于世界标准时间 (UTC) 偏移、时区名字和夏令时是否有效等信息。需要注意的是，只有一个具体的 `tzinfo` 类，即由 `datetime` 模块提供的 `timezone` 类。`timezone` 类可以代表相对于世界标准时间 (UTC) 固定偏移的简单时区，比如世界标准时间 (UTC) 自己或者北美东部时间或者东部夏令时。支持时区的详细程度取决于应用。世界各地的时间调

¹ If, that is, we ignore the effects of Relativity

整规则相比理性更加政治性，经常会变更。也没有一个基于世界标准时间（UTC）的标准套件适合用于所有应用。

The `datetime` module exports the following constants:

`datetime.MINYEAR`

`date` 或者 `datetime` 对象允许的最小年份。常量 `MINYEAR` 是 1。

`datetime.MAXYEAR`

`date` 或 `datetime` 对象允许最大的年份。常量 `MAXYEAR` 是 9999。

参见:

模块 `calendar` 日历相关函数

模块 `time` 时间的访问和转换

8.1.1 有效的类型

class `datetime.date`

理想的天真型日期假设公历是永远有效的。属性: `year`, `month`, and `day`。

class `datetime.time`

一个理想的时间独立于任何一天，假设每一天一共有 24*60*60 秒（这里没有”闰秒”的概念）。属性: `hour`, `minute`, `second`, `microsecond`, 和 `tzinfo`。

class `datetime.datetime`

日期和时间的结合。属性: `year`, `month`, `day`, `hour`, `minute`, `second`, `microsecond`, and `tzinfo`。

class `datetime.timedelta`

表示两个 `date` 对象或者 `time` 对象, 或者 `datetime` 对象之间的时间间隔，精确到微秒。

class `datetime.tzinfo`

一个描述时区信息的抽象基类。用于给 `datetime` 类和 `time` 类提供自定义的时间调整概念（例如，负责时区或者夏令时）。

class `datetime.timezone`

一个实现了 `tzinfo` 抽象基类的子类，用于表示相对于世界标准时间（UTC）的偏移量。

3.2 新版功能.

这些类型的对象都是不可变的。

`date` 类型的对象都是天真型的。

An object of type `time` or `datetime` may be naive or aware. A `datetime` object `d` is aware if `d.tzinfo` is not `None` and `d.tzinfo.utcoffset(d)` does not return `None`. If `d.tzinfo` is `None`, or if `d.tzinfo` is not `None` but `d.tzinfo.utcoffset(d)` returns `None`, `d` is naive. A `time` object `t` is aware if `t.tzinfo` is not `None` and `t.tzinfo.utcoffset(None)` does not return `None`. Otherwise, `t` is naive.

天真型和觉悟型之间的差别不适用于 `timedelta` 对象。

子类关系

```
object
├── timedelta
├── tzinfo
│   └── timezone
├── time
├── date
│   └── datetime
```

8.1.2 `timedelta` 类对象

`timedelta` 对象表示两个 `date` 或者 `time` 的时间间隔。

`class datetime.timedelta (days=0, seconds=0, microseconds=0, milliseconds=0, minutes=0, hours=0, weeks=0)`

所有的参数都是可选的并且默认为 0。这些参数可以是整数或者浮点数，也可以是正数或者负数。

只有 `days`, `*seconds*` 和 `microseconds` 会存储在内部，即 python 内部以 `days`, `*seconds*` 和 `microseconds` 三个单位作为存储的基本单位。参数单位转换规则如下：

- 1 毫秒会转换成 1000 微秒。
- 1 分钟会转换成 60 秒。
- 1 小时会转换成 3600 秒。
- 1 星期会转换成 7 天。

`days`, `seconds`, `microseconds` 本身也是标准化的，以保证表达方式的唯一性，例：

- `0 <= microseconds < 1000000`
- `0 <= seconds < 3600*24` (一天的秒数)
- `-999999999 <= days <= 999999999`

在有任何浮点型参数或者微秒为小数的情况下，所有小数均会按照前面的换算规则叠加到下一级，并使用 `round-half-to-even` 的方法对微秒进行取舍。没有浮点型参数情况下，转换的过程就是精确的（没有信息丢失）。

如果标准化后的 `days` 数值超过了指定范围，将会抛出 `OverflowError` 异常。

需要注意的是，负数被标准化后的结果会让你大吃一惊。例如，

```
>>> from datetime import timedelta
>>> d = timedelta(microseconds=-1)
>>> (d.days, d.seconds, d.microseconds)
(-1, 86399, 999999)
```

类属性：

`timedelta.min`

The most negative `timedelta` object, `timedelta(-999999999)`.

`timedelta.max`

The most positive `timedelta` object, `timedelta(days=999999999, hours=23, minutes=59, seconds=59, microseconds=999999)`.

`timedelta.resolution`

两个不相等的 `timedelta` 类对象最小的间隔为 `timedelta(microseconds=1)`。

需要注意的是，因为标准化的缘故，`timedelta.max > -timedelta.min`，`-timedelta.max` 不可以表示一个 `timedelta` 类对象。

实例属性（只读）：

属性	值
<code>days</code>	-999999999 至 999999999，含 999999999
<code>seconds</code>	0 至 86399，包含 86399
<code>microseconds</code>	0 至 999999，包含 999999

支持的运算：

运算	结果
<code>t1 = t2 + t3</code>	<i>t2</i> 和 <i>t3</i> 的总和。 <i>t1-t2 == t3</i> 和 <i>t1-t3 == t2</i> 的值永远为真。(1)
<code>t1 = t2 - t3</code>	<i>t2</i> 减 <i>t3</i> 的差。 <i>t1 == t2 - t3</i> 和 <i>t2 == t1 + t3</i> 值为永远为的真。(1)(6)
<code>t1 = t2 * i or t1 = i * t2</code>	乘以一个整数。假如 <i>i</i> != 0 则 <i>t1 // i == t2</i> 的值为真。
	In general, <i>t1 * i == t1 * (i-1) + t1</i> is true. (1)
<code>t1 = t2 * f or t1 = f * t2</code>	乘以一个浮点数，小数结果使用 round-half-to-even 方式进行取舍。
<code>f = t2 / t3</code>	Division (3) of <i>t2</i> by <i>t3</i> . Returns a <i>float</i> object.
<code>t1 = t2 / f or t1 = t2 / i</code>	除以一个浮点数或者整数。小数结果采用 round-half-to-even 方法进行取舍。
<code>t1 = t2 // i or t1 = t2 // t3</code>	取整除，余数部分 (如果有的话) 将被丢弃。在第二种情况下，返回一个整数。(3)
<code>t1 = t2 % t3</code>	余数为一个 <i>timedelta</i> 对象。(3)
<code>q, r = divmod(t1, t2)</code>	通过: <i>q</i> = <i>t1</i> // <i>t2</i> (3) and <i>r</i> = <i>t1</i> % <i>t2</i> 计算出商和余数。 <i>q</i> 是一个整数， <i>r</i> 是一个 <i>timedelta</i> 对象。
<code>+t1</code>	返回一个相同数值的 <i>timedelta</i> 对象。
<code>-t1</code>	等价于 <i>timedelta(-t1.days, -t1.seconds, -t1.microseconds)</i> , 和 <i>t1 * -1</i> . (1)(4)
<code>abs(t)</code>	当 <i>t.days</i> >= 0 时等于 <i>+t</i> , 当 <i>t.days</i> < 0 时 <i>-t</i> 。(2)
<code>str(t)</code>	返回一个形如 [D day[s],][H]H:MM:SS[.UUUUUU] 的字符串，当 <i>t</i> 为负数的时候，D 也为负数。(5)
<code>repr(t)</code>	Returns a string in the form <code>datetime.timedelta(D[, S[, U]])</code> , where D is negative for negative <i>t</i> . (5)

注释:

- (1) 精确但可能会溢出。
- (2) 精确且不会溢出。
- (3) 除以 0 将会抛出异常 `ZeroDivisionError`。
- (4) `-timedelta.max` 不是一个 *timedelta* 类对象。
- (5) String representations of *timedelta* objects are normalized similarly to their internal representation. This leads to somewhat unusual results for negative timedeltas. For example:

```
>>> timedelta(hours=-5)
datetime.timedelta(-1, 68400)
>>> print(_)
-1 day, 19:00:00
```

- (6) 表达式 `t2 - t3` 通常与 `t2 + (-t3)` 是等价的，除非 *t3* 等于 `timedelta.max`; 在这种情况下前者会返回结果，而后者则会溢出。

除了上面列举的操作以外 *timedelta* 对象还支持与 *date* 和 *datetime* 对象进行特定的相加和相减运算 (见下文)。

在 3.2 版更改: Floor division and true division of a *timedelta* object by another *timedelta* object are now supported, as are remainder operations and the `divmod()` function. True division and multiplication of a *timedelta* object by a *float* object are now supported.

Comparisons of *timedelta* objects are supported with the *timedelta* object representing the smaller duration considered to be the smaller *timedelta*. In order to stop mixed-type comparisons from falling back to the default comparison by object address, when a *timedelta* object is compared to an object of a different type, `TypeError` is raised unless the comparison is `==` or `!=`. The latter cases return `False` or `True`, respectively.

`timedelta` 对象是 *hashable* 类型的 (可以作为字典关键字), 支持高效获取, 在布尔上下文中, `timedelta` 对象大多数情况下都被视为真, 仅在不等于 `timedelta(0)` 时。

实例方法:

`timedelta.total_seconds()`

Return the total number of seconds contained in the duration. Equivalent to `td / timedelta(seconds=1)`.

需要注意的是, 时间间隔较大时, 这个方法的结果中的微秒将会失真 (大多数平台上大于 270 年视为一个较大的时间间隔)。

3.2 新版功能.

用法示例:

```
>>> from datetime import timedelta
>>> year = timedelta(days=365)
>>> another_year = timedelta(weeks=40, days=84, hours=23,
...                           minutes=50, seconds=600) # adds up to 365 days
>>> year.total_seconds()
31536000.0
>>> year == another_year
True
>>> ten_years = 10 * year
>>> ten_years, ten_years.days // 365
(datetime.timedelta(3650), 10)
>>> nine_years = ten_years - year
>>> nine_years, nine_years.days // 365
(datetime.timedelta(3285), 9)
>>> three_years = nine_years // 3;
>>> three_years, three_years.days // 365
(datetime.timedelta(1095), 3)
>>> abs(three_years - ten_years) == 2 * three_years + year
True
```

8.1.3 date 对象

A *date* object represents a date (year, month and day) in an idealized calendar, the current Gregorian calendar indefinitely extended in both directions. January 1 of year 1 is called day number 1, January 2 of year 1 is called day number 2, and so on. This matches the definition of the “proleptic Gregorian” calendar in Dershowitz and Reingold’s book *Calendrical Calculations*, where it’s the base calendar for all computations. See the book for algorithms for converting between proleptic Gregorian ordinals and many other calendar systems.

class `datetime.date` (*year, month, day*)

所有参数都是必选的。参数可以是在下面范围内的整数:

- `MINYEAR <= year <= MAXYEAR`
- `1 <= month <= 12`
- `1 <= 日期 <= 给定年月对应的天数`

如果参数不在这些范围内, 则抛出 `ValueError` 异常。

其它构造器, 所有的类方法:

classmethod `date.today()`

返回当地的当前日期。与 “`date.fromtimestamp(time.time())`” 等价。

classmethod `date.fromtimestamp` (*timestamp*)

Return the local date corresponding to the POSIX timestamp, such as is returned by `time.time()`. This

may raise `OverflowError`, if the timestamp is out of the range of values supported by the platform C `localtime()` function, and `OSError` on `localtime()` failure. It's common for this to be restricted to years from 1970 through 2038. Note that on non-POSIX systems that include leap seconds in their notion of a timestamp, leap seconds are ignored by `fromtimestamp()`.

在 3.3 版更改: Raise `OverflowError` instead of `ValueError` if the timestamp is out of the range of values supported by the platform C `localtime()` function. Raise `OSError` instead of `ValueError` on `localtime()` failure.

classmethod `date.fromordinal(ordinal)`

Return the date corresponding to the proleptic Gregorian ordinal, where January 1 of year 1 has ordinal 1. `ValueError` is raised unless $1 \leq \text{ordinal} \leq \text{date.max.toordinal}()$. For any date *d*, `date.fromordinal(d.toordinal()) == d`.

类属性:

`date.min`

最小的日期 `date(MINYEAR, 1, 1)`。

`date.max`

最大的日期, `date(MAXYEAR, 12, 31)`。

`date.resolution`

两个日期对象的最小间隔, `timedelta(days=1)`。

实例属性 (只读):

`date.year`

在 `MINYEAR` 和 `MAXYEAR` 之间, 包含边界。

`date.month`

1 至 12 (含)

`date.day`

返回 1 到指定年月的天数间的数字。

支持的运算:

运算	结果
<code>date2 = date1 + timedelta</code>	<i>date2</i> is <code>timedelta.days</code> days removed from <i>date1</i> . (1)
<code>date2 = date1 - timedelta</code>	Computes <i>date2</i> such that <code>date2 + timedelta == date1</code> . (2)
<code>timedelta = date1 - date2</code>	(3)
<code>date1 < date2</code>	<i>date1</i> is considered less than <i>date2</i> when <i>date1</i> precedes <i>date2</i> in time. (4)

注释:

- (1) *date2* is moved forward in time if `timedelta.days > 0`, or backward if `timedelta.days < 0`. Afterward `date2 - date1 == timedelta.days`. `timedelta.seconds` and `timedelta.microseconds` are ignored. `OverflowError` is raised if `date2.year` would be smaller than `MINYEAR` or larger than `MAXYEAR`.
- (2) `timedelta.seconds` 和 `timedelta.microseconds` 是被忽略的。
- (3) 精确且不会溢出。`timedelta.seconds` 和 `timedelta.microseconds` 均为 0, 同时 `date2 + timedelta == date1` 是成立的。
- (4) In other words, `date1 < date2` if and only if `date1.toordinal() < date2.toordinal()`. Date comparison raises `TypeError` if the other comparand isn't also a `date` object. However, `NotImplemented` is returned instead if the other comparand has a `timetuple()` attribute. This hook gives other kinds of date objects a chance at implementing mixed-type comparison. If not, when a `date` object is compared to an object

of a different type, `TypeError` is raised unless the comparison is `==` or `!=`. The latter cases return `False` or `True`, respectively.

日期可以作为字典的关键字。在布尔上下文中，所有的`date`对象都视为真。

实例方法：

`date.replace(year=self.year, month=self.month, day=self.day)`

Return a date with the same value, except for those parameters given new values by whichever keyword arguments are specified. For example, if `d == date(2002, 12, 31)`, then `d.replace(day=26) == date(2002, 12, 26)`.

`date.timetuple()`

Return a `time.struct_time` such as returned by `time.localtime()`. The hours, minutes and seconds are 0, and the DST flag is -1. `d.timetuple()` is equivalent to `time.struct_time((d.year, d.month, d.day, 0, 0, 0, d.weekday(), yday, -1))`, where `yday = d.toordinal() - date(d.year, 1, 1).toordinal() + 1` is the day number within the current year starting with 1 for January 1st.

`date.toordinal()`

Return the proleptic Gregorian ordinal of the date, where January 1 of year 1 has ordinal 1. For any `date` object `d`, `date.fromordinal(d.toordinal()) == d`.

`date.weekday()`

返回一个整数代表星期几，星期一为 0，星期天为 6。例如，`date(2002, 12, 4).weekday() == 2`，表示的是星期三。参阅 `isoweekday()`。

`date.isoweekday()`

返回一个整数代表星期几，星期一为 1，星期天为 7。例如：`date(2002, 12, 4).isoweekday() == 3`，表示星期三。参见 `weekday()`，`isocalendar()`。

`date.isocalendar()`

返回一个三元组，(ISO year, ISO week number, ISO weekday)。

ISO 日历是一个被广泛使用的公历。可以从 <https://www.staff.science.uu.nl/~gent0113/calendar/isocalendar.htm> 上查看更完整的说明。

The ISO year consists of 52 or 53 full weeks, and where a week starts on a Monday and ends on a Sunday. The first week of an ISO year is the first (Gregorian) calendar week of a year containing a Thursday. This is called week number 1, and the ISO year of that Thursday is the same as its Gregorian year.

For example, 2004 begins on a Thursday, so the first week of ISO year 2004 begins on Monday, 29 Dec 2003 and ends on Sunday, 4 Jan 2004, so that `date(2003, 12, 29).isocalendar() == (2004, 1, 1)` and `date(2004, 1, 4).isocalendar() == (2004, 1, 7)`.

`date.isoformat()`

返回一个 ISO 8601 格式的字符串，‘YYYY-MM-DD’。例如 `date(2002, 12, 4).isoformat() == '2002-12-04'`。

`date.__str__()`

对于日期对象 `d`，`str(d)` 等价于 `d.isoformat()`。

`date.ctime()`

Return a string representing the date, for example `date(2002, 12, 4).ctime() == 'Wed Dec 4 00:00:00 2002'`. `d.ctime()` is equivalent to `time.ctime(time.mktime(d.timetuple()))` on platforms where the native C `ctime()` function (which `time.ctime()` invokes, but which `date.ctime()` does not invoke) conforms to the C standard.

`date.strftime(format)`

Return a string representing the date, controlled by an explicit format string. Format codes referring to hours,

minutes or seconds will see 0 values. For a complete list of formatting directives, see *strftime()* and *strptime() Behavior*.

`date.__format__(format)`

Same as `date.strftime()`. This makes it possible to specify a format string for a `date` object in formatted string literals and when using `str.format()`. For a complete list of formatting directives, see *strftime()* and *strptime() Behavior*.

Example of counting days to an event:

```
>>> import time
>>> from datetime import date
>>> today = date.today()
>>> today
datetime.date(2007, 12, 5)
>>> today == date.fromtimestamp(time.time())
True
>>> my_birthday = date(today.year, 6, 24)
>>> if my_birthday < today:
...     my_birthday = my_birthday.replace(year=today.year + 1)
>>> my_birthday
datetime.date(2008, 6, 24)
>>> time_to_birthday = abs(my_birthday - today)
>>> time_to_birthday.days
202
```

Example of working with `date`:

```
>>> from datetime import date
>>> d = date.fromordinal(730920) # 730920th day after 1. 1. 0001
>>> d
datetime.date(2002, 3, 11)
>>> t = d.timetuple()
>>> for i in t:
...     print(i)
2002          # year
3             # month
11            # day
0
0
0
0             # weekday (0 = Monday)
70            # 70th day in the year
-1
>>> ic = d.isocalendar()
>>> for i in ic:
...     print(i)
2002          # ISO year
11            # ISO week number
1             # ISO day number ( 1 = Monday )
>>> d.isoformat()
'2002-03-11'
>>> d.strftime("%d/%m/%y")
'11/03/02'
>>> d.strftime("%A %d. %B %Y")
'Monday 11. March 2002'
>>> 'The {1} is {0:%d}, the {2} is {0:%B}.'.format(d, "day", "month")
'The day is 11, the month is March.'
```

8.1.4 `datetime` 对象

A `datetime` object is a single object containing all the information from a `date` object and a `time` object. Like a `date` object, `datetime` assumes the current Gregorian calendar extended in both directions; like a time object, `datetime` assumes there are exactly 3600×24 seconds in every day.

构造器:

class `datetime.datetime`(*year, month, day, hour=0, minute=0, second=0, microsecond=0, tzinfo=None, *, fold=0*)

`year, month, day` 三个参数是必选的。`tzinfo` 可以是 `None`, 或者是是一个 `tzinfo` 子类的实例。其余的参数可以是在以下范围内的整数:

- `MINYEAR <= year <= MAXYEAR`,
- `1 <= month <= 12`,
- `1 <= day <= 指定年月的天数`,
- `0 <= hour < 24`,
- `0 <= minute < 60`,
- `0 <= second < 60`,
- `0 <= microsecond < 1000000`,
- `fold` in `[0, 1]`.

如果参数不在这些范围内, 则抛出 `ValueError` 异常。

3.6 新版功能: 增加了 `fold` 参数。

其它构造器, 所有的类方法:

classmethod `datetime.today`()

Return the current local datetime, with `tzinfo` `None`. This is equivalent to `datetime.fromtimestamp(time.time())`. See also `now()`, `fromtimestamp()`.

classmethod `datetime.now`(*tz=None*)

Return the current local date and time. If optional argument `tz` is `None` or not specified, this is like `today()`, but, if possible, supplies more precision than can be gotten from going through a `time.time()` timestamp (for example, this may be possible on platforms supplying the C `gettimeofday()` function).

If `tz` is not `None`, it must be an instance of a `tzinfo` subclass, and the current date and time are converted to `tz`'s time zone. In this case the result is equivalent to `tz.fromutc(datetime.utcnow()).replace(tzinfo=tz)`. See also `today()`, `utcnow()`.

classmethod `datetime.utcnow`()

Return the current UTC date and time, with `tzinfo` `None`. This is like `now()`, but returns the current UTC date and time, as a naive `datetime` object. An aware current UTC datetime can be obtained by calling `datetime.now(timezone.utc)`. See also `now()`.

classmethod `datetime.fromtimestamp`(*timestamp, tz=None*)

Return the local date and time corresponding to the POSIX timestamp, such as is returned by `time.time()`. If optional argument `tz` is `None` or not specified, the timestamp is converted to the platform's local date and time, and the returned `datetime` object is naive.

If `tz` is not `None`, it must be an instance of a `tzinfo` subclass, and the timestamp is converted to `tz`'s time zone. In this case the result is equivalent to `tz.fromutc(datetime.utcfromtimestamp(timestamp)).replace(tzinfo=tz)`.

`fromtimestamp()` may raise `OverflowError`, if the timestamp is out of the range of values supported by the platform C `localtime()` or `gmtime()` functions, and `OSError` on `localtime()` or `gmtime()`

failure. It's common for this to be restricted to years in 1970 through 2038. Note that on non-POSIX systems that include leap seconds in their notion of a timestamp, leap seconds are ignored by `fromtimestamp()`, and then it's possible to have two timestamps differing by a second that yield identical `datetime` objects. See also `utcfromtimestamp()`.

在 3.3 版更改: Raise `OverflowError` instead of `ValueError` if the timestamp is out of the range of values supported by the platform C `localtime()` or `gmtime()` functions. Raise `OSError` instead of `ValueError` on `localtime()` or `gmtime()` failure.

在 3.6 版更改: `fromtimestamp()` may return instances with `fold` set to 1.

classmethod `datetime.utcnow(timestamp)`

Return the UTC `datetime` corresponding to the POSIX timestamp, with `tzinfo` None. This may raise `OverflowError`, if the timestamp is out of the range of values supported by the platform C `gmtime()` function, and `OSError` on `gmtime()` failure. It's common for this to be restricted to years in 1970 through 2038.

To get an aware `datetime` object, call `fromtimestamp()`:

```
datetime.fromtimestamp(timestamp, timezone.utc)
```

On the POSIX compliant platforms, it is equivalent to the following expression:

```
datetime(1970, 1, 1, tzinfo=timezone.utc) + timedelta(seconds=timestamp)
```

except the latter formula always supports the full years range: between `MINYEAR` and `MAXYEAR` inclusive.

在 3.3 版更改: Raise `OverflowError` instead of `ValueError` if the timestamp is out of the range of values supported by the platform C `gmtime()` function. Raise `OSError` instead of `ValueError` on `gmtime()` failure.

classmethod `datetime.fromordinal(ordinal)`

Return the `datetime` corresponding to the proleptic Gregorian ordinal, where January 1 of year 1 has ordinal 1. `ValueError` is raised unless $1 \leq \text{ordinal} \leq \text{datetime.max.toordinal}()$. The hour, minute, second and microsecond of the result are all 0, and `tzinfo` is None.

classmethod `datetime.combine(date, time, tzinfo=self.tzinfo)`

Return a new `datetime` object whose date components are equal to the given `date` object's, and whose time components are equal to the given `time` object's. If the `tzinfo` argument is provided, its value is used to set the `tzinfo` attribute of the result, otherwise the `tzinfo` attribute of the `time` argument is used.

For any `datetime` object `d`, `d == datetime.combine(d.date(), d.time(), d.tzinfo)`. If `date` is a `datetime` object, its time components and `tzinfo` attributes are ignored.

在 3.6 版更改: Added the `tzinfo` argument.

classmethod `datetime.strptime(date_string, format)`

Return a `datetime` corresponding to `date_string`, parsed according to `format`. This is equivalent to `datetime(*(time.strptime(date_string, format)[0:6]))`. `ValueError` is raised if the `date_string` and `format` can't be parsed by `time.strptime()` or if it returns a value which isn't a time tuple. For a complete list of formatting directives, see `strftime() and strptime() Behavior`.

类属性:

`datetime.min`

The earliest representable `datetime`, `datetime(MINYEAR, 1, 1, tzinfo=None)`.

`datetime.max`

The latest representable `datetime`, `datetime(MAXYEAR, 12, 31, 23, 59, 59, 999999, tzinfo=None)`.

`datetime.resolution`

The smallest possible difference between non-equal `datetime` objects, `timedelta(microseconds=1)`.

实例属性（只读）：

`datetime.year`

在`MINYEAR`和`MAXYEAR`之间，包含边界。

`datetime.month`

1 至 12（含）

`datetime.day`

返回 1 到指定年月的天数间的数字。

`datetime.hour`

In range(24).

`datetime.minute`

In range(60).

`datetime.second`

In range(60).

`datetime.microsecond`

In range(1000000).

`datetime.tzinfo`

The object passed as the `tzinfo` argument to the `datetime` constructor, or None if none was passed.

`datetime.fold`

In [0, 1]. Used to disambiguate wall times during a repeated interval. (A repeated interval occurs when clocks are rolled back at the end of daylight saving time or when the UTC offset for the current zone is decreased for political reasons.) The value 0 (1) represents the earlier (later) of the two moments with the same wall time representation.

3.6 新版功能.

支持的运算：

运算	结果
<code>datetime2 = datetime1 + timedelta</code>	(1)
<code>datetime2 = datetime1 - timedelta</code>	(2)
<code>timedelta = datetime1 - datetime2</code>	(3)
<code>datetime1 < datetime2</code>	Compares <code>datetime</code> to <code>datetime</code> . (4)

- (1) `datetime2` is a duration of `timedelta` removed from `datetime1`, moving forward in time if `timedelta.days > 0`, or backward if `timedelta.days < 0`. The result has the same `tzinfo` attribute as the input `datetime`, and `datetime2 - datetime1 == timedelta` after. `OverflowError` is raised if `datetime2.year` would be smaller than `MINYEAR` or larger than `MAXYEAR`. Note that no time zone adjustments are done even if the input is an aware object.
- (2) Computes the `datetime2` such that `datetime2 + timedelta == datetime1`. As for addition, the result has the same `tzinfo` attribute as the input `datetime`, and no time zone adjustments are done even if the input is aware.
- (3) Subtraction of a `datetime` from a `datetime` is defined only if both operands are naive, or if both are aware. If one is aware and the other is naive, `TypeError` is raised.

If both are naive, or both are aware and have the same `tzinfo` attribute, the `tzinfo` attributes are ignored, and the result is a `timedelta` object `t` such that `datetime2 + t == datetime1`. No time zone adjustments are done in this case.

If both are aware and have different `tzinfo` attributes, `a-b` acts as if `a` and `b` were first converted to naive UTC datetimes first. The result is `(a.replace(tzinfo=None) - a.utcoffset()) - (b.replace(tzinfo=None) - b.utcoffset())` except that the implementation never overflows.

- (4) *datetime1* is considered less than *datetime2* when *datetime1* precedes *datetime2* in time.

If one comparand is naive and the other is aware, *TypeError* is raised if an order comparison is attempted. For equality comparisons, naive instances are never equal to aware instances.

If both comparands are aware, and have the same *tzinfo* attribute, the common *tzinfo* attribute is ignored and the base datetimes are compared. If both comparands are aware and have different *tzinfo* attributes, the comparands are first adjusted by subtracting their UTC offsets (obtained from `self.utcoffset()`).

在 3.3 版更改: Equality comparisons between naive and aware *datetime* instances don't raise *TypeError*.

注解: In order to stop comparison from falling back to the default scheme of comparing object addresses, *datetime* comparison normally raises *TypeError* if the other comparand isn't also a *datetime* object. However, *NotImplemented* is returned instead if the other comparand has a `timetuple()` attribute. This hook gives other kinds of date objects a chance at implementing mixed-type comparison. If not, when a *datetime* object is compared to an object of a different type, *TypeError* is raised unless the comparison is `==` or `!=`. The latter cases return *False* or *True*, respectively.

datetime objects can be used as dictionary keys. In Boolean contexts, all *datetime* objects are considered to be true.

实例方法:

`datetime.date()`

Return *date* object with same year, month and day.

`datetime.time()`

Return *time* object with same hour, minute, second, microsecond and fold. *tzinfo* is *None*. See also method `timetz()`.

在 3.6 版更改: The fold value is copied to the returned *time* object.

`datetime.timetz()`

Return *time* object with same hour, minute, second, microsecond, fold, and *tzinfo* attributes. See also method `time()`.

在 3.6 版更改: The fold value is copied to the returned *time* object.

`datetime.replace(year=self.year, month=self.month, day=self.day, hour=self.hour, minute=self.minute, second=self.second, microsecond=self.microsecond, tzinfo=self.tzinfo, *fold=0)`

Return a *datetime* with the same attributes, except for those attributes given new values by whichever keyword arguments are specified. Note that *tzinfo=None* can be specified to create a naive *datetime* from an aware *datetime* with no conversion of date and time data.

3.6 新版功能: 增加了 *fold* 参数。

`datetime.astimezone(tz=None)`

Return a *datetime* object with new *tzinfo* attribute *tz*, adjusting the date and time data so the result is the same UTC time as *self*, but in *tz*'s local time.

If provided, *tz* must be an instance of a *tzinfo* subclass, and its `utcoffset()` and `dst()` methods must not return *None*. If *self* is naive, it is presumed to represent time in the system timezone.

If called without arguments (or with *tz=None*) the system local timezone is assumed for the target timezone. The *.tzinfo* attribute of the converted *datetime* instance will be set to an instance of *timezone* with the zone name and offset obtained from the OS.

If *self.tzinfo* is *tz*, `self.astimezone(tz)` is equal to *self*: no adjustment of date or time data is performed. Else the result is local time in the timezone *tz*, representing the same UTC time as *self*: after `astz = dt.astimezone(tz)`, `astz - astz.utcoffset()` will have the same date and time data as `dt - dt.utcoffset()`.

If you merely want to attach a time zone object *tz* to a datetime *dt* without adjustment of date and time data, use `dt.replace(tzinfo=tz)`. If you merely want to remove the time zone object from an aware datetime *dt* without conversion of date and time data, use `dt.replace(tzinfo=None)`.

Note that the default `tzinfo.fromutc()` method can be overridden in a `tzinfo` subclass to affect the result returned by `astimezone()`. Ignoring error cases, `astimezone()` acts like:

```
def astimezone(self, tz):
    if self.tzinfo is tz:
        return self
    # Convert self to UTC, and attach the new time zone object.
    utc = (self - self.utcoffset()).replace(tzinfo=tz)
    # Convert from UTC to tz's local time.
    return tz.fromutc(utc)
```

在 3.3 版更改: *tz* now can be omitted.

在 3.6 版更改: The `astimezone()` method can now be called on naive instances that are presumed to represent system local time.

`datetime.utcoffset()`

If `tzinfo` is `None`, returns `None`, else returns `self.tzinfo.utcoffset(self)`, and raises an exception if the latter doesn't return `None`, or a `timedelta` object representing a whole number of minutes with magnitude less than one day.

`datetime.dst()`

If `tzinfo` is `None`, returns `None`, else returns `self.tzinfo.dst(self)`, and raises an exception if the latter doesn't return `None`, or a `timedelta` object representing a whole number of minutes with magnitude less than one day.

`datetime.tzname()`

If `tzinfo` is `None`, returns `None`, else returns `self.tzinfo.tzname(self)`, raises an exception if the latter doesn't return `None` or a string object,

`datetime.timetuple()`

Return a `time.struct_time` such as returned by `time.localtime()`. `d.timetuple()` is equivalent to `time.struct_time((d.year, d.month, d.day, d.hour, d.minute, d.second, d.weekday(), yday, dst))`, where `yday = d.toordinal() - date(d.year, 1, 1).toordinal() + 1` is the day number within the current year starting with 1 for January 1st. The `tm_isdst` flag of the result is set according to the `dst()` method: `tzinfo` is `None` or `dst()` returns `None`, `tm_isdst` is set to -1; else if `dst()` returns a non-zero value, `tm_isdst` is set to 1; else `tm_isdst` is set to 0.

`datetime.utctimetuple()`

If `datetime` instance *d* is naive, this is the same as `d.timetuple()` except that `tm_isdst` is forced to 0 regardless of what `d.dst()` returns. DST is never in effect for a UTC time.

If *d* is aware, *d* is normalized to UTC time, by subtracting `d.utcoffset()`, and a `time.struct_time` for the normalized time is returned. `tm_isdst` is forced to 0. Note that an `OverflowError` may be raised if *d*.year was `MINYEAR` or `MAXYEAR` and UTC adjustment spills over a year boundary.

`datetime.toordinal()`

Return the proleptic Gregorian ordinal of the date. The same as `self.date().toordinal()`.

`datetime.timestamp()`

Return POSIX timestamp corresponding to the `datetime` instance. The return value is a `float` similar to that returned by `time.time()`.

Naive `datetime` instances are assumed to represent local time and this method relies on the platform `C mktime()` function to perform the conversion. Since `datetime` supports wider range of values than

`mktime()` on many platforms, this method may raise `OverflowError` for times far in the past or far in the future.

For aware `datetime` instances, the return value is computed as:

```
(dt - datetime(1970, 1, 1, tzinfo=timezone.utc)).total_seconds()
```

3.3 新版功能.

在 3.6 版更改: The `timestamp()` method uses the `fold` attribute to disambiguate the times during a repeated interval.

注解: There is no method to obtain the POSIX timestamp directly from a naive `datetime` instance representing UTC time. If your application uses this convention and your system timezone is not set to UTC, you can obtain the POSIX timestamp by supplying `tzinfo=timezone.utc`:

```
timestamp = dt.replace(tzinfo=timezone.utc).timestamp()
```

or by calculating the timestamp directly:

```
timestamp = (dt - datetime(1970, 1, 1)) / timedelta(seconds=1)
```

`datetime.weekday()`

Return the day of the week as an integer, where Monday is 0 and Sunday is 6. The same as `self.date().weekday()`. See also `isoweekday()`.

`datetime.isoweekday()`

Return the day of the week as an integer, where Monday is 1 and Sunday is 7. The same as `self.date().isoweekday()`. See also `weekday()`, `isocalendar()`.

`datetime.isocalendar()`

Return a 3-tuple, (ISO year, ISO week number, ISO weekday). The same as `self.date().isocalendar()`.

`datetime.isoformat(sep='T', timespec='auto')`

Return a string representing the date and time in ISO 8601 format, YYYY-MM-DDTHH:MM:SS.mmmmmmm or, if `microsecond` is 0, YYYY-MM-DDTHH:MM:SS

If `utcoffset()` does not return `None`, a 6-character string is appended, giving the UTC offset in (signed) hours and minutes: YYYY-MM-DDTHH:MM:SS.mmmmmmm+HH:MM or, if `microsecond` is 0 YYYY-MM-DDTHH:MM:SS+HH:MM

The optional argument `sep` (default 'T') is a one-character separator, placed between the date and time portions of the result. For example,

```
>>> from datetime import tzinfo, timedelta, datetime
>>> class TZ(tzinfo):
...     def utcoffset(self, dt): return timedelta(minutes=-399)
...
>>> datetime(2002, 12, 25, tzinfo=TZ()).isoformat(' ')
'2002-12-25 00:00:00-06:39'
```

The optional argument `timespec` specifies the number of additional components of the time to include (the default is 'auto'). It can be one of the following:

- 'auto': Same as 'seconds' if `microsecond` is 0, same as 'microseconds' otherwise.
- 'hours': Include the `hour` in the two-digit HH format.
- 'minutes': Include `hour` and `minute` in HH:MM format.

- 'seconds': Include *hour*, *minute*, and *second* in HH:MM:SS format.
- 'milliseconds': Include full time, but truncate fractional second part to milliseconds. HH:MM:SS.sss format.
- 'microseconds': Include full time in HH:MM:SS.mmmmmm format.

注解: Excluded time components are truncated, not rounded.

`ValueError` will be raised on an invalid *timespec* argument.

```
>>> from datetime import datetime
>>> datetime.now().isoformat(timespec='minutes')
'2002-12-25T00:00'
>>> dt = datetime(2015, 1, 1, 12, 30, 59, 0)
>>> dt.isoformat(timespec='microseconds')
'2015-01-01T12:30:59.000000'
```

3.6 新版功能: Added the *timespec* argument.

`datetime.__str__()`

For a *datetime* instance *d*, `str(d)` is equivalent to `d.isoformat(' ')`.

`datetime.ctime()`

Return a string representing the date and time, for example `datetime(2002, 12, 4, 20, 30, 40).ctime() == 'Wed Dec 4 20:30:40 2002'`. `d.ctime()` is equivalent to `time.ctime(time.mktime(d.timetuple()))` on platforms where the native C `ctime()` function (which *time.ctime()* invokes, but which *datetime.ctime()* does not invoke) conforms to the C standard.

`datetime.strftime(format)`

Return a string representing the date and time, controlled by an explicit format string. For a complete list of formatting directives, see *strftime()* and *strptime() Behavior*.

`datetime.__format__(format)`

Same as *datetime.strftime()*. This makes it possible to specify a format string for a *datetime* object in formatted string literals and when using *str.format()*. For a complete list of formatting directives, see *strftime()* and *strptime() Behavior*.

Examples of working with datetime objects:

```
>>> from datetime import datetime, date, time
>>> # Using datetime.combine()
>>> d = date(2005, 7, 14)
>>> t = time(12, 30)
>>> datetime.combine(d, t)
datetime.datetime(2005, 7, 14, 12, 30)
>>> # Using datetime.now() or datetime.utcnow()
>>> datetime.now()
datetime.datetime(2007, 12, 6, 16, 29, 43, 79043) # GMT +1
>>> datetime.utcnow()
datetime.datetime(2007, 12, 6, 15, 29, 43, 79060)
>>> # Using datetime.strptime()
>>> dt = datetime.strptime("21/11/06 16:30", "%d/%m/%y %H:%M")
>>> dt
datetime.datetime(2006, 11, 21, 16, 30)
>>> # Using datetime.timetuple() to get tuple of all attributes
>>> tt = dt.timetuple()
>>> for it in tt:
```

(下页继续)

(续上页)

```

...     print(it)
...
2006    # year
11      # month
21      # day
16      # hour
30      # minute
0       # second
1       # weekday (0 = Monday)
325     # number of days since 1st January
-1      # dst - method tzinfo.dst() returned None
>>> # Date in ISO format
>>> ic = dt.isocalendar()
>>> for it in ic:
...     print(it)
...
2006    # ISO year
47      # ISO week
2       # ISO weekday
>>> # Formatting datetime
>>> dt.strftime("%A, %d. %B %Y %I:%M%p")
'Tuesday, 21. November 2006 04:30PM'
>>> 'The {1} is {0:%d}, the {2} is {0:%B}, the {3} is {0:%I:%M%p}.'.format(dt, "day",
↪ "month", "time")
'The day is 21, the month is November, the time is 04:30PM.'

```

Using datetime with tzinfo:

```

>>> from datetime import timedelta, datetime, tzinfo
>>> class GMT1(tzinfo):
...     def utcoffset(self, dt):
...         return timedelta(hours=1) + self.dst(dt)
...     def dst(self, dt):
...         # DST starts last Sunday in March
...         d = datetime(dt.year, 4, 1) # ends last Sunday in October
...         self.dston = d - timedelta(days=d.weekday() + 1)
...         d = datetime(dt.year, 11, 1)
...         self.dstoff = d - timedelta(days=d.weekday() + 1)
...         if self.dston <= dt.replace(tzinfo=None) < self.dstoff:
...             return timedelta(hours=1)
...         else:
...             return timedelta(0)
...     def tzname(self, dt):
...         return "GMT +1"
...
>>> class GMT2(tzinfo):
...     def utcoffset(self, dt):
...         return timedelta(hours=2) + self.dst(dt)
...     def dst(self, dt):
...         d = datetime(dt.year, 4, 1)
...         self.dston = d - timedelta(days=d.weekday() + 1)
...         d = datetime(dt.year, 11, 1)
...         self.dstoff = d - timedelta(days=d.weekday() + 1)
...         if self.dston <= dt.replace(tzinfo=None) < self.dstoff:
...             return timedelta(hours=1)
...         else:

```

(下页继续)

(续上页)

```

...         return timedelta(0)
...     def tzname(self, dt):
...         return "GMT +2"
...
>>> gmt1 = GMT1()
>>> # Daylight Saving Time
>>> dt1 = datetime(2006, 11, 21, 16, 30, tzinfo=gmt1)
>>> dt1.dst()
datetime.timedelta(0)
>>> dt1.utcoffset()
datetime.timedelta(0, 3600)
>>> dt2 = datetime(2006, 6, 14, 13, 0, tzinfo=gmt1)
>>> dt2.dst()
datetime.timedelta(0, 3600)
>>> dt2.utcoffset()
datetime.timedelta(0, 7200)
>>> # Convert datetime to another time zone
>>> dt3 = dt2.astimezone(GMT2())
>>> dt3      # doctest: +ELLIPSIS
datetime.datetime(2006, 6, 14, 14, 0, tzinfo=<GMT2 object at 0x...>)
>>> dt2      # doctest: +ELLIPSIS
datetime.datetime(2006, 6, 14, 13, 0, tzinfo=<GMT1 object at 0x...>)
>>> dt2.utctimetuple() == dt3.utctimetuple()
True

```

8.1.5 time Objects

A time object represents a (local) time of day, independent of any particular day, and subject to adjustment via a *tzinfo* object.

class `datetime.time` (*hour=0, minute=0, second=0, microsecond=0, tzinfo=None, *, fold=0*)

All arguments are optional. *tzinfo* may be *None*, or an instance of a *tzinfo* subclass. The remaining arguments may be integers, in the following ranges:

- 0 <= hour < 24,
- 0 <= minute < 60,
- 0 <= second < 60,
- 0 <= microsecond < 1000000,
- fold in [0, 1].

If an argument outside those ranges is given, *ValueError* is raised. All default to 0 except *tzinfo*, which defaults to *None*.

类属性:

`time.min`

The earliest representable *time*, `time(0, 0, 0, 0)`.

`time.max`

The latest representable *time*, `time(23, 59, 59, 999999)`.

`time.resolution`

The smallest possible difference between non-equal *time* objects, `timedelta(microseconds=1)`, although note that arithmetic on *time* objects is not supported.

实例属性（只读）：

`time.hour`
In range(24).

`time.minute`
In range(60).

`time.second`
In range(60).

`time.microsecond`
In range(1000000).

`time.tzinfo`
The object passed as the `tzinfo` argument to the `time` constructor, or `None` if none was passed.

`time.fold`
In [0, 1]. Used to disambiguate wall times during a repeated interval. (A repeated interval occurs when clocks are rolled back at the end of daylight saving time or when the UTC offset for the current zone is decreased for political reasons.) The value 0 (1) represents the earlier (later) of the two moments with the same wall time representation.

3.6 新版功能.

支持的运算：

- comparison of `time` to `time`, where *a* is considered less than *b* when *a* precedes *b* in time. If one comparand is naive and the other is aware, `TypeError` is raised if an order comparison is attempted. For equality comparisons, naive instances are never equal to aware instances.

If both comparands are aware, and have the same `tzinfo` attribute, the common `tzinfo` attribute is ignored and the base times are compared. If both comparands are aware and have different `tzinfo` attributes, the comparands are first adjusted by subtracting their UTC offsets (obtained from `self.utcoffset()`). In order to stop mixed-type comparisons from falling back to the default comparison by object address, when a `time` object is compared to an object of a different type, `TypeError` is raised unless the comparison is `==` or `!=`. The latter cases return `False` or `True`, respectively.

在 3.3 版更改: Equality comparisons between naive and aware `time` instances don't raise `TypeError`.

- hash, use as dict key
- efficient pickling

In boolean contexts, a `time` object is always considered to be true.

在 3.5 版更改: Before Python 3.5, a `time` object was considered to be false if it represented midnight in UTC. This behavior was considered obscure and error-prone and has been removed in Python 3.5. See [bpo-13936](#) for full details.

实例方法：

`time.replace(hour=self.hour, minute=self.minute, second=self.second, microsecond=self.microsecond, tzinfo=self.tzinfo, *fold=0)`

Return a `time` with the same value, except for those attributes given new values by whichever keyword arguments are specified. Note that `tzinfo=None` can be specified to create a naive `time` from an aware `time`, without conversion of the time data.

3.6 新版功能: 增加了 `fold` 参数。

`time.isoformat(timespec='auto')`

Return a string representing the time in ISO 8601 format, HH:MM:SS.mmmmmmm or, if `microsecond` is 0, HH:MM:SS If `utcoffset()` does not return `None`, a 6-character string is appended, giving the UTC offset in (signed) hours and minutes: HH:MM:SS.mmmmmmm+HH:MM or, if `self.microsecond` is 0, HH:MM:SS+HH:MM

The optional argument *timespec* specifies the number of additional components of the time to include (the default is 'auto'). It can be one of the following:

- 'auto': Same as 'seconds' if *microsecond* is 0, same as 'microseconds' otherwise.
- 'hours': Include the *hour* in the two-digit HH format.
- 'minutes': Include *hour* and *minute* in HH:MM format.
- 'seconds': Include *hour*, *minute*, and *second* in HH:MM:SS format.
- 'milliseconds': Include full time, but truncate fractional second part to milliseconds. HH:MM:SS.sss format.
- 'microseconds': Include full time in HH:MM:SS.mmmmmm format.

注解: Excluded time components are truncated, not rounded.

ValueError will be raised on an invalid *timespec* argument.

```
>>> from datetime import time
>>> time(hour=12, minute=34, second=56, microsecond=123456).isoformat(timespec=
↳ 'minutes')
'12:34'
>>> dt = time(hour=12, minute=34, second=56, microsecond=0)
>>> dt.isoformat(timespec='microseconds')
'12:34:56.000000'
>>> dt.isoformat(timespec='auto')
'12:34:56'
```

3.6 新版功能: Added the *timespec* argument.

`time.__str__()`

For a time *t*, `str(t)` is equivalent to `t.isoformat()`.

`time.strftime(format)`

Return a string representing the time, controlled by an explicit format string. For a complete list of formatting directives, see *strftime()* and *strptime()* Behavior.

`time.__format__(format)`

Same as `time.strftime()`. This makes it possible to specify a format string for a *time* object in formatted string literals and when using `str.format()`. For a complete list of formatting directives, see *strftime()* and *strptime()* Behavior.

`time.utcoffset()`

If *tzinfo* is None, returns None, else returns `self.tzinfo.utcoffset(None)`, and raises an exception if the latter doesn't return None or a *timedelta* object representing a whole number of minutes with magnitude less than one day.

`time.dst()`

If *tzinfo* is None, returns None, else returns `self.tzinfo.dst(None)`, and raises an exception if the latter doesn't return None, or a *timedelta* object representing a whole number of minutes with magnitude less than one day.

`time.tzname()`

If *tzinfo* is None, returns None, else returns `self.tzinfo.tzname(None)`, or raises an exception if the latter doesn't return None or a string object.

示例:

```

>>> from datetime import time, tzinfo, timedelta
>>> class GMT1(tzinfo):
...     def utcoffset(self, dt):
...         return timedelta(hours=1)
...     def dst(self, dt):
...         return timedelta(0)
...     def tzname(self, dt):
...         return "Europe/Prague"
...
>>> t = time(12, 10, 30, tzinfo=GMT1())
>>> t                                     # doctest: +ELLIPSIS
datetime.time(12, 10, 30, tzinfo=<GMT1 object at 0x...>)
>>> gmt = GMT1()
>>> t.isoformat()
'12:10:30+01:00'
>>> t.dst()
datetime.timedelta(0)
>>> t.tzname()
'Europe/Prague'
>>> t.strftime("%H:%M:%S %Z")
'12:10:30 Europe/Prague'
>>> 'The {} is {:%H:%M}'.format("time", t)
'The time is 12:10.'

```

8.1.6 tzinfo 对象

class datetime.tzinfo

This is an abstract base class, meaning that this class should not be instantiated directly. You need to derive a concrete subclass, and (at least) supply implementations of the standard *tzinfo* methods needed by the *datetime* methods you use. The *datetime* module supplies a simple concrete subclass of *tzinfo*, *timezone*, which can represent timezones with fixed offset from UTC such as UTC itself or North American EST and EDT.

An instance of (a concrete subclass of) *tzinfo* can be passed to the constructors for *datetime* and *time* objects. The latter objects view their attributes as being in local time, and the *tzinfo* object supports methods revealing offset of local time from UTC, the name of the time zone, and DST offset, all relative to a date or time object passed to them.

Special requirement for pickling: A *tzinfo* subclass must have an `__init__()` method that can be called with no arguments, else it can be pickled but possibly not unpickled again. This is a technical requirement that may be relaxed in the future.

A concrete subclass of *tzinfo* may need to implement the following methods. Exactly which methods are needed depends on the uses made of aware *datetime* objects. If in doubt, simply implement all of them.

tzinfo.utcoffset(dt)

Return offset of local time from UTC, in minutes east of UTC. If local time is west of UTC, this should be negative. Note that this is intended to be the total offset from UTC; for example, if a *tzinfo* object represents both time zone and DST adjustments, *utcoffset()* should return their sum. If the UTC offset isn't known, return *None*. Else the value returned must be a *timedelta* object specifying a whole number of minutes in the range -1439 to 1439 inclusive (1440 = 24*60; the magnitude of the offset must be less than one day). Most implementations of *utcoffset()* will probably look like one of these two:

```

return CONSTANT                # fixed-offset class
return CONSTANT + self.dst(dt) # daylight-aware class

```

If *utcoffset()* does not return *None*, *dst()* should not return *None* either.

The default implementation of `utcoffset()` raises `NotImplementedError`.

`tzinfo.dst(dt)`

Return the daylight saving time (DST) adjustment, in minutes east of UTC, or `None` if DST information isn't known. Return `timedelta(0)` if DST is not in effect. If DST is in effect, return the offset as a `timedelta` object (see `utcoffset()` for details). Note that DST offset, if applicable, has already been added to the UTC offset returned by `utcoffset()`, so there's no need to consult `dst()` unless you're interested in obtaining DST info separately. For example, `datetime.datetime.timetuple()` calls its `tzinfo` attribute's `dst()` method to determine how the `tm_isdst` flag should be set, and `tzinfo.fromutc()` calls `dst()` to account for DST changes when crossing time zones.

An instance `tz` of a `tzinfo` subclass that models both standard and daylight times must be consistent in this sense:

```
tz.utcoffset(dt) - tz.dst(dt)
```

must return the same result for every `datetime dt` with `dt.tzinfo == tz`. For sane `tzinfo` subclasses, this expression yields the time zone's "standard offset", which should not depend on the date or the time, but only on geographic location. The implementation of `datetime.datetime.astimezone()` relies on this, but cannot detect violations; it's the programmer's responsibility to ensure it. If a `tzinfo` subclass cannot guarantee this, it may be able to override the default implementation of `tzinfo.fromutc()` to work correctly with `astimezone()` regardless.

Most implementations of `dst()` will probably look like one of these two:

```
def dst(self, dt):
    # a fixed-offset class: doesn't account for DST
    return timedelta(0)
```

or

```
def dst(self, dt):
    # Code to set dston and dstoff to the time zone's DST
    # transition times based on the input dt.year, and expressed
    # in standard local time. Then

    if dston <= dt.replace(tzinfo=None) < dstoff:
        return timedelta(hours=1)
    else:
        return timedelta(0)
```

The default implementation of `dst()` raises `NotImplementedError`.

`tzinfo.tzname(dt)`

Return the time zone name corresponding to the `datetime` object `dt`, as a string. Nothing about string names is defined by the `datetime` module, and there's no requirement that it mean anything in particular. For example, "GMT", "UTC", "-500", "-5:00", "EDT", "US/Eastern", "America/New York" are all valid replies. Return `None` if a string name isn't known. Note that this is a method rather than a fixed string primarily because some `tzinfo` subclasses will wish to return different names depending on the specific value of `dt` passed, especially if the `tzinfo` class is accounting for daylight time.

The default implementation of `tzname()` raises `NotImplementedError`.

These methods are called by a `datetime` or `time` object, in response to their methods of the same names. A `datetime` object passes itself as the argument, and a `time` object passes `None` as the argument. A `tzinfo` subclass's methods should therefore be prepared to accept a `dt` argument of `None`, or of class `datetime`.

When `None` is passed, it's up to the class designer to decide the best response. For example, returning `None` is appropriate if the class wishes to say that time objects don't participate in the `tzinfo` protocols. It may be more useful for `utcoffset(None)` to return the standard UTC offset, as there is no other convention for discovering the standard offset.

When a `datetime` object is passed in response to a `datetime` method, `dt.tzinfo` is the same object as `self`. `tzinfo` methods can rely on this, unless user code calls `tzinfo` methods directly. The intent is that the `tzinfo` methods interpret `dt` as being in local time, and not need worry about objects in other timezones.

There is one more `tzinfo` method that a subclass may wish to override:

`tzinfo.fromutc(dt)`

This is called from the default `datetime.astimezone()` implementation. When called from that, `dt.tzinfo` is `self`, and `dt`'s date and time data are to be viewed as expressing a UTC time. The purpose of `fromutc()` is to adjust the date and time data, returning an equivalent datetime in `self`'s local time.

Most `tzinfo` subclasses should be able to inherit the default `fromutc()` implementation without problems. It's strong enough to handle fixed-offset time zones, and time zones accounting for both standard and daylight time, and the latter even if the DST transition times differ in different years. An example of a time zone the default `fromutc()` implementation may not handle correctly in all cases is one where the standard offset (from UTC) depends on the specific date and time passed, which can happen for political reasons. The default implementations of `astimezone()` and `fromutc()` may not produce the result you want if the result is one of the hours straddling the moment the standard offset changes.

Skipping code for error cases, the default `fromutc()` implementation acts like:

```
def fromutc(self, dt):
    # raise ValueError error if dt.tzinfo is not self
    dtoff = dt.utcoffset()
    dtdst = dt.dst()
    # raise ValueError if dtoff is None or dtdst is None
    delta = dtoff - dtdst # this is self's standard offset
    if delta:
        dt += delta # convert to standard local time
        dtdst = dt.dst()
        # raise ValueError if dtdst is None
    if dtdst:
        return dt + dtdst
    else:
        return dt
```

Example `tzinfo` classes:

```
from datetime import tzinfo, timedelta, datetime, timezone

ZERO = timedelta(0)
HOUR = timedelta(hours=1)
SECOND = timedelta(seconds=1)

# A class capturing the platform's idea of local time.
# (May result in wrong values on historical times in
# timezones where UTC offset and/or the DST rules had
# changed in the past.)
import time as _time

STDOFFSET = timedelta(seconds = -_time.timezone)
if _time.daylight:
    DSTOFFSET = timedelta(seconds = -_time.altzone)
else:
    DSTOFFSET = STDOFFSET

DSTDIFF = DSTOFFSET - STDOFFSET
```

(下页继续)

(续上页)

```

class LocalTimezone(tzinfo):

    def fromutc(self, dt):
        assert dt.tzinfo is self
        stamp = (dt - datetime(1970, 1, 1, tzinfo=self)) // SECOND
        args = _time.localtime(stamp)[:6]
        dst_diff = DSTDIFF // SECOND
        # Detect fold
        fold = (args == _time.localtime(stamp - dst_diff))
        return datetime(*args, microsecond=dt.microsecond,
                        tzinfo=self, fold=fold)

    def utcoffset(self, dt):
        if self._isdst(dt):
            return DSTOFFSET
        else:
            return STDOFFSET

    def dst(self, dt):
        if self._isdst(dt):
            return DSTDIFF
        else:
            return ZERO

    def tzname(self, dt):
        return _time.tzname[self._isdst(dt)]

    def _isdst(self, dt):
        tt = (dt.year, dt.month, dt.day,
              dt.hour, dt.minute, dt.second,
              dt.weekday(), 0, 0)
        stamp = _time.mktime(tt)
        tt = _time.localtime(stamp)
        return tt.tm_isdst > 0

Local = LocalTimezone()

# A complete implementation of current DST rules for major US time zones.

def first_sunday_on_or_after(dt):
    days_to_go = 6 - dt.weekday()
    if days_to_go:
        dt += timedelta(days_to_go)
    return dt

# US DST Rules
#
# This is a simplified (i.e., wrong for a few cases) set of rules for US
# DST start and end times. For a complete and up-to-date set of DST rules
# and timezone definitions, visit the Olson Database (or try pytz):
# http://www.twinsun.com/tz/tz-link.htm
# http://sourceforge.net/projects/pytz/ (might not be up-to-date)
#
# In the US, since 2007, DST starts at 2am (standard time) on the second

```

(下页继续)

(续上页)

```

# Sunday in March, which is the first Sunday on or after Mar 8.
DSTSTART_2007 = datetime(1, 3, 8, 2)
# and ends at 2am (DST time) on the first Sunday of Nov.
DSTEND_2007 = datetime(1, 11, 1, 2)
# From 1987 to 2006, DST used to start at 2am (standard time) on the first
# Sunday in April and to end at 2am (DST time) on the last
# Sunday of October, which is the first Sunday on or after Oct 25.
DSTSTART_1987_2006 = datetime(1, 4, 1, 2)
DSTEND_1987_2006 = datetime(1, 10, 25, 2)
# From 1967 to 1986, DST used to start at 2am (standard time) on the last
# Sunday in April (the one on or after April 24) and to end at 2am (DST time)
# on the last Sunday of October, which is the first Sunday
# on or after Oct 25.
DSTSTART_1967_1986 = datetime(1, 4, 24, 2)
DSTEND_1967_1986 = DSTEND_1987_2006

def us_dst_range(year):
    # Find start and end times for US DST. For years before 1967, return
    # start = end for no DST.
    if 2006 < year:
        dststart, dstend = DSTSTART_2007, DSTEND_2007
    elif 1986 < year < 2007:
        dststart, dstend = DSTSTART_1987_2006, DSTEND_1987_2006
    elif 1966 < year < 1987:
        dststart, dstend = DSTSTART_1967_1986, DSTEND_1967_1986
    else:
        return (datetime(year, 1, 1), ) * 2

    start = first_sunday_on_or_after(dststart.replace(year=year))
    end = first_sunday_on_or_after(dstend.replace(year=year))
    return start, end

class USTimeZone(tzinfo):

    def __init__(self, hours, reprname, stdname, dstname):
        self.stdoffset = timedelta(hours=hours)
        self.reprname = reprname
        self.stdname = stdname
        self.dstname = dstname

    def __repr__(self):
        return self.reprname

    def tzname(self, dt):
        if self.dst(dt):
            return self.dstname
        else:
            return self.stdname

    def utcoffset(self, dt):
        return self.stdoffset + self.dst(dt)

    def dst(self, dt):
        if dt is None or dt.tzinfo is None:
            # An exception may be sensible here, in one or both cases.

```

(下页继续)

(续上页)

```

        # It depends on how you want to treat them. The default
        # fromutc() implementation (called by the default astimezone()
        # implementation) passes a datetime with dt.tzinfo is self.
        return ZERO
    assert dt.tzinfo is self
    start, end = us_dst_range(dt.year)
    # Can't compare naive to aware objects, so strip the timezone from
    # dt first.
    dt = dt.replace(tzinfo=None)
    if start + HOUR <= dt < end - HOUR:
        # DST is in effect.
        return HOUR
    if end - HOUR <= dt < end:
        # Fold (an ambiguous hour): use dt.fold to disambiguate.
        return ZERO if dt.fold else HOUR
    if start <= dt < start + HOUR:
        # Gap (a non-existent hour): reverse the fold rule.
        return HOUR if dt.fold else ZERO
    # DST is off.
    return ZERO

    def fromutc(self, dt):
        assert dt.tzinfo is self
        start, end = us_dst_range(dt.year)
        start = start.replace(tzinfo=self)
        end = end.replace(tzinfo=self)
        std_time = dt + self.stdoffset
        dst_time = std_time + HOUR
        if end <= dst_time < end + HOUR:
            # Repeated hour
            return std_time.replace(fold=1)
        if std_time < start or dst_time >= end:
            # Standard time
            return std_time
        if start <= std_time < end - HOUR:
            # Daylight saving time
            return dst_time

```

```

Eastern = USTimeZone(-5, "Eastern", "EST", "EDT")
Central = USTimeZone(-6, "Central", "CST", "CDT")
Mountain = USTimeZone(-7, "Mountain", "MST", "MDT")
Pacific = USTimeZone(-8, "Pacific", "PST", "PDT")

```

Note that there are unavoidable subtleties twice per year in a `tzinfo` subclass accounting for both standard and daylight time, at the DST transition points. For concreteness, consider US Eastern (UTC -0500), where EDT begins the minute after 1:59 (EST) on the second Sunday in March, and ends the minute after 1:59 (EDT) on the first Sunday in November:

UTC	3:MM	4:MM	5:MM	6:MM	7:MM	8:MM
EST	22:MM	23:MM	0:MM	1:MM	2:MM	3:MM
EDT	23:MM	0:MM	1:MM	2:MM	3:MM	4:MM
start	22:MM	23:MM	0:MM	1:MM	3:MM	4:MM
end	23:MM	0:MM	1:MM	1:MM	2:MM	3:MM

When DST starts (the “start” line), the local wall clock leaps from 1:59 to 3:00. A wall time of the form 2:MM doesn’t

`t` really make sense on that day, so `astimezone(Eastern)` won't deliver a result with `hour == 2` on the day DST begins. For example, at the Spring forward transition of 2016, we get

```
>>> u0 = datetime(2016, 3, 13, 5, tzinfo=timezone.utc)
>>> for i in range(4):
...     u = u0 + i*HOUR
...     t = u.astimezone(Eastern)
...     print(u.time(), 'UTC =', t.time(), t.tzname())
...
05:00:00 UTC = 00:00:00 EST
06:00:00 UTC = 01:00:00 EST
07:00:00 UTC = 03:00:00 EDT
08:00:00 UTC = 04:00:00 EDT
```

When DST ends (the “end” line), there's a potentially worse problem: there's an hour that can't be spelled unambiguously in local wall time: the last hour of daylight time. In Eastern, that's times of the form 5:MM UTC on the day daylight time ends. The local wall clock leaps from 1:59 (daylight time) back to 1:00 (standard time) again. Local times of the form 1:MM are ambiguous. `astimezone()` mimics the local clock's behavior by mapping two adjacent UTC hours into the same local hour then. In the Eastern example, UTC times of the form 5:MM and 6:MM both map to 1:MM when converted to Eastern, but earlier times have the *fold* attribute set to 0 and the later times have it set to 1. For example, at the Fall back transition of 2016, we get

```
>>> u0 = datetime(2016, 11, 6, 4, tzinfo=timezone.utc)
>>> for i in range(4):
...     u = u0 + i*HOUR
...     t = u.astimezone(Eastern)
...     print(u.time(), 'UTC =', t.time(), t.tzname(), t.fold)
...
04:00:00 UTC = 00:00:00 EDT 0
05:00:00 UTC = 01:00:00 EDT 0
06:00:00 UTC = 01:00:00 EST 1
07:00:00 UTC = 02:00:00 EST 0
```

Note that the *datetime* instances that differ only by the value of the *fold* attribute are considered equal in comparisons.

Applications that can't bear wall-time ambiguities should explicitly check the value of the *fold* attribute or avoid using hybrid *tzinfo* subclasses; there are no ambiguities when using *timezone*, or any other fixed-offset *tzinfo* subclass (such as a class representing only EST (fixed offset -5 hours), or only EDT (fixed offset -4 hours)).

参见:

dateutil.tz The standard library has *timezone* class for handling arbitrary fixed offsets from UTC and *timezone.utc* as UTC *timezone* instance.

dateutil.tz library brings the *IANA timezone database* (also known as the Olson database) to Python and its usage is recommended.

IANA timezone database The Time Zone Database (often called tz, tzdata or zoneinfo) contains code and data that represent the history of local time for many representative locations around the globe. It is updated periodically to reflect changes made by political bodies to time zone boundaries, UTC offsets, and daylight-saving rules.

8.1.7 timezone Objects

The *timezone* class is a subclass of *tzinfo*, each instance of which represents a timezone defined by a fixed offset from UTC. Note that objects of this class cannot be used to represent timezone information in the locations where different offsets are used in different days of the year or where historical changes have been made to civil time.

class `datetime.timezone` (*offset*, *name=None*)

The *offset* argument must be specified as a `timedelta` object representing the difference between the local time and UTC. It must be strictly between `-timedelta(hours=24)` and `timedelta(hours=24)` and represent a whole number of minutes, otherwise `ValueError` is raised.

The *name* argument is optional. If specified it must be a string that will be used as the value returned by the `datetime.tzname()` method.

3.2 新版功能.

`timezone.utcoffset(dt)`

Return the fixed value specified when the `timezone` instance is constructed. The *dt* argument is ignored. The return value is a `timedelta` instance equal to the difference between the local time and UTC.

`timezone.tzname(dt)`

Return the fixed value specified when the `timezone` instance is constructed. If *name* is not provided in the constructor, the name returned by `tzname(dt)` is generated from the value of the *offset* as follows. If *offset* is `timedelta(0)`, the name is “UTC”, otherwise it is a string ‘UTC±HH:MM’, where ± is the sign of *offset*, HH and MM are two digits of *offset.hours* and *offset.minutes* respectively.

在 3.6 版更改: Name generated from *offset*=`timedelta(0)` is now plain ‘UTC’, not ‘UTC+00:00’.

`timezone.dst(dt)`

Always returns None.

`timezone.fromutc(dt)`

Return *dt* + *offset*. The *dt* argument must be an aware `datetime` instance, with *tzinfo* set to *self*.

类属性:

`timezone.utc`

The UTC timezone, `timezone(timedelta(0))`.

8.1.8 strftime() and strptime() Behavior

`date`, `datetime`, and `time` objects all support a `strftime(format)` method, to create a string representing the time under the control of an explicit format string. Broadly speaking, `d.strftime(fmt)` acts like the `time` module’s `time.strftime(fmt, d.timetuple())` although not all objects support a `timetuple()` method.

Conversely, the `datetime.strptime()` class method creates a `datetime` object from a string representing a date and time and a corresponding format string. `datetime.strptime(date_string, format)` is equivalent to `datetime(*(time.strptime(date_string, format)[0:6]))`, except when the format includes sub-second components or timezone offset information, which are supported in `datetime.strptime` but are discarded by `time.strptime`.

For `time` objects, the format codes for year, month, and day should not be used, as time objects have no such values. If they’re used anyway, 1900 is substituted for the year, and 1 for the month and day.

For `date` objects, the format codes for hours, minutes, seconds, and microseconds should not be used, as `date` objects have no such values. If they’re used anyway, 0 is substituted for them.

The full set of format codes supported varies across platforms, because Python calls the platform C library’s `strftime()` function, and platform variations are common. To see the full set of format codes supported on your platform, consult the `strftime(3)` documentation.

The following is a list of all the format codes that the C standard (1989 version) requires, and these work on all platforms with a standard C implementation. Note that the 1999 version of the C standard added additional format codes.

指令	意义	示例	注释
%a	当地工作日的缩写。	Sun, Mon, ..., Sat (美国); So, Mo, ..., Sa (德国)	(1)
%A	当地工作日的全名。	Sunday, Monday, ..., Saturday (美国); Sonntag, Montag, ..., Samstag (德国)	(1)
%w	以十进制数显示的工作日，其中 0 表示星期日，6 表示星期六。	0, 1, ..., 6	
%d	补零后，以十进制数显示的月份中的一天。	01, 02, ..., 31	
%b	当地月份的缩写。	Jan, Feb, ..., Dec (美国); Jan, Feb, ..., Dez (德国)	(1)
%B	当地月份的全名。	January, February, ..., December (美国); Januar, Februar, ..., Dezember (德国)	(1)
%m	补零后，以十进制数显示的月份。	01, 02, ..., 12	
%y	补零后，以十进制数表示的，不带世纪的年份。	00, 01, ..., 99	
%Y	十进制数表示的带世纪的年份。	0001, 0002, ..., 2013, 2014, ..., 9998, 9999	(2)
%H	Hour (24-hour clock) as a zero-padded decimal number.	00, 01, ..., 23	
%I	Hour (12-hour clock) as a zero-padded decimal number.	01, 02, ..., 12	
%p	本地化的 AM 或 PM。	AM, PM (美国); am, pm (德国)	(1), (3)
%M	补零后，以十进制数显示的分钟。	00, 01, ..., 59	
%S	补零后，以十进制数显示的秒。	00, 01, ..., 59	(4)
%f	Microsecond as a decimal number, zero-padded on the left.	000000, 000001, ..., 999999	(5)
%z	UTC offset in the form +HHMM or -HHMM (empty string if the object is naive).	(empty), +0000, -0400, +1030	(6)
190			Chapter 8. 数据类型
%Z	Time zone name (empty string if the object is	(empty), UTC, EST, CST	

Several additional directives not required by the C89 standard are included for convenience. These parameters all correspond to ISO 8601 date values. These may not be available on all platforms when used with the `strptime()` method. The ISO 8601 year and ISO 8601 week directives are not interchangeable with the year and week number directives above. Calling `strptime()` with incomplete or ambiguous ISO 8601 directives will raise a `ValueError`.

指令	意义	示例	注释
%G	ISO 8601 year with century representing the year that contains the greater part of the ISO week (%V).	0001, 0002, ..., 2013, 2014, ..., 9998, 9999	(8)
%u	ISO 8601 weekday as a decimal number where 1 is Monday.	1, 2, ..., 7	
%V	ISO 8601 week as a decimal number with Monday as the first day of the week. Week 01 is the week containing Jan 4.	01, 02, ..., 53	(8)

3.6 新版功能: 增加了 %G, %u 和 %V。

注释:

- (1) Because the format depends on the current locale, care should be taken when making assumptions about the output value. Field orderings will vary (for example, “month/day/year” versus “day/month/year”), and the output may contain Unicode characters encoded using the locale’s default encoding (for example, if the current locale is `ja_JP`, the default encoding could be any one of `eucJP`, `SJIS`, or `utf-8`; use `locale.getlocale()` to determine the current locale’s encoding).
- (2) The `strptime()` method can parse years in the full [1, 9999] range, but years < 1000 must be zero-filled to 4-digit width.
在 3.2 版更改: In previous versions, `strptime()` method was restricted to years >= 1900.
在 3.3 版更改: 在版本 3.2 中, `strptime()` 方法只限于 years >= 1000。
- (3) When used with the `strptime()` method, the %p directive only affects the output hour field if the %I directive is used to parse the hour.
- (4) 与 `time` 模块不同的是, `datetime` 模块不支持闰秒。
- (5) When used with the `strptime()` method, the %f directive accepts from one to six digits and zero pads on the right. %f is an extension to the set of format characters in the C standard (but implemented separately in `datetime` objects, and therefore always available).
- (6) For a naive object, the %z and %Z format codes are replaced by empty strings.

对于一个觉悟型对象而言:

%z `utcoffset()` is transformed into a 5-character string of the form +HHMM or -HHMM, where HH is a 2-digit string giving the number of UTC offset hours, and MM is a 2-digit string giving the number of UTC offset minutes. For example, if `utcoffset()` returns `timedelta(hours=-3, minutes=-30)`, %z is replaced with the string '-0330'.

%Z If `tzname()` returns `None`, %Z is replaced by an empty string. Otherwise %Z is replaced by the returned value, which must be a string.

在 3.2 版更改: When the %z directive is provided to the `strptime()` method, an aware `datetime` object will be produced. The `tzinfo` of the result will be set to a `timezone` instance.

- (7) When used with the `strptime()` method, %U and %W are only used in calculations when the day of the week and the calendar year (%Y) are specified.
- (8) Similar to %U and %W, %V is only used in calculations when the day of the week and the ISO year (%G) are specified in a `strptime()` format string. Also note that %G and %Y are not interchangeable.

8.2 calendar — 日历相关函数

源代码: `Lib/calendar.py`

这个模块让你可以输出像 Unix `cal` 那样的日历, 它还提供了其它与日历相关的实用函数。默认情况下, 这些日历把星期一当作一周的第一天, 星期天为一周的最后一天 (按照欧洲惯例)。可以使用 `setfirstweekday()` 方法设置一周的第一天为星期天 (6) 或者其它任意一天。使用整数作为指定日期的参数。更多相关的函数, 参见 `datetime` 和 `time` 模块。

Most of these functions and classes rely on the `datetime` module which uses an idealized calendar, the current Gregorian calendar extended in both directions. This matches the definition of the “proleptic Gregorian” calendar in Dershowitz and Reingold’s book “Calendrical Calculations”, where it’s the base calendar for all computations.

class `calendar.Calendar` (*firstweekday=0*)

创建一个 `Calendar` 对象。*firstweekday* 是一个整数, 用于指定一周的第一天。0 是星期一 (默认值), 6 是星期天。

`Calendar` 对象提供了一些可被用于准备日历数据格式化的方法。这个类本身不执行任何格式化操作。这部分任务应由子类来完成。

`Calendar` 类的实例有下列方法:

iterweekdays ()

返回一个迭代器, 迭代器的内容为一星期的数字。迭代器的第一个值与 *firstweekday* 属性的值一至。

itermonthdates (*year, month*)

返回一个迭代器, 迭代器的内容为 *year* 年 *month* 月 (1-12) 的日期。这个迭代器返回当月的所有日期 (`datetime.date` 对象), 日期包含了本月头尾用于组成完整一周的日期。

itermonthdays2 (*year, month*)

Return an iterator for the month *month* in the year *year* similar to `itermonthdates()`. Days returned will be tuples consisting of a day number and a week day number.

itermonthdays (*year, month*)

Return an iterator for the month *month* in the year *year* similar to `itermonthdates()`. Days returned will simply be day numbers.

monthdatescalendar (*year, month*)

返回一个表示指定年月的周列表。周列表由七个 `datetime.date` 对象组成。

monthdays2calendar (*year, month*)

返回一个表示指定年月的周列表。周列表由七个代表日期的数字和代表周几的数字组成的二元元组。

monthdayscalendar (*year, month*)

返回一个表示指定年月的周列表。周列表由七个代表日期的数字组成。

yeardatescalendar (*year, width=3*)

返回可以用来格式化的指定年月的数据。返回的值是一个列表, 列表是月份组成的行。每一行包含了最多 *width* 个月 (默认为 3)。每个月包含了 4 到 6 周, 每周包含 1-7 天。每一天使用 `datetime.date` 对象。

yeardays2calendar (*year, width=3*)

返回可以用来模式化的指定年月的数据 (与 `yeardatescalendar()` 类似)。周列表的元素是由表示日期的数字和表示星期几的数字组成的元组。不在这个月的日子为 0。

yeardayscalendar (*year*, *width*=3)

返回可以用来模式化的指定年月的数据 (与 `yeardatescalendar()` 类似)。周列表的元素是表示日期的数字。不在这个月的日子为 0。

class `calendar.TextCalendar` (*firstweekday*=0)

可以使用这个类生成纯文本日历。

`TextCalendar` 实例有以下方法：

formatmonth (*theyear*, *themonth*, *w*=0, *l*=0)

返回一个多行字符串来表示指定年月的日历。*w* 为日期的宽度，但始终保持日期居中。*l* 指定了每星期占用的行数。以上这些还依赖于构造器或者 `setfirstweekday()` 方法指定的周的第一天是哪一天。

prmonth (*theyear*, *themonth*, *w*=0, *l*=0)

与 `formatmonth()` 方法一样，返回一个月的日历。

formatyear (*theyear*, *w*=2, *l*=1, *c*=6, *m*=3)

返回一个多行字符串，这个字符串为一个 *m* 列日历。可选参数 *w*, *l* 和 *c* 分别表示日期列数，周的行数，和月之间的间隔。同样，以上这些还依赖于构造器或者 `setfirstweekday()` 指定哪一天为一周的第一天。日历的第一年由平台依赖于使用的平台。

pryear (*theyear*, *w*=2, *l*=1, *c*=6, *m*=3)

与 `formatyear()` 方法一样，返回一整年的日历。

class `calendar.HTMLCalendar` (*firstweekday*=0)

可以使用这个类生成 HTML 日历。

`HTMLCalendar` instances have the following methods:

formatmonth (*theyear*, *themonth*, *withyear*=True)

返回一个 HTML 表格作为指定年月的日历。*withyear* 为真，则年份将会包含在表头，否则只显示月份。

formatyear (*theyear*, *width*=3)

返回一个 HTML 表格作为指定年份的日历。*width* (默认为 3) 用于规定每一行显示月份的数量。

formatyearpage (*theyear*, *width*=3, *css*='calendar.css', *encoding*=None)

返回一个完整的 HTML 页面作为指定年份的日历。*width**(默认为 3) 用于规定每一行显示的月份数量。**css* 为层叠样式表的名字。如果不使用任何层叠样式表，可以使用 `None`。*encoding* 为输出页面的编码 (默认为系统的默认编码)。

class `calendar.LocaleTextCalendar` (*firstweekday*=0, *locale*=None)

This subclass of `TextCalendar` can be passed a locale name in the constructor and will return month and weekday names in the specified locale. If this locale includes an encoding all strings containing month and weekday names will be returned as unicode.

class `calendar.LocaleHTMLCalendar` (*firstweekday*=0, *locale*=None)

This subclass of `HTMLCalendar` can be passed a locale name in the constructor and will return month and weekday names in the specified locale. If this locale includes an encoding all strings containing month and weekday names will be returned as unicode.

注解： The `formatweekday()` and `formatmonthname()` methods of these two classes temporarily change the current locale to the given *locale*. Because the current locale is a process-wide setting, they are not thread-safe.

对于简单的文本日历，这个模块提供了以下方法。

`calendar.setfirstweekday` (*weekday*)

设置每一周的开始 (0 表示星期一，6 表示星期天)。`calendar` 还提供了 MONDAY, TUESDAY, WEDNESDAY, THURSDAY, FRIDAY, SATURDAY 和 SUNDAY 几个常量方便使用。例如，设置每周的第一天为星期天

```
import calendar
calendar.setfirstweekday(calendar.SUNDAY)
```

`calendar.firstweekday()`

返回当前设置的每星期的第一天的数值。

`calendar.isleap(year)`

如果 *year* 是闰年则返回 *True* , 否则返回 *False*。

`calendar.leapdays(y1, y2)`

Returns the number of leap years in the range from *y1* to *y2* (exclusive), where *y1* and *y2* are years.

This function works for ranges spanning a century change.

`calendar.weekday(year, month, day)`

Returns the day of the week (0 is Monday) for *year* (1970–…), *month* (1–12), *day* (1–31).

`calendar.weekheader(n)`

Return a header containing abbreviated weekday names. *n* specifies the width in characters for one weekday.

`calendar.monthrange(year, month)`

Returns weekday of first day of the month and number of days in month, for the specified *year* and *month*.

`calendar.monthcalendar(year, month)`

Returns a matrix representing a month's calendar. Each row represents a week; days outside of the month are represented by zeros. Each week begins with Monday unless set by `setfirstweekday()`.

`calendar.prmonth(theyear, themonth, w=0, l=0)`

Prints a month's calendar as returned by `month()`.

`calendar.month(theyear, themonth, w=0, l=0)`

Returns a month's calendar in a multi-line string using the `formatmonth()` of the *TextCalendar* class.

`calendar.prcal(year, w=0, l=0, c=6, m=3)`

Prints the calendar for an entire year as returned by `calendar()`.

`calendar.calendar(year, w=2, l=1, c=6, m=3)`

Returns a 3-column calendar for an entire year as a multi-line string using the `formatyear()` of the *TextCalendar* class.

`calendar.timegm(tuple)`

An unrelated but handy function that takes a time tuple such as returned by the `gmtime()` function in the *time* module, and returns the corresponding Unix timestamp value, assuming an epoch of 1970, and the POSIX encoding. In fact, `time.gmtime()` and `timegm()` are each others' inverse.

The *calendar* module exports the following data attributes:

`calendar.day_name`

An array that represents the days of the week in the current locale.

`calendar.day_abbr`

An array that represents the abbreviated days of the week in the current locale.

`calendar.month_name`

An array that represents the months of the year in the current locale. This follows normal convention of January being month number 1, so it has a length of 13 and `month_name[0]` is the empty string.

`calendar.month_abbr`

An array that represents the abbreviated months of the year in the current locale. This follows normal convention of January being month number 1, so it has a length of 13 and `month_abbr[0]` is the empty string.

参见:

模块 `datetime` 为日期和时间提供与 `time` 模块相似功能的面向对象接口。

模块 `time` 底层时间相关函数。

8.3 collections — 容器数据类型

Source code: `Lib/collections/__init__.py`

这个模块实现了特定目标的容器，以提供 Python 标准内建容器 `dict` , `list` , `set` , 和 `tuple` 的替代选择。

<code>namedtuple()</code>	创建命名元组子类的工厂函数
<code>deque</code>	类似列表 (list) 的容器，实现了在两端快速添加 (append) 和弹出 (pop)
<code>ChainMap</code>	类似字典 (dict) 的容器类，将多个映射集合到一个视图里面
<code>Counter</code>	字典的子类，提供了可哈希对象的计数功能
<code>OrderedDict</code>	字典的子类，保存了他们被添加的顺序
<code>defaultdict</code>	字典的子类，提供了一个工厂函数，为字典查询提供一个默认值
<code>UserDict</code>	封装了字典对象，简化了字典子类化
<code>UserList</code>	封装了列表对象，简化了列表子类化
<code>UserString</code>	封装了列表对象，简化了字符串子类化

在 3.3 版更改: Moved *Collections Abstract Base Classes* to the `collections.abc` module. For backwards compatibility, they continue to be visible in this module as well.

8.3.1 ChainMap 对象

3.3 新版功能.

一个 `ChainMap` 类是为了将多个映射快速的链接到一起，这样它们就可以作为一个单元处理。它通常比创建一个新字典和多次调用 `update()` 要快很多。

这个类可以用于模拟嵌套作用域，并且在模版化的时候比较有用。

class `collections.ChainMap(*maps)`

一个 `ChainMap` 将多个字典或者其他映射组合在一起，创建一个单独的可更新的视图。如果没有 `maps` 被指定，就提供一个默认的空字典，这样一个新链至少有一个映射。

底层映射被存储在一个列表中。这个列表是公开的，可以通过 `maps` 属性存取和更新。没有其他的状态。搜索查询底层映射，直到一个键被找到。不同的是，写，更新和删除只操作第一个映射。

一个 `ChainMap` 通过引用合并底层映射。所以，如果一个底层映射更新了，这些更改会反映到 `ChainMap`。

支持所有常用字典方法。另外还有一个 `maps` 属性 (attribute)，一个创建子上下文的方法 (method)，一个存取它们首个映射的属性 (property):

maps

一个可以更新的映射列表。这个列表是按照第一次搜索到最后一次搜索的顺序组织的。它是仅有的存储状态，可以被修改。列表最少包含一个映射。

new_child(m=None)

返回一个新的 `ChainMap` 类，包含了一个新映射 (map)，后面跟随当前实例的全部映射 (map)。如果 `m` 被指定，它就成为不同新的实例，就是在所有映射前加上 `m`，如果没有指定，就加上一个空

字典，这样的话一个 `d.new_child()` 调用等价于 `ChainMap({}, *d.maps)`。这个方法用于创建子上下文，不改变任何父映射的值。

在 3.4 版更改: 添加了 `m` 可选参数。

parents

属性返回一个新的 *ChainMap* 包含所有的当前实例的映射，除了第一个。这样可以在搜索的时候跳过第一个映射。使用的场景类似在 *nested scopes* 嵌套作用域中使用 `nonlocal` 关键词。用例也可以类比内建函数 `super()`。一个 `d.parents` 的引用等价于 `ChainMap(*d.maps[1:])`。

参见:

- *MultiContext class* 在 *Enthought CodeTools package* 有支持写映射的选项。
- Django' s *Context class* for templating is a read-only chain of mappings. It also features pushing and popping of contexts similar to the *new_child()* method and the *parents()* property.
- *Nested Contexts recipe* 提供了是否对第一个映射或其他映射进行写和其他修改的选项。
- 一个 极简的只读版 *Chainmap*.

ChainMap 例子和方法

这一节提供了多个使用链映射的案例。

模拟 Python 内部 lookup 链的例子

```
import builtins
pylookup = ChainMap(locals(), globals(), vars(builtins))
```

让用户指定的命令行参数优先于环境变量，优先于默认值的例子

```
import os, argparse

defaults = {'color': 'red', 'user': 'guest'}

parser = argparse.ArgumentParser()
parser.add_argument('-u', '--user')
parser.add_argument('-c', '--color')
namespace = parser.parse_args()
command_line_args = {k:v for k, v in vars(namespace).items() if v}

combined = ChainMap(command_line_args, os.environ, defaults)
print(combined['color'])
print(combined['user'])
```

用 *ChainMap* 类模拟嵌套上下文的例子

```
c = ChainMap()           # Create root context
d = c.new_child()        # Create nested child context
e = c.new_child()        # Child of c, independent from d
e.maps[0]                # Current context dictionary -- like Python's locals()
e.maps[-1]               # Root context -- like Python's globals()
e.parents                # Enclosing context chain -- like Python's nonlocals

d['x']                   # Get first key in the chain of contexts
d['x'] = 1               # Set value in current context
del d['x']                # Delete from current context
list(d)                  # All nested values
```

(下页继续)

(续上页)

```

k in d          # Check all nested values
len(d)          # Number of nested values
d.items()       # All nested items
dict(d)         # Flatten into a regular dictionary

```

`ChainMap` 类只更新链中的第一个映射，但 `lookup` 会搜索整个链。然而，如果需要深度写和删除，也可以很容易的通过定义一个子类来实现它

```

class DeepChainMap(ChainMap):
    'Variant of ChainMap that allows direct updates to inner scopes'

    def __setitem__(self, key, value):
        for mapping in self.maps:
            if key in mapping:
                mapping[key] = value
                return
        self.maps[0][key] = value

    def __delitem__(self, key):
        for mapping in self.maps:
            if key in mapping:
                del mapping[key]
                return
        raise KeyError(key)

>>> d = DeepChainMap({'zebra': 'black'}, {'elephant': 'blue'}, {'lion': 'yellow'})
>>> d['lion'] = 'orange'          # update an existing key two levels down
>>> d['snake'] = 'red'           # new keys get added to the topmost dict
>>> del d['elephant']            # remove an existing key one level down
>>> d                            # display result
DeepChainMap({'zebra': 'black', 'snake': 'red'}, {}, {'lion': 'orange'})

```

8.3.2 Counter 对象

一个计数器工具提供快速和方便的计数。比如

```

>>> # Tally occurrences of words in a list
>>> cnt = Counter()
>>> for word in ['red', 'blue', 'red', 'green', 'blue', 'blue']:
...     cnt[word] += 1
>>> cnt
Counter({'blue': 3, 'red': 2, 'green': 1})

>>> # Find the ten most common words in Hamlet
>>> import re
>>> words = re.findall(r'\w+', open('hamlet.txt').read().lower())
>>> Counter(words).most_common(10)
[('the', 1143), ('and', 966), ('to', 762), ('of', 669), ('i', 631),
 ('you', 554), ('a', 546), ('my', 514), ('hamlet', 471), ('in', 451)]

```

class `collections.Counter` (`[iterable-or-mapping]`)

A `Counter` is a `dict` subclass for counting hashable objects. It is an unordered collection where elements are stored as dictionary keys and their counts are stored as dictionary values. Counts are allowed to be any integer value including zero or negative counts. The `Counter` class is similar to bags or multisets in other languages.

元素从一个 *iterable* 被计数或从其他的 *mapping* (or *counter*) 初始化:

```
>>> c = Counter()                    # a new, empty counter
>>> c = Counter('gallahad')         # a new counter from an iterable
>>> c = Counter({'red': 4, 'blue': 2}) # a new counter from a mapping
>>> c = Counter(cats=4, dogs=8)      # a new counter from keyword args
```

`Counter` 对象有一个字典接口, 如果引用的键没有任何记录, 就返回一个 0, 而不是弹出一个 `KeyError`:

```
>>> c = Counter(['eggs', 'ham'])
>>> c['bacon']                        # count of a missing element is zero
0
```

设置一个计数为 0 不会从计数器中移去一个元素。使用 `del` 来删除它:

```
>>> c['sausage'] = 0                 # counter entry with a zero count
>>> del c['sausage']                 # del actually removes the entry
```

3.1 新版功能.

计数器对象除了字典方法以外, 还提供了三个其他的方法:

`elements()`

返回一个迭代器, 每个元素重复计数的个数。元素顺序是任意的。如果一个元素的计数小于 1, `elements()` 就会忽略它。

```
>>> c = Counter(a=4, b=2, c=0, d=-2)
>>> sorted(c.elements())
['a', 'a', 'a', 'a', 'b', 'b']
```

`most_common([n])`

Return a list of the *n* most common elements and their counts from the most common to the least. If *n* is omitted or `None`, `most_common()` returns *all* elements in the counter. Elements with equal counts are ordered arbitrarily:

```
>>> Counter('abracadabra').most_common(3) # doctest: +SKIP
[('a', 5), ('r', 2), ('b', 2)]
```

`subtract([iterable-or-mapping])`

从迭代对象或映射对象减去元素。像 `dict.update()` 但是是减去, 而不是替换。输入和输出都可以是 0 或者负数。

```
>>> c = Counter(a=4, b=2, c=0, d=-2)
>>> d = Counter(a=1, b=2, c=3, d=4)
>>> c.subtract(d)
>>> c
Counter({'a': 3, 'b': 0, 'c': -3, 'd': -6})
```

3.2 新版功能.

通常字典方法都可用于 `Counter` 对象, 除了有两个方法工作方式与字典并不相同。

`fromkeys(iterable)`

这个类方法没有在 `Counter` 中实现。

`update([iterable-or-mapping])`

从迭代对象计数元素或者从另一个映射对象 (或计数器) 添加。像 `dict.update()` 但是是加上, 而不是替换。另外, 迭代对象应该是序列元素, 而不是一个 (key, value) 对。

`Counter` 对象的常用案例

```

sum(c.values())           # total of all counts
c.clear()                 # reset all counts
list(c)                   # list unique elements
set(c)                    # convert to a set
dict(c)                   # convert to a regular dictionary
c.items()                 # convert to a list of (elem, cnt) pairs
Counter(dict(list_of_pairs)) # convert from a list of (elem, cnt) pairs
c.most_common()[:n-1:-1]  # n least common elements
+c                        # remove zero and negative counts

```

提供了几个数学操作，可以结合 `Counter` 对象，以生产 **multisets** (计数器中大于 0 的元素)。加和减，结合计数器，通过加上或者减去元素的相应计数。交集和并集返回相应计数的最小或最大值。每种操作都可以接受带符号的计数，但是输出会忽略掉结果为零或者小于零的计数。

```

>>> c = Counter(a=3, b=1)
>>> d = Counter(a=1, b=2)
>>> c + d                               # add two counters together: c[x] + d[x]
Counter({'a': 4, 'b': 3})
>>> c - d                               # subtract (keeping only positive counts)
Counter({'a': 2})
>>> c & d                               # intersection: min(c[x], d[x]) # doctest: +SKIP
Counter({'a': 1, 'b': 1})
>>> c | d                               # union: max(c[x], d[x])
Counter({'a': 3, 'b': 2})

```

单目加和减（一元操作符）意思是从空计数器加或者减去。

```

>>> c = Counter(a=2, b=-4)
>>> +c
Counter({'a': 2})
>>> -c
Counter({'b': 4})

```

3.3 新版功能: 添加了对一元加，一元减和位置集合操作的支持。

注解：计数器主要是为了表达运行的正的计数而设计；但是，小心不要预先排除负数或者其他类型。为了帮助这些用例，这一节记录了最小范围和类型限制。

- `Counter` 类是一个字典的子类，不限制键和值。值用于表示计数，但你实际上可以存储任何其他值。
- The `most_common()` method requires only that the values be orderable.
- For in-place operations such as `c[key] += 1`, the value type need only support addition and subtraction. So fractions, floats, and decimals would work and negative values are supported. The same is also true for `update()` and `subtract()` which allow negative and zero values for both inputs and outputs.
- Multiset 多集合方法只为正值的使用情况设计。输入可以是负数或者 0，但只输出计数为正值。没有类型限制，但值类型需要支持加，减和比较操作。
- The `elements()` method requires integer counts. It ignores zero and negative counts.

参见：

- [Bag class](#) 在 Smalltalk。
- Wikipedia 链接 [Multisets](#)。
- [C++ multisets](#) 教程和例子。

- 数学操作和多集合用例，参考 *Knuth, Donald. The Art of Computer Programming Volume II, Section 4.6.3, Exercise 19*。
- To enumerate all distinct multisets of a given size over a given set of elements, see `itertools.combinations_with_replacement()`:

```
map(Counter, combinations_with_replacement('ABC', 2)) -> AA AB AC BB BC CC
```

8.3.3 deque 对象

class `collections.deque([iterable[, maxlen]])`

返回一个新的双向队列对象，从左到右初始化（用方法 `append()`），从 `iterable`（迭代对象）数据创建。如果 `iterable` 没有指定，新队列为空。

Deque 队列是由栈或者 `queue` 队列生成的（发音是“deck”，“double-ended queue”的简称）。Deque 支持线程安全，内存高效添加（`append`）和弹出（`pop`），从两端都可以，两个方向的大概开销都是 $O(1)$ 复杂度。

虽然 `list` 对象也支持类似操作，不过这里优化了定长操作和 `pop(0)` 和 `insert(0, v)` 的开销。它们引起 $O(n)$ 内存移动的操作，改变底层数据表达的大小和位置。

如果 `maxlen` 没有指定或者是 `None`，`deque`s 可以增长到任意长度。否则，`deque` 就限定到指定最大长度。一旦限定长度的 `deque` 满了，当新项加入时，同样数量的项就从另一端弹出。限定长度 `deque` 提供类似 Unix `filter tail` 的功能。它们同样可以用与追踪最近的交换和其他数据池活动。

双向队列（`deque`）对象支持以下方法：

append (*x*)

添加 *x* 到右端。

appendleft (*x*)

添加 *x* 到左端。

clear ()

移除所有元素，使其长度为 0。

copy ()

创建一份浅拷贝。

3.5 新版功能。

count (*x*)

计算 `deque` 中个数等于 *x* 的元素。

3.2 新版功能。

extend (*iterable*)

扩展 `deque` 的右侧，通过添加 `iterable` 参数中的元素。

extendleft (*iterable*)

扩展 `deque` 的左侧，通过添加 `iterable` 参数中的元素。注意，左添加时，在结果中 `iterable` 参数中的顺序将被反过来添加。

index (*x*[, *start*[, *stop*]])

返回第 *x* 个元素（从 *start* 开始计算，在 *stop* 之前）。返回第一个匹配，如果没找到的话，升起 `ValueError`。

3.5 新版功能。

insert (*i*, *x*)

在位置 *i* 插入 *x*。

如果插入会导致一个限长 `deque` 超出长度 `maxlen` 的话，就升起一个 `IndexError`。

3.5 新版功能.

pop()

移去并且返回一个元素，deque 最右侧的那一个。如果没有元素的话，就升起 `IndexError` 索引错误。

popleft()

移去并且返回一个元素，deque 最左侧的那一个。如果没有元素的话，就升起 `IndexError` 索引错误。

remove(value)

移去找到的第一个 `value`。如果没有的话就升起 `ValueError`。

reverse()

将 deque 逆序排列。返回 `None`。

3.2 新版功能.

rotate(n=1)

向右循环移动 `n` 步。如果 `n` 是负数，就向左循环。

如果 deque 不是空的，向右循环移动一步就等价于 `d.appendleft(d.pop())`，向左循环一步就等价于 `d.append(d.popleft())`。

Deque 对象同样提供了一个只读属性：

maxlen

Deque 的最大尺寸，如果没有限定的话就是 `None`。

3.1 新版功能.

除了以上，deque 还支持迭代，清洗，`len(d)`，`reversed(d)`，`copy.copy(d)`，`copy.deepcopy(d)`，成员测试 `in` 操作符，和下标引用 `d[-1]`。索引存取在两端的复杂度是 $O(1)$ ，在中间的复杂度比 $O(n)$ 略低。要快速存取，使用 `list` 来替代。

Deque 从版本 3.5 开始支持 `__add__()`，`__mul__()`，和 `__imul__()`。

示例：

```
>>> from collections import deque
>>> d = deque('ghi')                # make a new deque with three items
>>> for elem in d:                  # iterate over the deque's elements
...     print(elem.upper())
G
H
I

>>> d.append('j')                   # add a new entry to the right side
>>> d.appendleft('f')               # add a new entry to the left side
>>> d                               # show the representation of the deque
deque(['f', 'g', 'h', 'i', 'j'])

>>> d.pop()                         # return and remove the rightmost item
'j'
>>> d.popleft()                     # return and remove the leftmost item
'f'
>>> list(d)                         # list the contents of the deque
['g', 'h', 'i']
>>> d[0]                           # peek at leftmost item
'g'
>>> d[-1]                          # peek at rightmost item
'i'
```

(下页继续)

(续上页)

```

>>> list(reversed(d))           # list the contents of a deque in reverse
['i', 'h', 'g']
>>> 'h' in d                     # search the deque
True
>>> d.extend('jkl')             # add multiple elements at once
>>> d
deque(['g', 'h', 'i', 'j', 'k', 'l'])
>>> d.rotate(1)                 # right rotation
>>> d
deque(['l', 'g', 'h', 'i', 'j', 'k'])
>>> d.rotate(-1)                # left rotation
>>> d
deque(['g', 'h', 'i', 'j', 'k', 'l'])

>>> deque(reversed(d))          # make a new deque in reverse order
deque(['l', 'k', 'j', 'i', 'h', 'g'])
>>> d.clear()                   # empty the deque
>>> d.pop()                     # cannot pop from an empty deque
Traceback (most recent call last):
  File "<pyshell#6>", line 1, in <module>
    d.pop()
IndexError: pop from an empty deque

>>> d.extendleft('abc')         # extendleft() reverses the input order
>>> d
deque(['c', 'b', 'a'])

```

deque 用法

这一节展示了 deque 的多种用法。

限长 deque 提供了类似 Unix tail 过滤功能

```

def tail(filename, n=10):
    'Return the last n lines of a file'
    with open(filename) as f:
        return deque(f, n)

```

另一个用法是维护一个近期添加元素的序列，通过从右边添加和从左边弹出

```

def moving_average(iterable, n=3):
    # moving_average([40, 30, 50, 46, 39, 44]) --> 40.0 42.0 45.0 43.0
    # http://en.wikipedia.org/wiki/Moving_average
    it = iter(iterable)
    d = deque(itertools.islice(it, n-1))
    d.appendleft(0)
    s = sum(d)
    for elem in it:
        s += elem - d.popleft()
        d.append(elem)
        yield s / n

```

The rotate() method provides a way to implement deque slicing and deletion. For example, a pure Python implementation of del d[n] relies on the rotate() method to position elements to be popped:

```
def delete_nth(d, n):
    d.rotate(-n)
    d.popleft()
    d.rotate(n)
```

To implement *deque* slicing, use a similar approach applying `rotate()` to bring a target element to the left side of the deque. Remove old entries with `popleft()`, add new entries with `extend()`, and then reverse the rotation. With minor variations on that approach, it is easy to implement Forth style stack manipulations such as `dup`, `drop`, `swap`, `over`, `pick`, `rot`, and `roll`.

8.3.4 defaultdict 对象

class `collections.defaultdict` (`[default_factory[, ...]]`)

返回一个新的类似字典的对象。*defaultdict* 是内置 *dict* 类的子类。它重载了一个方法并添加了一个可写的实例变量。其余的功能与 *dict* 类相同，此处不再重复说明。

第一个参数 *default_factory* 提供了一个初始值。它默认为 `None`。所有的其他参数都等同与 *dict* 构建器中的参数对待，包括关键词参数。

defaultdict 对象除了支持 *dict* 的操作，还支持下面的方法作为扩展：

`__missing__` (*key*)

如果 *default_factory* 是 `None`，它就升起一个 *KeyError* 并将 *key* 作为参数。

如果 *default_factory* 不为 `None`，它就会被调用，不带参数，为 *key* 提供一个默认值，这个值和 *key* 作为一个对被插入到字典中，并返回。

如果调用 *default_factory* 升起了一个例外，这个例外就被扩散传递，不经过改变。

这个方法在查询键值失败时，会被 *dict* 中的 `__getitem__()` 调用。不管它是返回值或升起例外，都会被 `__getitem__()` 传递。

注意 `__missing__()` 不会被 `__getitem__()` 以外的其他方法调用。意思就是 `get()` 会向正常的 *dict* 那样返回 `None`，而不是使用 *default_factory*。

defaultdict 支持以下实例变量：

`default_factory`

这个属性被 `__missing__()` 方法使用；它从构建器的第一个参数初始化，如果提供了的话，否则就是 `None`。

defaultdict 例子

Using *list* as the *default_factory*, it is easy to group a sequence of key-value pairs into a dictionary of lists:

```
>>> s = [('yellow', 1), ('blue', 2), ('yellow', 3), ('blue', 4), ('red', 1)]
>>> d = defaultdict(list)
>>> for k, v in s:
...     d[k].append(v)
...
>>> sorted(d.items())
[('blue', [2, 4]), ('red', [1]), ('yellow', [1, 3])]
```

When each key is encountered for the first time, it is not already in the mapping; so an entry is automatically created using the *default_factory* function which returns an empty *list*. The `list.append()` operation then attaches the value to the new list. When keys are encountered again, the look-up proceeds normally (returning the list for that key) and

the `list.append()` operation adds another value to the list. This technique is simpler and faster than an equivalent technique using `dict.setdefault()`:

```
>>> d = {}
>>> for k, v in s:
...     d.setdefault(k, []).append(v)
...
>>> sorted(d.items())
[('blue', [2, 4]), ('red', [1]), ('yellow', [1, 3])]
```

Setting the `default_factory` to `int` makes the `defaultdict` useful for counting (like a bag or multiset in other languages):

```
>>> s = 'mississippi'
>>> d = defaultdict(int)
>>> for k in s:
...     d[k] += 1
...
>>> sorted(d.items())
[('i', 4), ('m', 1), ('p', 2), ('s', 4)]
```

When a letter is first encountered, it is missing from the mapping, so the `default_factory` function calls `int()` to supply a default count of zero. The increment operation then builds up the count for each letter.

函数 `int()` 总是返回 0，是常数函数的特殊情况。一个更快和灵活的方法是使用 `lambda` 函数，可以提供任何常量值 (不只是 0)：

```
>>> def constant_factory(value):
...     return lambda: value
>>> d = defaultdict(constant_factory('<missing>'))
>>> d.update(name='John', action='ran')
>>> '%(name)s %(action)s to %(object)s' % d
'John ran to <missing>'
```

Setting the `default_factory` to `set` makes the `defaultdict` useful for building a dictionary of sets:

```
>>> s = [('red', 1), ('blue', 2), ('red', 3), ('blue', 4), ('red', 1), ('blue', 4)]
>>> d = defaultdict(set)
>>> for k, v in s:
...     d[k].add(v)
...
>>> sorted(d.items())
[('blue', {2, 4}), ('red', {1, 3})]
```

8.3.5 `namedtuple()` 命名元组的工厂函数

命名元组赋予每个位置一个含义，提供可读性和自文档性。它们可以用于任何普通元组，并添加了通过名字获取值的能力，通过索引值也是可以的。

`collections.namedtuple` (*typename*, *field_names*, *, *verbose=False*, *rename=False*, *module=None*)

返回一个新的元组子类，名为 *typename*。这个新的子类用于创建类元组的对象，可以通过域名来获取属性值，同样也可以通过索引和迭代获取值。子类实例同样有文档字符串（类名和域名）另外一个有用的 `__repr__()` 方法，以 `name=value` 格式列明了元组内容。

field_names 是一个像 `['x', 'y']` 一样的字符串序列。另外 *field_names* 可以是一个纯字符串，用空白或逗号分隔开元素名，比如 `'x y'` 或者 `'x, y'`。

任何有效的 Python 标识符都可以作为域名，除了下划线开头的那些。有效标识符由字母，数字，下划线组成，但首字母不能是数字或下划线，另外不能是关键词 *keyword* 比如 *class*, *for*, *return*, *global*, *pass*, 或 *raise*。

如果 *rename* 为真，无效域名会自动转换成位置名。比如 ['abc', 'def', 'ghi', 'abc'] 转换成 ['abc', '_1', 'ghi', '_3']，消除关键词 *def* 和重复域名 *abc*。

If *verbose* is true, the class definition is printed after it is built. This option is outdated; instead, it is simpler to print the `__source__` attribute.

如果 *module* 值有定义，命名元组的 `__module__` 属性值就被设置。

命名元组实例没有字典，所以它们要更轻量，并且占用更小内存。

在 3.1 版更改: 添加了对 *rename* 的支持。

在 3.6 版更改: *verbose* 和 *rename* 参数成为仅限关键字参数。

在 3.6 版更改: 添加了 *module* 参数。

```
>>> # Basic example
>>> Point = namedtuple('Point', ['x', 'y'])
>>> p = Point(11, y=22)           # instantiate with positional or keyword arguments
>>> p[0] + p[1]                   # indexable like the plain tuple (11, 22)
33
>>> x, y = p                      # unpack like a regular tuple
>>> x, y
(11, 22)
>>> p.x + p.y                    # fields also accessible by name
33
>>> p                            # readable __repr__ with a name=value style
Point(x=11, y=22)
```

命名元组尤其有用于赋值 *csv* *sqlite3* 模块返回的元组

```
EmployeeRecord = namedtuple('EmployeeRecord', 'name, age, title, department, paygrade
→ ')

import csv
for emp in map(EmployeeRecord._make, csv.reader(open("employees.csv", "rb"))):
    print(emp.name, emp.title)

import sqlite3
conn = sqlite3.connect('/companydata')
cursor = conn.cursor()
cursor.execute('SELECT name, age, title, department, paygrade FROM employees')
for emp in map(EmployeeRecord._make, cursor.fetchall()):
    print(emp.name, emp.title)
```

除了继承元组的方法，命名元组还支持三个额外的方法和两个属性。为了防止域名冲突，方法和属性以下划线开始。

classmethod `somenamedtuple._make(iterable)`

类方法从存在的序列或迭代实例创建一个新实例。

```
>>> t = [11, 22]
>>> Point._make(t)
Point(x=11, y=22)
```

`somenamedtuple._asdict()`

Return a new *OrderedDict* which maps field names to their corresponding values:

```
>>> p = Point(x=11, y=22)
>>> p._asdict()
OrderedDict([('x', 11), ('y', 22)])
```

在 3.1 版更改: 返回一个 *OrderedDict* 而不是 *dict*。

`somenamedtuple._replace(**kwargs)`

返回一个新的命名元组实例, 并将指定域替换为新的值

```
>>> p = Point(x=11, y=22)
>>> p._replace(x=33)
Point(x=33, y=22)

>>> for partnum, record in inventory.items():
...     inventory[partnum] = record._replace(price=newprices[partnum],
↪ timestamp=time.now())
```

`somenamedtuple._source`

A string with the pure Python source code used to create the named tuple class. The source makes the named tuple self-documenting. It can be printed, executed using `exec()`, or saved to a file and imported.

3.3 新版功能.

`somenamedtuple._fields`

字符串元组列出了域名。用于提醒和从现有元组创建一个新的命名元组类型。

```
>>> p._fields          # view the field names
('x', 'y')

>>> Color = namedtuple('Color', 'red green blue')
>>> Pixel = namedtuple('Pixel', Point._fields + Color._fields)
>>> Pixel(11, 22, 128, 255, 0)
Pixel(x=11, y=22, red=128, green=255, blue=0)
```

要获取这个名字域的值, 使用 `getattr()` 函数:

```
>>> getattr(p, 'x')
11
```

To convert a dictionary to a named tuple, use the double-star-operator (as described in [tut-unpacking-arguments](#)):

```
>>> d = {'x': 11, 'y': 22}
>>> Point(**d)
Point(x=11, y=22)
```

因为一个命名元组是一个正常的 Python 类, 它可以很容易的通过子类更改功能。这里是如何添加一个计算域和定宽输出打印格式:

```
>>> class Point(namedtuple('Point', ['x', 'y'])):
...     __slots__ = ()
...     @property
...     def hypot(self):
...         return (self.x ** 2 + self.y ** 2) ** 0.5
...     def __str__(self):
...         return 'Point: x=%6.3f y=%6.3f hypot=%6.3f' % (self.x, self.y, self.
↪ hypot)

>>> for p in Point(3, 4), Point(14, 5/7):
```

(下页继续)

(续上页)

```
...     print(p)
Point: x= 3.000  y= 4.000  hypot= 5.000
Point: x=14.000  y= 0.714  hypot=14.018
```

上面的子类设置 `__slots__` 为一个空元组。通过阻止创建实例字典保持了较低的内存开销。

Subclassing is not useful for adding new, stored fields. Instead, simply create a new named tuple type from the `_fields` attribute:

```
>>> Point3D = namedtuple('Point3D', Point._fields + ('z',))
```

文档字符串可以自定义，通过直接赋值给 `__doc__` 属性：

```
>>> Book = namedtuple('Book', ['id', 'title', 'authors'])
>>> Book.__doc__ += ': Hardcover book in active collection'
>>> Book.id.__doc__ = '13-digit ISBN'
>>> Book.title.__doc__ = 'Title of first printing'
>>> Book.authors.__doc__ = 'List of authors sorted by last name'
```

在 3.5 版更改：文档字符串属性变成可写。

Default values can be implemented by using `_replace()` to customize a prototype instance:

```
>>> Account = namedtuple('Account', 'owner balance transaction_count')
>>> default_account = Account('<owner name>', 0.0, 0)
>>> johns_account = default_account._replace(owner='John')
>>> janes_account = default_account._replace(owner='Jane')
```

参见：

- [Recipe for named tuple abstract base class with a metaclass mix-in](#) by Jan Kaliszewski. Besides providing an *abstract base class* for named tuples, it also supports an alternate *metaclass*-based constructor that is convenient for use cases where named tuples are being subclassed.
- 对于以字典为底层的可变域名，参考 `types.SimpleNamespace()`。
- See `typing.NamedTuple()` for a way to add type hints for named tuples.

8.3.6 OrderedDict 对象

Ordered dictionaries are just like regular dictionaries but they remember the order that items were inserted. When iterating over an ordered dictionary, the items are returned in the order their keys were first added.

class `collections.OrderedDict` (`[items]`)

Return an instance of a dict subclass, supporting the usual *dict* methods. An *OrderedDict* is a dict that remembers the order that keys were first inserted. If a new entry overwrites an existing entry, the original insertion position is left unchanged. Deleting an entry and reinserting it will move it to the end.

3.1 新版功能.

popitem (`last=True`)

有序字典的 `popitem()` 方法移除并返回一个 (key, value) 键值对。如果 `last` 值为真，则按 LIFO (last-in, first-out) 后进先出的顺序返回键值对，否则就按 FIFO (first-in, first-out) 先进先出的顺序返回键值对。

move_to_end (`key`, `last=True`)

将现有 `key` 移动到有序字典的任一端。如果 `last` 为真值（默认）则将元素移至末尾；如果 `last` 为假值则将元素移至开头。如果 `key` 不存在则会触发 `KeyError`：

```
>>> d = OrderedDict.fromkeys('abcde')
>>> d.move_to_end('b')
>>> ''.join(d.keys())
'acdeb'
>>> d.move_to_end('b', last=False)
>>> ''.join(d.keys())
'bacde'
```

3.2 新版功能.

相对于通常的映射方法，有序字典还另外提供了逆序迭代的支持，通过 `reversed()`。

`OrderedDict` 之间的相等测试是顺序敏感的，实现为 `list(od1.items())==list(od2.items())`。`OrderedDict` 对象和其他的 `Mapping` 的相等测试，是顺序敏感的字典测试。这允许 `OrderedDict` 替换为任何字典可以使用的场所。

在 3.5 版更改：`OrderedDict` 的项 (item)，键 (key) 和值 (value) 视图现在支持逆序迭代，通过 `reversed()`。

在 3.6 版更改：**PEP 468** 赞成将关键词参数的顺序保留，通过传递给 `OrderedDict` 构造器和它的 `update()` 方法。

OrderedDict 例子和用法

Since an ordered dictionary remembers its insertion order, it can be used in conjunction with sorting to make a sorted dictionary:

```
>>> # regular unsorted dictionary
>>> d = {'banana': 3, 'apple': 4, 'pear': 1, 'orange': 2}

>>> # dictionary sorted by key
>>> OrderedDict(sorted(d.items(), key=lambda t: t[0]))
OrderedDict([('apple', 4), ('banana', 3), ('orange', 2), ('pear', 1)])

>>> # dictionary sorted by value
>>> OrderedDict(sorted(d.items(), key=lambda t: t[1]))
OrderedDict([('pear', 1), ('orange', 2), ('banana', 3), ('apple', 4)])

>>> # dictionary sorted by length of the key string
>>> OrderedDict(sorted(d.items(), key=lambda t: len(t[0])))
OrderedDict([('pear', 1), ('apple', 4), ('orange', 2), ('banana', 3)])
```

The new sorted dictionaries maintain their sort order when entries are deleted. But when new keys are added, the keys are appended to the end and the sort is not maintained.

It is also straight-forward to create an ordered dictionary variant that remembers the order the keys were *last* inserted. If a new entry overwrites an existing entry, the original insertion position is changed and moved to the end:

```
class LastUpdatedOrderedDict(OrderedDict):
    'Store items in the order the keys were last added'

    def __setitem__(self, key, value):
        if key in self:
            del self[key]
        OrderedDict.__setitem__(self, key, value)
```

An ordered dictionary can be combined with the `Counter` class so that the counter remembers the order elements are first encountered:

```
class OrderedCounter(Counter, OrderedDict):
    'Counter that remembers the order elements are first encountered'

    def __repr__(self):
        return '%s(%r)' % (self.__class__.__name__, OrderedDict(self))

    def __reduce__(self):
        return self.__class__, (OrderedDict(self),)
```

8.3.7 UserDict 对象

UserDict 类是用作字典对象的外包装。对这个类的需求已部分由直接创建 *dict* 的子类的功能所替代；不过，这个类处理起来更容易，因为底层的字典可以作为属性来访问。

```
class collections.UserDict([initialdata])
```

模拟一个字典类。这个实例的内容保存为一个正常字典，可以通过 *UserDict* 实例的 *data* 属性存取。如果提供了 *initialdata* 值，*data* 就被初始化为它的内容；注意一个 *initialdata* 的引用不会被保留作为其他用途。

UserDict 实例提供了以下属性作为扩展方法和操作的支持：

data

一个真实的字典，用于保存 *UserDict* 类的内容。

8.3.8 UserList 对象

这个类封装了列表对象。它是一个有用的基础类，对于你想自定义的类似列表的类，可以继承和覆盖现有的方法，也可以添加新的方法。这样我们可以对列表添加新的行为。

对这个类的需求已部分由直接创建 *list* 的子类的功能所替代；不过，这个类处理起来更容易，因为底层的列表可以作为属性来访问。

```
class collections.UserList([list])
```

模拟一个列表。这个实例的内容被保存为一个正常列表，通过 *UserList* 的 *data* 属性存取。实例内容被初始化为一个 *list* 的 *copy*，默认为 [] 空列表。*list* 可以是迭代对象，比如一个 Python 列表，或者一个 *UserList* 对象。

UserList 提供了以下属性作为可变序列的方法和操作的扩展：

data

一个 *list* 对象用于存储 *UserList* 的内容。

子类化的要求： *UserList* 的子类需要提供一个构造器，可以无参数调用，或者一个参数调用。返回一个新序列的列表操作需要创建一个实现类的实例。它假定了构造器可以以一个参数进行调用，这个参数是一个序列对象，作为数据源。

如果一个分离的类不希望依照这个需求，所有的特殊方法就必须重写；请参照源代码进行修改。

8.3.9 UserString 对象

UserString 类是用作字符串对象的外包装。对这个类的需求已部分由直接创建 *str* 的子类的功能所替代；不过，这个类处理起来更容易，因为底层的字符串可以作为属性来访问。

```
class collections.UserString([sequence])
```

Class that simulates a string or a Unicode string object. The instance's content is kept in a regular string object,

which is accessible via the `data` attribute of `UserString` instances. The instance's contents are initially set to a copy of `sequence`. The `sequence` can be an instance of `bytes`, `str`, `UserString` (or a subclass) or an arbitrary sequence which can be converted into a string using the built-in `str()` function.

在 3.5 版更改: 新方法 `__getnewargs__`, `__rmod__`, `casefold`, `format_map`, `isprintable`, 和 `maketrans`。

8.4 `collections.abc` —容器的抽象基类

3.3 新版功能: 该模块曾是 `collections` 模块的组成部分。

源代码: [Lib/_collections_abc.py](#)

该模块定义了一些抽象基类，它们可用于判断一个具体类是否具有某一特定的接口；例如，这个类是否可哈希，或其是否为映射类。

8.4.1 Collections Abstract Base Classes

The `collections` module offers the following ABCs:

ABC	Inherits from	Abstract Methods	Mixin Methods
<i>Container</i>		<code>__contains__</code>	
<i>Hashable</i>		<code>__hash__</code>	
<i>Iterable</i>		<code>__iter__</code>	
<i>Iterator</i>	<i>Iterable</i>	<code>__next__</code>	<code>__iter__</code>
<i>Reversible</i>	<i>Iterable</i>	<code>__reversed__</code>	
<i>Generator</i>	<i>Iterator</i>	<code>send</code> , <code>throw</code>	<code>close</code> , <code>__iter__</code> , <code>__next__</code>
<i>Sized</i>		<code>__len__</code>	
<i>Callable</i>		<code>__call__</code>	
<i>Collection</i>	<i>Sized</i> , <i>Iterable</i> , <i>Container</i>	<code>__contains__</code> , <code>__iter__</code> , <code>__len__</code>	
<i>Sequence</i>	<i>Reversible</i> , <i>Collection</i>	<code>__getitem__</code> , <code>__len__</code>	<code>__contains__</code> , <code>__iter__</code> , <code>__reversed__</code> , <code>index</code> , and <code>count</code>
<i>MutableSequence</i>	<i>Sequence</i>	<code>__getitem__</code> , <code>__setitem__</code> , <code>__delitem__</code> , <code>__len__</code> , <code>insert</code>	Inherited <i>Sequence</i> methods and <code>append</code> , <code>reverse</code> , <code>extend</code> , <code>pop</code> , <code>remove</code> , and <code>__iadd__</code>
<i>ByteString</i>	<i>Sequence</i>	<code>__getitem__</code> , <code>__len__</code>	Inherited <i>Sequence</i> methods
<i>Set</i>	<i>Collection</i>	<code>__contains__</code> , <code>__iter__</code> , <code>__len__</code>	<code>__le__</code> , <code>__lt__</code> , <code>__eq__</code> , <code>__ne__</code> , <code>__gt__</code> , <code>__ge__</code> , <code>__and__</code> , <code>__or__</code> , <code>__sub__</code> , <code>__xor__</code> , and <code>isdisjoint</code>
<i>MutableSet</i>	<i>Set</i>	<code>__contains__</code> , <code>__iter__</code> , <code>__len__</code> , <code>add</code> , <code>discard</code>	Inherited <i>Set</i> methods and <code>clear</code> , <code>pop</code> , <code>remove</code> , <code>__ior__</code> , <code>__iand__</code> , <code>__ixor__</code> , and <code>__isub__</code>
<i>Mapping</i>	<i>Collection</i>	<code>__getitem__</code> , <code>__iter__</code> , <code>__len__</code>	<code>__contains__</code> , <code>keys</code> , <code>items</code> , <code>values</code> , <code>get</code> , <code>__eq__</code> , and <code>__ne__</code>
<i>MutableMapping</i>	<i>Mapping</i>	<code>__getitem__</code> , <code>__setitem__</code> , <code>__delitem__</code> , <code>__iter__</code> , <code>__len__</code>	Inherited <i>Mapping</i> methods and <code>pop</code> , <code>popitem</code> , <code>clear</code> , <code>update</code> , and <code>setdefault</code>
<i>MappingView</i>	<i>Sized</i>		<code>__len__</code>
<i>ItemsView</i>	<i>MappingView</i> , <i>Set</i>		<code>__contains__</code> , <code>__iter__</code>
<i>KeysView</i>	<i>MappingView</i> , <i>Set</i>		<code>__contains__</code> , <code>__iter__</code>
<i>ValuesView</i>	<i>MappingView</i>		<code>__contains__</code> , <code>__iter__</code>
<i>Awaitable</i>		<code>__await__</code>	
<i>Coroutine</i>	<i>Awaitable</i>	<code>send</code> , <code>throw</code>	<code>close</code>
<i>AsyncIterable</i>		<code>__aiter__</code>	
<i>AsyncIterator</i>	<i>AsyncIterable</i>	<code>__anext__</code>	<code>__aiter__</code>
<i>AsyncGenerator</i>	<i>AsyncIterator</i>	<code>asend</code> , <code>athrow</code>	<code>aclose</code> , <code>__aiter__</code> , <code>__anext__</code>

```
class collections.abc.Container
```

```
class collections.abc.Hashable
```

```
class collections.abc.Sized
```

```
class collections.abc.Callable
```

ABCs for classes that provide respectively the methods `__contains__()`, `__hash__()`, `__len__()`, and `__call__()`.

```
class collections.abc.Iterable
```

ABC for classes that provide the `__iter__()` method.

Checking `isinstance(obj, Iterable)` detects classes that are registered as *Iterable* or that have an `__iter__()` method, but it does not detect classes that iterate with the `__getitem__()` method. The only reliable way to determine whether an object is *iterable* is to call `iter(obj)`.

class `collections.abc.Collection`

ABC for sized iterable container classes.

3.6 新版功能.

class `collections.abc.Iterator`

ABC for classes that provide the `__iter__()` and `__next__()` methods. See also the definition of *iterator*.

class `collections.abc.Reversible`

ABC for iterable classes that also provide the `__reversed__()` method.

3.6 新版功能.

class `collections.abc.Generator`

ABC for generator classes that implement the protocol defined in [PEP 342](#) that extends iterators with the `send()`, `throw()` and `close()` methods. See also the definition of *generator*.

3.5 新版功能.

class `collections.abc.Sequence`

class `collections.abc.MutableSequence`

class `collections.abc.ByteString`

ABCs for read-only and mutable *sequences*.

Implementation note: Some of the mixin methods, such as `__iter__()`, `__reversed__()` and `index()`, make repeated calls to the underlying `__getitem__()` method. Consequently, if `__getitem__()` is implemented with constant access speed, the mixin methods will have linear performance; however, if the underlying method is linear (as it would be with a linked list), the mixins will have quadratic performance and will likely need to be overridden.

在 3.5 版更改: The `index()` method added support for *stop* and *start* arguments.

class `collections.abc.Set`

class `collections.abc.MutableSet`

ABCs for read-only and mutable sets.

class `collections.abc.Mapping`

class `collections.abc.MutableMapping`

ABCs for read-only and mutable *mappings*.

class `collections.abc.MappingView`

class `collections.abc.ItemsView`

class `collections.abc.KeysView`

class `collections.abc.ValuesView`

ABCs for mapping, items, keys, and values *views*.

class `collections.abc.Awaitable`

ABC for *awaitable* objects, which can be used in `await` expressions. Custom implementations must provide the `__await__()` method.

Coroutine objects and instances of the *Coroutine* ABC are all instances of this ABC.

注解: In CPython, generator-based coroutines (generators decorated with `types.coroutine()` or `asyncio.coroutine()`) are *awaitables*, even though they do not have an `__await__()` method. Using

`isinstance(gencoro, Awaitable)` for them will return `False`. Use `inspect.isawaitable()` to detect them.

3.5 新版功能.

class `collections.abc.Coroutine`

ABC for coroutine compatible classes. These implement the following methods, defined in `coroutine-objects`: `send()`, `throw()`, and `close()`. Custom implementations must also implement `__await__()`. All *Coroutine* instances are also instances of *Awaitable*. See also the definition of *coroutine*.

注解: In CPython, generator-based coroutines (generators decorated with `types.coroutine()` or `asyncio.coroutine()`) are *awaitables*, even though they do not have an `__await__()` method. Using `isinstance(gencoro, Coroutine)` for them will return `False`. Use `inspect.isawaitable()` to detect them.

3.5 新版功能.

class `collections.abc.AsyncIterable`

ABC for classes that provide `__aiter__` method. See also the definition of *asynchronous iterable*.

3.5 新版功能.

class `collections.abc.AsyncIterator`

ABC for classes that provide `__aiter__` and `__anext__` methods. See also the definition of *asynchronous iterator*.

3.5 新版功能.

class `collections.abc.AsyncGenerator`

ABC for asynchronous generator classes that implement the protocol defined in **PEP 525** and **PEP 492**.

3.6 新版功能.

These ABCs allow us to ask classes or instances if they provide particular functionality, for example:

```
size = None
if isinstance(myvar, collections.abc.Sized):
    size = len(myvar)
```

Several of the ABCs are also useful as mixins that make it easier to develop classes supporting container APIs. For example, to write a class supporting the full *Set* API, it is only necessary to supply the three underlying abstract methods: `__contains__()`, `__iter__()`, and `__len__()`. The ABC supplies the remaining methods such as `__and__()` and `isdisjoint()`:

```
class ListBasedSet(collections.abc.Set):
    ''' Alternate set implementation favoring space over speed
        and not requiring the set elements to be hashable. '''
    def __init__(self, iterable):
        self.elements = lst = []
        for value in iterable:
            if value not in lst:
                lst.append(value)

    def __iter__(self):
        return iter(self.elements)

    def __contains__(self, value):
```

(下页继续)

(续上页)

```

        return value in self.elements

    def __len__(self):
        return len(self.elements)

s1 = ListBasedSet('abcdef')
s2 = ListBasedSet('defghi')
overlap = s1 & s2           # The __and__() method is supported automatically

```

Notes on using *Set* and *MutableSet* as a mixin:

- (1) Since some set operations create new sets, the default mixin methods need a way to create new instances from an iterable. The class constructor is assumed to have a signature in the form `ClassName(iterable)`. That assumption is factored-out to an internal classmethod called `_from_iterable()` which calls `cls(iterable)` to produce a new set. If the *Set* mixin is being used in a class with a different constructor signature, you will need to override `_from_iterable()` with a classmethod that can construct new instances from an iterable argument.
- (2) To override the comparisons (presumably for speed, as the semantics are fixed), redefine `__le__()` and `__ge__()`, then the other operations will automatically follow suit.
- (3) The *Set* mixin provides a `_hash()` method to compute a hash value for the set; however, `__hash__()` is not defined because not all sets are hashable or immutable. To add set hashability using mixins, inherit from both *Set()* and *Hashable()*, then define `__hash__ = Set._hash`.

参见:

- [OrderedSet recipe](#) for an example built on *MutableSet*.
- For more about ABCs, see the [abc](#) module and [PEP 3119](#).

8.5 heapq — 堆队列算法

源码: [Lib/heapq.py](#)

这个模块提供了堆队列算法的实现，也称为优先队列算法。

堆是一个二叉树，它的每个父节点的值都只会小于或大于所有孩子节点（的值）。它使用了数组来实现：从零开始计数，对于所有的 k ，都有 “ $\text{heap}[k] \leq \text{heap}[2*k+1]$ ” 和 $\text{heap}[k] \leq \text{heap}[2*k+2]$ 。为了便于比较，不存在的元素被认为是无限大。堆最有趣的特性在于最小的元素总是在根结点： $\text{heap}[0]$ 。

这个 API 与教材的堆算法实现不太一样，在于两方面：(a) 我们使用了基于零开始进行索引。这使得节点和其孩子节点索引之间的关系不太直观，但是更加适合

基于这两方面，把堆看作原生的 Python list 也没什么奇怪的： $\text{heap}[0]$ 表示最小的元素，同时 `heap.sort()` 维护了堆的不变性！

要创建一个堆，可以使用 list 来初始化为 `[]`，或者你可以通过一个函数 `heapify()`，来把一个 list 转换成堆。

定义了以下函数：

`heapq.heappush(heap, item)`

将 `item` 的值加入 `heap` 中，保持堆的不变性。

`heapq.heappop(heap)`

弹出并返回 `heap` 的最小的元素，保持堆的不变性。如果堆为空，抛出 `IndexError`。使用 `heap[0]`，可以只访问最小的元素而不弹出它。

`heapq.heappushpop(heap, item)`

将 *item* 放入堆中，然后弹出并返回 *heap* 的最小元素。该组合操作比先调用 `heappush()` 再调用 `heappop()` 运行起来更有效率。

`heapq.heapify(x)`

将 list *x* 转换成堆，原地，线性时间内。

`heapq.heapreplace(heap, item)`

Pop and return the smallest item from the *heap*, and also push the new *item*. The heap size doesn't change. If the heap is empty, `IndexError` is raised.

This one step operation is more efficient than a `heappop()` followed by `heappush()` and can be more appropriate when using a fixed-size heap. The pop/push combination always returns an element from the heap and replaces it with *item*.

The value returned may be larger than the *item* added. If that isn't desired, consider using `heappushpop()` instead. Its push/pop combination returns the smaller of the two values, leaving the larger value on the heap.

The module also offers three general purpose functions based on heaps.

`heapq.merge(*iterables, key=None, reverse=False)`

Merge multiple sorted inputs into a single sorted output (for example, merge timestamped entries from multiple log files). Returns an *iterator* over the sorted values.

Similar to `sorted(itertools.chain(*iterables))` but returns an iterable, does not pull the data into memory all at once, and assumes that each of the input streams is already sorted (smallest to largest).

具有两个可选参数，它们都必须指定为关键字参数。

key specifies a *key function* of one argument that is used to extract a comparison key from each input element. The default value is `None` (compare the elements directly).

reverse is a boolean value. If set to `True`, then the input elements are merged as if each comparison were reversed.

在 3.5 版更改: Added the optional *key* and *reverse* parameters.

`heapq.nlargest(n, iterable, key=None)`

Return a list with the *n* largest elements from the dataset defined by *iterable*. *key*, if provided, specifies a function of one argument that is used to extract a comparison key from each element in *iterable* (for example, `key=str.lower`). Equivalent to: `sorted(iterable, key=key, reverse=True)[:n]`.

`heapq.nsmallest(n, iterable, key=None)`

Return a list with the *n* smallest elements from the dataset defined by *iterable*. *key*, if provided, specifies a function of one argument that is used to extract a comparison key from each element in *iterable* (for example, `key=str.lower`). Equivalent to: `sorted(iterable, key=key)[:n]`.

The latter two functions perform best for smaller values of *n*. For larger values, it is more efficient to use the `sorted()` function. Also, when *n*==1, it is more efficient to use the built-in `min()` and `max()` functions. If repeated usage of these functions is required, consider turning the iterable into an actual heap.

8.5.1 基本示例

A *heapsort* can be implemented by pushing all values onto a heap and then popping off the smallest values one at a time:

```
>>> def heapsort(iterable):
...     h = []
...     for value in iterable:
...         heappush(h, value)
...     return [heappop(h) for i in range(len(h))]
... 
```

(下页继续)

(续上页)

```
>>> heapsort([1, 3, 5, 7, 9, 2, 4, 6, 8, 0])
[0, 1, 2, 3, 4, 5, 6, 7, 8, 9]
```

This is similar to `sorted(iterable)`, but unlike `sorted()`, this implementation is not stable.

Heap elements can be tuples. This is useful for assigning comparison values (such as task priorities) alongside the main record being tracked:

```
>>> h = []
>>> heappush(h, (5, 'write code'))
>>> heappush(h, (7, 'release product'))
>>> heappush(h, (1, 'write spec'))
>>> heappush(h, (3, 'create tests'))
>>> heappop(h)
(1, 'write spec')
```

8.5.2 Priority Queue Implementation Notes

A **priority queue** is common use for a heap, and it presents several implementation challenges:

- Sort stability: how do you get two tasks with equal priorities to be returned in the order they were originally added?
- Tuple comparison breaks for (priority, task) pairs if the priorities are equal and the tasks do not have a default comparison order.
- If the priority of a task changes, how do you move it to a new position in the heap?
- Or if a pending task needs to be deleted, how do you find it and remove it from the queue?

A solution to the first two challenges is to store entries as 3-element list including the priority, an entry count, and the task. The entry count serves as a tie-breaker so that two tasks with the same priority are returned in the order they were added. And since no two entry counts are the same, the tuple comparison will never attempt to directly compare two tasks.

The remaining challenges revolve around finding a pending task and making changes to its priority or removing it entirely. Finding a task can be done with a dictionary pointing to an entry in the queue.

Removing the entry or changing its priority is more difficult because it would break the heap structure invariants. So, a possible solution is to mark the entry as removed and add a new entry with the revised priority:

```
pq = []                                # list of entries arranged in a heap
entry_finder = {}                       # mapping of tasks to entries
REMOVED = '<removed-task>'              # placeholder for a removed task
counter = itertools.count()             # unique sequence count

def add_task(task, priority=0):
    'Add a new task or update the priority of an existing task'
    if task in entry_finder:
        remove_task(task)
    count = next(counter)
    entry = [priority, count, task]
    entry_finder[task] = entry
    heappush(pq, entry)

def remove_task(task):
    'Mark an existing task as REMOVED. Raise KeyError if not found.'
    entry = entry_finder.pop(task)
    entry[-1] = REMOVED
```

(下页继续)

(续上页)

```
def pop_task():
    'Remove and return the lowest priority task. Raise KeyError if empty.'
    while pq:
        priority, count, task = heappop(pq)
        if task is not REMOVED:
            del entry_finder[task]
            return task
    raise KeyError('pop from an empty priority queue')
```

8.5.3 Theory

Heaps are arrays for which $a[k] \leq a[2k+1]$ and $a[k] \leq a[2k+2]$ for all k , counting elements from 0. For the sake of comparison, non-existing elements are considered to be infinite. The interesting property of a heap is that $a[0]$ is always its smallest element.

The strange invariant above is meant to be an efficient memory representation for a tournament. The numbers below are k , not $a[k]$:



In the tree above, each cell k is topping $2k+1$ and $2k+2$. In a usual binary tournament we see in sports, each cell is the winner over the two cells it tops, and we can trace the winner down the tree to see all opponents s/he had. However, in many computer applications of such tournaments, we do not need to trace the history of a winner. To be more memory efficient, when a winner is promoted, we try to replace it by something else at a lower level, and the rule becomes that a cell and the two cells it tops contain three different items, but the top cell “wins” over the two topped cells.

If this heap invariant is protected at all time, index 0 is clearly the overall winner. The simplest algorithmic way to remove it and find the “next” winner is to move some loser (let’s say cell 30 in the diagram above) into the 0 position, and then percolate this new 0 down the tree, exchanging values, until the invariant is re-established. This is clearly logarithmic on the total number of items in the tree. By iterating over all items, you get an $O(n \log n)$ sort.

A nice feature of this sort is that you can efficiently insert new items while the sort is going on, provided that the inserted items are not “better” than the last 0th element you extracted. This is especially useful in simulation contexts, where the tree holds all incoming events, and the “win” condition means the smallest scheduled time. When an event schedules other events for execution, they are scheduled into the future, so they can easily go into the heap. So, a heap is a good structure for implementing schedulers (this is what I used for my MIDI sequencer :-).

Various structures for implementing schedulers have been extensively studied, and heaps are good for this, as they are reasonably speedy, the speed is almost constant, and the worst case is not much different than the average case. However, there are other representations which are more efficient overall, yet the worst cases might be terrible.

Heaps are also very useful in big disk sorts. You most probably all know that a big sort implies producing “runs” (which are pre-sorted sequences, whose size is usually related to the amount of CPU memory), followed by a merging passes for these runs, which merging is often very cleverly organised¹. It is very important that the initial sort produces the longest

¹ The disk balancing algorithms which are current, nowadays, are more annoying than clever, and this is a consequence of the seeking capabilities of the disks. On devices which cannot seek, like big tape drives, the story was quite different, and one had to be very clever to ensure (far in advance)

runs possible. Tournaments are a good way to achieve that. If, using all the memory available to hold a tournament, you replace and percolate items that happen to fit the current run, you'll produce runs which are twice the size of the memory for random input, and much better for input fuzzily ordered.

Moreover, if you output the 0th item on disk and get an input which may not fit in the current tournament (because the value “wins” over the last output value), it cannot fit in the heap, so the size of the heap decreases. The freed memory could be cleverly reused immediately for progressively building a second heap, which grows at exactly the same rate the first heap is melting. When the first heap completely vanishes, you switch heaps and start a new run. Clever and quite effective!

In a word, heaps are useful memory structures to know. I use them in a few applications, and I think it is good to keep a ‘heap’ module around. :-)

8.6 bisect —Array bisection algorithm

Source code: [Lib/bisect.py](#)

This module provides support for maintaining a list in sorted order without having to sort the list after each insertion. For long lists of items with expensive comparison operations, this can be an improvement over the more common approach. The module is called *bisect* because it uses a basic bisection algorithm to do its work. The source code may be most useful as a working example of the algorithm (the boundary conditions are already right!).

定义了以下函数：

`bisect.bisect_left(a, x, lo=0, hi=len(a))`

Locate the insertion point for *x* in *a* to maintain sorted order. The parameters *lo* and *hi* may be used to specify a subset of the list which should be considered; by default the entire list is used. If *x* is already present in *a*, the insertion point will be before (to the left of) any existing entries. The return value is suitable for use as the first parameter to `list.insert()` assuming that *a* is already sorted.

The returned insertion point *i* partitions the array *a* into two halves so that `all(val < x for val in a[lo:i])` for the left side and `all(val >= x for val in a[i:hi])` for the right side.

`bisect.bisect_right(a, x, lo=0, hi=len(a))`

`bisect.bisect(a, x, lo=0, hi=len(a))`

Similar to `bisect_left()`, but returns an insertion point which comes after (to the right of) any existing entries of *x* in *a*.

The returned insertion point *i* partitions the array *a* into two halves so that `all(val <= x for val in a[lo:i])` for the left side and `all(val > x for val in a[i:hi])` for the right side.

`bisect.insort_left(a, x, lo=0, hi=len(a))`

Insert *x* in *a* in sorted order. This is equivalent to `a.insert(bisect.bisect_left(a, x, lo, hi), x)` assuming that *a* is already sorted. Keep in mind that the $O(\log n)$ search is dominated by the slow $O(n)$ insertion step.

`bisect.insort_right(a, x, lo=0, hi=len(a))`

`bisect.insort(a, x, lo=0, hi=len(a))`

Similar to `insort_left()`, but inserting *x* in *a* after any existing entries of *x*.

参见：

[SortedCollection](#) recipe that uses `bisect` to build a full-featured collection class with straight-forward search methods and support for a key-function. The keys are precomputed to save unnecessary calls to the key function during searches.

that each tape movement will be the most effective possible (that is, will best participate at “progressing” the merge). Some tapes were even able to read backwards, and this was also used to avoid the rewinding time. Believe me, real good tape sorts were quite spectacular to watch! From all times, sorting has always been a Great Art! :-)

8.6.1 Searching Sorted Lists

The above `bisect()` functions are useful for finding insertion points but can be tricky or awkward to use for common searching tasks. The following five functions show how to transform them into the standard lookups for sorted lists:

```
def index(a, x):
    'Locate the leftmost value exactly equal to x'
    i = bisect_left(a, x)
    if i != len(a) and a[i] == x:
        return i
    raise ValueError

def find_lt(a, x):
    'Find rightmost value less than x'
    i = bisect_left(a, x)
    if i:
        return a[i-1]
    raise ValueError

def find_le(a, x):
    'Find rightmost value less than or equal to x'
    i = bisect_right(a, x)
    if i:
        return a[i-1]
    raise ValueError

def find_gt(a, x):
    'Find leftmost value greater than x'
    i = bisect_right(a, x)
    if i != len(a):
        return a[i]
    raise ValueError

def find_ge(a, x):
    'Find leftmost item greater than or equal to x'
    i = bisect_left(a, x)
    if i != len(a):
        return a[i]
    raise ValueError
```

8.6.2 Other Examples

The `bisect()` function can be useful for numeric table lookups. This example uses `bisect()` to look up a letter grade for an exam score (say) based on a set of ordered numeric breakpoints: 90 and up is an 'A', 80 to 89 is a 'B', and so on:

```
>>> def grade(score, breakpoints=[60, 70, 80, 90], grades='FDCBA'):
...     i = bisect(breakpoints, score)
...     return grades[i]
...
>>> [grade(score) for score in [33, 99, 77, 70, 89, 90, 100]]
['F', 'A', 'C', 'C', 'B', 'A', 'A']
```

Unlike the `sorted()` function, it does not make sense for the `bisect()` functions to have `key` or `reversed` arguments because that would lead to an inefficient design (successive calls to `bisect` functions would not “remember” all of the previous key lookups).

Instead, it is better to search a list of precomputed keys to find the index of the record in question:

```
>>> data = [('red', 5), ('blue', 1), ('yellow', 8), ('black', 0)]
>>> data.sort(key=lambda r: r[1])
>>> keys = [r[1] for r in data]           # precomputed list of keys
>>> data[bisect_left(keys, 0)]
('black', 0)
>>> data[bisect_left(keys, 1)]
('blue', 1)
>>> data[bisect_left(keys, 5)]
('red', 5)
>>> data[bisect_left(keys, 8)]
('yellow', 8)
```

8.7 array — Efficient arrays of numeric values

This module defines an object type which can compactly represent an array of basic values: characters, integers, floating point numbers. Arrays are sequence types and behave very much like lists, except that the type of objects stored in them is constrained. The type is specified at object creation time by using a *type code*, which is a single character. The following type codes are defined:

Type code	C 类型	Python Type	Minimum size in bytes	注释
'b'	signed char	int	1	
'B'	unsigned char	int	1	
'u'	Py_UNICODE	Unicode character	2	(1)
'h'	signed short	int	2	
'H'	unsigned short	int	2	
'i'	signed int	int	2	
'I'	无符号整型	int	2	
'l'	signed long	int	4	
'L'	无符号长整型	int	4	
'q'	signed long long	int	8	(2)
'Q'	无符号 long long	int	8	(2)
'f'	float	float	4	
'd'	double	float	8	

注释:

- (1) The 'u' type code corresponds to Python's obsolete unicode character (Py_UNICODE which is wchar_t). Depending on the platform, it can be 16 bits or 32 bits.

'u' will be removed together with the rest of the Py_UNICODE API.

Deprecated since version 3.3, will be removed in version 4.0.

- (2) The 'q' and 'Q' type codes are available only if the platform C compiler used to build Python supports C long long, or, on Windows, __int64.

3.3 新版功能.

The actual representation of values is determined by the machine architecture (strictly speaking, by the C implementation). The actual size can be accessed through the `itemsize` attribute.

The module defines the following type:

class `array.array`(*typecode*[, *initializer*])

A new array whose items are restricted by *typecode*, and initialized from the optional *initializer* value, which must be a list, a *bytes-like object*, or iterable over elements of the appropriate type.

If given a list or string, the initializer is passed to the new array's `fromlist()`, `frombytes()`, or `fromunicode()` method (see below) to add initial items to the array. Otherwise, the iterable initializer is passed to the `extend()` method.

`array.typecodes`

A string with all available type codes.

Array objects support the ordinary sequence operations of indexing, slicing, concatenation, and multiplication. When using slice assignment, the assigned value must be an array object with the same type code; in all other cases, `TypeError` is raised. Array objects also implement the buffer interface, and may be used wherever *bytes-like objects* are supported.

The following data items and methods are also supported:

`array.typecode`

The typecode character used to create the array.

`array.itemsize`

The length in bytes of one array item in the internal representation.

`array.append(x)`

Append a new item with value *x* to the end of the array.

`array.buffer_info()`

Return a tuple (*address*, *length*) giving the current memory address and the length in elements of the buffer used to hold array's contents. The size of the memory buffer in bytes can be computed as `array.buffer_info()[1] * array.itemsize`. This is occasionally useful when working with low-level (and inherently unsafe) I/O interfaces that require memory addresses, such as certain `ioctl()` operations. The returned numbers are valid as long as the array exists and no length-changing operations are applied to it.

注解: When using array objects from code written in C or C++ (the only way to effectively make use of this information), it makes more sense to use the buffer interface supported by array objects. This method is maintained for backward compatibility and should be avoided in new code. The buffer interface is documented in `bufferobjects`.

`array.byteswap()`

“Byteswap” all items of the array. This is only supported for values which are 1, 2, 4, or 8 bytes in size; for other types of values, `RuntimeError` is raised. It is useful when reading data from a file written on a machine with a different byte order.

`array.count(x)`

Return the number of occurrences of *x* in the array.

`array.extend(iterable)`

Append items from *iterable* to the end of the array. If *iterable* is another array, it must have *exactly* the same type code; if not, `TypeError` will be raised. If *iterable* is not an array, it must be iterable and its elements must be the right type to be appended to the array.

`array.frombytes(s)`

Appends items from the string, interpreting the string as an array of machine values (as if it had been read from a file using the `fromfile()` method).

3.2 新版功能: `fromstring()` is renamed to `frombytes()` for clarity.

`array.fromfile(f, n)`

Read *n* items (as machine values) from the *file object* *f* and append them to the end of the array. If less than *n*

items are available, `EOFError` is raised, but the items that were available are still inserted into the array. `f` must be a real built-in file object; something else with a `read()` method won't do.

`array.fromlist(list)`

Append items from the list. This is equivalent to `for x in list: a.append(x)` except that if there is a type error, the array is unchanged.

`array.fromstring()`

Deprecated alias for `frombytes()`.

`array.fromunicode(s)`

Extends this array with data from the given unicode string. The array must be a type 'u' array; otherwise a `ValueError` is raised. Use `array.frombytes(unicodestring.encode(enc))` to append Unicode data to an array of some other type.

`array.index(x)`

Return the smallest *i* such that *i* is the index of the first occurrence of *x* in the array.

`array.insert(i, x)`

Insert a new item with value *x* in the array before position *i*. Negative values are treated as being relative to the end of the array.

`array.pop([i])`

Removes the item with the index *i* from the array and returns it. The optional argument defaults to `-1`, so that by default the last item is removed and returned.

`array.remove(x)`

Remove the first occurrence of *x* from the array.

`array.reverse()`

Reverse the order of the items in the array.

`array.tobytes()`

Convert the array to an array of machine values and return the bytes representation (the same sequence of bytes that would be written to a file by the `tofile()` method.)

3.2 新版功能: `tostring()` is renamed to `tobytes()` for clarity.

`array.tofile(f)`

Write all items (as machine values) to the *file object* *f*.

`array.tolist()`

Convert the array to an ordinary list with the same items.

`array.tostring()`

Deprecated alias for `tobytes()`.

`array.tounicode()`

Convert the array to a unicode string. The array must be a type 'u' array; otherwise a `ValueError` is raised. Use `array.tobytes().decode(enc)` to obtain a unicode string from an array of some other type.

When an array object is printed or converted to a string, it is represented as `array(typecode, initializer)`. The *initializer* is omitted if the array is empty, otherwise it is a string if the *typecode* is 'u', otherwise it is a list of numbers. The string is guaranteed to be able to be converted back to an array with the same type and value using `eval()`, so long as the `array` class has been imported using `from array import array`. Examples:

```
array('l')
array('u', 'hello \u2641')
array('l', [1, 2, 3, 4, 5])
array('d', [1.0, 2.0, 3.14])
```

参见:

Module `struct` Packing and unpacking of heterogeneous binary data.

Module `xdrlib` Packing and unpacking of External Data Representation (XDR) data as used in some remote procedure call systems.

The Numerical Python Documentation The Numeric Python extension (NumPy) defines another array type; see <http://www.numpy.org/> for further information about Numerical Python.

8.8 weakref — 弱引用

源码: [Lib/weakref.py](#)

`weakref` 模块允许 Python 程序员创建对象的 *weak references*。

在下文中，术语 *referent* 表示由弱引用引用的对象。

对对象的弱引用不能保证对象存活：当对象的引用只剩弱引用时，*garbage collection* 可以销毁引用并将其内存重用于其他内容。但是，在实际销毁对象之前，即使没有强引用，弱引用也一直能返回该对象。

弱引用的主要用途是实现保存大对象的高速缓存或映射，但又并希望大对象仅仅因为它出现在高速缓存或映射中而保持存活。

例如，如果您有许多大型二进制图像对象，则可能希望将名称与每个对象关联起来。如果您使用 Python 字典将名称映射到图像，或将图像映射到名称，则图像对象将保持活动状态，因为它们在字典中显示为值或键。`weakref` 模块提供的 `WeakKeyDictionary` 和 `WeakValueDictionary` 类可以替代 Python 字典，使用弱引用来构造映射，这些映射不会仅仅因为它们出现在映射对象中而使对象保持存活。例如，如果一个图像对象是 `WeakValueDictionary` 中的值，那么当对该图像对象的剩余引用是弱映射对象所持有的弱引用时，垃圾回收可以回收该对象并将其在弱映射对象中相应的条目删除。

`WeakKeyDictionary` 和 `WeakValueDictionary` 在它们的实现中使用弱引用，在弱引用上设置回调函数，当键或值被垃圾回收回收时通知弱字典。`WeakSet` 实现了 `set` 接口，但像 `WeakKeyDictionary` 一样，只持有其元素的弱引用。

`finalize` 提供了注册一个对象被垃圾收集时要调用的清理函数的方式。这比在原始弱引用上设置回调函数更简单，因为模块会自动确保对象被回收前终结器一直保持存活。

这些弱容器类型之一或者 `finalize` 就是大多数程序所需要的 - 通常不需要直接创建自己的弱引用。`weakref` 模块暴露了低级机制，以便于高级用途。

Not all objects can be weakly referenced; those objects which can include class instances, functions written in Python (but not in C), instance methods, sets, frozensets, some *file objects*, *generators*, type objects, sockets, arrays, dequeues, regular expression pattern objects, and code objects.

在 3.2 版更改: 添加了对 `thread.lock`，`threading.Lock` 和代码对象的支持。

几个内建类型如 `list` 和 `dict` 不直接支持弱引用，但可以通过子类化添加支持:

```
class Dict(dict):
    pass

obj = Dict(red=1, green=2, blue=3)  # this object is weak referenceable
```

Other built-in types such as `tuple` and `int` do not support weak references even when subclassed (This is an implementation detail and may be different across various Python implementations.).

Extension types can easily be made to support weak references; see `weakref-support`.

class `weakref.ref(object[, callback])`

返回对 对象的弱引用。如果原始对象仍然存活，则可以通过调用引用对象来检索原始对象；如果引用的原始对象不再存在，则调用引用对象将得到`None`。如果提供了回调而且值不是`None`，并且返回的弱引用对象仍然存活，则在对象即将终结时将调用回调；弱引用对象将作为回调的唯一参数传递；指示物将不再可用。

It is allowable for many weak references to be constructed for the same object. Callbacks registered for each weak reference will be called from the most recently registered callback to the oldest registered callback.

Exceptions raised by the callback will be noted on the standard error output, but cannot be propagated; they are handled in exactly the same way as exceptions raised from an object's `__del__()` method.

Weak references are *hashable* if the *object* is hashable. They will maintain their hash value even after the *object* was deleted. If `hash()` is called the first time only after the *object* was deleted, the call will raise `TypeError`.

Weak references support tests for equality, but not ordering. If the referents are still alive, two references have the same equality relationship as their referents (regardless of the *callback*). If either referent has been deleted, the references are equal only if the reference objects are the same object.

This is a subclassable type rather than a factory function.

__callback__

This read-only attribute returns the callback currently associated to the weakref. If there is no callback or if the referent of the weakref is no longer alive then this attribute will have value `None`.

在 3.4 版更改: Added the `__callback__` attribute.

`weakref.proxy(object[, callback])`

Return a proxy to *object* which uses a weak reference. This supports use of the proxy in most contexts instead of requiring the explicit dereferencing used with weak reference objects. The returned object will have a type of either `ProxyType` or `CallableProxyType`, depending on whether *object* is callable. Proxy objects are not *hashable* regardless of the referent; this avoids a number of problems related to their fundamentally mutable nature, and prevent their use as dictionary keys. *callback* is the same as the parameter of the same name to the `ref()` function.

`weakref.getweakrefcount(object)`

Return the number of weak references and proxies which refer to *object*.

`weakref.getweakrefs(object)`

Return a list of all weak reference and proxy objects which refer to *object*.

class `weakref.WeakKeyDictionary([dict])`

Mapping class that references keys weakly. Entries in the dictionary will be discarded when there is no longer a strong reference to the key. This can be used to associate additional data with an object owned by other parts of an application without adding attributes to those objects. This can be especially useful with objects that override attribute accesses.

注解: Caution: Because a `WeakKeyDictionary` is built on top of a Python dictionary, it must not change size when iterating over it. This can be difficult to ensure for a `WeakKeyDictionary` because actions performed by the program during iteration may cause items in the dictionary to vanish “by magic” (as a side effect of garbage collection).

`WeakKeyDictionary` objects have an additional method that exposes the internal references directly. The references are not guaranteed to be “live” at the time they are used, so the result of calling the references needs to be checked before being used. This can be used to avoid creating references that will cause the garbage collector to keep the keys around longer than needed.

`WeakKeyDictionary.keyrefs()`

Return an iterable of the weak references to the keys.

class weakref.**WeakValueDictionary** (*[dict]*)

Mapping class that references values weakly. Entries in the dictionary will be discarded when no strong reference to the value exists any more.

注解: Caution: Because a *WeakValueDictionary* is built on top of a Python dictionary, it must not change size when iterating over it. This can be difficult to ensure for a *WeakValueDictionary* because actions performed by the program during iteration may cause items in the dictionary to vanish “by magic” (as a side effect of garbage collection).

WeakValueDictionary objects have an additional method that has the same issues as the *keyrefs()* method of *WeakKeyDictionary* objects.

WeakValueDictionary.**valuerefs**()

Return an iterable of the weak references to the values.

class weakref.**WeakSet** (*[elements]*)

Set class that keeps weak references to its elements. An element will be discarded when no strong reference to it exists any more.

class weakref.**WeakMethod** (*method*)

A custom *ref* subclass which simulates a weak reference to a bound method (i.e., a method defined on a class and looked up on an instance). Since a bound method is ephemeral, a standard weak reference cannot keep hold of it. *WeakMethod* has special code to recreate the bound method until either the object or the original function dies:

```
>>> class C:
...     def method(self):
...         print("method called!")
...
>>> c = C()
>>> r = weakref.ref(c.method)
>>> r()
>>> r = weakref.WeakMethod(c.method)
>>> r()
<bound method C.method of <__main__.C object at 0x7fc859830220>>
>>> r() ()
method called!
>>> del c
>>> gc.collect()
0
>>> r()
>>>
```

3.4 新版功能.

class weakref.**finalize** (*obj, func, *args, **kwargs*)

Return a callable finalizer object which will be called when *obj* is garbage collected. Unlike an ordinary weak reference, a finalizer will always survive until the reference object is collected, greatly simplifying lifecycle management.

A finalizer is considered *alive* until it is called (either explicitly or at garbage collection), and after that it is *dead*. Calling a live finalizer returns the result of evaluating *func(*arg, **kwargs)*, whereas calling a dead finalizer returns *None*.

Exceptions raised by finalizer callbacks during garbage collection will be shown on the standard error output, but cannot be propagated. They are handled in the same way as exceptions raised from an object's *__del__()* method or a weak reference's callback.

When the program exits, each remaining live finalizer is called unless its `atexit` attribute has been set to false. They are called in reverse order of creation.

A finalizer will never invoke its callback during the later part of the *interpreter shutdown* when module globals are liable to have been replaced by `None`.

`__call__()`

If *self* is alive then mark it as dead and return the result of calling `func(*args, **kwargs)`. If *self* is dead then return `None`.

`detach()`

If *self* is alive then mark it as dead and return the tuple `(obj, func, args, kwargs)`. If *self* is dead then return `None`.

`peek()`

If *self* is alive then return the tuple `(obj, func, args, kwargs)`. If *self* is dead then return `None`.

`alive`

Property which is true if the finalizer is alive, false otherwise.

`atexit`

A writable boolean property which by default is true. When the program exits, it calls all remaining live finalizers for which `atexit` is true. They are called in reverse order of creation.

注解: It is important to ensure that *func*, *args* and *kwargs* do not own any references to *obj*, either directly or indirectly, since otherwise *obj* will never be garbage collected. In particular, *func* should not be a bound method of *obj*.

3.4 新版功能.

`weakref.ReferenceType`

The type object for weak references objects.

`weakref.ProxyType`

The type object for proxies of objects which are not callable.

`weakref.CallableProxyType`

The type object for proxies of callable objects.

`weakref.ProxyTypes`

Sequence containing all the type objects for proxies. This can make it simpler to test if an object is a proxy without being dependent on naming both proxy types.

exception `weakref.ReferenceError`

Exception raised when a proxy object is used but the underlying object has been collected. This is the same as the standard `ReferenceError` exception.

参见:

PEP 205 - Weak References The proposal and rationale for this feature, including links to earlier implementations and information about similar features in other languages.

8.8.1 弱引用对象

Weak reference objects have no methods and no attributes besides `ref.__callback__`. A weak reference object allows the referent to be obtained, if it still exists, by calling it:

```
>>> import weakref
>>> class Object:
...     pass
...
>>> o = Object()
>>> r = weakref.ref(o)
>>> o2 = r()
>>> o is o2
True
```

If the referent no longer exists, calling the reference object returns *None*:

```
>>> del o, o2
>>> print(r())
None
```

Testing that a weak reference object is still live should be done using the expression `ref() is not None`. Normally, application code that needs to use a reference object should follow this pattern:

```
# r is a weak reference object
o = r()
if o is None:
    # referent has been garbage collected
    print("Object has been deallocated; can't frobnicate.")
else:
    print("Object is still live!")
    o.do_something_useful()
```

Using a separate test for “liveness” creates race conditions in threaded applications; another thread can cause a weak reference to become invalidated before the weak reference is called; the idiom shown above is safe in threaded applications as well as single-threaded applications.

Specialized versions of *ref* objects can be created through subclassing. This is used in the implementation of the *WeakValueDictionary* to reduce the memory overhead for each entry in the mapping. This may be most useful to associate additional information with a reference, but could also be used to insert additional processing on calls to retrieve the referent.

This example shows how a subclass of *ref* can be used to store additional information about an object and affect the value that’s returned when the referent is accessed:

```
import weakref

class ExtendedRef(weakref.ref):
    def __init__(self, ob, callback=None, **annotations):
        super(ExtendedRef, self).__init__(ob, callback)
        self.__counter = 0
        for k, v in annotations.items():
            setattr(self, k, v)

    def __call__(self):
        """Return a pair containing the referent and the number of
        times the reference has been called.
        """
        ob = super(ExtendedRef, self).__call__()
        if ob is not None:
            self.__counter += 1
            ob = (ob, self.__counter)
        return ob
```

8.8.2 示例

This simple example shows how an application can use object IDs to retrieve objects that it has seen before. The IDs of the objects can then be used in other data structures without forcing the objects to remain alive, but the objects can still be retrieved by ID if they do.

```
import weakref

_id2obj_dict = weakref.WeakValueDictionary()

def remember(obj):
    oid = id(obj)
    _id2obj_dict[oid] = obj
    return oid

def id2obj(oid):
    return _id2obj_dict[oid]
```

8.8.3 Finalizer Objects

The main benefit of using `finalize` is that it makes it simple to register a callback without needing to preserve the returned finalizer object. For instance

```
>>> import weakref
>>> class Object:
...     pass
...
>>> kenny = Object()
>>> weakref.finalize(kenny, print, "You killed Kenny!") #doctest:+ELLIPSIS
<finalize object at ...; for 'Object' at ...>
>>> del kenny
You killed Kenny!
```

The finalizer can be called directly as well. However the finalizer will invoke the callback at most once.

```
>>> def callback(x, y, z):
...     print("CALLBACK")
...     return x + y + z
...
>>> obj = Object()
>>> f = weakref.finalize(obj, callback, 1, 2, z=3)
>>> assert f.alive
>>> assert f() == 6
CALLBACK
>>> assert not f.alive
>>> f() # callback not called because finalizer dead
>>> del obj # callback not called because finalizer dead
```

You can unregister a finalizer using its `detach()` method. This kills the finalizer and returns the arguments passed to the constructor when it was created.

```
>>> obj = Object()
>>> f = weakref.finalize(obj, callback, 1, 2, z=3)
>>> f.detach() #doctest:+ELLIPSIS
(<__main__.Object object ...>, <function callback ...>, (1, 2), {'z': 3})
```

(下页继续)

(续上页)

```
>>> newobj, func, args, kwargs = _
>>> assert not f.alive
>>> assert newobj is obj
>>> assert func(*args, **kwargs) == 6
CALLBACK
```

Unless you set the `atexit` attribute to `False`, a finalizer will be called when the program exits if it is still alive. For instance

```
>>> obj = Object()
>>> weakref.finalize(obj, print, "obj dead or exiting") #doctest:+ELLIPSIS
<finalize object at ...; for 'Object' at ...>
>>> exit() #doctest:+SKIP
obj dead or exiting
```

8.8.4 Comparing finalizers with `__del__()` methods

Suppose we want to create a class whose instances represent temporary directories. The directories should be deleted with their contents when the first of the following events occurs:

- the object is garbage collected,
- the object's `remove()` method is called, or
- the program exits.

We might try to implement the class using a `__del__()` method as follows:

```
class TempDir:
    def __init__(self):
        self.name = tempfile.mkdtemp()

    def remove(self):
        if self.name is not None:
            shutil.rmtree(self.name)
            self.name = None

    @property
    def removed(self):
        return self.name is None

    def __del__(self):
        self.remove()
```

Starting with Python 3.4, `__del__()` methods no longer prevent reference cycles from being garbage collected, and module globals are no longer forced to `None` during *interpreter shutdown*. So this code should work without any issues on CPython.

However, handling of `__del__()` methods is notoriously implementation specific, since it depends on internal details of the interpreter's garbage collector implementation.

A more robust alternative can be to define a finalizer which only references the specific functions and objects that it needs, rather than having access to the full state of the object:

```
class TempDir:
    def __init__(self):
```

(下页继续)

(续上页)

```
self.name = tempfile.mkdtemp()
self._finalizer = weakref.finalize(self, shutil.rmtree, self.name)

def remove(self):
    self._finalizer()

@property
def removed(self):
    return not self._finalizer.alive
```

Defined like this, our finalizer only receives a reference to the details it needs to clean up the directory appropriately. If the object never gets garbage collected the finalizer will still be called at exit.

The other advantage of weakref based finalizers is that they can be used to register finalizers for classes where the definition is controlled by a third party, such as running code when a module is unloaded:

```
import weakref, sys
def unloading_module():
    # implicit reference to the module globals from the function body
    weakref.finalize(sys.modules[__name__], unloading_module)
```

注解: If you create a finalizer object in a daemon thread just as the program exits then there is the possibility that the finalizer does not get called at exit. However, in a daemon thread `atexit.register()`, `try: ... finally:` ... and `with: ...` do not guarantee that cleanup occurs either.

8.9 types —Dynamic type creation and names for built-in types

Source code: [Lib/types.py](#)

This module defines utility function to assist in dynamic creation of new types.

It also defines names for some object types that are used by the standard Python interpreter, but not exposed as builtins like `int` or `str` are.

Finally, it provides some additional type-related utility classes and functions that are not fundamental enough to be builtins.

8.9.1 Dynamic Type Creation

`types.new_class(name, bases=(), kwds=None, exec_body=None)`
Creates a class object dynamically using the appropriate metaclass.

The first three arguments are the components that make up a class definition header: the class name, the base classes (in order), the keyword arguments (such as `metaclass`).

The `exec_body` argument is a callback that is used to populate the freshly created class namespace. It should accept the class namespace as its sole argument and update the namespace directly with the class contents. If no callback is provided, it has the same effect as passing in `lambda ns: ns`.

3.3 新版功能.

`types.prepare_class` (*name*, *bases*=(), *kwds*=None)

Calculates the appropriate metaclass and creates the class namespace.

The arguments are the components that make up a class definition header: the class name, the base classes (in order) and the keyword arguments (such as `metaclass`).

The return value is a 3-tuple: `metaclass`, `namespace`, `kwds`

metaclass is the appropriate metaclass, *namespace* is the prepared class namespace and *kwds* is an updated copy of the passed in *kwds* argument with any 'metaclass' entry removed. If no *kwds* argument is passed in, this will be an empty dict.

3.3 新版功能.

在 3.6 版更改: The default value for the `namespace` element of the returned tuple has changed. Now an insertion-order-preserving mapping is used when the metaclass does not have a `__prepare__` method,

参见:

metaclasses Full details of the class creation process supported by these functions

PEP 3115 - Python 3000 中的元类 引入 `__prepare__` 命名空间钩子

8.9.2 Standard Interpreter Types

This module provides names for many of the types that are required to implement a Python interpreter. It deliberately avoids including some of the types that arise only incidentally during processing such as the `listiterator` type.

Typical use of these names is for `isinstance()` or `issubclass()` checks.

Standard names are defined for the following types:

`types.FunctionType`

`types.LambdaType`

用户自定义函数以及由 `lambda` 表达式所创建函数的类型。

`types.GeneratorType`

generator 迭代器对象的类型，由生成器函数创建。

`types.CoroutineType`

coroutine 对象的类型，由 `async def` 函数创建。

3.5 新版功能.

`types.AsyncGeneratorType`

asynchronous generator 迭代器对象的类型，由异步生成器函数创建。

3.6 新版功能.

`types.CodeType`

The type for code objects such as returned by `compile()`.

`types.MethodType`

The type of methods of user-defined class instances.

`types.BuiltinFunctionType`

`types.BuiltinMethodType`

The type of built-in functions like `len()` or `sys.exit()`, and methods of built-in classes. (Here, the term “built-in” means “written in C”.)

class `types.ModuleType` (*name*, *doc*=None)

The type of *modules*. Constructor takes the name of the module to be created and optionally its *docstring*.

注解: Use `importlib.util.module_from_spec()` to create a new module if you wish to set the various import-controlled attributes.

__doc__

The *docstring* of the module. Defaults to `None`.

__loader__

The *loader* which loaded the module. Defaults to `None`.

在 3.4 版更改: Defaults to `None`. Previously the attribute was optional.

__name__

模块的名字

__package__

Which *package* a module belongs to. If the module is top-level (i.e. not a part of any specific package) then the attribute should be set to `' '`, else it should be set to the name of the package (which can be `__name__` if the module is a package itself). Defaults to `None`.

在 3.4 版更改: Defaults to `None`. Previously the attribute was optional.

types.TracebackType

The type of traceback objects such as found in `sys.exc_info()[2]`.

types.FrameType

The type of frame objects such as found in `tb.tb_frame` if `tb` is a traceback object.

types.GetSetDescriptorType

The type of objects defined in extension modules with `PyGetSetDef`, such as `FrameType.f_locals` or `array.array.typecode`. This type is used as descriptor for object attributes; it has the same purpose as the *property* type, but for classes defined in extension modules.

types.MemberDescriptorType

The type of objects defined in extension modules with `PyMemberDef`, such as `datetime.timedelta.days`. This type is used as descriptor for simple C data members which use standard conversion functions; it has the same purpose as the *property* type, but for classes defined in extension modules.

CPython implementation detail: In other implementations of Python, this type may be identical to `GetSetDescriptorType`.

class types.MappingProxyType(mapping)

Read-only proxy of a mapping. It provides a dynamic view on the mapping's entries, which means that when the mapping changes, the view reflects these changes.

3.3 新版功能.

key in proxy

Return `True` if the underlying mapping has a key *key*, else `False`.

proxy[key]

Return the item of the underlying mapping with key *key*. Raises a *KeyError* if *key* is not in the underlying mapping.

iter(proxy)

Return an iterator over the keys of the underlying mapping. This is a shortcut for `iter(proxy.keys())`.

len(proxy)

Return the number of items in the underlying mapping.

copy()

Return a shallow copy of the underlying mapping.

get (*key* [, *default*])

Return the value for *key* if *key* is in the underlying mapping, else *default*. If *default* is not given, it defaults to `None`, so that this method never raises a `KeyError`.

items ()

Return a new view of the underlying mapping's items ((*key*, *value*) pairs).

keys ()

Return a new view of the underlying mapping's keys.

values ()

Return a new view of the underlying mapping's values.

8.9.3 Additional Utility Classes and Functions

`class types.SimpleNamespace`

A simple *object* subclass that provides attribute access to its namespace, as well as a meaningful repr.

Unlike *object*, with `SimpleNamespace` you can add and remove attributes. If a `SimpleNamespace` object is initialized with keyword arguments, those are directly added to the underlying namespace.

The type is roughly equivalent to the following code:

```
class SimpleNamespace:
    def __init__(self, **kwargs):
        self.__dict__.update(kwargs)

    def __repr__(self):
        keys = sorted(self.__dict__)
        items = ("{}={!r}".format(k, self.__dict__[k]) for k in keys)
        return "{}({})".format(type(self).__name__, ", ".join(items))

    def __eq__(self, other):
        return self.__dict__ == other.__dict__
```

`SimpleNamespace` may be useful as a replacement for `class NS: pass`. However, for a structured record type use `namedtuple()` instead.

3.3 新版功能.

`types.DynamicClassAttribute` (*fget=None, fset=None, fdel=None, doc=None*)

Route attribute access on a class to `__getattr__`.

This is a descriptor, used to define attributes that act differently when accessed through an instance and through a class. Instance access remains normal, but access to an attribute through a class will be routed to the class's `__getattr__` method; this is done by raising `AttributeError`.

This allows one to have properties active on an instance, and have virtual attributes on the class with the same name (see `Enum` for an example).

3.4 新版功能.

8.9.4 Coroutine Utility Functions

`types.coroutine` (*gen_func*)

This function transforms a *generator* function into a *coroutine function* which returns a generator-based coroutine. The generator-based coroutine is still a *generator iterator*, but is also considered to be a *coroutine* object and is *awaitable*. However, it may not necessarily implement the `__await__()` method.

If `gen_func` is a generator function, it will be modified in-place.

If `gen_func` is not a generator function, it will be wrapped. If it returns an instance of `collections.abc.Generator`, the instance will be wrapped in an *awaitable* proxy object. All other types of objects will be returned as is.

3.5 新版功能.

8.10 copy — 浅层 (shallow) 和深层 (deep) 复制操作

源代码: [Lib/copy.py](#)

Python 中赋值语句不复制对象，而是在目标和对象之间创建绑定 (bindings) 关系。对于自身可变或者包含可变项的集合对象，开发者有时需要生成其副本用于改变操作，进而避免改变原对象。本模块提供了通用的浅层复制和深层复制操作（如下所述）。

接口摘要：

`copy.copy(x)`

返回 `x` 的浅层复制。

`copy.deepcopy(x[, memo])`

返回 `x` 的深层复制。

exception `copy.error`

针对模块特定错误引发。

浅层复制和深层复制之间的区别仅与复合对象（即包含其他对象的对象，如列表或类的实例）相关：

- 一个 浅层复制会构造一个新的复合对象，然后（在可能的范围内）将原对象中找到的 引用插入其中。
- 一个 深层复制会构造一个新的复合对象，然后递归地将原始对象中所找到的对象的 副本插入。

深度复制操作通常存在两个问题，而浅层复制操作并不存在这些问题：

- 递归对象（直接或间接包含对自身引用的复合对象）可能会导致递归循环。
- 由于深层复制会复制所有内容，因此可能会过多复制（例如本应该在副本之间共享的数据）。

The `deepcopy()` function avoids these problems by:

- 保留在当前复制过程中已复制的对象的“备忘录” (memo) 字典；以及
- 允许用户定义的类型重载复制操作或复制的组件集合。

该模块不复制模块、方法、栈追踪 (stack trace)、栈帧 (stack frame)、文件、套接字、窗口、数组以及任何类似的类型。它通过不改变地返回原始对象来（浅层或深层地）“复制”函数和类；这与 `pickle` 模块处理这类问题的方式是相似的。

制作字典的浅层复制可以使用 `dict.copy()` 方法，而制作列表的浅层复制可以通过赋值整个列表的切片完成，例如，`copied_list = original_list[:]`。

类可以使用与控制序列化 (pickling) 操作相同的接口来控制复制操作，关于这些方法的描述信息请参考 `pickle` 模块。实际上，`copy` 模块使用的正是从 `copyreg` 模块中注册的 `pickle` 函数。

想要给一个类定义它自己的拷贝操作实现，可以通过定义特殊方法 `__copy__()` 和 `__deepcopy__()`。调用前者以实现浅层拷贝操作，该方法不用传入额外参数。调用后者以实现深层拷贝操作；它应传入一个参数即 memo 字典。如果 `__deepcopy__()` 实现需要创建一个组件的深层拷贝，它应当调用 `deepcopy()` 函数并以该组件作为第一个参数，而将 memo 字典作为第二个参数。

参见：

模块 `pickle` 讨论了支持对象状态检索和恢复的特殊方法。

8.11 pprint — 数据美化输出

源代码: `Lib/pprint.py`

The `pprint` module provides a capability to “pretty-print” arbitrary Python data structures in a form which can be used as input to the interpreter. If the formatted structures include objects which are not fundamental Python types, the representation may not be loadable. This may be the case if objects such as files, sockets or classes are included, as well as many other objects which are not representable as Python literals.

The formatted representation keeps objects on a single line if it can, and breaks them onto multiple lines if they don't fit within the allowed width. Construct `PrettyPrinter` objects explicitly if you need to adjust the width constraint.

Dictionaries are sorted by key before the display is computed. 字典在被计算前通过键值来排序。

`pprint` 模块定义了一个类:

class `pprint.PrettyPrinter` (`indent=1`, `width=80`, `depth=None`, `stream=None`, *, `compact=False`)

Construct a `PrettyPrinter` instance. This constructor understands several keyword parameters. An output stream may be set using the `stream` keyword; the only method used on the stream object is the file protocol's `write()` method. If not specified, the `PrettyPrinter` adopts `sys.stdout`. The amount of indentation added for each recursive level is specified by `indent`; the default is one. Other values can cause output to look a little odd, but can make nesting easier to spot. The number of levels which may be printed is controlled by `depth`; if the data structure being printed is too deep, the next contained level is replaced by `...`. By default, there is no constraint on the depth of the objects being formatted. The desired output width is constrained using the `width` parameter; the default is 80 characters. If a structure cannot be formatted within the constrained width, a best effort will be made. If `compact` is false (the default) each item of a long sequence will be formatted on a separate line. If `compact` is true, as many items as will fit within the `width` will be formatted on each output line.

在 3.4 版更改: 加入 `*compact*` 参数。

```
>>> import pprint
>>> stuff = ['spam', 'eggs', 'lumberjack', 'knights', 'ni']
>>> stuff.insert(0, stuff[:])
>>> pp = pprint.PrettyPrinter(indent=4)
>>> pp.pprint(stuff)
[  ['spam', 'eggs', 'lumberjack', 'knights', 'ni'],
   ['spam',
    'eggs',
    'lumberjack',
    'knights',
    'ni']]
>>> pp = pprint.PrettyPrinter(width=41, compact=True)
>>> pp.pprint(stuff)
[['spam', 'eggs', 'lumberjack',
  'knights', 'ni'],
 ['spam', 'eggs', 'lumberjack', 'knights',
  'ni']]
>>> tup = ('spam', ('eggs', ('lumberjack', ('knights', ('ni', ('dead',
... ('parrot', ('fresh fruit',))))))))))
>>> pp = pprint.PrettyPrinter(depth=6)
>>> pp.pprint(tup)
('spam', ('eggs', ('lumberjack', ('knights', ('ni', ('dead', (...)))))))
```

The `pprint` module also provides several shortcut functions:

`pprint.pformat(object, indent=1, width=80, depth=None, *, compact=False)`

Return the formatted representation of *object* as a string. *indent*, *width*, *depth* and *compact* will be passed to the `PrettyPrinter` constructor as formatting parameters.

在 3.4 版更改: 加入 **compact** 参数。

`pprint.pprint(object, stream=None, indent=1, width=80, depth=None, *, compact=False)`

Prints the formatted representation of *object* on *stream*, followed by a newline. If *stream* is `None`, `sys.stdout` is used. This may be used in the interactive interpreter instead of the `print()` function for inspecting values (you can even reassign `print = pprint.pprint` for use within a scope). *indent*, *width*, *depth* and *compact* will be passed to the `PrettyPrinter` constructor as formatting parameters.

在 3.4 版更改: 加入 **compact** 参数。

```
>>> import pprint
>>> stuff = ['spam', 'eggs', 'lumberjack', 'knights', 'ni']
>>> stuff.insert(0, stuff)
>>> pprint.pprint(stuff)
[<Recursion on list with id=...>,
 'spam',
 'eggs',
 'lumberjack',
 'knights',
 'ni']
```

`pprint.isreadable(object)`

Determine if the formatted representation of *object* is “readable,” or can be used to reconstruct the value using `eval()`. This always returns `False` for recursive objects.

```
>>> pprint.isreadable(stuff)
False
```

`pprint.isrecursive(object)`

Determine if *object* requires a recursive representation.

One more support function is also defined:

`pprint.saferepr(object)`

Return a string representation of *object*, protected against recursive data structures. If the representation of *object* exposes a recursive entry, the recursive reference will be represented as `<Recursion on typename with id=number>`. The representation is not otherwise formatted.

```
>>> pprint.saferepr(stuff)
"[<Recursion on list with id=...>, 'spam', 'eggs', 'lumberjack', 'knights', 'ni']"
```

8.11.1 PrettyPrinter Objects

`PrettyPrinter` instances have the following methods:

`PrettyPrinter.pformat(object)`

Return the formatted representation of *object*. This takes into account the options passed to the `PrettyPrinter` constructor.

`PrettyPrinter.pprint(object)`

Print the formatted representation of *object* on the configured stream, followed by a newline.

The following methods provide the implementations for the corresponding functions of the same names. Using these methods on an instance is slightly more efficient since new `PrettyPrinter` objects don't need to be created.

`PrettyPrinter.isreadable(object)`

Determine if the formatted representation of the object is “readable,” or can be used to reconstruct the value using `eval()`. Note that this returns `False` for recursive objects. If the `depth` parameter of the `PrettyPrinter` is set and the object is deeper than allowed, this returns `False`.

`PrettyPrinter.isrecursive(object)`

Determine if the object requires a recursive representation.

This method is provided as a hook to allow subclasses to modify the way objects are converted to strings. The default implementation uses the internals of the `saferepr()` implementation.

`PrettyPrinter.format(object, context, maxlevels, level)`

Returns three values: the formatted version of `object` as a string, a flag indicating whether the result is readable, and a flag indicating whether recursion was detected. The first argument is the object to be presented. The second is a dictionary which contains the `id()` of objects that are part of the current presentation context (direct and indirect containers for `object` that are affecting the presentation) as the keys; if an object needs to be presented which is already represented in `context`, the third return value should be `True`. Recursive calls to the `format()` method should add additional entries for containers to this dictionary. The third argument, `maxlevels`, gives the requested limit to recursion; this will be 0 if there is no requested limit. This argument should be passed unmodified to recursive calls. The fourth argument, `level`, gives the current level; recursive calls should be passed a value less than that of the current call.

8.11.2 示例

To demonstrate several uses of the `pprint()` function and its parameters, let’s fetch information about a project from PyPI:

```
>>> import json
>>> import pprint
>>> from urllib.request import urlopen
>>> with urlopen('https://pypi.org/pypi/sampleproject/json') as resp:
...     project_info = json.load(resp)['info']
```

In its basic form, `pprint()` shows the whole object:

```
>>> pprint.pprint(project_info)
{'author': 'The Python Packaging Authority',
 'author_email': 'pypa-dev@googlegroups.com',
 'bugtrack_url': None,
 'classifiers': ['Development Status :: 3 - Alpha',
                  'Intended Audience :: Developers',
                  'License :: OSI Approved :: MIT License',
                  'Programming Language :: Python :: 2',
                  'Programming Language :: Python :: 2.6',
                  'Programming Language :: Python :: 2.7',
                  'Programming Language :: Python :: 3',
                  'Programming Language :: Python :: 3.2',
                  'Programming Language :: Python :: 3.3',
                  'Programming Language :: Python :: 3.4',
                  'Topic :: Software Development :: Build Tools'],
 'description': 'A sample Python project\n'
                '=====\n'
                '\n'
                'This is the description file for the project.\n'
                '\n'
                'The file should use UTF-8 encoding and be written using '}
```

(下页继续)

(续上页)

```

'ReStructured Text. It\n'
'will be used to generate the project webpage on PyPI, and '
'should be written for\n'
'that purpose.\n'
'\n'
'Typical contents for this file would include an overview of '
'the project, basic\n'
'usage examples, etc. Generally, including the project '
'changelog in here is not\n'
'a good idea, although a simple "What\'s New" section for the '
'most recent version\n'
'may be appropriate.',
'description_content_type': None,
'docs_url': None,
'download_url': 'UNKNOWN',
'downloads': {'last_day': -1, 'last_month': -1, 'last_week': -1},
'home_page': 'https://github.com/pypa/sampleproject',
'keywords': 'sample setuptools development',
'license': 'MIT',
'maintainer': None,
'maintainer_email': None,
'name': 'sampleproject',
'package_url': 'https://pypi.org/project/sampleproject/',
'platform': 'UNKNOWN',
'project_url': 'https://pypi.org/project/sampleproject/',
'project_urls': {'Download': 'UNKNOWN',
                  'Homepage': 'https://github.com/pypa/sampleproject'},
'release_url': 'https://pypi.org/project/sampleproject/1.2.0/',
'requires_dist': None,
'requires_python': None,
'summary': 'A sample Python project',
'version': '1.2.0'}

```

The result can be limited to a certain *depth* (ellipsis is used for deeper contents):

```

>>> pprint.pprint(project_info, depth=1)
{'author': 'The Python Packaging Authority',
 'author_email': 'pypa-dev@googlegroups.com',
 'bugtrack_url': None,
 'classifiers': [...],
 'description': 'A sample Python project\n'
               '=====\n'
               '\n'
               'This is the description file for the project.\n'
               '\n'
               'The file should use UTF-8 encoding and be written using '
               'ReStructured Text. It\n'
               'will be used to generate the project webpage on PyPI, and '
               'should be written for\n'
               'that purpose.\n'
               '\n'
               'Typical contents for this file would include an overview of '
               'the project, basic\n'
               'usage examples, etc. Generally, including the project '
               'changelog in here is not\n'
               'a good idea, although a simple "What\'s New" section for the '

```

(下页继续)

(续上页)

```

        'most recent version\n'
        'may be appropriate.',
'description_content_type': None,
'docs_url': None,
'download_url': 'UNKNOWN',
'downloads': {...},
'home_page': 'https://github.com/pypa/sampleproject',
'keywords': 'sample setuptools development',
'license': 'MIT',
'maintainer': None,
'maintainer_email': None,
'name': 'sampleproject',
'package_url': 'https://pypi.org/project/sampleproject/',
'platform': 'UNKNOWN',
'project_url': 'https://pypi.org/project/sampleproject/',
'project_urls': {...},
'release_url': 'https://pypi.org/project/sampleproject/1.2.0/',
'requires_dist': None,
'requires_python': None,
'summary': 'A sample Python project',
'version': '1.2.0'}

```

Additionally, maximum character *width* can be suggested. If a long object cannot be split, the specified width will be exceeded:

```

>>> pprint.pprint(project_info, depth=1, width=60)
{'author': 'The Python Packaging Authority',
 'author_email': 'pypa-dev@googlegroups.com',
 'bugtrack_url': None,
 'classifiers': [...],
 'description': 'A sample Python project\n'
               '=====\n'
               '\n'
               'This is the description file for the '
               'project.\n'
               '\n'
               'The file should use UTF-8 encoding and be '
               'written using ReStructured Text. It\n'
               'will be used to generate the project '
               'webpage on PyPI, and should be written '
               'for\n'
               'that purpose.\n'
               '\n'
               'Typical contents for this file would '
               'include an overview of the project, '
               'basic\n'
               'usage examples, etc. Generally, including '
               'the project changelog in here is not\n'
               'a good idea, although a simple "What\'s '
               'New" section for the most recent version\n'
               'may be appropriate.',
'description_content_type': None,
'docs_url': None,
'download_url': 'UNKNOWN',
'downloads': {...},
'home_page': 'https://github.com/pypa/sampleproject',

```

(下页继续)

(续上页)

```
'keywords': 'sample setuptools development',
'license': 'MIT',
'maintainer': None,
'maintainer_email': None,
'name': 'sampleproject',
'package_url': 'https://pypi.org/project/sampleproject/',
'platform': 'UNKNOWN',
'project_url': 'https://pypi.org/project/sampleproject/',
'project_urls': {...},
'release_url': 'https://pypi.org/project/sampleproject/1.2.0/',
'requires_dist': None,
'requires_python': None,
'summary': 'A sample Python project',
'version': '1.2.0'}
```

8.12 reprlib — Alternate repr() implementation

Source code: [Lib/reprlib.py](#)

The `reprlib` module provides a means for producing object representations with limits on the size of the resulting strings. This is used in the Python debugger and may be useful in other contexts as well.

This module provides a class, an instance, and a function:

class reprlib.Repr

Class which provides formatting services useful in implementing functions similar to the built-in `repr()`; size limits for different object types are added to avoid the generation of representations which are excessively long.

`reprlib.aRepr`

This is an instance of `Repr` which is used to provide the `repr()` function described below. Changing the attributes of this object will affect the size limits used by `repr()` and the Python debugger.

`reprlib.repr(obj)`

This is the `repr()` method of `aRepr`. It returns a string similar to that returned by the built-in function of the same name, but with limits on most sizes.

In addition to size-limiting tools, the module also provides a decorator for detecting recursive calls to `__repr__()` and substituting a placeholder string instead.

`@reprlib.recursive_repr(fillvalue="...")`

Decorator for `__repr__()` methods to detect recursive calls within the same thread. If a recursive call is made, the `fillvalue` is returned, otherwise, the usual `__repr__()` call is made. For example:

```
>>> class MyList(list):
...     @recursive_repr()
...     def __repr__(self):
...         return '<' + '|'.join(map(repr, self)) + '>'
...
>>> m = MyList('abc')
>>> m.append(m)
>>> m.append('x')
>>> print(m)
<'a'|'b'|'c'|...|'x'>
```

3.2 新版功能.

8.12.1 Repr Objects

Repr instances provide several attributes which can be used to provide size limits for the representations of different object types, and methods which format specific object types.

`Repr.maxlevel`

Depth limit on the creation of recursive representations. The default is 6.

`Repr.maxdict`

`Repr.maxlist`

`Repr.maxtuple`

`Repr.maxset`

`Repr.maxfrozenset`

`Repr.maxdeque`

`Repr.maxarray`

Limits on the number of entries represented for the named object type. The default is 4 for *maxdict*, 5 for *maxarray*, and 6 for the others.

`Repr.maxlong`

Maximum number of characters in the representation for an integer. Digits are dropped from the middle. The default is 40.

`Repr.maxstring`

Limit on the number of characters in the representation of the string. Note that the “normal” representation of the string is used as the character source: if escape sequences are needed in the representation, these may be mangled when the representation is shortened. The default is 30.

`Repr.maxother`

This limit is used to control the size of object types for which no specific formatting method is available on the *Repr* object. It is applied in a similar manner as *maxstring*. The default is 20.

`Repr.repr(obj)`

The equivalent to the built-in *repr()* that uses the formatting imposed by the instance.

`Repr.repr1(obj, level)`

Recursive implementation used by *repr()*. This uses the type of *obj* to determine which formatting method to call, passing it *obj* and *level*. The type-specific methods should call *repr1()* to perform recursive formatting, with *level - 1* for the value of *level* in the recursive call.

`Repr.repr_TYPE(obj, level)`

Formatting methods for specific types are implemented as methods with a name based on the type name. In the method name, **TYPE** is replaced by `'_'.join(type(obj).__name__.split())`. Dispatch to these methods is handled by *repr1()*. Type-specific methods which need to recursively format a value should call `self.repr1(subobj, level - 1)`.

8.12.2 Subclassing Repr Objects

The use of dynamic dispatching by *Repr.repr1()* allows subclasses of *Repr* to add support for additional built-in object types or to modify the handling of types already supported. This example shows how special support for file objects could be added:

```
import reprlib
import sys

class MyRepr(reprlib.Repr):

    def repr_TextIOWrapper(self, obj, level):
```

(下页继续)

(续上页)

```

    if obj.name in {'<stdin>', '<stdout>', '<stderr>'}:
        return obj.name
    return repr(obj)

aRepr = MyRepr()
print(aRepr.repr(sys.stdin))           # prints '<stdin>'

```

8.13 enum — Support for enumerations

3.4 新版功能.

Source code: [Lib/enum.py](#)

An enumeration is a set of symbolic names (members) bound to unique, constant values. Within an enumeration, the members can be compared by identity, and the enumeration itself can be iterated over.

8.13.1 模块内容

This module defines four enumeration classes that can be used to define unique sets of names and values: *Enum*, *IntEnum*, *Flag*, and *IntFlag*. It also defines one decorator, *unique()*, and one helper, *auto*.

class `enum.Enum`

Base class for creating enumerated constants. See section *Functional API* for an alternate construction syntax.

class `enum.IntEnum`

Base class for creating enumerated constants that are also subclasses of *int*.

class `enum.IntFlag`

Base class for creating enumerated constants that can be combined using the bitwise operators without losing their *IntFlag* membership. *IntFlag* members are also subclasses of *int*.

class `enum.Flag`

Base class for creating enumerated constants that can be combined using the bitwise operations without losing their *Flag* membership.

`enum.unique()`

Enum class decorator that ensures only one name is bound to any one value.

class `enum.auto`

Instances are replaced with an appropriate value for Enum members.

3.6 新版功能: `Flag`, `IntFlag`, `auto`

8.13.2 Creating an Enum

Enumerations are created using the `class` syntax, which makes them easy to read and write. An alternative creation method is described in *Functional API*. To define an enumeration, subclass *Enum* as follows:

```

>>> from enum import Enum
>>> class Color(Enum):
...     RED = 1

```

(下页继续)

(续上页)

```
...     GREEN = 2
...     BLUE = 3
...
```

注解: Enum member values

Member values can be anything: *int*, *str*, etc.. If the exact value is unimportant you may use *auto* instances and an appropriate value will be chosen for you. Care must be taken if you mix *auto* with other values.

注解: Nomenclature

- The class `Color` is an *enumeration* (or *enum*)
- The attributes `Color.RED`, `Color.GREEN`, etc., are *enumeration members* (or *enum members*) and are functionally constants.
- The enum members have *names* and *values* (the name of `Color.RED` is `RED`, the value of `Color.BLUE` is `3`, etc.)

注解: Even though we use the `class` syntax to create Enums, Enums are not normal Python classes. See [How are Enums different?](#) for more details.

Enumeration members have human readable string representations:

```
>>> print(Color.RED)
Color.RED
```

...while their `repr` has more information:

```
>>> print(repr(Color.RED))
<Color.RED: 1>
```

The *type* of an enumeration member is the enumeration it belongs to:

```
>>> type(Color.RED)
<enum 'Color'>
>>> isinstance(Color.GREEN, Color)
True
>>>
```

Enum members also have a property that contains just their item name:

```
>>> print(Color.RED.name)
RED
```

Enumerations support iteration, in definition order:

```
>>> class Shake(Enum):
...     VANILLA = 7
...     CHOCOLATE = 4
...     COOKIES = 9
...     MINT = 3
```

(下页继续)

(续上页)

```

...
>>> for shake in Shake:
...     print(shake)
...
Shake.VANILLA
Shake.CHOCOLATE
Shake.COOKIES
Shake.MINT

```

Enumeration members are hashable, so they can be used in dictionaries and sets:

```

>>> apples = {}
>>> apples[Color.RED] = 'red delicious'
>>> apples[Color.GREEN] = 'granny smith'
>>> apples == {Color.RED: 'red delicious', Color.GREEN: 'granny smith'}
True

```

8.13.3 Programmatic access to enumeration members and their attributes

Sometimes it's useful to access members in enumerations programmatically (i.e. situations where `Color.RED` won't do because the exact color is not known at program-writing time). Enum allows such access:

```

>>> Color(1)
<Color.RED: 1>
>>> Color(3)
<Color.BLUE: 3>

```

If you want to access enum members by *name*, use item access:

```

>>> Color['RED']
<Color.RED: 1>
>>> Color['GREEN']
<Color.GREEN: 2>

```

If you have an enum member and need its name or value:

```

>>> member = Color.RED
>>> member.name
'RED'
>>> member.value
1

```

8.13.4 Duplicating enum members and values

Having two enum members with the same name is invalid:

```

>>> class Shape(Enum):
...     SQUARE = 2
...     SQUARE = 3
...
Traceback (most recent call last):
...
TypeError: Attempted to reuse key: 'SQUARE'

```


(续上页)

```
>>> list(Color)
[<Color.RED: 1>, <Color.BLUE: 2>, <Color.GREEN: 3>]
```

The values are chosen by `_generate_next_value_()`, which can be overridden:

```
>>> class AutoName(Enum):
...     def _generate_next_value_(name, start, count, last_values):
...         return name
...
>>> class Ordinal(AutoName):
...     NORTH = auto()
...     SOUTH = auto()
...     EAST = auto()
...     WEST = auto()
...
>>> list(Ordinal)
[<Ordinal.NORTH: 'NORTH'>, <Ordinal.SOUTH: 'SOUTH'>, <Ordinal.EAST: 'EAST'>, <Ordinal.
↪WEST: 'WEST'>]
```

注解: The goal of the default `_generate_next_value_()` methods is to provide the next `int` in sequence with the last `int` provided, but the way it does this is an implementation detail and may change.

8.13.7 Iteration

Iterating over the members of an enum does not provide the aliases:

```
>>> list(Shape)
[<Shape.SQUARE: 2>, <Shape.DIAMOND: 1>, <Shape.CIRCLE: 3>]
```

The special attribute `__members__` is an ordered dictionary mapping names to members. It includes all names defined in the enumeration, including the aliases:

```
>>> for name, member in Shape.__members__.items():
...     name, member
...
('SQUARE', <Shape.SQUARE: 2>)
('DIAMOND', <Shape.DIAMOND: 1>)
('CIRCLE', <Shape.CIRCLE: 3>)
('ALIAS_FOR_SQUARE', <Shape.SQUARE: 2>)
```

The `__members__` attribute can be used for detailed programmatic access to the enumeration members. For example, finding all the aliases:

```
>>> [name for name, member in Shape.__members__.items() if member.name != name]
['ALIAS_FOR_SQUARE']
```

8.13.8 比较

Enumeration members are compared by identity:

```
>>> Color.RED is Color.RED
True
>>> Color.RED is Color.BLUE
False
>>> Color.RED is not Color.BLUE
True
```

Ordered comparisons between enumeration values are *not* supported. Enum members are not integers (but see [IntEnum](#) below):

```
>>> Color.RED < Color.BLUE
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
TypeError: '<' not supported between instances of 'Color' and 'Color'
```

Equality comparisons are defined though:

```
>>> Color.BLUE == Color.RED
False
>>> Color.BLUE != Color.RED
True
>>> Color.BLUE == Color.BLUE
True
```

Comparisons against non-enumeration values will always compare not equal (again, [IntEnum](#) was explicitly designed to behave differently, see below):

```
>>> Color.BLUE == 2
False
```

8.13.9 Allowed members and attributes of enumerations

The examples above use integers for enumeration values. Using integers is short and handy (and provided by default by the [Functional API](#)), but not strictly enforced. In the vast majority of use-cases, one doesn't care what the actual value of an enumeration is. But if the value *is* important, enumerations can have arbitrary values.

Enumerations are Python classes, and can have methods and special methods as usual. If we have this enumeration:

```
>>> class Mood(Enum):
...     FUNKY = 1
...     HAPPY = 3
...
...     def describe(self):
...         # self is the member here
...         return self.name, self.value
...
...     def __str__(self):
...         return 'my custom str! {0}'.format(self.value)
...
...     @classmethod
...     def favorite_mood(cls):
...         # cls here is the enumeration
...         return cls.HAPPY
...
... 
```

Then:

```
>>> Mood.favorite_mood()
<Mood.HAPPY: 3>
>>> Mood.HAPPY.describe()
('HAPPY', 3)
>>> str(Mood.FUNKY)
'my custom str! 1'
```

The rules for what is allowed are as follows: names that start and end with a single underscore are reserved by enum and cannot be used; all other attributes defined within an enumeration will become members of this enumeration, with the exception of special methods (`__str__()`, `__add__()`, etc.) and descriptors (methods are also descriptors).

Note: if your enumeration defines `__new__()` and/or `__init__()` then whatever value(s) were given to the enum member will be passed into those methods. See [Planet](#) for an example.

8.13.10 Restricted subclassing of enumerations

Subclassing an enumeration is allowed only if the enumeration does not define any members. So this is forbidden:

```
>>> class MoreColor(Color):
...     PINK = 17
...
Traceback (most recent call last):
...
TypeError: Cannot extend enumerations
```

But this is allowed:

```
>>> class Foo(Enum):  
...     def some_behavior(self):  
...         pass  
...  
>>> class Bar(Foo):  
...     HAPPY = 1  
...     SAD = 2  
...
```

Allowing subclassing of enums that define members would lead to a violation of some important invariants of types and instances. On the other hand, it makes sense to allow sharing some common behavior between a group of enumerations. (See *OrderedEnum* for an example.)

8.13.11 Pickling

Enumerations can be pickled and unpickled:

```
>>> from test.test_enum import Fruit
>>> from pickle import dumps, loads
>>> Fruit.TOMATO is loads(dumps(Fruit.TOMATO))
True
```

The usual restrictions for pickling apply: picklable enums must be defined in the top level of a module, since unpickling requires them to be importable from that module.

注解: With pickle protocol version 4 it is possible to easily pickle enums nested in other classes.

It is possible to modify how Enum members are pickled/unpickled by defining `__reduce_ex__()` in the enumeration class.

8.13.12 Functional API

The `Enum` class is callable, providing the following functional API:

```
>>> Animal = Enum('Animal', 'ANT BEE CAT DOG')
>>> Animal
<enum 'Animal'>
>>> Animal.ANT
<Animal.ANT: 1>
>>> Animal.ANT.value
1
>>> list(Animal)
[<Animal.ANT: 1>, <Animal.BEE: 2>, <Animal.CAT: 3>, <Animal.DOG: 4>]
```

The semantics of this API resemble *namedtuple*. The first argument of the call to `Enum` is the name of the enumeration.

The second argument is the *source* of enumeration member names. It can be a whitespace-separated string of names, a sequence of names, a sequence of 2-tuples with key/value pairs, or a mapping (e.g. dictionary) of names to values. The last two options enable assigning arbitrary values to enumerations; the others auto-assign increasing integers starting with 1 (use the `start` parameter to specify a different starting value). A new class derived from `Enum` is returned. In other words, the above assignment to `Animal` is equivalent to:

```
>>> class Animal(Enum):
...     ANT = 1
...     BEE = 2
...     CAT = 3
...     DOG = 4
... 
```

The reason for defaulting to 1 as the starting number and not 0 is that 0 is `False` in a boolean sense, but enum members all evaluate to `True`.

Pickling enums created with the functional API can be tricky as frame stack implementation details are used to try and figure out which module the enumeration is being created in (e.g. it will fail if you use a utility function in separate module, and also may not work on IronPython or Jython). The solution is to specify the module name explicitly as follows:

```
>>> Animal = Enum('Animal', 'ANT BEE CAT DOG', module=__name__)
```

警告: If module is not supplied, and Enum cannot determine what it is, the new Enum members will not be unpicklable; to keep errors closer to the source, pickling will be disabled.

The new pickle protocol 4 also, in some circumstances, relies on `__qualname__` being set to the location where pickle will be able to find the class. For example, if the class was made available in class `SomeData` in the global scope:

```
>>> Animal = Enum('Animal', 'ANT BEE CAT DOG', qualname='SomeData.Animal')
```

The complete signature is:

```
Enum(value='NewEnumName', names=<...>, *, module='...', qualname='...', type=<mixed-
↳ in class>, start=1)
```

值 What the new Enum class will record as its name.

names The Enum members. This can be a whitespace or comma separated string (values will start at 1 unless otherwise specified):

```
'RED GREEN BLUE' | 'RED, GREEN, BLUE' | 'RED, GREEN, BLUE'
```

or an iterator of names:

```
['RED', 'GREEN', 'BLUE']
```

or an iterator of (name, value) pairs:

```
[('CYAN', 4), ('MAGENTA', 5), ('YELLOW', 6)]
```

or a mapping:

```
{'CHARTREUSE': 7, 'SEA_GREEN': 11, 'ROSEMARY': 42}
```

模块 name of module where new Enum class can be found.

qualname where in module new Enum class can be found.

类型 type to mix in to new Enum class.

start number to start counting at if only names are passed in.

在 3.5 版更改: The *start* parameter was added.

8.13.13 Derived Enumerations

IntEnum

The first variation of *Enum* that is provided is also a subclass of *int*. Members of an *IntEnum* can be compared to integers; by extension, integer enumerations of different types can also be compared to each other:

```
>>> from enum import IntEnum
>>> class Shape(IntEnum):
...     CIRCLE = 1
...     SQUARE = 2
...
>>> class Request(IntEnum):
...     POST = 1
...     GET = 2
...
>>> Shape == 1
False
>>> Shape.CIRCLE == 1
True
>>> Shape.CIRCLE == Request.POST
True
```

However, they still can't be compared to standard *Enum* enumerations:

```
>>> class Shape(IntEnum):
...     CIRCLE = 1
...     SQUARE = 2
...
```

(下页继续)

(续上页)

```
>>> class Color(Enum):
...     RED = 1
...     GREEN = 2
...
>>> Shape.CIRCLE == Color.RED
False
```

IntEnum values behave like integers in other ways you'd expect:

```
>>> int(Shape.CIRCLE)
1
>>> ['a', 'b', 'c'][Shape.CIRCLE]
'b'
>>> [i for i in range(Shape.SQUARE)]
[0, 1]
```

IntFlag

The next variation of *Enum* provided, *IntFlag*, is also based on *int*. The difference being *IntFlag* members can be combined using the bitwise operators (&, |, ^, ~) and the result is still an *IntFlag* member. However, as the name implies, *IntFlag* members also subclass *int* and can be used wherever an *int* is used. Any operation on an *IntFlag* member besides the bit-wise operations will lose the *IntFlag* membership.

3.6 新版功能.

Sample *IntFlag* class:

```
>>> from enum import IntFlag
>>> class Perm(IntFlag):
...     R = 4
...     W = 2
...     X = 1
...
>>> Perm.R | Perm.W
<Perm.R|W: 6>
>>> Perm.R + Perm.W
6
>>> RW = Perm.R | Perm.W
>>> Perm.R in RW
True
```

It is also possible to name the combinations:

```
>>> class Perm(IntFlag):
...     R = 4
...     W = 2
...     X = 1
...     RWX = 7
>>> Perm.RWX
<Perm.RWX: 7>
>>> ~Perm.RWX
<Perm.-8: -8>
```

Another important difference between *IntFlag* and *Enum* is that if no flags are set (the value is 0), its boolean evaluation is *False*:

```
>>> Perm.R & Perm.X
<Perm.0: 0>
>>> bool(Perm.R & Perm.X)
False
```

Because *IntFlag* members are also subclasses of *int* they can be combined with them:

```
>>> Perm.X | 8
<Perm.8|X: 9>
```

标志

The last variation is *Flag*. Like *IntFlag*, *Flag* members can be combined using the bitwise operators (&, |, ^, ~). Unlike *IntFlag*, they cannot be combined with, nor compared against, any other *Flag* enumeration, nor *int*. While it is possible to specify the values directly it is recommended to use *auto* as the value and let *Flag* select an appropriate value.

3.6 新版功能.

Like *IntFlag*, if a combination of *Flag* members results in no flags being set, the boolean evaluation is *False*:

```
>>> from enum import Flag, auto
>>> class Color(Flag):
...     RED = auto()
...     BLUE = auto()
...     GREEN = auto()
...
>>> Color.RED & Color.GREEN
<Color.0: 0>
>>> bool(Color.RED & Color.GREEN)
False
```

Individual flags should have values that are powers of two (1, 2, 4, 8, ...), while combinations of flags won't:

```
>>> class Color(Flag):
...     RED = auto()
...     BLUE = auto()
...     GREEN = auto()
...     WHITE = RED | BLUE | GREEN
...
>>> Color.WHITE
<Color.WHITE: 7>
```

Giving a name to the “no flags set” condition does not change its boolean value:

```
>>> class Color(Flag):
...     BLACK = 0
...     RED = auto()
...     BLUE = auto()
...     GREEN = auto()
...
>>> Color.BLACK
<Color.BLACK: 0>
>>> bool(Color.BLACK)
False
```

注解: For the majority of new code, `Enum` and `Flag` are strongly recommended, since `IntEnum` and `IntFlag` break some semantic promises of an enumeration (by being comparable to integers, and thus by transitivity to other unrelated enumerations). `IntEnum` and `IntFlag` should be used only in cases where `Enum` and `Flag` will not do; for example, when integer constants are replaced with enumerations, or for interoperability with other systems.

Others

While `IntEnum` is part of the `enum` module, it would be very simple to implement independently:

```
class IntEnum(int, Enum):
    pass
```

This demonstrates how similar derived enumerations can be defined; for example a `StrEnum` that mixes in `str` instead of `int`.

Some rules:

1. When subclassing `Enum`, mix-in types must appear before `Enum` itself in the sequence of bases, as in the `IntEnum` example above.
2. While `Enum` can have members of any type, once you mix in an additional type, all the members must have values of that type, e.g. `int` above. This restriction does not apply to mix-ins which only add methods and don't specify another data type such as `int` or `str`.
3. When another data type is mixed in, the `value` attribute is *not the same* as the enum member itself, although it is equivalent and will compare equal.
4. %-style formatting: `%s` and `%r` call the `Enum` class's `__str__()` and `__repr__()` respectively; other codes (such as `%i` or `%h` for `IntEnum`) treat the enum member as its mixed-in type.
5. Formatted string literals, `str.format()`, and `format()` will use the mixed-in type's `__format__()`. If the `Enum` class's `str()` or `repr()` is desired, use the `!s` or `!r` format codes.

8.13.14 Interesting examples

While `Enum`, `IntEnum`, `IntFlag`, and `Flag` are expected to cover the majority of use-cases, they cannot cover them all. Here are recipes for some different types of enumerations that can be used directly, or as examples for creating one's own.

Omitting values

In many use-cases one doesn't care what the actual value of an enumeration is. There are several ways to define this type of simple enumeration:

- use instances of `auto` for the value
- use instances of `object` as the value
- use a descriptive string as the value
- use a tuple as the value and a custom `__new__()` to replace the tuple with an `int` value

Using any of these methods signifies to the user that these values are not important, and also enables one to add, remove, or reorder members without having to renumber the remaining members.

Whichever method you choose, you should provide a `repr()` that also hides the (unimportant) value:

```
>>> class NoValue(Enum):
...     def __repr__(self):
...         return '<%s.%s>' % (self.__class__.__name__, self.name)
... 
```

Using `auto`

Using `auto` would look like:

```
>>> class Color(NoValue):
...     RED = auto()
...     BLUE = auto()
...     GREEN = auto()
...
>>> Color.GREEN
<Color.GREEN>
```

Using `object`

Using `object` would look like:

```
>>> class Color(NoValue):
...     RED = object()
...     GREEN = object()
...     BLUE = object()
...
>>> Color.GREEN
<Color.GREEN>
```

Using a descriptive string

Using a string as the value would look like:

```
>>> class Color(NoValue):
...     RED = 'stop'
...     GREEN = 'go'
...     BLUE = 'too fast!'
...
>>> Color.GREEN
<Color.GREEN>
>>> Color.GREEN.value
'go'
```

Using a custom `__new__()`

Using an auto-numbering `__new__()` would look like:

```
>>> class AutoNumber(NoValue):
...     def __new__(cls):
...         value = len(cls.__members__) + 1
```

(下页继续)

(续上页)

```

...     obj = object.__new__(cls)
...     obj._value_ = value
...     return obj
...
>>> class Color(AutoNumber):
...     RED = ()
...     GREEN = ()
...     BLUE = ()
...
>>> Color.GREEN
<Color.GREEN>
>>> Color.GREEN.value
2

```

注解: The `__new__()` method, if defined, is used during creation of the Enum members; it is then replaced by Enum's `__new__()` which is used after class creation for lookup of existing members.

OrderedEnum

An ordered enumeration that is not based on *IntEnum* and so maintains the normal *Enum* invariants (such as not being comparable to other enumerations):

```

>>> class OrderedEnum(Enum):
...     def __ge__(self, other):
...         if self.__class__ is other.__class__:
...             return self.value >= other.value
...         return NotImplemented
...     def __gt__(self, other):
...         if self.__class__ is other.__class__:
...             return self.value > other.value
...         return NotImplemented
...     def __le__(self, other):
...         if self.__class__ is other.__class__:
...             return self.value <= other.value
...         return NotImplemented
...     def __lt__(self, other):
...         if self.__class__ is other.__class__:
...             return self.value < other.value
...         return NotImplemented
...
>>> class Grade(OrderedEnum):
...     A = 5
...     B = 4
...     C = 3
...     D = 2
...     F = 1
...
>>> Grade.C < Grade.A
True

```

DuplicateFreeEnum

Raises an error if a duplicate member name is found instead of creating an alias:

```
>>> class DuplicateFreeEnum(Enum):
...     def __init__(self, *args):
...         cls = self.__class__
...         if any(self.value == e.value for e in cls):
...             a = self.name
...             e = cls(self.value).name
...             raise ValueError(
...                 "aliases not allowed in DuplicateFreeEnum: %r --> %r"
...                 % (a, e))
...
>>> class Color(DuplicateFreeEnum):
...     RED = 1
...     GREEN = 2
...     BLUE = 3
...     GRENE = 2
...
Traceback (most recent call last):
...
ValueError: aliases not allowed in DuplicateFreeEnum: 'GRENE' --> 'GREEN'
```

注解: This is a useful example for subclassing Enum to add or change other behaviors as well as disallowing aliases. If the only desired change is disallowing aliases, the `unique()` decorator can be used instead.

Planet

If `__new__()` or `__init__()` is defined the value of the enum member will be passed to those methods:

```
>>> class Planet(Enum):
...     MERCURY = (3.303e+23, 2.4397e6)
...     VENUS   = (4.869e+24, 6.0518e6)
...     EARTH   = (5.976e+24, 6.37814e6)
...     MARS    = (6.421e+23, 3.3972e6)
...     JUPITER = (1.9e+27, 7.1492e7)
...     SATURN  = (5.688e+26, 6.0268e7)
...     URANUS  = (8.686e+25, 2.5559e7)
...     NEPTUNE = (1.024e+26, 2.4746e7)
...     def __init__(self, mass, radius):
...         self.mass = mass          # in kilograms
...         self.radius = radius      # in meters
...     @property
...     def surface_gravity(self):
...         # universal gravitational constant (m3 kg-1 s-2)
...         G = 6.67300E-11
...         return G * self.mass / (self.radius * self.radius)
...
>>> Planet.EARTH.value
(5.976e+24, 6378140.0)
>>> Planet.EARTH.surface_gravity
9.802652743337129
```

8.13.15 How are Enums different?

Enums have a custom metaclass that affects many aspects of both derived Enum classes and their instances (members).

Enum Classes

The `EnumMeta` metaclass is responsible for providing the `__contains__()`, `__dir__()`, `__iter__()` and other methods that allow one to do things with an `Enum` class that fail on a typical class, such as `list(Color)` or `some_var in Color`. `EnumMeta` is responsible for ensuring that various other methods on the final `Enum` class are correct (such as `__new__()`, `__getnewargs__()`, `__str__()` and `__repr__()`).

Enum Members (aka instances)

The most interesting thing about Enum members is that they are singletons. `EnumMeta` creates them all while it is creating the `Enum` class itself, and then puts a custom `__new__()` in place to ensure that no new ones are ever instantiated by returning only the existing member instances.

Finer Points

Supported `__dunder__` names

`__members__` is an `OrderedDict` of `member_name:member` items. It is only available on the class.

`__new__()`, if specified, must create and return the enum members; it is also a very good idea to set the member's `_value_` appropriately. Once all the members are created it is no longer used.

Supported `_sunder_` names

- `_name_` – name of the member
- `_value_` – value of the member; can be set / modified in `__new__`
- `_missing_` – a lookup function used when a value is not found; may be overridden
- `_order_` – used in Python 2/3 code to ensure member order is consistent (class attribute, removed during class creation)
- `_generate_next_value_` – used by the *Functional API* and by *auto* to get an appropriate value for an enum member; may be overridden

3.6 新版功能: `_missing_`, `_order_`, `_generate_next_value_`

To help keep Python 2 / Python 3 code in sync an `_order_` attribute can be provided. It will be checked against the actual order of the enumeration and raise an error if the two do not match:

```
>>> class Color(Enum):
...     _order_ = 'RED GREEN BLUE'
...     RED = 1
...     BLUE = 3
...     GREEN = 2
...
Traceback (most recent call last):
...
TypeError: member order does not match _order_
```

注解: In Python 2 code the `_order_` attribute is necessary as definition order is lost before it can be recorded.

Enum member type

Enum members are instances of their *Enum* class, and are normally accessed as `EnumClass.member`. Under certain circumstances they can also be accessed as `EnumClass.member.member`, but you should never do this as that lookup may fail or, worse, return something besides the *Enum* member you are looking for (this is another good reason to use all-uppercase names for members):

```
>>> class FieldTypes(Enum):
...     name = 0
...     value = 1
...     size = 2
...
>>> FieldTypes.value.size
<FieldTypes.size: 2>
>>> FieldTypes.size.value
2
```

在 3.5 版更改.

Boolean value of Enum classes and members

Enum members that are mixed with non-*Enum* types (such as *int*, *str*, etc.) are evaluated according to the mixed-in type's rules; otherwise, all members evaluate as *True*. To make your own Enum's boolean evaluation depend on the member's value add the following to your class:

```
def __bool__(self):
    return bool(self.value)
```

Enum classes always evaluate as *True*.

Enum classes with methods

If you give your *Enum* subclass extra methods, like the *Planet* class above, those methods will show up in a `dir()` of the member, but not of the class:

```
>>> dir(Planet)
['EARTH', 'JUPITER', 'MARS', 'MERCURY', 'NEPTUNE', 'SATURN', 'URANUS', 'VENUS', '__class__', '__doc__', '__members__', '__module__']
>>> dir(Planet.EARTH)
['__class__', '__doc__', '__module__', 'name', 'surface_gravity', 'value']
```

Combining members of Flag

If a combination of Flag members is not named, the `repr()` will include all named flags and all named combinations of flags that are in the value:

```
>>> class Color(Flag):
...     RED = auto()
...     GREEN = auto()
...     BLUE = auto()
...     MAGENTA = RED | BLUE
...     YELLOW = RED | GREEN
...     CYAN = GREEN | BLUE
...
>>> Color(3) # named combination
<Color.YELLOW: 3>
>>> Color(7) # not named combination
<Color.CYAN|MAGENTA|BLUE|YELLOW|GREEN|RED: 7>
```

数字和数学模块

本章介绍的模块提供与数字和数学相关的函数和数据类型。*numbers* 模块定义了数字类型的抽象层次结构。*math* 和 *cmath* 模块包含浮点数和复数的各种数学函数。*decimal* 模块支持使用任意精度算术的十进制数的精确表示。

本章记录以下模块：

9.1 *numbers* — 数字的抽象基类

源代码： [Lib/numbers.py](#)

numbers 模块 (PEP 3141) 定义了数字抽象基类的层次结构，其中逐级定义了更多操作。此模块中所定义的类型都不可被实例化。

class *numbers.Number*

数字的层次结构的基础。如果你只想确认参数 *x* 是不是数字而不关心其类型，则使用 “*isinstance(x, Number)*”。

9.1.1 数字的层次

class *numbers.Complex*

Subclasses of this type describe complex numbers and include the operations that work on the built-in *complex* type. These are: conversions to *complex* and *bool*, *real*, *imag*, *+*, *-*, ***, */*, *abs()*, *conjugate()*, *==*, and *!=*. All except *-* and *!=* are abstract.

real

Abstract. Retrieves the real component of this number.

imag

Abstract. Retrieves the imaginary component of this number.

abstractmethod conjugate()

Abstract. Returns the complex conjugate. For example, `(1+3j).conjugate() == (1-3j)`.

class numbers.Real

To *Complex*, *Real* adds the operations that work on real numbers.

In short, those are: a conversion to *float*, *math.trunc()*, *round()*, *math.floor()*, *math.ceil()*, *divmod()*, *//*, *%*, *<*, *<=*, *>*, and *>=*.

Real also provides defaults for *complex()*, *real*, *imag*, and *conjugate()*.

class numbers.Rational

Subtypes *Real* and adds *numerator* and *denominator* properties, which should be in lowest terms. With these, it provides a default for *float()*.

numerator

Abstract.

denominator

Abstract.

class numbers.Integral

Subtypes *Rational* and adds a conversion to *int*. Provides defaults for *float()*, *numerator*, and *denominator*. Adds abstract methods for **** and bit-string operations: *<<*, *>>*, *&*, *^*, *|*, *~*.

9.1.2 Notes for type implementors

Implementors should be careful to make equal numbers equal and hash them to the same values. This may be subtle if there are two different extensions of the real numbers. For example, *fractions.Fraction* implements *hash()* as follows:

```
def __hash__(self):
    if self.denominator == 1:
        # Get integers right.
        return hash(self.numerator)
    # Expensive check, but definitely correct.
    if self == float(self):
        return hash(float(self))
    else:
        # Use tuple's hash to avoid a high collision rate on
        # simple fractions.
        return hash((self.numerator, self.denominator))
```

Adding More Numeric ABCs

There are, of course, more possible ABCs for numbers, and this would be a poor hierarchy if it precluded the possibility of adding those. You can add *MyFoo* between *Complex* and *Real* with:

```
class MyFoo(Complex): ...
MyFoo.register(Real)
```

Implementing the arithmetic operations

We want to implement the arithmetic operations so that mixed-mode operations either call an implementation whose author knew about the types of both arguments, or convert both to the nearest built in type and do the operation there. For subtypes of *Integral*, this means that *__add__()* and *__radd__()* should be defined as:

```

class MyIntegral(Integral):

    def __add__(self, other):
        if isinstance(other, MyIntegral):
            return do_my_adding_stuff(self, other)
        elif isinstance(other, OtherTypeIKnowAbout):
            return do_my_other_adding_stuff(self, other)
        else:
            return NotImplemented

    def __radd__(self, other):
        if isinstance(other, MyIntegral):
            return do_my_adding_stuff(other, self)
        elif isinstance(other, OtherTypeIKnowAbout):
            return do_my_other_adding_stuff(other, self)
        elif isinstance(other, Integral):
            return int(other) + int(self)
        elif isinstance(other, Real):
            return float(other) + float(self)
        elif isinstance(other, Complex):
            return complex(other) + complex(self)
        else:
            return NotImplemented

```

There are 5 different cases for a mixed-type operation on subclasses of `Complex`. I'll refer to all of the above code that doesn't refer to `MyIntegral` and `OtherTypeIKnowAbout` as “boilerplate”. `a` will be an instance of `A`, which is a subtype of `Complex` (`a : A <: Complex`), and `b : B <: Complex`. I'll consider `a + b`:

1. If `A` defines an `__add__()` which accepts `b`, all is well.
2. If `A` falls back to the boilerplate code, and it were to return a value from `__add__()`, we'd miss the possibility that `B` defines a more intelligent `__radd__()`, so the boilerplate should return `NotImplemented` from `__add__()`. (Or `A` may not implement `__add__()` at all.)
3. Then `B`'s `__radd__()` gets a chance. If it accepts `a`, all is well.
4. If it falls back to the boilerplate, there are no more possible methods to try, so this is where the default implementation should live.
5. If `B <: A`, Python tries `B.__radd__` before `A.__add__`. This is ok, because it was implemented with knowledge of `A`, so it can handle those instances before delegating to `Complex`.

If `A <: Complex` and `B <: Real` without sharing any other knowledge, then the appropriate shared operation is the one involving the built in `complex`, and both `__radd__()`s land there, so `a+b == b+a`.

Because most of the operations on any given type will be very similar, it can be useful to define a helper function which generates the forward and reverse instances of any given operator. For example, `fractions.Fraction` uses:

```

def _operator_fallbacks(monomorphic_operator, fallback_operator):
    def forward(a, b):
        if isinstance(b, (int, Fraction)):
            return monomorphic_operator(a, b)
        elif isinstance(b, float):
            return fallback_operator(float(a), b)
        elif isinstance(b, complex):
            return fallback_operator(complex(a), b)
        else:
            return NotImplemented
    forward.__name__ = '__' + fallback_operator.__name__ + '__'

```

(下页继续)

(续上页)

```

forward.__doc__ = monomorphic_operator.__doc__

def reverse(b, a):
    if isinstance(a, Rational):
        # Includes ints.
        return monomorphic_operator(a, b)
    elif isinstance(a, numbers.Real):
        return fallback_operator(float(a), float(b))
    elif isinstance(a, numbers.Complex):
        return fallback_operator(complex(a), complex(b))
    else:
        return NotImplemented
reverse.__name__ = '__r' + fallback_operator.__name__ + '__'
reverse.__doc__ = monomorphic_operator.__doc__

return forward, reverse

def _add(a, b):
    """a + b"""
    return Fraction(a.numerator * b.denominator +
                    b.numerator * a.denominator,
                    a.denominator * b.denominator)

__add__, __radd__ = _operator_fallbacks(_add, operator.add)

# ...

```

9.2 math — 数学函数

This module is always available. It provides access to the mathematical functions defined by the C standard.

这些函数不适用于复数；如果你需要计算复数，请使用 `cmath` 模块中的同名函数。将支持计算复数的函数区分开的目的，来自于大多数开发者并不愿意像数学家一样需要学习复数的概念。得到一个异常而不是一个复数结果使得开发者能够更早地监测到传递给这些函数的参数中包含复数，进而调查其产生的原因。

该模块提供了以下函数。除非另有明确说明，否则所有返回值均为浮点数。

9.2.1 数论与表示函数

`math.ceil(x)`

返回 x 的上限，即大于或者等于 x 的最小整数。如果 x 不是一个浮点数，则委托 `x.__ceil__()`，返回一个 *Integral* 类的值。

`math.copysign(x, y)`

返回一个基于 x 的绝对值和 y 的符号的浮点数。在支持带符号零的平台上，`copysign(1.0, -0.0)` 返回 `-1.0`。

`math.fabs(x)`

返回 x 的绝对值。

`math.factorial(x)`

Return x factorial. Raises *ValueError* if x is not integral or is negative.

`math.floor(x)`

返回 x 的向下取整，小于或等于 x 的最大整数。如果 x 不是浮点数，则委托 `x.__floor__()`，它应返回 *Integral* 值。

`math.fmod(x, y)`

返回 `fmod(x, y)`，由平台 C 库定义。请注意，Python 表达式 `x % y` 可能不会返回相同的结果。C 标准的目的是 `fmod(x, y)` 完全（数学上；到无限精度）等于 $x - n*y$ 对于某个整数 n ，使得结果具有与 x 相同的符号和小于 `abs(y)` 的幅度。Python 的 `x % y` 返回带有 y 符号的结果，并且可能不能完全计算浮点参数。例如，`fmod(-1e-100, 1e100)` 是 $-1e-100$ ，但 Python 的 `-1e-100 % 1e100` 的结果是 $1e100-1e-100$ ，它不能完全表示为浮点数，并且取整为令人惊讶的 $1e100$ 。出于这个原因，函数 `fmod()` 在使用浮点数时通常是首选，而 Python 的 `x % y` 在使用整数时是首选。

`math.frexp(x)`

返回 x 的尾数和指数作为对“(m, e)”。 m 是一个浮点数， e 是一个整数，正好是 $x == m * 2**e$ 。如果 x 为零，则返回 $(0.0, 0)$ ，否则返回 $0.5 <= \text{abs}(m) < 1$ 。这用于以可移植方式“分离”浮点数的内部表示。

`math.fsum(iterable)`

返回迭代中的精确浮点值。通过跟踪多个中间部分和来避免精度损失：

```
>>> sum([.1, .1, .1, .1, .1, .1, .1, .1, .1, .1])
0.9999999999999999
>>> fsum([.1, .1, .1, .1, .1, .1, .1, .1, .1, .1])
1.0
```

该算法的准确性取决于 IEEE-754 算术保证和舍入模式为半偶的典型情况。在某些非 Windows 版本中，底层 C 库使用扩展精度添加，并且有时可能会使中间和加倍，导致它在最低有效位中关闭。

有关待进一步讨论和两种替代方法，参见 [ASPN cookbook recipes for accurate floating point summation](#)。

`math.gcd(a, b)`

返回整数 a 和 b 的最大公约数。如果 a 或 b 之一非零，则 `gcd(a, b)` 的值是能同时整除 a 和 b 的最大正整数。`gcd(0, 0)` 返回 0。

3.5 新版功能。

`math.isclose(a, b, *, rel_tol=1e-09, abs_tol=0.0)`

若 a 和 b 的值比较接近则返回 True，否则返回 False。

根据给定的绝对和相对容差确定两个值是否被认为是接近的。

`rel_tol` 是相对容差——它是 a 和 b 之间允许的最大差值，相对于 a 或 b 的较大绝对值。例如，要设置 5% 的容差，请传递 `rel_tol=0.05`。默认容差为 $1e-09$ ，确保两个值在大约 9 位十进制数字内相同。`rel_tol` 必须大于零。

`abs_tol` 是最小绝对容差——对于接近零的比较很有用。`abs_tol` 必须至少为零。

如果没有错误发生，结果将是：`abs(a-b) <= max(rel_tol * max(abs(a), abs(b)), abs_tol)`。

IEEE 754 特殊值 NaN，`inf` 和 `-inf` 将根据 IEEE 规则处理。具体来说，NaN 不被认为接近任何其他值，包括 NaN。`inf` 和 `-inf` 只被认为接近自己。

3.5 新版功能。

参见：

[PEP 485](#) ——用于测试近似相等的函数

`math.isfinite(x)`

如果 x 既不是无穷大也不是 NaN，则返回 True，否则返回 False。（注意 0.0 被认为是有限的。）

3.2 新版功能。

`math.isinf(x)`

如果 x 是正或负无穷大，则返回 `True`，否则返回 `False`。

`math.isnan(x)`

如果 x 是 NaN (不是数字)，则返回 `True`，否则返回 `False`。

`math.ldexp(x, i)`

返回 $x * (2^{**i})$ 。这基本上是函数 `frexp()` 的反函数。

`math.modf(x)`

返回 x 的小数和整数部分。两个结果都带有 x 的符号并且是浮点数。

`math.trunc(x)`

返回 *Real* 值 x 截断为 *Integral* (通常是整数)。委托给 `x.__trunc__()`。

注意 `frexp()` 和 `modf()` 具有与它们的 C 等价函数不同的调用/返回模式：它们采用单个参数并返回一对值，而不是通过‘输出形参’返回它们的第二个返回参数 (Python 中没有这样的东西)。

对于 `ceil()`，`floor()` 和 `modf()` 函数，请注意 所有足够大的浮点数都是精确整数。Python 浮点数通常不超过 53 位的精度 (与平台 C `double` 类型相同)，在这种情况下，任何浮点 x 与 `abs(x) >= 2**52` 必然没有小数位。

9.2.2 幂函数与对数函数

`math.exp(x)`

Return e^{**x} 。

`math.expm1(x)`

Return $e^{**x} - 1$ 。For small floats x , the subtraction in `exp(x) - 1` can result in a significant loss of precision; the `expm1()` function provides a way to compute this quantity to full precision:

```
>>> from math import exp, expm1
>>> exp(1e-5) - 1 # gives result accurate to 11 places
1.0000050000069649e-05
>>> expm1(1e-5) # result accurate to full precision
1.0000050000166668e-05
```

3.2 新版功能.

`math.log(x[, base])`

使用一个参数，返回 x 的自然对数 (底为 e)。

使用两个参数，返回给定的 *base* 的对数 x ，计算为 `log(x) / log(base)`。

`math.log1p(x)`

返回 $1+x$ (base e) 的自然对数。以对于接近零的 x 精确的方式计算结果。

`math.log2(x)`

返回 x 以 2 为底的对数。这通常比 `log(x, 2)` 更准确。

3.3 新版功能.

参见:

`int.bit_length()` 返回表示二进制整数所需的位数，不包括符号和前导零。

`math.log10(x)`

返回 x 底为 10 的对数。这通常比 `log(x, 10)` 更准确。

`math.pow(x, y)`

将返回 x 的 y 次幂。特殊情况尽可能遵循 C99 标准的附录‘F’。特别是，`pow(1.0, x)` 和 `pow(x,`

0.0) 总是返回 1.0，即使 x 是零或 NaN。如果 x 和 y 都是有限的， x 是负数， y 不是整数那么 $\text{pow}(x, y)$ 是未定义的，并且引发 `ValueError`。

与内置的 `**` 运算符不同，`math.pow()` 将其参数转换为 `float` 类型。使用 `**` 或内置的 `pow()` 函数来计算精确的整数幂。

`math.sqrt(x)`
返回 x 的平方根。

9.2.3 三角函数

`math.acos(x)`
以弧度为单位返回 x 的反余弦值。

`math.asin(x)`
以弧度为单位返回 x 的反正弦值。

`math.atan(x)`
以弧度为单位返回 x 的反正切值。

`math.atan2(y, x)`
以弧度为单位返回 $\text{atan}(y / x)$ 。结果是在 $-\pi$ 和 π 之间。从原点到点 (x, y) 的平面矢量使该角度与正 X 轴成正比。`atan2()` 的点的两个输入的符号都是已知的，因此它可以计算角度的正确象限。例如， $\text{atan}(1)$ 和 $\text{atan2}(1, 1)$ 都是 $\pi/4$ ，但 $\text{atan2}(-1, -1)$ 是 $-3\pi/4$ 。

`math.cos(x)`
返回 x 弧度的余弦值。

`math.hypot(x, y)`
返回欧几里德范数， $\sqrt{x^2 + y^2}$ 。这是从原点到点 (x, y) 的向量长度。

`math.sin(x)`
返回 x 弧度的正弦值。

`math.tan(x)`
返回 x 弧度的正切值。

9.2.4 角度转换

`math.degrees(x)`
将角度 x 从弧度转换为度数。

`math.radians(x)`
将角度 x 从度数转换为弧度。

9.2.5 双曲函数

双曲函数是基于双曲线而非圆来对三角函数进行模拟。

`math.acosh(x)`
返回 x 的反双曲余弦值。

`math.asinh(x)`
返回 x 的反双曲正弦值。

`math.atanh(x)`
返回 x 的反双曲正切值。

`math.cosh(x)`
返回 x 的双曲余弦值。

`math.sinh(x)`
返回 x 的双曲正弦值。

`math.tanh(x)`
返回 x 的双曲正切值。

9.2.6 特殊函数

`math.erf(x)`
返回 x 处的 `error function` 。

`erf()` 函数可用于计算传统的统计函数，如 累积标准正态分布

```
def phi(x):  
    'Cumulative distribution function for the standard normal distribution'  
    return (1.0 + erf(x / sqrt(2.0))) / 2.0
```

3.2 新版功能.

`math.erfc(x)`
返回 x 处的互补误差函数。互补误差函数 定义为 $1.0 - \text{erf}(x)$ 。它用于 x 的大值，从其中减去一个会导致 有效位数损失。

3.2 新版功能.

`math.gamma(x)`
返回 x 处的 伽马函数 值。

3.2 新版功能.

`math.lgamma(x)`
返回 Gamma 函数在 x 绝对值的自然对数。

3.2 新版功能.

9.2.7 常数

`math.pi`
The mathematical constant $\pi = 3.141592\dots$, to available precision.

`math.e`
The mathematical constant $e = 2.718281\dots$, to available precision.

`math.tau`
The mathematical constant $\tau = 6.283185\dots$, to available precision. Tau is a circle constant equal to 2π , the ratio of a circle's circumference to its radius. To learn more about Tau, check out Vi Hart's video [Pi is \(still\) Wrong](#), and start celebrating [Tau day](#) by eating twice as much pie!

3.6 新版功能.

`math.inf`
浮点正无穷大。(对于负无穷大，使用 `-math.inf`。) 相当于 “`float('inf')`” 的输出。

3.5 新版功能.

`math.nan`

浮点“非数字”(NaN)值。相当于 `float('nan')` 的输出。

3.5 新版功能.

CPython implementation detail: `math` 模块主要包含围绕平台 C 数学库函数的简单包装器。特殊情况下的行为在适当情况下遵循 C99 标准的附录 F。当前的实现将引发 `ValueError` 用于无效操作, 如 `sqrt(-1.0)` 或 `log(0.0)` (其中 C99 附件 F 建议发出无效操作信号或被零除), 和 `OverflowError` 用于溢出的结果 (例如, `exp(1000.0)`)。除非一个或多个输入参数是 NaN, 否则不会从上述任何函数返回 NaN; 在这种情况下, 大多数函数将返回一个 NaN, 但是 (再次遵循 C99 附件 F) 这个规则有一些例外, 例如 `pow(float('nan'), 0.0)` 或 `hypot(float('nan'), float('inf'))`。

请注意, Python 不会将显式 NaN 与静默 NaN 区分开来, 并且显式 NaN 的行为仍未明确。典型的行为是将所有 NaN 视为静默的。

参见:

`cmath` 模块 这里很多函数的复数版本。

9.3 cmath —— 关于复数的数学函数

This module is always available. It provides access to mathematical functions for complex numbers. The functions in this module accept integers, floating-point numbers or complex numbers as arguments. They will also accept any Python object that has either a `__complex__()` or a `__float__()` method: these methods are used to convert the object to a complex or floating-point number, respectively, and the function is then applied to the result of the conversion.

注解: On platforms with hardware and system-level support for signed zeros, functions involving branch cuts are continuous on *both* sides of the branch cut: the sign of the zero distinguishes one side of the branch cut from the other. On platforms that do not support signed zeros the continuity is as specified below.

9.3.1 到极坐标和从极坐标的转换

A Python complex number `z` is stored internally using *rectangular* or *Cartesian* coordinates. It is completely determined by its *real part* `z.real` and its *imaginary part* `z.imag`. In other words:

```
z == z.real + z.imag*1j
```

Polar coordinates give an alternative way to represent a complex number. In polar coordinates, a complex number `z` is defined by the modulus `r` and the phase angle `phi`. The modulus `r` is the distance from `z` to the origin, while the phase `phi` is the counterclockwise angle, measured in radians, from the positive x-axis to the line segment that joins the origin to `z`.

下面的函数可用于原生直角坐标与极坐标的相互转换。

`cmath.phase(x)`

Return the phase of `x` (also known as the *argument* of `x`), as a float. `phase(x)` is equivalent to `math.atan2(x.imag, x.real)`. The result lies in the range $[-\pi, \pi]$, and the branch cut for this operation lies along the negative real axis, continuous from above. On systems with support for signed zeros (which includes most systems in current use), this means that the sign of the result is the same as the sign of `x.imag`, even when `x.imag` is zero:

```
>>> phase(complex(-1.0, 0.0))
3.141592653589793
>>> phase(complex(-1.0, -0.0))
-3.141592653589793
```

注解: The modulus (absolute value) of a complex number x can be computed using the built-in `abs()` function. There is no separate `cmath` module function for this operation.

`cmath.polar(x)`

Return the representation of x in polar coordinates. Returns a pair (r, phi) where r is the modulus of x and phi is the phase of x . `polar(x)` is equivalent to `(abs(x), phase(x))`.

`cmath.rect(r, phi)`

Return the complex number x with polar coordinates r and phi . Equivalent to $r * (\text{math.cos(phi)} + \text{math.sin(phi)} * 1j)$.

9.3.2 幂函数与对数函数

`cmath.exp(x)`

Return the exponential value e^{**x} .

`cmath.log(x[, base])`

Returns the logarithm of x to the given *base*. If the *base* is not specified, returns the natural logarithm of x . There is one branch cut, from 0 along the negative real axis to $-\infty$, continuous from above.

`cmath.log10(x)`

Return the base-10 logarithm of x . This has the same branch cut as `log()`.

`cmath.sqrt(x)`

Return the square root of x . This has the same branch cut as `log()`.

9.3.3 三角函数

`cmath.acos(x)`

Return the arc cosine of x . There are two branch cuts: One extends right from 1 along the real axis to ∞ , continuous from below. The other extends left from -1 along the real axis to $-\infty$, continuous from above.

`cmath.asin(x)`

Return the arc sine of x . This has the same branch cuts as `acos()`.

`cmath.atan(x)`

Return the arc tangent of x . There are two branch cuts: One extends from $1j$ along the imaginary axis to ∞j , continuous from the right. The other extends from $-1j$ along the imaginary axis to $-\infty j$, continuous from the left.

`cmath.cos(x)`

返回 $*x*$ 的余弦。

`cmath.sin(x)`

返回 $*x*$ 的正弦。

`cmath.tan(x)`

返回 $*x*$ 的正切。

9.3.4 双曲函数

`cmath.acosh(x)`

Return the inverse hyperbolic cosine of x . There is one branch cut, extending left from 1 along the real axis to $-\infty$, continuous from above.

`cmath.asinh(x)`

Return the inverse hyperbolic sine of x . There are two branch cuts: One extends from $1j$ along the imaginary axis to ∞j , continuous from the right. The other extends from $-1j$ along the imaginary axis to $-\infty j$, continuous from the left.

`cmath.atanh(x)`

Return the inverse hyperbolic tangent of x . There are two branch cuts: One extends from 1 along the real axis to ∞ , continuous from below. The other extends from -1 along the real axis to $-\infty$, continuous from above.

`cmath.cosh(x)`

返回 x 的双曲余弦值。

`cmath.sinh(x)`

返回 x 的双曲正弦值。

`cmath.tanh(x)`

返回 x 的双曲正切值。

9.3.5 Classification functions

`cmath.isfinite(x)`

Return `True` if both the real and imaginary parts of x are finite, and `False` otherwise.

3.2 新版功能。

`cmath.isinf(x)`

Return `True` if either the real or the imaginary part of x is an infinity, and `False` otherwise.

`cmath.isnan(x)`

Return `True` if either the real or the imaginary part of x is a NaN, and `False` otherwise.

`cmath.isclose(a, b, *, rel_tol=1e-09, abs_tol=0.0)`

若 a 和 b 的值比较接近则返回 `True`，否则返回 `False`。

根据给定的绝对和相对容差确定两个值是否被认为是接近的。

`rel_tol` 是相对容差——它是 a 和 b 之间允许的最大差值，相对于 a 或 b 的较大绝对值。例如，要设置 5% 的容差，请传递 `rel_tol=0.05`。默认容差为 `1e-09`，确保两个值在大约 9 位十进制数字内相同。`rel_tol` 必须大于零。

`abs_tol` 是最小绝对容差——对于接近零的比较很有用。`abs_tol` 必须至少为零。

如果没有错误发生，结果将是：`abs(a-b) <= max(rel_tol * max(abs(a), abs(b)), abs_tol)`。

IEEE 754 特殊值 NaN，`inf` 和 `-inf` 将根据 IEEE 规则处理。具体来说，NaN 不被认为接近任何其他值，包括 NaN。`inf` 和 `-inf` 只被认为接近自己。

3.5 新版功能。

参见：

[PEP 485](#) ——用于测试近似相等的函数

9.3.6 常数

`cmath.pi`

The mathematical constant π , as a float.

`cmath.e`

The mathematical constant e , as a float.

`cmath.tau`

The mathematical constant τ , as a float.

3.6 新版功能.

`cmath.inf`

Floating-point positive infinity. Equivalent to `float('inf')`.

3.6 新版功能.

`cmath.infj`

Complex number with zero real part and positive infinity imaginary part. Equivalent to `complex(0.0, float('inf'))`.

3.6 新版功能.

`cmath.nan`

A floating-point “not a number” (NaN) value. Equivalent to `float('nan')`.

3.6 新版功能.

`cmath.nanj`

Complex number with zero real part and NaN imaginary part. Equivalent to `complex(0.0, float('nan'))`.

3.6 新版功能.

Note that the selection of functions is similar, but not identical, to that in module `math`. The reason for having two modules is that some users aren't interested in complex numbers, and perhaps don't even know what they are. They would rather have `math.sqrt(-1)` raise an exception than return a complex number. Also note that the functions defined in `cmath` always return a complex number, even if the answer can be expressed as a real number (in which case the complex number has an imaginary part of zero).

A note on branch cuts: They are curves along which the given function fails to be continuous. They are a necessary feature of many complex functions. It is assumed that if you need to compute with complex functions, you will understand about branch cuts. Consult almost any (not too elementary) book on complex variables for enlightenment. For information of the proper choice of branch cuts for numerical purposes, a good reference should be the following:

参见:

Kahan, W: Branch cuts for complex elementary functions; or, Much ado about nothing's sign bit. In Iserles, A., and Powell, M. (eds.), The state of the art in numerical analysis. Clarendon Press (1987) pp165–211.

9.4 decimal — 十进制定点和浮点运算

源码: [Lib/decimal.py](#)

`decimal` 模块为快速正确舍入的十进制浮点运算提供支持。它提供了 `float` 数据类型以外的几个优点:

- Decimal “基于一个浮点模型，它是为人们设计的，并且必然具有最重要的指导原则——计算机必须提供与人们在学校学习的算法相同的算法。”——摘自十进制算术规范。

- 十进制数字可以准确表示。相比之下，数字如 1.1 和 2.2 在二进制浮点中没有精确的表示。最终用户通常不希望“1.1 + 2.2”显示为 3.3000000000000003，就像二进制浮点一样。
- 精确性延续到算术中。在十进制浮点数中， $0.1 + 0.1 + 0.1 - 0.3$ 恰好等于零。在二进制浮点数中，结果为 $5.5511151231257827\text{e-}017$ 。虽然接近于零，但差异妨碍了可靠的相等性检验，并且差异可能会累积。因此，在具有严格相等不变量的会计应用程序中，`decimal` 是首选。
- 十进制模块包含一个重要位置的概念，因此 1.30 + 1.20 是 2.50。保留尾随零以表示重要性。这是货币申请的惯常陈述。对于乘法，“教科书”方法使用被乘数中的所有数字。例如， $1.3 * 1.2$ 给出 1.56 而 $1.30 * 1.20$ 给出 1.5600。
- 与基于硬件的二进制浮点不同，十进制模块具有用户可更改的精度（默认为 28 个位置），可以与给定问题所需的一样大：

```
>>> from decimal import *
>>> getcontext().prec = 6
>>> Decimal(1) / Decimal(7)
Decimal('0.142857')
>>> getcontext().prec = 28
>>> Decimal(1) / Decimal(7)
Decimal('0.1428571428571428571428571428571429')
```

- 二进制和十进制浮点都是根据已发布的标准实现的。虽然内置浮点类型只公开其功能的一小部分，但十进制模块公开了标准的所有必需部分。在需要时，程序员可以完全控制舍入和信号处理。这包括通过使用异常来阻止任何不精确操作来强制执行精确算术的选项。
- 十进制模块旨在支持“无偏见，精确的非连续十进制算术（有时称为定点算术）和舍入浮点算术”。——摘自十进制算术规范。

模块设计以三个概念为中心：十进制数，算术上下文和信号。

十进制数是不可变的。它有一个符号，系数数字和一个指数。为了保持重要性，系数数字不会截断尾随零。十进制数也包括特殊值，例如 `Infinity`，`-Infinity`，和 `NaN`。该标准还区分 `-0` 和 `+0`。

算术的上下文是指定精度、舍入规则、指数限制、指示操作结果的标志以及确定符号是否被视为异常的陷阱启用器的环境。舍入选项包括 `ROUND_CEILING`、`ROUND_DOWN`、`ROUND_FLOOR`、`ROUND_HALF_DOWN`、`ROUND_HALF_EVEN`、`ROUND_HALF_UP`、`ROUND_UP` 以及 `ROUND_05UP`。

信号是在计算过程中出现的异常条件组。根据应用程序的需要，信号可能会被忽略，被视为信息，或被视为异常。十进制模块中的信号有：`Clamped`、`InvalidOperation`、`DivisionByZero`、`Inexact`、`Rounded`、`Subnormal`、`Overflow`、`Underflow` 以及 `FloatOperation`。

对于每个信号，都有一个标志和一个陷阱启动器。遇到信号时，其标志设置为 1，然后，如果陷阱启用器设置为 1，则引发异常。标志是粘性的，因此用户需要在监控计算之前重置它们。

参见：

- IBM 的通用十进制算术规范，[The General Decimal Arithmetic Specification](#).

9.4.1 快速入门教程

通常使用小数的开始是导入模块，使用 `getcontext()` 查看当前上下文，并在必要时为精度、舍入或启用的陷阱设置新值：

```
>>> from decimal import *
>>> getcontext()
Context(prec=28, rounding=ROUND_HALF_EVEN, Emin=-999999, Emax=999999,
        capitals=1, clamp=0, flags=[], traps=[Overflow, DivisionByZero,
        InvalidOperation])
```

(下页继续)

(续上页)

```
>>> getcontext().prec = 7          # Set a new precision
```

可以从整数、字符串、浮点数或元组构造十进制实例。从整数或浮点构造将执行该整数或浮点值的精确转换。十进制数包括特殊值，例如 NaN 代表“非数字”，正的和负的 Infinity，和 -0

```
>>> getcontext().prec = 28
>>> Decimal(10)
Decimal('10')
>>> Decimal('3.14')
Decimal('3.14')
>>> Decimal(3.14)
Decimal('3.1400000000000000124344978758017532527446746826171875')
>>> Decimal((0, (3, 1, 4), -2))
Decimal('3.14')
>>> Decimal(str(2.0 ** 0.5))
Decimal('1.4142135623730951')
>>> Decimal(2) ** Decimal('0.5')
Decimal('1.414213562373095048801688724')
>>> Decimal('NaN')
Decimal('NaN')
>>> Decimal('-Infinity')
Decimal('-Infinity')
```

如果`FloatOperation`信号被捕获，构造函数中的小数和浮点数的意外混合或排序比较会引发异常

```
>>> c = getcontext()
>>> c.traps[FloatOperation] = True
>>> Decimal(3.14)
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
decimal.FloatOperation: [
>>> Decimal('3.5') < 3.7
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
decimal.FloatOperation: [
>>> Decimal('3.5') == 3.5
True
```

3.3 新版功能.

新 **Decimal** 的重要性仅由输入的位数决定。上下文精度和舍入仅在算术运算期间发挥作用。

```
>>> getcontext().prec = 6
>>> Decimal('3.0')
Decimal('3.0')
>>> Decimal('3.1415926535')
Decimal('3.1415926535')
>>> Decimal('3.1415926535') + Decimal('2.7182818285')
Decimal('5.85987')
>>> getcontext().rounding = ROUND_UP
>>> Decimal('3.1415926535') + Decimal('2.7182818285')
Decimal('5.85988')
```

如果超出了 C 版本的内部限制，则构造一个十进制将引发 *InvalidOperation*

```
>>> Decimal("1e999999999999999999")
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
decimal.InvalidOperation: [
```

在 3.3 版更改.

小数与 Python 的其余部分很好地交互。这是一个小的十进制浮点飞行杂技团：

```
>>> data = list(map(Decimal, '1.34 1.87 3.45 2.35 1.00 0.03 9.25'.split()))
>>> max(data)
Decimal('9.25')
>>> min(data)
Decimal('0.03')
>>> sorted(data)
[Decimal('0.03'), Decimal('1.00'), Decimal('1.34'), Decimal('1.87'),
 Decimal('2.35'), Decimal('3.45'), Decimal('9.25')]
>>> sum(data)
Decimal('19.29')
>>> a,b,c = data[:3]
>>> str(a)
'1.34'
>>> float(a)
1.34
>>> round(a, 1)
Decimal('1.3')
>>> int(a)
1
>>> a * 5
Decimal('6.70')
>>> a * b
Decimal('2.5058')
>>> c % a
Decimal('0.77')
```

Decimal 也可以使用一些数学函数：

```
>>> getcontext().prec = 28
>>> Decimal(2).sqrt()
Decimal('1.414213562373095048801688724')
>>> Decimal(1).exp()
Decimal('2.718281828459045235360287471')
>>> Decimal('10').ln()
Decimal('2.302585092994045684017991455')
>>> Decimal('10').log10()
Decimal('1')
```

`quantize()` 方法将数字四舍五入为固定指数。此方法对于将结果舍入到固定的位置的货币应用程序非常有用：

```
>>> Decimal('7.325').quantize(Decimal('.01'), rounding=ROUND_DOWN)
Decimal('7.32')
>>> Decimal('7.325').quantize(Decimal('1.'), rounding=ROUND_UP)
Decimal('8')
```

如上所示，`getcontext()` 函数访问当前上下文并允许更改设置。这种方法满足大多数应用程序的需求。

对于更高级的工作，使用 `Context()` 构造函数创建备用上下文可能很有用。要使用备用活动，请使用 `setcontext()` 函数。

根据标准，`decimal` 模块提供了两个现成的标准上下文 `BasicContext` 和 `ExtendedContext`。前者对调试特别有用，因为许多陷阱都已启用：

```
>>> myothercontext = Context(prec=60, rounding=ROUND_HALF_DOWN)
>>> setcontext(myothercontext)
>>> Decimal(1) / Decimal(7)
Decimal('0.142857142857142857142857142857142857142857142857142857')

>>> ExtendedContext
Context(prec=9, rounding=ROUND_HALF_EVEN, Emin=-999999, Emax=999999,
        capitals=1, clamp=0, flags=[], traps=[])
>>> setcontext(ExtendedContext)
>>> Decimal(1) / Decimal(7)
Decimal('0.142857143')
>>> Decimal(42) / Decimal(0)
Decimal('Infinity')

>>> setcontext(BasicContext)
>>> Decimal(42) / Decimal(0)
Traceback (most recent call last):
  File "<pyshell#143>", line 1, in -toplevel-
    Decimal(42) / Decimal(0)
DivisionByZero: x / 0
```

上下文还具有用于监视计算期间遇到的异常情况的信号标志。标志保持设置直到明确清除，因此最好通过使用 `clear_flags()` 方法清除每组受监控计算之前的标志。：

```
>>> setcontext(ExtendedContext)
>>> getcontext().clear_flags()
>>> Decimal(355) / Decimal(113)
Decimal('3.14159292')
>>> getcontext()
Context(prec=9, rounding=ROUND_HALF_EVEN, Emin=-999999, Emax=999999,
        capitals=1, clamp=0, flags=[Inexact, Rounded], traps=[])
```

`flags` 条目显示对 π 的有理逼近被舍入（超出上下文精度的数字被抛弃）并且结果是不精确的（一些丢弃的数字不为零）。

使用上下文的 `traps` 字段中的字典设置单个陷阱：

```
>>> setcontext(ExtendedContext)
>>> Decimal(1) / Decimal(0)
Decimal('Infinity')
>>> getcontext().traps[DivisionByZero] = 1
>>> Decimal(1) / Decimal(0)
Traceback (most recent call last):
  File "<pyshell#112>", line 1, in -toplevel-
    Decimal(1) / Decimal(0)
DivisionByZero: x / 0
```

大多数程序仅在程序开始时调整当前上下文一次。并且，在许多应用程序中，数据在循环内单个强制转换为 `Decimal`。通过创建上下文集和小数，程序的大部分操作数据与其他 Python 数字类型没有区别。

9.4.2 Decimal 对象

class `decimal.Decimal` (*value="0", context=None*)

根据 *value* 构造一个新的 `Decimal` 对象。

value 可以是整数，字符串，元组，*float*，或另一个 *Decimal* 对象。如果没有给出 *value*，则返回 `Decimal('0')`。如果 *value* 是一个字符串，它应该在前导和尾随空格字符以及下划线被删除之后符合十进制数字字符串语法：

```
sign      ::= '+' | '-'
digit     ::= '0' | '1' | '2' | '3' | '4' | '5' | '6' | '7' | '8' | '9'
indicator ::= 'e' | 'E'
digits    ::= digit [digit]...
decimal-part ::= digits '.' [digits] | ['.'] digits
exponent-part ::= indicator [sign] digits
infinity  ::= 'Infinity' | 'Inf'
nan       ::= 'NaN' [digits] | 'sNaN' [digits]
numeric-value ::= decimal-part [exponent-part] | infinity
numeric-string ::= [sign] numeric-value | [sign] nan
```

当上面出现 *digit* 时也允许其他十进制数码。其中包括来自各种其他语言系统的十进制数码（例如阿拉伯-印地语和天城文的数码）以及全宽数码 `'\uff10'` 到 `'\uff19'`。

如果 *value* 是一个 *tuple*，它应该有三个组件，一个符号（0 表示正数或 1 表示负数），一个数字的 *tuple* 和整数指数。例如，`Decimal((0, (1, 4, 1, 4), -3))` 返回 `Decimal('1.414')`。

如果 *value* 是 *float*，则二进制浮点值无损地转换为其精确的十进制等效值。此转换通常需要 53 位或更多位数的精度。例如，`Decimal(float('1.1'))` 转换为 `Decimal('1.100000000000000088817841970012523233890533447265625')`。

context 精度不会影响存储的位数。这完全由 *value* 中的位数决定。例如，`Decimal('3.00000')` 记录所有五个零，即使上下文精度只有三。

context 参数的目的是确定 *value* 是格式错误的字符串时该怎么做。如果上下文陷阱 *InvalidOperation*，则引发异常；否则，构造函数返回一个新的 *Decimal*，其值为 NaN。

构造完成后，*Decimal* 对象是不可变的。

在 3.2 版更改：现在允许构造函数的参数为 *float* 实例。

在 3.3 版更改：*float* 参数在设置 *FloatOperation* 陷阱时引发异常。默认情况下，陷阱已关闭。

在 3.6 版更改：允许下划线进行分组，就像代码中的整数和浮点文字一样。

十进制浮点对象与其他内置数值类型共享许多属性，例如 *float* 和 *int*。所有常用的数学运算和特殊方法都适用。同样，十进制对象可以复制、pickle、打印、用作字典键、用作集合元素、比较、排序和强制转换为另一种类型（例如 *float* 或 *int*）。

算术对十进制对象和算术对整数和浮点数有一些小的差别。当余数运算符 `%` 应用于 *Decimal* 对象时，结果的符号是 被除数的符号，而不是除数的符号：

```
>>> (-7) % 4
1
>>> Decimal(-7) % Decimal(4)
Decimal('-3')
```

整数除法运算符 `//` 的行为类似，返回真商的整数部分（截断为零）而不是它的向下取整，以便保留通常的标识 `x == (x // y) * y + x % y`：

```
>>> -7 // 4
-2
>>> Decimal(-7) // Decimal(4)
Decimal('-1')
```

`%` 和 `//` 运算符实现了 *remainder* 和 *divide-integer* 操作（分别），如规范中所述。

十进制对象通常不能与浮点数或 `fractions.Fraction` 实例在算术运算中结合使用：例如，尝试将 `Decimal` 加到 `float`，将引发 `TypeError`。但是，可以使用 Python 的比较运算符来比较 `Decimal` 实例 `x` 和另一个数字 `y`。这样可以避免在对不同类型的数字进行相等比较时混淆结果。

在 3.2 版更改：现在完全支持 `Decimal` 实例和其他数字类型之间的混合类型比较。

除了标准的数字属性，十进制浮点对象还有许多专门的方法：

`adjusted()`

在移出系数最右边的数字之后返回调整后的指数，直到只剩下前导数字：`Decimal('321e+5').adjusted()` 返回 7。用于确定最高有效位相对于小数点的位置。

`as_integer_ratio()`

返回一对 `(n, d)` 整数，表示给定的 `Decimal` 实例作为分数、最简形式项并带有正分母：

```
>>> Decimal('-3.14').as_integer_ratio()
(-157, 50)
```

转换是精确的。在 `Infinity` 上引发 `OverflowError`，在 `NaN` 上引起 `ValueError`。

3.6 新版功能。

`as_tuple()`

返回一个 *named tuple* 表示的数字：`DecimalTuple(sign, digits, exponent)`。

`canonical()`

返回参数的规范编码。目前，一个 `Decimal` 实例的编码始终是规范的，因此该操作返回其参数不变。

`compare(other, context=None)`

比较两个 `Decimal` 实例的值。`compare()` 返回一个 `Decimal` 实例，如果任一操作数是 `NaN`，那么结果是 `NaN`

```
a or b is a NaN ==> Decimal('NaN')
a < b           ==> Decimal('-1')
a == b          ==> Decimal('0')
a > b           ==> Decimal('1')
```

`compare_signal(other, context=None)`

除了所有 `NaN` 信号之外，此操作与 `compare()` 方法相同。也就是说，如果两个操作数都不是信号 `NaN`，那么任何静默的 `NaN` 操作数都被视为信号 `NaN`。

`compare_total(other, context=None)`

使用它们的抽象表示而不是它们的数值来比较两个操作数。类似于 `compare()` 方法，但结果给出了一个总排序 `Decimal` 实例。两个 `Decimal` 实例具有相同的数值但不同的表示形式在此排序中比较不相等：

```
>>> Decimal('12.0').compare_total(Decimal('12'))
Decimal('-1')
```

静默和发出信号的 `NaN` 也包括在总排序中。这个函数的结果是 `Decimal('0')` 如果两个操作数具有相同的表示，或是 `Decimal('-1')` 如果第一个操作数的总顺序低于第二个操作数，或是 `Decimal('1')` 如果第一个操作数在总顺序中高于第二个操作数。有关总排序的详细信息，请参阅规范。

此操作不受上下文影响且静默：不更改任何标志且不执行舍入。作为例外，如果无法准确转换第二个操作数，则 C 版本可能会引发 `InvalidOperation`。

`compare_total_mag(other, context=None)`

比较两个操作数使用它们的抽象表示而不是它们的值，如 `compare_total()`，但忽略每个

操作数的符号。`x.compare_total_mag(y)` 相当于 `x.copy_abs().compare_total(y.copy_abs())`。

此操作不受上下文影响且静默：不更改任何标志且不执行舍入。作为例外，如果无法准确转换第二个操作数，则 C 版本可能会引发 `InvalidOperation`。

conjugate()

只返回 `self`，这种方法只符合 `Decimal` 规范。

copy_abs()

返回参数的绝对值。此操作不受上下文影响并且是静默的：没有更改标志且不执行舍入。

copy_negate()

回到参数的否定。此操作不受上下文影响并且是静默的：没有标志更改且不执行舍入。

copy_sign(other, context=None)

返回第一个操作数的副本，其符号设置为与第二个操作数的符号相同。例如：

```
>>> Decimal('2.3').copy_sign(Decimal('-1.5'))
Decimal('-2.3')
```

此操作不受上下文影响且静默：不更改任何标志且不执行舍入。作为例外，如果无法准确转换第二个操作数，则 C 版本可能会引发 `InvalidOperation`。

exp(context=None)

返回给定数字的（自然）指数函数“ e^x ”的值。结果使用 `ROUND_HALF_EVEN` 舍入模式正确舍入。

```
>>> Decimal(1).exp()
Decimal('2.718281828459045235360287471')
>>> Decimal(321).exp()
Decimal('2.561702493119680037517373933E+139')
```

from_float(f)

将浮点数转换为十进制数的类方法。

注意，`Decimal.from_float(0.1)` 与 `Decimal('0.1')` 不同。由于 0.1 在二进制浮点中不能精确表示，因此该值存储为最接近的可表示值，即 `0x1.999999999999ap-4`。十进制的等效值是 `'0.1000000000000000055511151231257827021181583404541015625'`。

注解：从 Python 3.2 开始，`Decimal` 实例也可以直接从 `float` 构造。

```
>>> Decimal.from_float(0.1)
Decimal('0.1000000000000000055511151231257827021181583404541015625')
>>> Decimal.from_float(float('nan'))
Decimal('NaN')
>>> Decimal.from_float(float('inf'))
Decimal('Infinity')
>>> Decimal.from_float(float('-inf'))
Decimal('-Infinity')
```

3.1 新版功能.

fma(other, third, context=None)

混合乘法加法。返回 `self*other+third`，中间乘积 `self*other` 没有四舍五入。

```
>>> Decimal(2).fma(3, 5)
Decimal('11')
```

is_canonical()

如果参数是规范的，则为返回 *True*，否则为 *False*。目前，*Decimal* 实例总是规范的，所以这个操作总是返回 *True*。

is_finite()

如果参数是一个有限的数，则返回为 *True*；如果参数为无穷大或 NaN，则返回为 *False*。

is_infinite()

如果参数为正负无穷大，则返回为 *True*，否则为 *False*。

is_nan()

如果参数为 NaN（无论是否静默），则返回为 *True*，否则为 *False*。

is_normal(context=None)

如果参数是一个有限正规数，返回 *True*，如果参数是 0、次正规数、无穷大或是 NaN，返回 *False*。

is_qnan()

如果参数为静默 NaN，返回 *True*，否则返回 *False*。

is_signed()

如果参数带有负号，则返回为 *True*，否则返回 *False*。注意，0 和 NaN 都可带有符号。

is_snan()

如果参数为显式 NaN，则返回 *True*，否则返回 *False*。

is_subnormal(context=None)

如果参数为次正规数，则返回 *True*，否则返回 *False*。

is_zero()

如果参数是 0（正负皆可），则返回 *True*，否则返回 *False*。

ln(context=None)

Return the natural (base e) logarithm of the operand. The result is correctly rounded using the *ROUND_HALF_EVEN* rounding mode.

log10(context=None)

Return the base ten logarithm of the operand. The result is correctly rounded using the *ROUND_HALF_EVEN* rounding mode.

logb(context=None)

For a nonzero number, return the adjusted exponent of its operand as a *Decimal* instance. If the operand is a zero then *Decimal('-Infinity')* is returned and the *DivisionByZero* flag is raised. If the operand is an infinity then *Decimal('Infinity')* is returned.

logical_and(other, context=None)

logical_and() is a logical operation which takes two *logical operands* (see *Logical operands*). The result is the digit-wise and of the two operands.

logical_invert(context=None)

logical_invert() is a logical operation. The result is the digit-wise inversion of the operand.

logical_or(other, context=None)

logical_or() is a logical operation which takes two *logical operands* (see *Logical operands*). The result is the digit-wise or of the two operands.

logical_xor(other, context=None)

logical_xor() is a logical operation which takes two *logical operands* (see *Logical operands*). The result is the digit-wise exclusive or of the two operands.

max(other, context=None)

Like *max(self, other)* except that the context rounding rule is applied before returning and that NaN values are either signaled or ignored (depending on the context and whether they are signaling or quiet).

max_mag (*other*, *context=None*)

Similar to the `max()` method, but the comparison is done using the absolute values of the operands.

min (*other*, *context=None*)

Like `min(self, other)` except that the context rounding rule is applied before returning and that NaN values are either signaled or ignored (depending on the context and whether they are signaling or quiet).

min_mag (*other*, *context=None*)

Similar to the `min()` method, but the comparison is done using the absolute values of the operands.

next_minus (*context=None*)

Return the largest number representable in the given context (or in the current thread's context if no context is given) that is smaller than the given operand.

next_plus (*context=None*)

Return the smallest number representable in the given context (or in the current thread's context if no context is given) that is larger than the given operand.

next_toward (*other*, *context=None*)

If the two operands are unequal, return the number closest to the first operand in the direction of the second operand. If both operands are numerically equal, return a copy of the first operand with the sign set to be the same as the sign of the second operand.

normalize (*context=None*)

Normalize the number by stripping the rightmost trailing zeros and converting any result equal to `Decimal('0')` to `Decimal('0e0')`. Used for producing canonical values for attributes of an equivalence class. For example, `Decimal('32.100')` and `Decimal('0.321000e+2')` both normalize to the equivalent value `Decimal('32.1')`.

number_class (*context=None*)

Return a string describing the *class* of the operand. The returned value is one of the following ten strings.

- "-Infinity", indicating that the operand is negative infinity.
- "-Normal", indicating that the operand is a negative normal number.
- "-Subnormal", indicating that the operand is negative and subnormal.
- "-Zero", indicating that the operand is a negative zero.
- "+Zero", indicating that the operand is a positive zero.
- "+Subnormal", indicating that the operand is positive and subnormal.
- "+Normal", indicating that the operand is a positive normal number.
- "+Infinity", indicating that the operand is positive infinity.
- "NaN", indicating that the operand is a quiet NaN (Not a Number).
- "sNaN", indicating that the operand is a signaling NaN.

quantize (*exp*, *rounding=None*, *context=None*)

Return a value equal to the first operand after rounding and having the exponent of the second operand.

```
>>> Decimal('1.41421356').quantize(Decimal('1.000'))
Decimal('1.414')
```

Unlike other operations, if the length of the coefficient after the quantize operation would be greater than precision, then an *InvalidOperation* is signaled. This guarantees that, unless there is an error condition, the quantized exponent is always equal to that of the right-hand operand.

Also unlike other operations, quantize never signals Underflow, even if the result is subnormal and inexact.

If the exponent of the second operand is larger than that of the first then rounding may be necessary. In this case, the rounding mode is determined by the `rounding` argument if given, else by the given `context` argument; if neither argument is given the rounding mode of the current thread's context is used.

An error is returned whenever the resulting exponent is greater than `Emax` or less than `Etiny`.

radix()

Return `Decimal(10)`, the radix (base) in which the `Decimal` class does all its arithmetic. Included for compatibility with the specification.

remainder_near (*other*, *context=None*)

Return the remainder from dividing *self* by *other*. This differs from `self % other` in that the sign of the remainder is chosen so as to minimize its absolute value. More precisely, the return value is `self - n * other` where *n* is the integer nearest to the exact value of `self / other`, and if two integers are equally near then the even one is chosen.

If the result is zero then its sign will be the sign of *self*.

```
>>> Decimal(18).remainder_near(Decimal(10))
Decimal('-2')
>>> Decimal(25).remainder_near(Decimal(10))
Decimal('5')
>>> Decimal(35).remainder_near(Decimal(10))
Decimal('-5')
```

rotate (*other*, *context=None*)

Return the result of rotating the digits of the first operand by an amount specified by the second operand. The second operand must be an integer in the range `-precision` through `precision`. The absolute value of the second operand gives the number of places to rotate. If the second operand is positive then rotation is to the left; otherwise rotation is to the right. The coefficient of the first operand is padded on the left with zeros to length `precision` if necessary. The sign and exponent of the first operand are unchanged.

same_quantum (*other*, *context=None*)

Test whether *self* and *other* have the same exponent or whether both are NaN.

此操作不受上下文影响且静默：不更改任何标志且不执行舍入。作为例外，如果无法准确转换第二个操作数，则 C 版本可能会引发 `InvalidOperation`。

scaleb (*other*, *context=None*)

Return the first operand with exponent adjusted by the second. Equivalently, return the first operand multiplied by `10**other`. The second operand must be an integer.

shift (*other*, *context=None*)

Return the result of shifting the digits of the first operand by an amount specified by the second operand. The second operand must be an integer in the range `-precision` through `precision`. The absolute value of the second operand gives the number of places to shift. If the second operand is positive then the shift is to the left; otherwise the shift is to the right. Digits shifted into the coefficient are zeros. The sign and exponent of the first operand are unchanged.

sqrt (*context=None*)

Return the square root of the argument to full precision.

to_eng_string (*context=None*)

Convert to a string, using engineering notation if an exponent is needed.

Engineering notation has an exponent which is a multiple of 3. This can leave up to 3 digits to the left of the decimal place and may require the addition of either one or two trailing zeros.

For example, this converts `Decimal('123E+1')` to `Decimal('1.23E+3')`.

to_integral (*rounding=None, context=None*)

Identical to the `to_integral_value()` method. The `to_integral` name has been kept for compatibility with older versions.

to_integral_exact (*rounding=None, context=None*)

Round to the nearest integer, signaling *Inexact* or *Rounded* as appropriate if rounding occurs. The rounding mode is determined by the `rounding` parameter if given, else by the given `context`. If neither parameter is given then the rounding mode of the current context is used.

to_integral_value (*rounding=None, context=None*)

Round to the nearest integer without signaling *Inexact* or *Rounded*. If given, applies *rounding*; otherwise, uses the rounding method in either the supplied *context* or the current context.

Logical operands

The `logical_and()`, `logical_invert()`, `logical_or()`, and `logical_xor()` methods expect their arguments to be *logical operands*. A *logical operand* is a *Decimal* instance whose exponent and sign are both zero, and whose digits are all either 0 or 1.

9.4.3 Context objects

Contexts are environments for arithmetic operations. They govern precision, set rules for rounding, determine which signals are treated as exceptions, and limit the range for exponents.

Each thread has its own current context which is accessed or changed using the `getcontext()` and `setcontext()` functions:

`decimal.getcontext()`

Return the current context for the active thread.

`decimal.setcontext(c)`

Set the current context for the active thread to *c*.

You can also use the `with` statement and the `localcontext()` function to temporarily change the active context.

`decimal.localcontext(ctx=None)`

Return a context manager that will set the current context for the active thread to a copy of *ctx* on entry to the `with`-statement and restore the previous context when exiting the `with`-statement. If no context is specified, a copy of the current context is used.

For example, the following code sets the current decimal precision to 42 places, performs a calculation, and then automatically restores the previous context:

```
from decimal import localcontext

with localcontext() as ctx:
    ctx.prec = 42    # Perform a high precision calculation
    s = calculate_something()
s = +s    # Round the final result back to the default precision
```

New contexts can also be created using the *Context* constructor described below. In addition, the module provides three pre-made contexts:

class `decimal.BasicContext`

This is a standard context defined by the General Decimal Arithmetic Specification. Precision is set to nine. Rounding is set to *ROUND_HALF_UP*. All flags are cleared. All traps are enabled (treated as exceptions) except *Inexact*, *Rounded*, and *Subnormal*.

Because many of the traps are enabled, this context is useful for debugging.

class decimal.**ExtendedContext**

This is a standard context defined by the General Decimal Arithmetic Specification. Precision is set to nine. Rounding is set to `ROUND_HALF_EVEN`. All flags are cleared. No traps are enabled (so that exceptions are not raised during computations).

Because the traps are disabled, this context is useful for applications that prefer to have result value of NaN or Infinity instead of raising exceptions. This allows an application to complete a run in the presence of conditions that would otherwise halt the program.

class decimal.**DefaultContext**

This context is used by the `Context` constructor as a prototype for new contexts. Changing a field (such as precision) has the effect of changing the default for new contexts created by the `Context` constructor.

This context is most useful in multi-threaded environments. Changing one of the fields before threads are started has the effect of setting system-wide defaults. Changing the fields after threads have started is not recommended as it would require thread synchronization to prevent race conditions.

In single threaded environments, it is preferable to not use this context at all. Instead, simply create contexts explicitly as described below.

The default values are `prec=28`, `rounding=ROUND_HALF_EVEN`, and enabled traps for `Overflow`, `InvalidOperation`, and `DivisionByZero`.

In addition to the three supplied contexts, new contexts can be created with the `Context` constructor.

class decimal.**Context** (*prec=None, rounding=None, Emin=None, Emax=None, capitals=None, clamp=None, flags=None, traps=None*)

Creates a new context. If a field is not specified or is `None`, the default values are copied from the `DefaultContext`. If the `flags` field is not specified or is `None`, all flags are cleared.

`prec` is an integer in the range `[1, MAX_PREC]` that sets the precision for arithmetic operations in the context.

The `rounding` option is one of the constants listed in the section *Rounding Modes*.

The `traps` and `flags` fields list any signals to be set. Generally, new contexts should only set traps and leave the flags clear.

The `Emin` and `Emax` fields are integers specifying the outer limits allowable for exponents. `Emin` must be in the range `[MIN_EMIN, 0]`, `Emax` in the range `[0, MAX_EMAX]`.

The `capitals` field is either 0 or 1 (the default). If set to 1, exponents are printed with a capital E; otherwise, a lowercase e is used: `Decimal('6.02e+23')`.

The `clamp` field is either 0 (the default) or 1. If set to 1, the exponent `e` of a `Decimal` instance representable in this context is strictly limited to the range `Emin - prec + 1 <= e <= Emax - prec + 1`. If `clamp` is 0 then a weaker condition holds: the adjusted exponent of the `Decimal` instance is at most `Emax`. When `clamp` is 1, a large normal number will, where possible, have its exponent reduced and a corresponding number of zeros added to its coefficient, in order to fit the exponent constraints; this preserves the value of the number but loses information about significant trailing zeros. For example:

```
>>> Context(prec=6, Emax=999, clamp=1).create_decimal('1.23e999')
Decimal('1.23000E+999')
```

A `clamp` value of 1 allows compatibility with the fixed-width decimal interchange formats specified in IEEE 754.

The `Context` class defines several general purpose methods as well as a large number of methods for doing arithmetic directly in a given context. In addition, for each of the `Decimal` methods described above (with the exception of the `adjusted()` and `as_tuple()` methods) there is a corresponding `Context` method. For example, for a `Context` instance `C` and `Decimal` instance `x`, `C.exp(x)` is equivalent to `x.exp(context=C)`. Each `Context` method accepts a Python integer (an instance of `int`) anywhere that a `Decimal` instance is accepted.

clear_flags()

Resets all of the flags to 0.

clear_traps()

Resets all of the traps to 0.

3.3 新版功能.

copy()

Return a duplicate of the context.

copy_decimal(num)

Return a copy of the Decimal instance num.

create_decimal(num)

Creates a new Decimal instance from *num* but using *self* as context. Unlike the *Decimal* constructor, the context precision, rounding method, flags, and traps are applied to the conversion.

This is useful because constants are often given to a greater precision than is needed by the application. Another benefit is that rounding immediately eliminates unintended effects from digits beyond the current precision. In the following example, using unrounded inputs means that adding zero to a sum can change the result:

```
>>> getcontext().prec = 3
>>> Decimal('3.4445') + Decimal('1.0023')
Decimal('4.45')
>>> Decimal('3.4445') + Decimal(0) + Decimal('1.0023')
Decimal('4.44')
```

This method implements the to-number operation of the IBM specification. If the argument is a string, no leading or trailing whitespace or underscores are permitted.

create_decimal_from_float(f)

Creates a new Decimal instance from a float *f* but rounding using *self* as the context. Unlike the *Decimal.from_float()* class method, the context precision, rounding method, flags, and traps are applied to the conversion.

```
>>> context = Context(prec=5, rounding=ROUND_DOWN)
>>> context.create_decimal_from_float(math.pi)
Decimal('3.1415')
>>> context = Context(prec=5, traps=[Inexact])
>>> context.create_decimal_from_float(math.pi)
Traceback (most recent call last):
...
decimal.Inexact: None
```

3.1 新版功能.

Etiny()

Returns a value equal to $E_{\min} - \text{prec} + 1$ which is the minimum exponent value for subnormal results. When underflow occurs, the exponent is set to *Etiny*.

Etop()

Returns a value equal to $E_{\max} - \text{prec} + 1$.

The usual approach to working with decimals is to create *Decimal* instances and then apply arithmetic operations which take place within the current context for the active thread. An alternative approach is to use context methods for calculating within a specific context. The methods are similar to those for the *Decimal* class and are only briefly recounted here.

abs (*x*)
Returns the absolute value of *x*.

add (*x*, *y*)
Return the sum of *x* and *y*.

canonical (*x*)
Returns the same Decimal object *x*.

compare (*x*, *y*)
Compares *x* and *y* numerically.

compare_signal (*x*, *y*)
Compares the values of the two operands numerically.

compare_total (*x*, *y*)
Compares two operands using their abstract representation.

compare_total_mag (*x*, *y*)
Compares two operands using their abstract representation, ignoring sign.

copy_abs (*x*)
Returns a copy of *x* with the sign set to 0.

copy_negate (*x*)
Returns a copy of *x* with the sign inverted.

copy_sign (*x*, *y*)
Copies the sign from *y* to *x*.

divide (*x*, *y*)
Return *x* divided by *y*.

divide_int (*x*, *y*)
Return *x* divided by *y*, truncated to an integer.

divmod (*x*, *y*)
Divides two numbers and returns the integer part of the result.

exp (*x*)
Returns e^{**x} .

fma (*x*, *y*, *z*)
Returns *x* multiplied by *y*, plus *z*.

is_canonical (*x*)
Returns `True` if *x* is canonical; otherwise returns `False`.

is_finite (*x*)
Returns `True` if *x* is finite; otherwise returns `False`.

is_infinite (*x*)
Returns `True` if *x* is infinite; otherwise returns `False`.

is_nan (*x*)
Returns `True` if *x* is a qNaN or sNaN; otherwise returns `False`.

is_normal (*x*)
Returns `True` if *x* is a normal number; otherwise returns `False`.

is_qnan (*x*)
Returns `True` if *x* is a quiet NaN; otherwise returns `False`.

is_signed (*x*)
Returns True if *x* is negative; otherwise returns False.

is_snan (*x*)
Returns True if *x* is a signaling NaN; otherwise returns False.

is_subnormal (*x*)
Returns True if *x* is subnormal; otherwise returns False.

is_zero (*x*)
Returns True if *x* is a zero; otherwise returns False.

ln (*x*)
Returns the natural (base e) logarithm of *x*.

log10 (*x*)
Returns the base 10 logarithm of *x*.

logb (*x*)
Returns the exponent of the magnitude of the operand's MSD.

logical_and (*x*, *y*)
Applies the logical operation *and* between each operand's digits.

logical_invert (*x*)
Invert all the digits in *x*.

logical_or (*x*, *y*)
Applies the logical operation *or* between each operand's digits.

logical_xor (*x*, *y*)
Applies the logical operation *xor* between each operand's digits.

max (*x*, *y*)
Compares two values numerically and returns the maximum.

max_mag (*x*, *y*)
Compares the values numerically with their sign ignored.

min (*x*, *y*)
Compares two values numerically and returns the minimum.

min_mag (*x*, *y*)
Compares the values numerically with their sign ignored.

minus (*x*)
Minus corresponds to the unary prefix minus operator in Python.

multiply (*x*, *y*)
Return the product of *x* and *y*.

next_minus (*x*)
Returns the largest representable number smaller than *x*.

next_plus (*x*)
Returns the smallest representable number larger than *x*.

next_toward (*x*, *y*)
Returns the number closest to *x*, in direction towards *y*.

normalize (*x*)
Reduces *x* to its simplest form.

number_class (*x*)

Returns an indication of the class of *x*.

plus (*x*)

Plus corresponds to the unary prefix plus operator in Python. This operation applies the context precision and rounding, so it is *not* an identity operation.

power (*x*, *y*, *modulo*=None)

Return *x* to the power of *y*, reduced modulo *modulo* if given.

With two arguments, compute $x**y$. If *x* is negative then *y* must be integral. The result will be inexact unless *y* is integral and the result is finite and can be expressed exactly in ‘precision’ digits. The rounding mode of the context is used. Results are always correctly-rounded in the Python version.

在 3.3 版更改: The C module computes `power()` in terms of the correctly-rounded `exp()` and `ln()` functions. The result is well-defined but only “almost always correctly-rounded”.

With three arguments, compute $(x**y) \% modulo$. For the three argument form, the following restrictions on the arguments hold:

- all three arguments must be integral
- *y* must be nonnegative
- at least one of *x* or *y* must be nonzero
- *modulo* must be nonzero and have at most ‘precision’ digits

The value resulting from `Context.power(x, y, modulo)` is equal to the value that would be obtained by computing $(x**y) \% modulo$ with unbounded precision, but is computed more efficiently. The exponent of the result is zero, regardless of the exponents of *x*, *y* and *modulo*. The result is always exact.

quantize (*x*, *y*)

Returns a value equal to *x* (rounded), having the exponent of *y*.

radix ()

Just returns 10, as this is Decimal, :)

remainder (*x*, *y*)

Returns the remainder from integer division.

The sign of the result, if non-zero, is the same as that of the original dividend.

remainder_near (*x*, *y*)

Returns $x - y * n$, where *n* is the integer nearest the exact value of x / y (if the result is 0 then its sign will be the sign of *x*).

rotate (*x*, *y*)

Returns a rotated copy of *x*, *y* times.

same_quantum (*x*, *y*)

Returns `True` if the two operands have the same exponent.

scaleb (*x*, *y*)

Returns the first operand after adding the second value its exp.

shift (*x*, *y*)

Returns a shifted copy of *x*, *y* times.

sqrt (*x*)

Square root of a non-negative number to context precision.

subtract (*x*, *y*)Return the difference between *x* and *y*.**to_eng_string** (*x*)

Convert to a string, using engineering notation if an exponent is needed.

Engineering notation has an exponent which is a multiple of 3. This can leave up to 3 digits to the left of the decimal place and may require the addition of either one or two trailing zeros.

to_integral_exact (*x*)

Rounds to an integer.

to_sci_string (*x*)

Converts a number to a string using scientific notation.

9.4.4 常数

The constants in this section are only relevant for the C module. They are also included in the pure Python version for compatibility.

	32-bit	64-bit
<code>decimal.MAX_PREC</code>	425000000	999999999999999999
<code>decimal.MAX_EMAX</code>	425000000	999999999999999999
<code>decimal.MIN_EMIN</code>	-425000000	-999999999999999999
<code>decimal.MIN_ETINY</code>	-849999999	-1999999999999999997

`decimal.HAVE_THREADS`

The default value is `True`. If Python is compiled without threads, the C version automatically disables the expensive thread local context machinery. In this case, the value is `False`.

9.4.5 Rounding modes

`decimal.ROUND_CEILING`

Round towards Infinity.

`decimal.ROUND_DOWN`

Round towards zero.

`decimal.ROUND_FLOOR`

Round towards -Infinity.

`decimal.ROUND_HALF_DOWN`

Round to nearest with ties going towards zero.

`decimal.ROUND_HALF_EVEN`

Round to nearest with ties going to nearest even integer.

`decimal.ROUND_HALF_UP`

Round to nearest with ties going away from zero.

`decimal.ROUND_UP`

Round away from zero.

`decimal.ROUND_05UP`

Round away from zero if last digit after rounding towards zero would have been 0 or 5; otherwise round towards zero.

9.4.6 Signals

Signals represent conditions that arise during computation. Each corresponds to one context flag and one context trap enabler.

The context flag is set whenever the condition is encountered. After the computation, flags may be checked for informational purposes (for instance, to determine whether a computation was exact). After checking the flags, be sure to clear all flags before starting the next computation.

If the context's trap enabler is set for the signal, then the condition causes a Python exception to be raised. For example, if the *DivisionByZero* trap is set, then a *DivisionByZero* exception is raised upon encountering the condition.

class `decimal.Clamped`

Altered an exponent to fit representation constraints.

Typically, clamping occurs when an exponent falls outside the context's `Emin` and `Emax` limits. If possible, the exponent is reduced to fit by adding zeros to the coefficient.

class `decimal.DecimalException`

Base class for other signals and a subclass of *ArithmeticError*.

class `decimal.DivisionByZero`

Signals the division of a non-infinite number by zero.

Can occur with division, modulo division, or when raising a number to a negative power. If this signal is not trapped, returns `Infinity` or `-Infinity` with the sign determined by the inputs to the calculation.

class `decimal.Inexact`

Indicates that rounding occurred and the result is not exact.

Signals when non-zero digits were discarded during rounding. The rounded result is returned. The signal flag or trap is used to detect when results are inexact.

class `decimal.InvalidOperation`

An invalid operation was performed.

Indicates that an operation was requested that does not make sense. If not trapped, returns `NaN`. Possible causes include:

```
Infinity - Infinity
0 * Infinity
Infinity / Infinity
x % 0
Infinity % x
sqrt(-x) and x > 0
0 ** 0
x ** (non-integer)
x ** Infinity
```

class `decimal.Overflow`

Numerical overflow.

Indicates the exponent is larger than `Emax` after rounding has occurred. If not trapped, the result depends on the rounding mode, either pulling inward to the largest representable finite number or rounding outward to `Infinity`. In either case, *Inexact* and *Rounded* are also signaled.

class `decimal.Rounded`

Rounding occurred though possibly no information was lost.

Signaled whenever rounding discards digits; even if those digits are zero (such as rounding `5.00` to `5.0`). If not trapped, returns the result unchanged. This signal is used to detect loss of significant digits.

class `decimal.Subnormal`

Exponent was lower than `Emin` prior to rounding.

Occurs when an operation result is subnormal (the exponent is too small). If not trapped, returns the result unchanged.

class `decimal.Underflow`

Numerical underflow with result rounded to zero.

Occurs when a subnormal result is pushed to zero by rounding. *Inexact* and *Subnormal* are also signaled.

class `decimal.FloatOperation`

Enable stricter semantics for mixing floats and Decimals.

If the signal is not trapped (default), mixing floats and Decimals is permitted in the *Decimal* constructor, *create_decimal()* and all comparison operators. Both conversion and comparisons are exact. Any occurrence of a mixed operation is silently recorded by setting *FloatOperation* in the context flags. Explicit conversions with *from_float()* or *create_decimal_from_float()* do not set the flag.

Otherwise (the signal is trapped), only equality comparisons and explicit conversions are silent. All other mixed operations raise *FloatOperation*.

The following table summarizes the hierarchy of signals:

```
exceptions.ArithmeticError(exceptions.Exception)
    DecimalException
        Clamped
        DivisionByZero(DecimalException, exceptions.ZeroDivisionError)
        Inexact
            Overflow(Inexact, Rounded)
            Underflow(Inexact, Rounded, Subnormal)
        InvalidOperation
        Rounded
        Subnormal
        FloatOperation(DecimalException, exceptions.TypeError)
```

9.4.7 Floating Point Notes

Mitigating round-off error with increased precision

The use of decimal floating point eliminates decimal representation error (making it possible to represent `0.1` exactly); however, some operations can still incur round-off error when non-zero digits exceed the fixed precision.

The effects of round-off error can be amplified by the addition or subtraction of nearly offsetting quantities resulting in loss of significance. Knuth provides two instructive examples where rounded floating point arithmetic with insufficient precision causes the breakdown of the associative and distributive properties of addition:

```
# Examples from Seminumerical Algorithms, Section 4.2.2.
>>> from decimal import Decimal, getcontext
>>> getcontext().prec = 8

>>> u, v, w = Decimal(11111113), Decimal(-11111111), Decimal('7.51111111')
>>> (u + v) + w
Decimal('9.5111111')
>>> u + (v + w)
Decimal('10')

>>> u, v, w = Decimal(20000), Decimal(-6), Decimal('6.0000003')
>>> (u*v) + (u*w)
Decimal('0.01')
>>> u * (v+w)
Decimal('0.0060000')
```

The `decimal` module makes it possible to restore the identities by expanding the precision sufficiently to avoid loss of significance:

```
>>> getcontext().prec = 20
>>> u, v, w = Decimal(11111113), Decimal(-11111111), Decimal('7.51111111')
>>> (u + v) + w
Decimal('9.51111111')
>>> u + (v + w)
Decimal('9.51111111')
>>>
>>> u, v, w = Decimal(20000), Decimal(-6), Decimal('6.0000003')
>>> (u*v) + (u*w)
Decimal('0.0060000')
>>> u * (v+w)
Decimal('0.0060000')
```

Special values

The number system for the `decimal` module provides special values including NaN, sNaN, -Infinity, Infinity, and two zeros, +0 and -0.

Infinities can be constructed directly with: `Decimal('Infinity')`. Also, they can arise from dividing by zero when the `DivisionByZero` signal is not trapped. Likewise, when the `Overflow` signal is not trapped, infinity can result from rounding beyond the limits of the largest representable number.

The infinities are signed (affine) and can be used in arithmetic operations where they get treated as very large, indeterminate numbers. For instance, adding a constant to infinity gives another infinite result.

Some operations are indeterminate and return NaN, or if the `InvalidOperation` signal is trapped, raise an exception. For example, `0/0` returns NaN which means “not a number”. This variety of NaN is quiet and, once created, will flow through other computations always resulting in another NaN. This behavior can be useful for a series of computations that occasionally have missing inputs—it allows the calculation to proceed while flagging specific results as invalid.

A variant is sNaN which signals rather than remaining quiet after every operation. This is a useful return value when an invalid result needs to interrupt a calculation for special handling.

The behavior of Python’s comparison operators can be a little surprising where a NaN is involved. A test for equality where one of the operands is a quiet or signaling NaN always returns `False` (even when doing `Decimal('NaN')==Decimal('NaN')`), while a test for inequality always returns `True`. An attempt to compare two Decimals using any of the `<`, `<=`, `>` or `>=` operators will raise the `InvalidOperation` signal if either

operand is a NaN, and return *False* if this signal is not trapped. Note that the General Decimal Arithmetic specification does not specify the behavior of direct comparisons; these rules for comparisons involving a NaN were taken from the IEEE 854 standard (see Table 3 in section 5.7). To ensure strict standards-compliance, use the `compare()` and `compare-signal()` methods instead.

The signed zeros can result from calculations that underflow. They keep the sign that would have resulted if the calculation had been carried out to greater precision. Since their magnitude is zero, both positive and negative zeros are treated as equal and their sign is informational.

In addition to the two signed zeros which are distinct yet equal, there are various representations of zero with differing precisions yet equivalent in value. This takes a bit of getting used to. For an eye accustomed to normalized floating point representations, it is not immediately obvious that the following calculation returns a value equal to zero:

```
>>> 1 / Decimal('Infinity')
Decimal('0E-1000026')
```

9.4.8 Working with threads

The `getcontext()` function accesses a different *Context* object for each thread. Having separate thread contexts means that threads may make changes (such as `getcontext().prec=10`) without interfering with other threads.

Likewise, the `setcontext()` function automatically assigns its target to the current thread.

If `setcontext()` has not been called before `getcontext()`, then `getcontext()` will automatically create a new context for use in the current thread.

The new context is copied from a prototype context called *DefaultContext*. To control the defaults so that each thread will use the same values throughout the application, directly modify the *DefaultContext* object. This should be done *before* any threads are started so that there won't be a race condition between threads calling `getcontext()`. For example:

```
# Set applicationwide defaults for all threads about to be launched
DefaultContext.prec = 12
DefaultContext.rounding = ROUND_DOWN
DefaultContext.traps = ExtendedContext.traps.copy()
DefaultContext.traps[InvalidOperation] = 1
setcontext(DefaultContext)

# Afterwards, the threads can be started
t1.start()
t2.start()
t3.start()
...
```

9.4.9 Recipes

Here are a few recipes that serve as utility functions and that demonstrate ways to work with the *Decimal* class:

```
def moneyfmt(value, places=2, curr='', sep=',', dp='.',
             pos='', neg='-', trailneg=''):
    """Convert Decimal to a money formatted string.

    places:  required number of places after the decimal point
    curr:    optional currency symbol before the sign (may be blank)
    sep:     optional grouping separator (comma, period, space, or blank)
    dp:      decimal point indicator (comma or period)
    pos:     sign of positive number: "" (default) or "+"
    neg:     sign of negative number: "-" (default) or "(-)"
    trailneg: trailing minus sign: "" (default) or "(-)"
```

(下页继续)

(续上页)

```

        only specify as blank when places is zero
pos:      optional sign for positive numbers: '+', space or blank
neg:      optional sign for negative numbers: '-', '(', space or blank
trailneg: optional trailing minus indicator: '-', ')', space or blank

>>> d = Decimal('-1234567.8901')
>>> moneyfmt(d, curr='$')
'-$1,234,567.89'
>>> moneyfmt(d, places=0, sep='.', dp='', neg='', trailneg='-')
'1.234.568-'
>>> moneyfmt(d, curr='$', neg='(', trailneg='')
'($1,234,567.89)'
>>> moneyfmt(Decimal(123456789), sep=' ')
'123 456 789.00'
>>> moneyfmt(Decimal('-0.02'), neg='<', trailneg='>')
'<0.02>'

"""
q = Decimal(10) ** -places      # 2 places --> '0.01'
sign, digits, exp = value.quantize(q).as_tuple()
result = []
digits = list(map(str, digits))
build, next = result.append, digits.pop
if sign:
    build(trailneg)
for i in range(places):
    build(next() if digits else '0')
if places:
    build(dp)
if not digits:
    build('0')
i = 0
while digits:
    build(next())
    i += 1
    if i == 3 and digits:
        i = 0
        build(sep)
build(curr)
build(neg if sign else pos)
return ''.join(reversed(result))

def pi():
    """Compute Pi to the current precision.

>>> print(pi())
3.141592653589793238462643383

"""
getcontext().prec += 2 # extra digits for intermediate steps
three = Decimal(3)     # substitute "three=3.0" for regular floats
lasts, t, s, n, na, d, da = 0, three, 3, 1, 0, 0, 24
while s != lasts:
    lasts = s
    n, na = n+na, na+8
    d, da = d+da, da+32

```

(下页继续)

(续上页)

```

        t = (t * n) / d
        s += t
    getcontext().prec -= 2
    return +s                # unary plus applies the new precision

def exp(x):
    """Return e raised to the power of x.  Result type matches input type.

    >>> print(exp(Decimal(1)))
    2.718281828459045235360287471
    >>> print(exp(Decimal(2)))
    7.389056098930650227230427461
    >>> print(exp(2.0))
    7.38905609893
    >>> print(exp(2+0j))
    (7.38905609893+0j)

    """
    getcontext().prec += 2
    i, lasts, s, fact, num = 0, 0, 1, 1, 1
    while s != lasts:
        lasts = s
        i += 1
        fact *= i
        num *= x
        s += num / fact
    getcontext().prec -= 2
    return +s

def cos(x):
    """Return the cosine of x as measured in radians.

    The Taylor series approximation works best for a small value of x.
    For larger values, first compute x = x % (2 * pi).

    >>> print(cos(Decimal('0.5')))
    0.8775825618903727161162815826
    >>> print(cos(0.5))
    0.87758256189
    >>> print(cos(0.5+0j))
    (0.87758256189+0j)

    """
    getcontext().prec += 2
    i, lasts, s, fact, num, sign = 0, 0, 1, 1, 1, 1
    while s != lasts:
        lasts = s
        i += 2
        fact *= i * (i-1)
        num *= x * x
        sign *= -1
        s += num / fact * sign
    getcontext().prec -= 2
    return +s

def sin(x):

```

(下页继续)

(续上页)

```

"""Return the sine of x as measured in radians.

The Taylor series approximation works best for a small value of x.
For larger values, first compute x = x % (2 * pi).

>>> print(sin(Decimal('0.5')))
0.4794255386042030002732879352
>>> print(sin(0.5))
0.479425538604
>>> print(sin(0.5+0j))
(0.479425538604+0j)

"""
getcontext().prec += 2
i, lasts, s, fact, num, sign = 1, 0, x, 1, x, 1
while s != lasts:
    lasts = s
    i += 2
    fact *= i * (i-1)
    num *= x * x
    sign *= -1
    s += num / fact * sign
getcontext().prec -= 2
return +s

```

9.4.10 Decimal FAQ

Q. It is cumbersome to type `decimal.Decimal('1234.5')`. Is there a way to minimize typing when using the interactive interpreter?

A. Some users abbreviate the constructor to just a single letter:

```

>>> D = decimal.Decimal
>>> D('1.23') + D('3.45')
Decimal('4.68')

```

Q. In a fixed-point application with two decimal places, some inputs have many places and need to be rounded. Others are not supposed to have excess digits and need to be validated. What methods should be used?

A. The `quantize()` method rounds to a fixed number of decimal places. If the *Inexact* trap is set, it is also useful for validation:

```

>>> TWOPLACES = Decimal(10) ** -2           # same as Decimal('0.01')

```

```

>>> # Round to two places
>>> Decimal('3.214').quantize(TWOPLACES)
Decimal('3.21')

```

```

>>> # Validate that a number does not exceed two places
>>> Decimal('3.21').quantize(TWOPLACES, context=Context(traps=[Inexact]))
Decimal('3.21')

```

```

>>> Decimal('3.214').quantize(TWOPLACES, context=Context(traps=[Inexact]))
Traceback (most recent call last):

```

(下页继续)

(续上页)

```
...
Inexact: None
```

Q. Once I have valid two place inputs, how do I maintain that invariant throughout an application?

A. Some operations like addition, subtraction, and multiplication by an integer will automatically preserve fixed point. Others operations, like division and non-integer multiplication, will change the number of decimal places and need to be followed-up with a `quantize()` step:

```
>>> a = Decimal('102.72')          # Initial fixed-point values
>>> b = Decimal('3.17')
>>> a + b                          # Addition preserves fixed-point
Decimal('105.89')
>>> a - b
Decimal('99.55')
>>> a * 42                         # So does integer multiplication
Decimal('4314.24')
>>> (a * b).quantize(TWOPLACES)    # Must quantize non-integer multiplication
Decimal('325.62')
>>> (b / a).quantize(TWOPLACES)    # And quantize division
Decimal('0.03')
```

In developing fixed-point applications, it is convenient to define functions to handle the `quantize()` step:

```
>>> def mul(x, y, fp=TWOPLACES):
...     return (x * y).quantize(fp)
>>> def div(x, y, fp=TWOPLACES):
...     return (x / y).quantize(fp)
```

```
>>> mul(a, b)                      # Automatically preserve fixed-point
Decimal('325.62')
>>> div(b, a)
Decimal('0.03')
```

Q. There are many ways to express the same value. The numbers 200, 200.000, 2E2, and 02E+4 all have the same value at various precisions. Is there a way to transform them to a single recognizable canonical value?

A. The `normalize()` method maps all equivalent values to a single representative:

```
>>> values = map(Decimal, '200 200.000 2E2 .02E+4'.split())
>>> [v.normalize() for v in values]
[Decimal('2E+2'), Decimal('2E+2'), Decimal('2E+2'), Decimal('2E+2')]
```

Q. Some decimal values always print with exponential notation. Is there a way to get a non-exponential representation?

A. For some values, exponential notation is the only way to express the number of significant places in the coefficient. For example, expressing 5.0E+3 as 5000 keeps the value constant but cannot show the original's two-place significance.

If an application does not care about tracking significance, it is easy to remove the exponent and trailing zeroes, losing significance, but keeping the value unchanged:

```
>>> def remove_exponent(d):
...     return d.quantize(Decimal(1)) if d == d.to_integral() else d.normalize()
```

```
>>> remove_exponent(Decimal('5E+3'))
Decimal('5000')
```

Q. Is there a way to convert a regular float to a *Decimal*?

A. Yes, any binary floating point number can be exactly expressed as a *Decimal* though an exact conversion may take more precision than intuition would suggest:

```
>>> Decimal(math.pi)
Decimal('3.141592653589793115997963468544185161590576171875')
```

Q. Within a complex calculation, how can I make sure that I haven't gotten a spurious result because of insufficient precision or rounding anomalies.

A. The decimal module makes it easy to test results. A best practice is to re-run calculations using greater precision and with various rounding modes. Widely differing results indicate insufficient precision, rounding mode issues, ill-conditioned inputs, or a numerically unstable algorithm.

Q. I noticed that context precision is applied to the results of operations but not to the inputs. Is there anything to watch out for when mixing values of different precisions?

A. Yes. The principle is that all values are considered to be exact and so is the arithmetic on those values. Only the results are rounded. The advantage for inputs is that “what you type is what you get”. A disadvantage is that the results can look odd if you forget that the inputs haven't been rounded:

```
>>> getcontext().prec = 3
>>> Decimal('3.104') + Decimal('2.104')
Decimal('5.21')
>>> Decimal('3.104') + Decimal('0.000') + Decimal('2.104')
Decimal('5.20')
```

The solution is either to increase precision or to force rounding of inputs using the unary plus operation:

```
>>> getcontext().prec = 3
>>> +Decimal('1.23456789')      # unary plus triggers rounding
Decimal('1.23')
```

Alternatively, inputs can be rounded upon creation using the *Context.create_decimal()* method:

```
>>> Context(prec=5, rounding=ROUND_DOWN).create_decimal('1.2345678')
Decimal('1.2345')
```

9.5 fractions — 分数

源代码 [Lib/fractions.py](#)

fractions 模块支持分数运算。

分数实例可以由一对整数，一个分数，或者一个字符串构建而成。

```
class fractions.Fraction(numerator=0, denominator=1)
class fractions.Fraction(other_fraction)
class fractions.Fraction(float)
class fractions.Fraction(decimal)
class fractions.Fraction(string)
```

The first version requires that *numerator* and *denominator* are instances of *numbers.Rational* and returns a new *Fraction* instance with value *numerator/denominator*. If *denominator* is 0, it raises

a `ZeroDivisionError`. The second version requires that `other_fraction` is an instance of `numbers.Rational` and returns a `Fraction` instance with the same value. The next two versions accept either a `float` or a `decimal.Decimal` instance, and return a `Fraction` instance with exactly the same value. Note that due to the usual issues with binary floating-point (see [tut-fp-issues](#)), the argument to `Fraction(1.1)` is not exactly equal to $11/10$, and so `Fraction(1.1)` does *not* return `Fraction(11, 10)` as one might expect. (But see the documentation for the `limit_denominator()` method below.) The last version of the constructor expects a string or unicode instance. The usual form for this instance is:

```
[sign] numerator ['/' denominator]
```

where the optional sign may be either '+' or '-' and numerator and denominator (if present) are strings of decimal digits. In addition, any string that represents a finite value and is accepted by the `float` constructor is also accepted by the `Fraction` constructor. In either form the input string may also have leading and/or trailing whitespace. Here are some examples:

```
>>> from fractions import Fraction
>>> Fraction(16, -10)
Fraction(-8, 5)
>>> Fraction(123)
Fraction(123, 1)
>>> Fraction()
Fraction(0, 1)
>>> Fraction('3/7')
Fraction(3, 7)
>>> Fraction('-3/7 ')
Fraction(-3, 7)
>>> Fraction('1.414213 \t\n')
Fraction(1414213, 1000000)
>>> Fraction('-.125')
Fraction(-1, 8)
>>> Fraction('7e-6')
Fraction(7, 1000000)
>>> Fraction(2.25)
Fraction(9, 4)
>>> Fraction(1.1)
Fraction(2476979795053773, 2251799813685248)
>>> from decimal import Decimal
>>> Fraction(Decimal('1.1'))
Fraction(11, 10)
```

The `Fraction` class inherits from the abstract base class `numbers.Rational`, and implements all of the methods and operations from that class. `Fraction` instances are hashable, and should be treated as immutable. In addition, `Fraction` has the following properties and methods:

在 3.2 版更改: The `Fraction` constructor now accepts `float` and `decimal.Decimal` instances.

numerator

最简分数形式的分子。

denominator

最简分数形式的分母。

from_float (flt)

This class method constructs a `Fraction` representing the exact value of `flt`, which must be a `float`. Beware that `Fraction.from_float(0.3)` is not the same value as `Fraction(3, 10)`.

注解: From Python 3.2 onwards, you can also construct a `Fraction` instance directly from a `float`.

from_decimal (*dec*)

This class method constructs a *Fraction* representing the exact value of *dec*, which must be a *decimal.Decimal* instance.

注解: From Python 3.2 onwards, you can also construct a *Fraction* instance directly from a *decimal.Decimal* instance.

limit_denominator (*max_denominator=1000000*)

Finds and returns the closest *Fraction* to *self* that has denominator at most *max_denominator*. This method is useful for finding rational approximations to a given floating-point number:

```
>>> from fractions import Fraction
>>> Fraction('3.1415926535897932').limit_denominator(1000)
Fraction(355, 113)
```

or for recovering a rational number that's represented as a float:

```
>>> from math import pi, cos
>>> Fraction(cos(pi/3))
Fraction(4503599627370497, 9007199254740992)
>>> Fraction(cos(pi/3)).limit_denominator()
Fraction(1, 2)
>>> Fraction(1.1).limit_denominator()
Fraction(11, 10)
```

__floor__ ()

Returns the greatest *int* \leq *self*. This method can also be accessed through the *math.floor()* function:

```
>>> from math import floor
>>> floor(Fraction(355, 113))
3
```

__ceil__ ()

Returns the least *int* \geq *self*. This method can also be accessed through the *math.ceil()* function.

__round__ ()**__round__** (*ndigits*)

The first version returns the nearest *int* to *self*, rounding half to even. The second version rounds *self* to the nearest multiple of *Fraction(1, 10**ndigits)* (logically, if *ndigits* is negative), again rounding half toward even. This method can also be accessed through the *round()* function.

fractions.gcd (*a, b*)

返回整数 *a* 和 *b* 的最大公约数。如果 *a* 或 *b* 之一非零，则 *gcd(a, b)* 的绝对值是能同时整除 *a* 和 *b* 的最大整数。若 *b* 非零，则 *gcd(a, b)* 与 *b* 同号；否则返回值与 *a* 同号。*gcd(0, 0)* 返回 0。

3.5 版后已移除: 由 *math.gcd()* 取代。

参见:

numbers 模块 The abstract base classes making up the numeric tower.

9.6 random — 生成伪随机数

源码: [Lib/random.py](#)

该模块实现了各种分布的伪随机数生成器。

对于整数，从范围中有统一的选择。对于序列，存在随机元素的统一选择、用于生成列表的随机排列的函数、以及用于随机抽样而无需替换的函数。

在实数轴上，有计算均匀、正态（高斯）、对数正态、负指数、伽马和贝塔分布的函数。为了生成角度分布，可以使用 von Mises 分布。

几乎所有模块函数都依赖于基本函数 `random()`，它在半开区间 `[0.0,1.0)` 内均匀生成随机浮点数。Python 使用 Mersenne Twister 作为核心生成器。它产生 53 位精度浮点数，周期为 $2^{19937}-1$ ，其在 C 中的底层实现既快又线程安全。Mersenne Twister 是现存最广泛测试的随机数发生器之一。但是，因为完全确定性，它不适用于所有目的，并且完全不适合加密目的。

这个模块提供的函数实际上是 `random.Random` 类的隐藏实例的绑定方法。你可以实例化自己的 `Random` 类实例以获取不共享状态的生成器。

如果你想使用自己设计的不同基础生成器，类 `Random` 也可以作为子类：在这种情况下，重载 `random()`、`seed()`、`getstate()` 以及 `setstate()` 方法。可选地，新生成器可以提供 `getrandbits()` 方法——这允许 `randrange()` 在任意大的范围内产生选择。

`random` 模块还提供 `SystemRandom` 类，它使用系统函数 `os.urandom()` 从操作系统提供的源生成随机数。

警告： 不应将此模块的伪随机生成器用于安全目的。有关安全性或加密用途，请参阅 `secrets` 模块。

参见：

M. Matsumoto and T. Nishimura, “Mersenne Twister: A 623-dimensionally equidistributed uniform pseudorandom number generator”, ACM Transactions on Modeling and Computer Simulation Vol. 8, No. 1, January pp.3–30 1998.

[Complementary-Multiply-with-Carry recipe](#) 用于兼容的替代随机数发生器，具有长周期和相对简单的更新操作。

9.6.1 簿记功能

`random.seed(a=None, version=2)`

初始化随机数生成器。

如果 `a` 被省略或为 `None`，则使用当前系统时间。如果操作系统提供随机源，则使用它们而不是系统时间（有关可用性的详细信息，请参阅 `os.urandom()` 函数）。

如果 `a` 是 `int` 类型，则直接使用。

对于版本 2（默认的），`str`、`bytes` 或 `bytearray` 对象转换为 `int` 并使用它的所有位。

对于版本 1（用于从旧版本的 Python 再现随机序列），用于 `str` 和 `bytes` 的算法生成更窄的种子范围。

在 3.2 版更改：已移至版本 2 方案，该方案使用字符串种子中的所有位。

`random.getstate()`

返回捕获生成器当前内部状态的对象。这个对象可以传递给 `setstate()` 来恢复状态。

`random.setstate(state)`

`state` 应该是从之前调用 `getstate()` 获得的，并且 `setstate()` 将生成器的内部状态恢复到 `getstate()` 被调用时的状态。

`random.getrandbits(k)`

返回带有 `k` 位随机的 Python 整数。此方法随 MersenneTwister 生成器一起提供，其他一些生成器也可以将其作为 API 的可选部分提供。如果可用，`getrandbits()` 启用 `randrange()` 来处理任意大范围。

9.6.2 整数用函数

`random.randrange(stop)`

`random.randrange(start, stop[, step])`

从 `range(start, stop, step)` 返回一个随机选择的元素。这相当于 `choice(range(start, stop, step))`，但实际上并没有构建一个 `range` 对象。

位置参数模式匹配 `range()`。不应使用关键字参数，因为该函数可能以意外的方式使用它们。

在 3.2 版更改：`randrange()` 在生成均匀分布的值方面更为复杂。以前它使用了像 `int(random()*n)`“这样的形式，它可以产生稍微不均匀的分布。

`random.randint(a, b)`

返回随机整数 N 满足 $a \leq N \leq b$ 。相当于 `randrange(a, b+1)`。

9.6.3 序列用函数

`random.choice(seq)`

从非空序列 `seq` 返回一个随机元素。如果 `seq` 为空，则引发 `IndexError`。

`random.choices(population, weights=None, *, cum_weights=None, k=1)`

从 `*population*` 中选择替换，返回大小为 `k` 的元素列表。如果 `population` 为空，则引发 `IndexError`。

如果指定了 `weight` 序列，则根据相对权重进行选择。或者，如果给出 `cum_weights` 序列，则根据累积权重（可能使用 `itertools.accumulate()` 计算）进行选择。例如，相对权重 “[10, 5, 30, 5]” 相当于累积权重 “[10, 15, 45, 50]”。在内部，相对权重在进行选择之前会转换为累积权重，因此提供累积权重可以节省工作量。

如果既未指定 `weight` 也未指定 `cum_weights`，则以相等的概率进行选择。如果提供了权重序列，则它必须与 `population` 序列的长度相同。一个 `TypeError` 指定了 `weights` 和 `*cum_weights*`。

`weights` 或 `cum_weights` 可以使用任何与 `random()` 返回的 `float` 值互操作的数值类型（包括整数，浮点数和分数但不包括十进制小数）。

3.6 新版功能。

`random.shuffle(x[, random])`

将序列 `x` 随机打乱位置。

可选参数 `random` 是一个 0 参数函数，在 [0.0, 1.0) 中返回随机浮点数；默认情况下，这是函数 `random()`。

要改变一个不可变的序列并返回一个新的打乱列表，请使用 `“sample(x, k=len(x))”`。

请注意，即使对于小的 `len(x)`，`x` 的排列总数也可以快速增长，大于大多数随机数生成器的周期。这意味着长序列的大多数排列永远不会产生。例如，长度为 2080 的序列是可以在 Mersenne Twister 随机数生成器的周期内拟合的最大序列。

`random.sample(population, k)`

返回从总体序列或集合中选择的唯一元素的 `k` 长度列表。用于无重复的随机抽样。

返回包含来自总体的元素的新列表，同时保持原始总体不变。结果列表按选择顺序排列，因此所有子切片也将是有效的随机样本。这允许抽奖获奖者（样本）被划分为大奖和第二名获胜者（子切片）。

总体成员不必是 `hashable` 或 `unique`。如果总体包含重复，则每次出现都是样本中可能的选择。

要从一系列整数中选择样本，请使用 `range()` 对象作为参数。对于从大量人群中采样，这种方法特别快速且节省空间：`sample(range(10000000), k=60)`。

如果样本大小大于总体大小，则引发 `ValueError`。

9.6.4 实值分布

以下函数生成特定的实值分布。如常用数学实践中所使用的那样，函数参数以分布方程中的相应变量命名；大多数这些方程都可以在任何统计学教材中找到。

`random.random()`

返回 [0.0, 1.0) 范围内的下一个随机浮点数。

`random.uniform(a, b)`

返回一个随机浮点数 N ，当 $a \leq b$ 时 $a \leq N \leq b$ ，当 $b < a$ 时 $b \leq N \leq a$ 。

取决于等式 $a + (b-a) * \text{random}()$ 中的浮点舍入，终点 b 可以包括或不包括在该范围内。

`random.triangular(low, high, mode)`

返回一个随机浮点数 N ，使得 $low \leq N \leq high$ 并在这些边界之间使用指定的 *mode*。*low* 和 *high* 边界默认为零和一。*mode* 参数默认为边界之间的中点，给出对称分布。

`random.betavariate(alpha, beta)`

Beta 分布。参数的条件是 $\alpha > 0$ 和 $\beta > 0$ 。返回值的范围介于 0 和 1 之间。

`random.expovariate(lambd)`

指数分布。*lambd* 是 1.0 除以所需的平均值，它应该非零的。（该参数本应命名为 “lambda”，但这是 Python 中的保留字。）如果 *lambd* 为正，则返回值的范围为 0 到正无穷大；如果 *lambd* 为负，则返回值从负无穷大到 0。

`random.gammavariate(alpha, beta)`

Gamma 分布。（不是 gamma 函数！）参数的条件是 $\alpha > 0$ 和 $\beta > 0$ 。

概率分布函数是：

$$\text{pdf}(x) = \frac{x^{(\alpha - 1)} * \text{math.exp}(-x / \beta)}{\text{math.gamma}(\alpha) * \beta ** \alpha}$$

`random.gauss(mu, sigma)`

高斯分布。*mu* 是平均值，*sigma* 是标准差。这比下面定义的 `normalvariate()` 函数略快。

`random.lognormvariate(mu, sigma)`

对数正态分布。如果你采用这个分布的自然对数，你将得到一个正态分布，平均值为 *mu* 和标准差为 *sigma*。*mu* 可以是任何值，*sigma* 必须大于零。

`random.normalvariate(mu, sigma)`

正态分布。*mu* 是平均值，*sigma* 是标准差。

`random.vonmisesvariate(mu, kappa)`

mu 是平均角度，以弧度表示，介于 0 和 2π 之间，*kappa* 是浓度参数，必须大于或等于零。如果 *kappa* 等于零，则该分布在 0 到 2π 的范围内减小到均匀的随机角度。

`random.paretovariate(alpha)`

帕累托分布。*alpha* 是形状参数。

`random.weibullvariate(alpha, beta)`

威布尔分布。*alpha* 是比例参数，*beta* 是形状参数。

9.6.5 替代生成器

`class random.SystemRandom([seed])`

使用 `os.urandom()` 函数的类，用从操作系统提供的源生成随机数。这并非适用于所有系统。也不依赖于软件状态，序列不可重现。因此，`seed()` 方法没有效果而被忽略。`getstate()` 和 `setstate()` 方法如果被调用则引发 `NotImplementedError`。

9.6.6 关于再现性的说明

有时能够重现伪随机数生成器给出的序列是有用的。通过重新使用种子值，只要多个线程没有运行，相同的序列就可以在两次不同运行之间重现。

大多数随机模块的算法和种子函数都会在 Python 版本中发生变化，但保证两个方面不会改变：

- 如果添加了新的播种方法，则将提供向后兼容的播种机。
- 当兼容的播种机被赋予相同的种子时，生成器的 `random()` 方法将继续产生相同的序列。

9.6.7 例子和配方

基本示例:

```
>>> random()                                # Random float:  0.0 <= x < 1.0
0.37444887175646646

>>> uniform(2.5, 10.0)                      # Random float:  2.5 <= x < 10.0
3.1800146073117523

>>> expovariate(1 / 5)                      # Interval between arrivals averaging 5.
↪seconds
5.148957571865031

>>> randrange(10)                           # Integer from 0 to 9 inclusive
7

>>> randrange(0, 101, 2)                    # Even integer from 0 to 100 inclusive
26

>>> choice(['win', 'lose', 'draw'])          # Single random element from a sequence
'draw'

>>> deck = 'ace two three four'.split()
>>> shuffle(deck)                           # Shuffle a list
>>> deck
['four', 'two', 'ace', 'three']

>>> sample([10, 20, 30, 40, 50], k=4)        # Four samples without replacement
[40, 10, 50, 30]
```

模拟:

```
>>> # Six roulette wheel spins (weighted sampling with replacement)
>>> choices(['red', 'black', 'green'], [18, 18, 2], k=6)
['red', 'green', 'black', 'black', 'red', 'black']

>>> # Deal 20 cards without replacement from a deck of 52 playing cards
>>> # and determine the proportion of cards with a ten-value
>>> # (a ten, jack, queen, or king).
>>> deck = collections.Counter(tens=16, low_cards=36)
>>> seen = sample(list(deck.elements()), k=20)
>>> seen.count('tens') / 20
0.15

>>> # Estimate the probability of getting 5 or more heads from 7 spins
```

(下页继续)

(续上页)

```
>>> # of a biased coin that settles on heads 60% of the time.
>>> trial = lambda: choices('HT', cum_weights=(0.60, 1.00), k=7).count('H') >= 5
>>> sum(trial() for i in range(10000)) / 10000
0.4169

>>> # Probability of the median of 5 samples being in middle two quartiles
>>> trial = lambda: 2500 <= sorted(choices(range(10000), k=5))[2] < 7500
>>> sum(trial() for i in range(10000)) / 10000
0.7958
```

statistical bootstrapping 使用重采样和替换来估计大小为五的样本的均值的置信区间的示例:

```
# http://statistics.about.com/od/Applications/a/Example-Of-Bootstrapping.htm
from statistics import mean
from random import choices

data = 1, 2, 4, 4, 10
means = sorted(mean(choices(data, k=5)) for i in range(20))
print(f'The sample mean of {mean(data):.1f} has a 90% confidence '
      f'interval from {means[1]:.1f} to {means[-2]:.1f}')
```

使用 重新采样排列测试 来确定统计学显著性或者使用 p-值 来观察药物与安慰剂的作用之间差异的示例:

```
# Example from "Statistics is Easy" by Dennis Shasha and Manda Wilson
from statistics import mean
from random import shuffle

drug = [54, 73, 53, 70, 73, 68, 52, 65, 65]
placebo = [54, 51, 58, 44, 55, 52, 42, 47, 58, 46]
observed_diff = mean(drug) - mean(placebo)

n = 10000
count = 0
combined = drug + placebo
for i in range(n):
    shuffle(combined)
    new_diff = mean(combined[:len(drug)]) - mean(combined[len(drug):])
    count += (new_diff >= observed_diff)

print(f'{n} label reshufflings produced only {count} instances with a difference')
print(f'at least as extreme as the observed difference of {observed_diff:.1f}.')
print(f'The one-sided p-value of {count / n:.4f} leads us to reject the null')
print(f'hypothesis that there is no difference between the drug and the placebo.')
```

模拟单个服务器队列中的到达时间和服务交付:

```
from random import expovariate, gauss
from statistics import mean, median, stdev

average_arrival_interval = 5.6
average_service_time = 5.0
stdev_service_time = 0.5

num_waiting = 0
arrivals = []
starts = []
```

(下页继续)

(续上页)

```

arrival = service_end = 0.0
for i in range(20000):
    if arrival <= service_end:
        num_waiting += 1
        arrival += expovariate(1.0 / average_arrival_interval)
        arrivals.append(arrival)
    else:
        num_waiting -= 1
        service_start = service_end if num_waiting else arrival
        service_time = gauss(average_service_time, stdev_service_time)
        service_end = service_start + service_time
        starts.append(service_start)

waits = [start - arrival for arrival, start in zip(arrivals, starts)]
print(f'Mean wait: {mean(waits):.1f}. Stdev wait: {stdev(waits):.1f}.')
print(f'Median wait: {median(waits):.1f}. Max wait: {max(waits):.1f}.')

```

参见:

[Statistics for Hackers](#) Jake Vanderplas 撰写的视频教程，使用一些基本概念进行统计分析，包括模拟、抽样、改组和交叉验证。

[Economics Simulation](#) Peter Norvig 编写的市场模拟，显示了该模块提供的许多工具和分布的有效使用（高斯、均匀、样本、beta 变量、选择、三角和随机范围等）。

[A Concrete Introduction to Probability \(using Python\)](#) Peter Norvig 撰写的教程，涵盖了概率论基础知识，如何编写模拟，以及如何使用 Python 进行数据分析。

9.7 statistics —Mathematical statistics functions

3.4 新版功能.

Source code: [Lib/statistics.py](#)

This module provides functions for calculating mathematical statistics of numeric (Real-valued) data.

注解: Unless explicitly noted otherwise, these functions support *int*, *float*, *decimal.Decimal* and *fractions.Fraction*. Behaviour with other types (whether in the numeric tower or not) is currently unsupported. Mixed types are also undefined and implementation-dependent. If your input data consists of mixed types, you may be able to use *map()* to ensure a consistent result, e.g. *map(float, input_data)*.

9.7.1 Averages and measures of central location

These functions calculate an average or typical value from a population or sample.

<code>mean()</code>	Arithmetic mean (“average”) of data.
<code>harmonic_mean()</code>	Harmonic mean of data.
<code>median()</code>	Median (middle value) of data.
<code>median_low()</code>	Low median of data.
<code>median_high()</code>	High median of data.
<code>median_grouped()</code>	Median, or 50th percentile, of grouped data.
<code>mode()</code>	Mode (most common value) of discrete data.

9.7.2 Measures of spread

These functions calculate a measure of how much the population or sample tends to deviate from the typical or average values.

<code>pstdev()</code>	Population standard deviation of data.
<code>pvariance()</code>	Population variance of data.
<code>stdev()</code>	Sample standard deviation of data.
<code>variance()</code>	Sample variance of data.

9.7.3 Function details

Note: The functions do not require the data given to them to be sorted. However, for reading convenience, most of the examples show sorted sequences.

`statistics.mean(data)`

Return the sample arithmetic mean of *data* which can be a sequence or iterator.

The arithmetic mean is the sum of the data divided by the number of data points. It is commonly called “the average”, although it is only one of many different mathematical averages. It is a measure of the central location of the data.

If *data* is empty, `StatisticsError` will be raised.

Some examples of use:

```
>>> mean([1, 2, 3, 4, 4])
2.8
>>> mean([-1.0, 2.5, 3.25, 5.75])
2.625

>>> from fractions import Fraction as F
>>> mean([F(3, 7), F(1, 21), F(5, 3), F(1, 3)])
Fraction(13, 21)

>>> from decimal import Decimal as D
>>> mean([D("0.5"), D("0.75"), D("0.625"), D("0.375")])
Decimal('0.5625')
```

注解: The mean is strongly affected by outliers and is not a robust estimator for central location: the mean is not necessarily a typical example of the data points. For more robust, although less efficient, measures of central location, see `median()` and `mode()`. (In this case, “efficient” refers to statistical efficiency rather than computational efficiency.)

The sample mean gives an unbiased estimate of the true population mean, which means that, taken on average over all the possible samples, `mean(sample)` converges on the true mean of the entire population. If *data* represents the entire population rather than a sample, then `mean(data)` is equivalent to calculating the true population mean μ .

`statistics.harmonic_mean(data)`

Return the harmonic mean of *data*, a sequence or iterator of real-valued numbers.

The harmonic mean, sometimes called the subcontrary mean, is the reciprocal of the arithmetic `mean()` of the reciprocals of the data. For example, the harmonic mean of three values *a*, *b* and *c* will be equivalent to $3 / (1/a + 1/b + 1/c)$.

The harmonic mean is a type of average, a measure of the central location of the data. It is often appropriate when averaging quantities which are rates or ratios, for example speeds. For example:

Suppose an investor purchases an equal value of shares in each of three companies, with P/E (price/earning) ratios of 2.5, 3 and 10. What is the average P/E ratio for the investor's portfolio?

```
>>> harmonic_mean([2.5, 3, 10]) # For an equal investment portfolio.
3.6
```

Using the arithmetic mean would give an average of about 5.167, which is too high.

`StatisticsError` is raised if *data* is empty, or any element is less than zero.

3.6 新版功能.

`statistics.median(data)`

Return the median (middle value) of numeric data, using the common “mean of middle two” method. If *data* is empty, `StatisticsError` is raised. *data* can be a sequence or iterator.

The median is a robust measure of central location, and is less affected by the presence of outliers in your data. When the number of data points is odd, the middle data point is returned:

```
>>> median([1, 3, 5])
3
```

When the number of data points is even, the median is interpolated by taking the average of the two middle values:

```
>>> median([1, 3, 5, 7])
4.0
```

This is suited for when your data is discrete, and you don't mind that the median may not be an actual data point.

If your data is ordinal (supports order operations) but not numeric (doesn't support addition), you should use `median_low()` or `median_high()` instead.

参见:

`median_low()`, `median_high()`, `median_grouped()`

`statistics.median_low(data)`

Return the low median of numeric data. If *data* is empty, `StatisticsError` is raised. *data* can be a sequence or iterator.

The low median is always a member of the data set. When the number of data points is odd, the middle value is returned. When it is even, the smaller of the two middle values is returned.

```
>>> median_low([1, 3, 5])
3
>>> median_low([1, 3, 5, 7])
3
```

Use the low median when your data are discrete and you prefer the median to be an actual data point rather than interpolated.

`statistics.median_high(data)`

Return the high median of data. If *data* is empty, *StatisticsError* is raised. *data* can be a sequence or iterator.

The high median is always a member of the data set. When the number of data points is odd, the middle value is returned. When it is even, the larger of the two middle values is returned.

```
>>> median_high([1, 3, 5])
3
>>> median_high([1, 3, 5, 7])
5
```

Use the high median when your data are discrete and you prefer the median to be an actual data point rather than interpolated.

`statistics.median_grouped(data, interval=1)`

Return the median of grouped continuous data, calculated as the 50th percentile, using interpolation. If *data* is empty, *StatisticsError* is raised. *data* can be a sequence or iterator.

```
>>> median_grouped([52, 52, 53, 54])
52.5
```

In the following example, the data are rounded, so that each value represents the midpoint of data classes, e.g. 1 is the midpoint of the class 0.5–1.5, 2 is the midpoint of 1.5–2.5, 3 is the midpoint of 2.5–3.5, etc. With the data given, the middle value falls somewhere in the class 3.5–4.5, and interpolation is used to estimate it:

```
>>> median_grouped([1, 2, 2, 3, 4, 4, 4, 4, 4, 5])
3.7
```

Optional argument *interval* represents the class interval, and defaults to 1. Changing the class interval naturally will change the interpolation:

```
>>> median_grouped([1, 3, 3, 5, 7], interval=1)
3.25
>>> median_grouped([1, 3, 3, 5, 7], interval=2)
3.5
```

This function does not check whether the data points are at least *interval* apart.

CPython implementation detail: Under some circumstances, *median_grouped()* may coerce data points to floats. This behaviour is likely to change in the future.

参见:

- “Statistics for the Behavioral Sciences” , Frederick J Gravetter and Larry B Wallnau (8th Edition).
- Calculating the [median](#).
- The [SSMEDIAN](#) function in the Gnome Gnumeric spreadsheet, including [this discussion](#).

`statistics.mode(data)`

Return the most common data point from discrete or nominal *data*. The mode (when it exists) is the most typical value, and is a robust measure of central location.

If *data* is empty, or if there is not exactly one most common value, `StatisticsError` is raised.

mode assumes discrete data, and returns a single value. This is the standard treatment of the mode as commonly taught in schools:

```
>>> mode([1, 1, 2, 3, 3, 3, 3, 4])
3
```

The mode is unique in that it is the only statistic which also applies to nominal (non-numeric) data:

```
>>> mode(["red", "blue", "blue", "red", "green", "red", "red"])
'red'
```

`statistics.pstdev(data, mu=None)`

Return the population standard deviation (the square root of the population variance). See `pvariance()` for arguments and other details.

```
>>> pstdev([1.5, 2.5, 2.5, 2.75, 3.25, 4.75])
0.986893273527251
```

`statistics.pvariance(data, mu=None)`

Return the population variance of *data*, a non-empty iterable of real-valued numbers. Variance, or second moment about the mean, is a measure of the variability (spread or dispersion) of data. A large variance indicates that the data is spread out; a small variance indicates it is clustered closely around the mean.

If the optional second argument *mu* is given, it should be the mean of *data*. If it is missing or `None` (the default), the mean is automatically calculated.

Use this function to calculate the variance from the entire population. To estimate the variance from a sample, the `variance()` function is usually a better choice.

Raises `StatisticsError` if *data* is empty.

示例:

```
>>> data = [0.0, 0.25, 0.25, 1.25, 1.5, 1.75, 2.75, 3.25]
>>> pvariance(data)
1.25
```

If you have already calculated the mean of your data, you can pass it as the optional second argument *mu* to avoid recalculation:

```
>>> mu = mean(data)
>>> pvariance(data, mu)
1.25
```

This function does not attempt to verify that you have passed the actual mean as *mu*. Using arbitrary values for *mu* may lead to invalid or impossible results.

Decimals and Fractions are supported:

```
>>> from decimal import Decimal as D
>>> pvariance([D("27.5"), D("30.25"), D("30.25"), D("34.5"), D("41.75")])
Decimal('24.815')
```

(下页继续)

(续上页)

```
>>> from fractions import Fraction as F
>>> pvariance([F(1, 4), F(5, 4), F(1, 2)])
Fraction(13, 72)
```

注解： When called with the entire population, this gives the population variance σ^2 . When called on a sample instead, this is the biased sample variance s^2 , also known as variance with N degrees of freedom.

If you somehow know the true population mean μ , you may use this function to calculate the variance of a sample, giving the known population mean as the second argument. Provided the data points are representative (e.g. independent and identically distributed), the result will be an unbiased estimate of the population variance.

`statistics.stdev(data, xbar=None)`

Return the sample standard deviation (the square root of the sample variance). See `variance()` for arguments and other details.

```
>>> stdev([1.5, 2.5, 2.5, 2.75, 3.25, 4.75])
1.0810874155219827
```

`statistics.variance(data, xbar=None)`

Return the sample variance of *data*, an iterable of at least two real-valued numbers. Variance, or second moment about the mean, is a measure of the variability (spread or dispersion) of data. A large variance indicates that the data is spread out; a small variance indicates it is clustered closely around the mean.

If the optional second argument *xbar* is given, it should be the mean of *data*. If it is missing or `None` (the default), the mean is automatically calculated.

Use this function when your data is a sample from a population. To calculate the variance from the entire population, see `pvariance()`.

Raises `StatisticsError` if *data* has fewer than two values.

示例：

```
>>> data = [2.75, 1.75, 1.25, 0.25, 0.5, 1.25, 3.5]
>>> variance(data)
1.3720238095238095
```

If you have already calculated the mean of your data, you can pass it as the optional second argument *xbar* to avoid recalculation:

```
>>> m = mean(data)
>>> variance(data, m)
1.3720238095238095
```

This function does not attempt to verify that you have passed the actual mean as *xbar*. Using arbitrary values for *xbar* can lead to invalid or impossible results.

Decimal and Fraction values are supported:

```
>>> from decimal import Decimal as D
>>> variance([D("27.5"), D("30.25"), D("30.25"), D("34.5"), D("41.75")])
Decimal('31.01875')

>>> from fractions import Fraction as F
>>> variance([F(1, 6), F(1, 2), F(5, 3)])
Fraction(67, 108)
```

注解: This is the sample variance s^2 with Bessel's correction, also known as variance with N-1 degrees of freedom. Provided that the data points are representative (e.g. independent and identically distributed), the result should be an unbiased estimate of the true population variance.

If you somehow know the actual population mean μ you should pass it to the `pvariance()` function as the *mu* parameter to get the variance of a sample.

9.7.4 异常

A single exception is defined:

exception `statistics.StatisticsError`

Subclass of `ValueError` for statistics-related exceptions.

本章里描述的模块提供了函数和类，以支持函数式编程风格和在可调用对象上的通用操作。

本章对下列模块进行说明：

10.1 `itertools` — 为高效循环而创建迭代器的函数

本模块实现一系列 *iterator*，这些迭代器受到 APL，Haskell 和 SML 的启发。为了适用于 Python，它们都被重新写过。

本模块标准化了一个快速、高效利用内存的核心工具集，这些工具本身或组合都很有用。它们一起形成了“迭代器代数”，这使得在纯 Python 中有可能创建简洁又高效的专用工具。

例如，SML 有一个制表工具： `tabulate(f)`，它可产生一个序列 `f(0)`， `f(1)`，...。在 Python 中可以组合 `map()` 和 `count()` 实现： `map(f, count())`。

这些内置工具同时也能很好地与 *operator* 模块中的高效函数配合使用。例如，我们可以将两个向量的点积映射到乘法运算符： `sum(map(operator.mul, vector1, vector2))`。

无穷迭代器：

迭代器	实参	结果	示例
<code>count()</code>	start, [step]	start, start+step, start+2*step, ...	<code>count(10) --> 10 11 12 13 14 ...</code>
<code>cycle()</code>	p	p0, p1, ...,plast, p0, p1, ...	<code>cycle('ABCD') --> A B C D A B C D ...</code>
<code>repeat()</code>	elem [,n]	elem, elem, elem, ... 重复无限次或 n 次	<code>repeat(10, 3) --> 10 10 10</code>

根据最短输入序列长度停止的迭代器：

迭代器	实参	结果	示例
<code>accumulate()</code>	<code>p [,func]</code>	<code>p0, p0+p1, p0+p1+p2, ...</code>	<code>accumulate([1,2,3,4,5]) --> 1 3 6 10 15</code>
<code>chain()</code>	<code>p, q, ...</code>	<code>p0, p1, ...plast, q0, q1, ...</code>	<code>chain('ABC', 'DEF') --> A B C D E F</code>
<code>chain.from_iterable()</code>	iterable	<code>p0, p1, ...plast, q0, q1, ...</code>	<code>chain.from_iterable(['ABC', 'DEF']) --> A B C D E F</code>
<code>compress()</code>	<code>data, selectors</code>	<code>(d[0] if s[0]), (d[1] if s[1]), ...</code>	<code>compress('ABCDEF', [1,0,1,0,1,1]) --> A C E F</code>
<code>dropwhile()</code>	<code>pred, seq</code>	<code>seq[n], seq[n+1], ...</code> 从 <code>pred</code> 首次真值测试失败开始	<code>dropwhile(lambda x: x<5, [1,4,6,4,1]) --> 6 4 1</code>
<code>filterfalse()</code>	<code>pred, seq</code>	<code>seq</code> 中 <code>pred(x)</code> 为假值的元素, <code>x</code> 是 <code>seq</code> 中的元素。	<code>filterfalse(lambda x: x%2, range(10)) --> 0 2 4 6 8</code>
<code>groupby()</code>	iterable[, key]	根据 <code>key(v)</code> 值分组的迭代器	
<code>islice()</code>	<code>seq, [start,] stop [, step]</code>	<code>seq[start:stop:step]</code> 中的元素	<code>islice('ABCDEFG', 2, None) --> C D E F G</code>
<code>starmap()</code>	<code>func, seq</code>	<code>func(*seq[0]), func(*seq[1]), ...</code>	<code>starmap(pow, [(2,5), (3,2), (10,3)]) --> 32 9 1000</code>
<code>takewhile()</code>	<code>pred, seq</code>	<code>seq[0], seq[1], ...</code> , 直到 <code>pred</code> 真值测试失败	<code>takewhile(lambda x: x<5, [1,4,6,4,1]) --> 1 4</code>
<code>tee()</code>	<code>it, n</code>	<code>it1, it2, ...itn</code> 将一个迭代器拆分为 <code>n</code> 个迭代器	
<code>zip_longest()</code>	<code>p, q, ...</code>	<code>(p[0], q[0]), (p[1], q[1]), ...</code>	<code>zip_longest('ABCD', 'xy', fillvalue='-') --> Ax By C- D-</code>

排列组合迭代器:

迭代器	实参	结果
<code>product()</code>	<code>p, q, ...[repeat=1]</code>	笛卡尔积, 相当于嵌套的 <code>for</code> 循环
<code>permutations()</code>	<code>p[, r]</code>	长度 <code>r</code> 元组, 所有可能的排列, 无重复元素
<code>combinations()</code>	<code>p, r</code>	长度 <code>r</code> 元组, 有序, 无重复元素
<code>combinations_with_replacement()</code>	<code>p, r</code>	长度 <code>r</code> 元组, 有序, 元素可重复
<code>product('ABCD', repeat=2)</code>		AA AB AC AD BA BB BC BD CA CB CC CD DA DB DC DD
<code>permutations('ABCD', 2)</code>		AB AC AD BA BC BD CA CB CD DA DB DC
<code>combinations('ABCD', 2)</code>		AB AC AD BC BD CD
<code>combinations_with_replacement('ABCD', 2)</code>		AA AB AC AD BB BC BD CC CD DD

10.1.1 Itertool 函数

下列模块函数均创建并返回迭代器。有些迭代器不限制输出流长度, 所以它们只应在能截断输出流的函数或循环中使用。

`itertools.accumulate(iterable[, func])`

创建一个迭代器, 返回累加和或其他二元函数的累加结果 (通过可选参数 `func` 指定)。如果提供了 `func`, 它应是 2 个参数的函数。输入 `iterable` 元素类型应是 `func` 能支持的任意类型。(例如, 对于默认

法操作，元素可以是任一支持加法的类型，包括`Decimal` 或 `Fraction`)。如果可迭代对象的输入为空，输出也为空。

大致相当于：

```
def accumulate(iterable, func=operator.add):
    'Return running totals'
    # accumulate([1,2,3,4,5]) --> 1 3 6 10 15
    # accumulate([1,2,3,4,5], operator.mul) --> 1 2 6 24 120
    it = iter(iterable)
    try:
        total = next(it)
    except StopIteration:
        return
    yield total
    for element in it:
        total = func(total, element)
        yield total
```

`func` 参数有几种用法。它可以被设为`min()` 最终得到一个最小值，或者设为`max()` 最终得到一个最大值，或设为`operator.mul()` 最终得到一个乘积。摊销表可通过累加利息和支付款项得到。给 `iterable` 设置初始值并只将参数 `func` 设为累加总数可以对一阶 递归关系 建模。

```
>>> data = [3, 4, 6, 2, 1, 9, 0, 7, 5, 8]
>>> list(accumulate(data, operator.mul))      # running product
[3, 12, 72, 144, 144, 1296, 0, 0, 0, 0]
>>> list(accumulate(data, max))              # running maximum
[3, 4, 6, 6, 6, 9, 9, 9, 9, 9]

# Amortize a 5% loan of 1000 with 4 annual payments of 90
>>> cashflows = [1000, -90, -90, -90, -90]
>>> list(accumulate(cashflows, lambda bal, pmt: bal*1.05 + pmt))
[1000, 960.0, 918.0, 873.9000000000001, 827.5950000000001]

# Chaotic recurrence relation https://en.wikipedia.org/wiki/Logistic_map
>>> logistic_map = lambda x, _: r * x * (1 - x)
>>> r = 3.8
>>> x0 = 0.4
>>> inputs = repeat(x0, 36)                  # only the initial value is used
>>> [format(x, '.2f') for x in accumulate(inputs, logistic_map)]
['0.40', '0.91', '0.30', '0.81', '0.60', '0.92', '0.29', '0.79', '0.63',
 '0.88', '0.39', '0.90', '0.33', '0.84', '0.52', '0.95', '0.18', '0.57',
 '0.93', '0.25', '0.71', '0.79', '0.63', '0.88', '0.39', '0.91', '0.32',
 '0.83', '0.54', '0.95', '0.20', '0.60', '0.91', '0.30', '0.80', '0.60']
```

参考一个类似函数`functools.reduce()`，它只返回一个最终累积值。

3.2 新版功能.

在 3.3 版更改: 增加可选参数 `func`。

`itertools.chain(*iterables)`

创建一个迭代器，它首先返回第一个可迭代对象中所有元素，接着返回下一个可迭代对象中所有元素，直到耗尽所有可迭代对象中的元素。可将多个序列处理为单个序列。大致相当于：

```
def chain(*iterables):
    # chain('ABC', 'DEF') --> A B C D E F
    for it in iterables:
```

(下页继续)

(续上页)

```

for element in it:
    yield element

```

classmethod `chain.from_iterable(iterable)`

构建类似`chain()` 迭代器的另一个选择。从一个单独的可迭代参数中得到链式输入，该参数是延迟计算的。大致相当于：

```

def from_iterable(iterables):
    # chain.from_iterable(['ABC', 'DEF']) --> A B C D E F
    for it in iterables:
        for element in it:
            yield element

```

itertools.combinations(iterable, r)

返回由输入 *iterable* 中元素组成长度为 *r* 的子序列。

组合按照字典序返回。所以如果输入 *iterable* 是有序的，生成的组合元组也是有序的。

即使元素的值相同，不同位置的元素也被认为是不同的。如果元素各自不同，那么每个组合中没有重复元素。

大致相当于：

```

def combinations(iterable, r):
    # combinations('ABCD', 2) --> AB AC AD BC BD CD
    # combinations(range(4), 3) --> 012 013 023 123
    pool = tuple(iterable)
    n = len(pool)
    if r > n:
        return
    indices = list(range(r))
    yield tuple(pool[i] for i in indices)
    while True:
        for i in reversed(range(r)):
            if indices[i] != i + n - r:
                break
        else:
            return
        indices[i] += 1
        for j in range(i+1, r):
            indices[j] = indices[j-1] + 1
        yield tuple(pool[i] for i in indices)

```

`combinations()` 的代码可被改写为 `permutations()` 过滤后的子序列，（相对于元素在输入中的位置）元素不是有序的。

```

def combinations(iterable, r):
    pool = tuple(iterable)
    n = len(pool)
    for indices in permutations(range(n), r):
        if sorted(indices) == list(indices):
            yield tuple(pool[i] for i in indices)

```

当 $0 \leq r \leq n$ 时，返回项的个数是 $n! / r! / (n-r)!$ ；当 $r > n$ 时，返回项个数为 0。

itertools.combinations_with_replacement(iterable, r)

返回由输入 *iterable* 中元素组成的长度为 *r* 的子序列，允许每个元素可重复出现。

组合按照字典序返回。所以如果输入 *iterable* 是有序的，生成的组合元组也是有序的。

不同位置的元素是不同的，即使它们的值相同。因此如果输入中的元素都是不同的话，返回的组合中元素也都会不同。

大致相当于：

```
def combinations_with_replacement(iterable, r):
    # combinations_with_replacement('ABC', 2) --> AA AB AC BB BC CC
    pool = tuple(iterable)
    n = len(pool)
    if not n and r:
        return
    indices = [0] * r
    yield tuple(pool[i] for i in indices)
    while True:
        for i in reversed(range(r)):
            if indices[i] != n - 1:
                break
        else:
            return
        indices[i:] = [indices[i] + 1] * (r - i)
        yield tuple(pool[i] for i in indices)
```

`combinations_with_replacement()` 的代码可被改写为 `production()` 过滤后的子序列，（相对于元素在输入中的位置）元素不是有序的。

```
def combinations_with_replacement(iterable, r):
    pool = tuple(iterable)
    n = len(pool)
    for indices in product(range(n), repeat=r):
        if sorted(indices) == list(indices):
            yield tuple(pool[i] for i in indices)
```

当 $n > 0$ 时，返回项个数为 $(n+r-1)! / r! / (n-1)!$ 。

3.1 新版功能.

`itertools.compress(data, selectors)`

创建一个迭代器，它返回 `data` 中经 `selectors` 真值测试为 `True` 的元素。迭代器在两者较短的长度处停止。大致相当于：

```
def compress(data, selectors):
    # compress('ABCDEF', [1,0,1,0,1,1]) --> A C E F
    return (d for d, s in zip(data, selectors) if s)
```

3.1 新版功能.

`itertools.count(start=0, step=1)`

创建一个迭代器，它从 `start` 值开始，返回均匀间隔的值。常用于 `map()` 中的实参来生成连续的数据点。此外，还用于 `zip()` 来添加序列号。大致相当于：

```
def count(start=0, step=1):
    # count(10) --> 10 11 12 13 14 ...
    # count(2.5, 0.5) -> 2.5 3.0 3.5 ...
    n = start
    while True:
        yield n
        n += step
```

当对浮点数计数时，替换为乘法代码有时精度会更好，例如：`(start + step * i for i in count())`。

在 3.1 版更改: 增加参数 *step* , 允许非整型。

`itertools.cycle(iterable)`

创建一个迭代器, 返回 *iterable* 中所有元素并保存一个副本。当取完 *iterable* 中所有元素, 返回副本中的所有元素。无限重复。大致相当于:

```
def cycle(iterable):
    # cycle('ABCD') --> A B C D A B C D A B C D ...
    saved = []
    for element in iterable:
        yield element
        saved.append(element)
    while saved:
        for element in saved:
            yield element
```

注意, 该函数可能需要相当大的辅助空间 (取决于 *iterable* 的长度)。

`itertools.dropwhile(predicate, iterable)`

创建一个迭代器, 如果 *predicate* 为 `true`, 迭代器丢弃这些元素, 然后返回其他元素。注意, 迭代器在 *predicate* 首次为 `false` 之前不会产生任何输出, 所以可能需要一定长度的启动时间。大致相当于:

```
def dropwhile(predicate, iterable):
    # dropwhile(lambda x: x<5, [1,4,6,4,1]) --> 6 4 1
    iterable = iter(iterable)
    for x in iterable:
        if not predicate(x):
            yield x
            break
    for x in iterable:
        yield x
```

`itertools.filterfalse(predicate, iterable)`

创建一个迭代器, 只返回 *iterable* 中 *predicate* 为 `False` 的元素。如果 *predicate* 是 `None`, 返回真值测试为 `false` 的元素。大致相当于:

```
def filterfalse(predicate, iterable):
    # filterfalse(lambda x: x%2, range(10)) --> 0 2 4 6 8
    if predicate is None:
        predicate = bool
    for x in iterable:
        if not predicate(x):
            yield x
```

`itertools.groupby(iterable, key=None)`

创建一个迭代器, 返回 *iterable* 中连续的键和组。*key* 是一个计算元素键值函数。如果未指定或为 `None`, *key* 缺省为恒等函数 (identity function), 返回元素不变。一般来说, *iterable* 需用同一个键值函数预先排序。

`groupby()` 操作类似于 Unix 中的 `uniq`。当每次 *key* 函数产生的键值改变时, 迭代器会分组或生成一个新组 (这就是为什么通常需要使用同一个键值函数先对数据进行排序)。这种行为与 SQL 的 `GROUP BY` 操作不同, SQL 的操作会忽略输入的顺序将相同键值的元素分在同组中。

返回的组本身也是一个迭代器, 它与 `groupby()` 共享底层的可迭代对象。因为源是共享的, 当 `groupby()` 对象向后迭代时, 前一个组将消失。因此如果稍后还需要返回结果, 可保存为列表:

```

groups = []
uniquekeys = []
data = sorted(data, key=keyfunc)
for k, g in groupby(data, keyfunc):
    groups.append(list(g))      # Store group iterator as a list
    uniquekeys.append(k)

```

`groupby()` 大致相当于:

```

class groupby:
    # [k for k, g in groupby('AAAABBBCCDAABBB')] --> A B C D A B
    # [list(g) for k, g in groupby('AAAABBBCCD')] --> AAAA BBB CC D
    def __init__(self, iterable, key=None):
        if key is None:
            key = lambda x: x
        self.keyfunc = key
        self.it = iter(iterable)
        self.tgtkey = self.currkey = self.currvalue = object()
    def __iter__(self):
        return self
    def __next__(self):
        while self.currkey == self.tgtkey:
            self.currvalue = next(self.it)      # Exit on StopIteration
            self.currkey = self.keyfunc(self.currvalue)
        self.tgtkey = self.currkey
        return (self.currkey, self._grouper(self.tgtkey))
    def _grouper(self, tgtkey):
        while self.currkey == tgtkey:
            yield self.currvalue
            try:
                self.currvalue = next(self.it)
            except StopIteration:
                return
            self.currkey = self.keyfunc(self.currvalue)

```

`itertools.islice(iterable, stop)`

`itertools.islice(iterable, start, stop[, step])`

创建一个迭代器，返回从 `iterable` 里选中的元素。如果 `start` 不是 0，跳过 `iterable` 中的元素，直到到达 `start` 这个位置。之后迭代器连续返回元素，除非 `step` 设置的值很高导致被跳过。如果 `stop` 为 `None`，迭代器耗光为止；否则，在指定的位置停止。与普通的切片不同，`islice()` 不支持将 `start`，`stop`，或 `step` 设为负值。可用来从内部数据结构被压平的数据中提取相关字段（例如一个多行报告，它的名称字段出现在每三行上）。大致相当于：

```

def islice(iterable, *args):
    # islice('ABCDEFGH', 2) --> A B
    # islice('ABCDEFGH', 2, 4) --> C D
    # islice('ABCDEFGH', 2, None) --> C D E F G
    # islice('ABCDEFGH', 0, None, 2) --> A C E G
    s = slice(*args)
    start, stop, step = s.start or 0, s.stop or sys.maxsize, s.step or 1
    it = iter(range(start, stop, step))
    try:
        nexti = next(it)
    except StopIteration:
        # Consume *iterable* up to the *start* position.
        for i, element in zip(range(start), iterable):
            pass

```

(下页继续)

(续上页)

```

    return
try:
    for i, element in enumerate(iterable):
        if i == nexti:
            yield element
            nexti = next(it)
except StopIteration:
    # Consume to *stop*.
    for i, element in zip(range(i + 1, stop), iterable):
        pass

```

如果 *start* 为 `None`，迭代从 0 开始。如果 *step* 为 `None`，步长缺省为 1。

`itertools.permutations(iterable, r=None)`

连续返回由 *iterable* 元素生成长度为 *r* 的排列。

如果 *r* 未指定或为 `None`，*r* 默认设置为 *iterable* 的长度，这种情况下，生成所有全长排列。

排列依字典序发出。因此，如果 *iterable* 是已排序的，排列元组将有序地产出。

即使元素的值相同，不同位置的元素也被认为是不同的。如果元素值都不同，每个排列中的元素值不会重复。

大致相当于：

```

def permutations(iterable, r=None):
    # permutations('ABCD', 2) --> AB AC AD BA BC BD CA CB CD DA DB DC
    # permutations(range(3)) --> 012 021 102 120 201 210
    pool = tuple(iterable)
    n = len(pool)
    r = n if r is None else r
    if r > n:
        return
    indices = list(range(n))
    cycles = list(range(n, n-r, -1))
    yield tuple(pool[i] for i in indices[:r])
    while n:
        for i in reversed(range(r)):
            cycles[i] -= 1
            if cycles[i] == 0:
                indices[i:] = indices[i+1:] + indices[i:i+1]
                cycles[i] = n - i
            else:
                j = cycles[i]
                indices[i], indices[-j] = indices[-j], indices[i]
                yield tuple(pool[i] for i in indices[:r])
                break
        else:
            return

```

`permutations()` 的代码也可被改写为 `product()` 的子序列，只要将含有重复元素（来自输入中同一位置的）的项排除。

```

def permutations(iterable, r=None):
    pool = tuple(iterable)
    n = len(pool)
    r = n if r is None else r
    for indices in product(range(n), repeat=r):

```

(下页继续)

(续上页)

```
if len(set(indices)) == r:
    yield tuple(pool[i] for i in indices)
```

当 $0 \leq r \leq n$ ，返回项个数为 $n! / (n-r)!$ ；当 $r > n$ ，返回项个数为 0。

`itertools.product(*iterables, repeat=1)`

可迭代对象输入的笛卡儿积。

大致相当于生成器表达式中的嵌套循环。例如，`product(A, B)` 和 `((x,y) for x in A for y in B)` 返回结果一样。

嵌套循环像里程表那样循环变动，每次迭代时将最右侧的元素向后迭代。这种模式形成了一种字典序，因此如果输入的可迭代对象是已排序的，笛卡尔积元组依次序发出。

要计算可迭代对象自身的笛卡尔积，将可选参数 `repeat` 设定为要重复的次数。例如，`product(A, repeat=4)` 和 `product(A, A, A, A)` 是一样的。

该函数大致相当于下面的代码，只不过实际实现方案不会在内存中创建中间结果。

```
def product(*args, repeat=1):
    # product('ABCD', 'xy') --> Ax Ay Bx By Cx Cy Dx Dy
    # product(range(2), repeat=3) --> 000 001 010 011 100 101 110 111
    pools = [tuple(pool) for pool in args] * repeat
    result = [[]]
    for pool in pools:
        result = [x+[y] for x in result for y in pool]
    for prod in result:
        yield tuple(prod)
```

`itertools.repeat(object[, times])`

创建一个迭代器，不断重复 `object`。除非设定参数 `times`，否则将无限重复。可用于 `map()` 函数中的参数，被调用函数可得到一个不变参数。也可用于 `zip()` 的参数以在元组记录中创建一个不变的部分。

大致相当于：

```
def repeat(object, times=None):
    # repeat(10, 3) --> 10 10 10
    if times is None:
        while True:
            yield object
    else:
        for i in range(times):
            yield object
```

`repeat` 最常见的用途就是在 `map` 或 `zip` 提供一个常量流：

```
>>> list(map(pow, range(10), repeat(2)))
[0, 1, 4, 9, 16, 25, 36, 49, 64, 81]
```

`itertools.starmap(function, iterable)`

创建一个迭代器，使用从可迭代对象中获取的参数来计算该函数。当参数对应的形参已从一个单独可迭代对象组合为元组时（数据已被“预组对”）可用此函数代替 `map()`。`map()` 与 `starmap()` 之间的区别可以类比 `function(a,b)` 与 `function(*c)` 的区别。大致相当于：

```
def starmap(function, iterable):
    # starmap(pow, [(2,5), (3,2), (10,3)]) --> 32 9 1000
    for args in iterable:
        yield function(*args)
```

`itertools takewhile` (*predicate, iterable*)

创建一个迭代器，只要 `predicate` 为真就从可迭代对象中返回元素。大致相当于：

```
def takewhile(predicate, iterable):
    # takewhile(lambda x: x<5, [1,4,6,4,1]) --> 1 4
    for x in iterable:
        if predicate(x):
            yield x
        else:
            break
```

`itertools tee` (*iterable, n=2*)

从一个可迭代对象中返回 *n* 个独立的迭代器。

下面的 Python 代码能帮助解释 *tee* 做了什么（尽管实际的实现更复杂，而且仅使用了一个底层的 FIFO 队列）。

大致相当于：

```
def tee(iterable, n=2):
    it = iter(iterable)
    deque = [collections.deque() for i in range(n)]
    def gen(mydeque):
        while True:
            if not mydeque:
                # when the local deque is empty
                try:
                    newval = next(it)
                    # fetch a new value and
                except StopIteration:
                    return
                for d in deque:
                    # load it to all the deque
                    d.append(newval)
            yield mydeque.popleft()
    return tuple(gen(d) for d in deque)
```

一旦 *tee()* 实施了一次分裂，原有的 *iterable* 不应再被使用；否则 *tee* 对象无法得知 *iterable* 可能已向后迭代。

该迭代工具可能需要相当大的辅助存储空间（这取决于要保存多少临时数据）。通常，如果一个迭代器在另一个迭代器开始之前就要使用大部份或全部数据，使用 *list()* 会比 *tee()* 更快。

`itertools zip_longest` (**iterables, fillvalue=None*)

创建一个迭代器，从每个可迭代对象中收集元素。如果可迭代对象的长度未对齐，将根据 *fillvalue* 填充缺失值。迭代持续到耗光最长的可迭代对象。大致相当于：

```
class ZipExhausted(Exception):
    pass

def zip_longest(*args, **kwargs):
    # zip_longest('ABCD', 'xy', fillvalue='-') --> Ax By C- D-
    fillvalue = kwargs.get('fillvalue')
    counter = len(args) - 1
    def sentinel():
        nonlocal counter
        if not counter:
            raise ZipExhausted
        counter -= 1
        yield fillvalue
    fillers = repeat(fillvalue)
    iterators = [chain(it, sentinel(), fillers) for it in args]
```

(下页继续)

(续上页)

```

try:
    while iterators:
        yield tuple(map(next, iterators))
except ZipExhausted:
    pass

```

如果其中一个可迭代对象有无限长度，`zip_longest()` 函数应封装在限制调用次数的场景中（例如 `islice()` 或 `takewhile()`）。除非指定，`fillvalue` 默认为 `None`。

10.1.2 Itertools 食谱

本节将展示如何使用现有的 `itertools` 作为基础构件来创建扩展的工具集。

扩展的工具提供了与底层工具集相同的高性能。保持了超棒的内存利用率，因为一次只处理一个元素，而不是将整个可迭代对象加载到内存。代码量保持得很小，以函数式风格将这些工具连接在一起，有助于消除临时变量。速度依然很快，因为倾向于使用“矢量化”构件来取代解释器开销大的 `for` 循环和 *generator*。

```

def take(n, iterable):
    "Return first n items of the iterable as a list"
    return list(islice(iterable, n))

def prepend(value, iterator):
    "Prepend a single value in front of an iterator"
    # prepend(1, [2, 3, 4]) -> 1 2 3 4
    return chain([value], iterator)

def tabulate(function, start=0):
    "Return function(0), function(1), ..."
    return map(function, count(start))

def tail(n, iterable):
    "Return an iterator over the last n items"
    # tail(3, 'ABCDEFG') --> E F G
    return iter(collections.deque(iterable, maxlen=n))

def consume(iterator, n=None):
    "Advance the iterator n-steps ahead. If n is None, consume entirely."
    # Use functions that consume iterators at C speed.
    if n is None:
        # feed the entire iterator into a zero-length deque
        collections.deque(iterator, maxlen=0)
    else:
        # advance to the empty slice starting at position n
        next(islice(iterator, n, n), None)

def nth(iterable, n, default=None):
    "Returns the nth item or a default value"
    return next(islice(iterable, n, None), default)

def all_equal(iterable):
    "Returns True if all the elements are equal to each other"
    g = groupby(iterable)
    return next(g, True) and not next(g, False)

def quantify(iterable, pred=bool):

```

(下页继续)

(续上页)

```

    "Count how many times the predicate is true"
    return sum(map(pred, iterable))

def padnone(iterable):
    """Returns the sequence elements and then returns None indefinitely.

    Useful for emulating the behavior of the built-in map() function.
    """
    return chain(iterable, repeat(None))

def ncycles(iterable, n):
    "Returns the sequence elements n times"
    return chain.from_iterable(repeat(tuple(iterable), n))

def dotproduct(vec1, vec2):
    return sum(map(operator.mul, vec1, vec2))

def flatten(listOfLists):
    "Flatten one level of nesting"
    return chain.from_iterable(listOfLists)

def repeatfunc(func, times=None, *args):
    """Repeat calls to func with specified arguments.

    Example:  repeatfunc(random.random)
    """
    if times is None:
        return starmap(func, repeat(args))
    return starmap(func, repeat(args, times))

def pairwise(iterable):
    "s -> (s0,s1), (s1,s2), (s2, s3), ..."
    a, b = tee(iterable)
    next(b, None)
    return zip(a, b)

def grouper(iterable, n, fillvalue=None):
    "Collect data into fixed-length chunks or blocks"
    # grouper('ABCDEFG', 3, 'x') --> ABC DEF Gxx"
    args = [iter(iterable)] * n
    return zip_longest(*args, fillvalue=fillvalue)

def roundrobin(*iterables):
    "roundrobin('ABC', 'D', 'EF') --> A D E B F C"
    # Recipe credited to George Sakkis
    num_active = len(iterables)
    nexts = cycle(iter(it).__next__ for it in iterables)
    while num_active:
        try:
            for next in nexts:
                yield next()
        except StopIteration:
            # Remove the iterator we just exhausted from the cycle.
            num_active -= 1
            nexts = cycle(islice(nexts, num_active))

```

(下页继续)

(续上页)

```

def partition(pred, iterable):
    'Use a predicate to partition entries into false entries and true entries'
    # partition(is_odd, range(10)) --> 0 2 4 6 8   and   1 3 5 7 9
    t1, t2 = tee(iterable)
    return filterfalse(pred, t1), filter(pred, t2)

def powerset(iterable):
    "powerset([1,2,3]) --> () (1,) (2,) (3,) (1,2) (1,3) (2,3) (1,2,3)"
    s = list(iterable)
    return chain.from_iterable(combinations(s, r) for r in range(len(s)+1))

def unique_everseen(iterable, key=None):
    "List unique elements, preserving order. Remember all elements ever seen."
    # unique_everseen('AAAABBBCCDAABBB') --> A B C D
    # unique_everseen('ABBCcAD', str.lower) --> A B C D
    seen = set()
    seen_add = seen.add
    if key is None:
        for element in filterfalse(seen.__contains__, iterable):
            seen_add(element)
            yield element
    else:
        for element in iterable:
            k = key(element)
            if k not in seen:
                seen_add(k)
                yield element

def unique_justseen(iterable, key=None):
    "List unique elements, preserving order. Remember only the element just seen."
    # unique_justseen('AAAABBBCCDAABBB') --> A B C D A B
    # unique_justseen('ABBCcAD', str.lower) --> A B C A D
    return map(next, map(itemgetter(1), groupby(iterable, key)))

def iter_except(func, exception, first=None):
    """ Call a function repeatedly until an exception is raised.

    Converts a call-until-exception interface to an iterator interface.
    Like builtins.iter(func, sentinel) but uses an exception instead
    of a sentinel to end the loop.

    Examples:
        iter_except(functools.partial(heappop, h), IndexError)   # priority queue_
    ↪ iterator
        iter_except(d.popitem, KeyError)                         # non-blocking dict_
    ↪ iterator
        iter_except(d.popleft, IndexError)                       # non-blocking deque_
    ↪ iterator
        iter_except(q.get_nowait, Queue.Empty)                  # loop over a_
    ↪ producer Queue
        iter_except(s.pop, KeyError)                             # non-blocking set_
    ↪ iterator

    """
    try:
        if first is not None:

```

(下页继续)

(续上页)

```

        yield first()                # For database APIs needing an initial cast to
↪db.first()
        while True:
            yield func()
        except exception:
            pass

def first_true(iterable, default=False, pred=None):
    """Returns the first true value in the iterable.

    If no true value is found, returns *default*

    If *pred* is not None, returns the first item
    for which pred(item) is true.

    """
    # first_true([a,b,c], x) --> a or b or c or x
    # first_true([a,b], x, f) --> a if f(a) else b if f(b) else x
    return next(filter(pred, iterable), default)

def random_product(*args, repeat=1):
    "Random selection from itertools.product(*args, **kwargs)"
    pools = [tuple(pool) for pool in args] * repeat
    return tuple(random.choice(pool) for pool in pools)

def random_permutation(iterable, r=None):
    "Random selection from itertools.permutations(iterable, r)"
    pool = tuple(iterable)
    r = len(pool) if r is None else r
    return tuple(random.sample(pool, r))

def random_combination(iterable, r):
    "Random selection from itertools.combinations(iterable, r)"
    pool = tuple(iterable)
    n = len(pool)
    indices = sorted(random.sample(range(n), r))
    return tuple(pool[i] for i in indices)

def random_combination_with_replacement(iterable, r):
    "Random selection from itertools.combinations_with_replacement(iterable, r)"
    pool = tuple(iterable)
    n = len(pool)
    indices = sorted(random.randrange(n) for i in range(r))
    return tuple(pool[i] for i in indices)

def nth_combination(iterable, r, index):
    'Equivalent to list(combinations(iterable, r))[index]'
    pool = tuple(iterable)
    n = len(pool)
    if r < 0 or r > n:
        raise ValueError
    c = 1
    k = min(r, n-r)
    for i in range(1, k+1):
        c = c * (n - k + i) // i
    if index < 0:

```

(下页继续)

(续上页)

```

    index += c
    if index < 0 or index >= c:
        raise IndexError
    result = []
    while r:
        c, n, r = c*r//n, n-1, r-1
        while index >= c:
            index -= c
            c, n = c*(n-r)//n, n-1
        result.append(pool[-1-n])
    return tuple(result)

```

注意，通过将全局查找替换为局部变量的缺省值，上述配方中有很多可以这样优化。例如，*dotproduct* 配方可以这样写：

```

def dotproduct(vec1, vec2, sum=sum, map=map, mul=operator.mul):
    return sum(map(mul, vec1, vec2))

```

10.2 functools — 高阶函数和可调用对象上的操作

源代码: [Lib/functools.py](#)

functools 模块应用于高阶函数，即——参数或（和）返回值为其他函数的函数。通常来说，此模块的功能适用于所有可调用对象。

functools 模块定义了以下函数：

`functools.cmp_to_key(func)`

将（旧式的）比较函数转换为新式的 *key function*。在类似于 `sorted()`，`min()`，`max()`，`heapq.nlargest()`，`heapq.nsmallest()`，`itertools.groupby()` 等函数的 *key* 参数中使用。此函数主要用作将 Python 2 程序转换至新版的转换工具，以保持对比较函数的兼容。

比较函数意为一个可调用对象，该对象接受两个参数并比较它们，结果为小于则返回一个负数，相等则返回零，大于则返回一个正数。*key function* 则是一个接受一个参数，并返回另一个用以排序的值的可调用对象。

示例：

```
sorted(iterable, key=cmp_to_key(locale.strcoll)) # locale-aware sort order
```

有关排序示例和简要排序教程，请参阅 [sortinghowto](#)。

3.2 新版功能。

`@functools.lru_cache(maxsize=128, typed=False)`

一个为函数提供缓存功能的装饰器，缓存 *maxsize* 组传入参数，在下次以相同参数调用时直接返回上一次的结果。用以节约高开销或 I/O 函数的调用时间。

由于使用了字典存储缓存，所以该函数的固定参数和关键字参数必须是可哈希的。

如果 *maxsize* 设置为 `None`，LRU 功能将被禁用且缓存数量无上限。*maxsize* 设置为 2 的幂时可获得最佳性能。

如果 *typed* 设置为 `true`，不同类型的函数参数将被分别缓存。例如，`f(3)` 和 `f(3.0)` 将被视为不同而分别缓存。

为了衡量缓存的有效性以便调整 *maxsize* 形参，被装饰的函数带有一个 `cache_info()` 函数。当调用 `cache_info()` 函数时，返回一个具名元组，包含命中次数 *hits*，未命中次数 *misses*，最大缓存数量 *maxsize* 和当前缓存大小 *currsz*。在多线程环境中，命中数与未命中数是不完全准确的。

该装饰器也提供了一个用于清理/使缓存失效的函数 `cache_clear()`。

原始的未经装饰的函数可以通过 `__wrapped__` 属性访问。它可以用于检查、绕过缓存，或使用不同的缓存再次装饰原始函数。

“最久未使用算法” (LRU) 缓存 在“最近的调用是即将到来的调用的最佳预测因子”时性能最好（比如，新闻服务器上最受欢迎的文章倾向于每天更改）。“缓存大小限制”参数保证缓存不会在长时间运行的进程比如说网站服务器上无限制的增加自身的大小。

一般来说，LRU 缓存只在当你想要重用之前计算的结果时使用。因此，用它缓存具有副作用的函数、需要在每次调用时创建不同、易变的对象的函数或者诸如 `time()` 或 `random()` 之类的不纯函数是没有意义的。

静态 Web 内容的 LRU 缓存示例:

```
@lru_cache(maxsize=32)
def get_pep(num):
    'Retrieve text of a Python Enhancement Proposal'
    resource = 'http://www.python.org/dev/peps/pep-%04d/' % num
    try:
        with urllib.request.urlopen(resource) as s:
            return s.read()
    except urllib.error.HTTPError:
        return 'Not Found'

>>> for n in 8, 290, 308, 320, 8, 218, 320, 279, 289, 320, 9991:
...     pep = get_pep(n)
...     print(n, len(pep))

>>> get_pep.cache_info()
CacheInfo(hits=3, misses=8, maxsize=32, currsz=8)
```

以下是使用缓存通过 动态规划 计算 斐波那契数列 的例子。

```
@lru_cache(maxsize=None)
def fib(n):
    if n < 2:
        return n
    return fib(n-1) + fib(n-2)

>>> [fib(n) for n in range(16)]
[0, 1, 1, 2, 3, 5, 8, 13, 21, 34, 55, 89, 144, 233, 377, 610]

>>> fib.cache_info()
CacheInfo(hits=28, misses=16, maxsize=None, currsz=16)
```

3.2 新版功能.

在 3.3 版更改: 添加 *typed* 选项。

@functools.total_ordering

给定一个声明一个或多个全比较排序方法的类，这个类装饰器实现剩余的方法。这减轻了指定所有可能的全比较操作的工作。

此类必须包含以下方法之一: `__lt__()`、`__le__()`、`__gt__()` 或 `__ge__()`。另外，此类必须支持 `__eq__()` 方法。

例如:

```
@total_ordering
class Student:
    def __is_valid_operand(self, other):
        return (hasattr(other, "lastname") and
                hasattr(other, "firstname"))
    def __eq__(self, other):
        if not self.__is_valid_operand(other):
            return NotImplemented
        return ((self.lastname.lower(), self.firstname.lower()) ==
                (other.lastname.lower(), other.firstname.lower()))
    def __lt__(self, other):
        if not self.__is_valid_operand(other):
            return NotImplemented
        return ((self.lastname.lower(), self.firstname.lower()) <
                (other.lastname.lower(), other.firstname.lower()))
```

注解: While this decorator makes it easy to create well behaved totally ordered types, it *does* come at the cost of slower execution and more complex stack traces for the derived comparison methods. If performance benchmarking indicates this is a bottleneck for a given application, implementing all six rich comparison methods instead is likely to provide an easy speed boost.

3.2 新版功能.

在 3.4 版更改: Returning NotImplemented from the underlying comparison function for unrecognised types is now supported.

`functools.partial(func, *args, **keywords)`

Return a new *partial object* which when called will behave like *func* called with the positional arguments *args* and keyword arguments *keywords*. If more arguments are supplied to the call, they are appended to *args*. If additional keyword arguments are supplied, they extend and override *keywords*. Roughly equivalent to:

```
def partial(func, *args, **keywords):
    def newfunc(*fargs, **fkeywords):
        newkeywords = keywords.copy()
        newkeywords.update(fkeywords)
        return func(*args, *fargs, **newkeywords)
    newfunc.func = func
    newfunc.args = args
    newfunc.keywords = keywords
    return newfunc
```

The *partial()* is used for partial function application which “freezes” some portion of a function’s arguments and/or keywords resulting in a new object with a simplified signature. For example, *partial()* can be used to create a callable that behaves like the *int()* function where the *base* argument defaults to two:

```
>>> from functools import partial
>>> basetwo = partial(int, base=2)
>>> basetwo.__doc__ = 'Convert base 2 string to an int.'
>>> basetwo('10010')
18
```

`class functools.partialmethod(func, *args, **keywords)`

Return a new *partialmethod* descriptor which behaves like *partial* except that it is designed to be used as a method definition rather than being directly callable.

func must be a *descriptor* or a callable (objects which are both, like normal functions, are handled as descriptors).

When *func* is a descriptor (such as a normal Python function, *classmethod()*, *staticmethod()*, *abstractmethod()* or another instance of *partialmethod*), calls to *__get__* are delegated to the underlying descriptor, and an appropriate *partial object* returned as the result.

When *func* is a non-descriptor callable, an appropriate bound method is created dynamically. This behaves like a normal Python function when used as a method: the *self* argument will be inserted as the first positional argument, even before the *args* and *keywords* supplied to the *partialmethod* constructor.

示例:

```
>>> class Cell(object):
...     def __init__(self):
...         self._alive = False
...     @property
...     def alive(self):
...         return self._alive
...     def set_state(self, state):
...         self._alive = bool(state)
...     set_alive = partialmethod(set_state, True)
...     set_dead = partialmethod(set_state, False)
...
>>> c = Cell()
>>> c.alive
False
>>> c.set_alive()
>>> c.alive
True
```

3.4 新版功能.

`functools.reduce(function, iterable[, initializer])`

Apply *function* of two arguments cumulatively to the items of *sequence*, from left to right, so as to reduce the sequence to a single value. For example, `reduce(lambda x, y: x+y, [1, 2, 3, 4, 5])` calculates `(((1+2)+3)+4)+5`. The left argument, *x*, is the accumulated value and the right argument, *y*, is the update value from the *sequence*. If the optional *initializer* is present, it is placed before the items of the sequence in the calculation, and serves as a default when the sequence is empty. If *initializer* is not given and *sequence* contains only one item, the first item is returned.

大致相当于:

```
def reduce(function, iterable, initializer=None):
    it = iter(iterable)
    if initializer is None:
        value = next(it)
    else:
        value = initializer
    for element in it:
        value = function(value, element)
    return value
```

`@functools.singledispatch`

Transform a function into a *single-dispatch generic function*.

To define a generic function, decorate it with the `@singledispatch` decorator. Note that the dispatch happens on the type of the first argument, create your function accordingly:

```
>>> from functools import singledispatch
>>> @singledispatch
... def fun(arg, verbose=False):
...     if verbose:
...         print("Let me just say,", end=" ")
...     print(arg)
```

To add overloaded implementations to the function, use the `register()` attribute of the generic function. It is a decorator, taking a type parameter and decorating a function implementing the operation for that type:

```
>>> @fun.register(int)
... def _(arg, verbose=False):
...     if verbose:
...         print("Strength in numbers, eh?", end=" ")
...     print(arg)
...
>>> @fun.register(list)
... def _(arg, verbose=False):
...     if verbose:
...         print("Enumerate this:")
...     for i, elem in enumerate(arg):
...         print(i, elem)
```

To enable registering lambdas and pre-existing functions, the `register()` attribute can be used in a functional form:

```
>>> def nothing(arg, verbose=False):
...     print("Nothing.")
...
>>> fun.register(type(None), nothing)
```

The `register()` attribute returns the undecorated function which enables decorator stacking, pickling, as well as creating unit tests for each variant independently:

```
>>> @fun.register(float)
... @fun.register(Decimal)
... def fun_num(arg, verbose=False):
...     if verbose:
...         print("Half of your number:", end=" ")
...     print(arg / 2)
...
>>> fun_num is fun
False
```

When called, the generic function dispatches on the type of the first argument:

```
>>> fun("Hello, world.")
Hello, world.
>>> fun("test.", verbose=True)
Let me just say, test.
>>> fun(42, verbose=True)
Strength in numbers, eh? 42
>>> fun(['spam', 'spam', 'eggs', 'spam'], verbose=True)
Enumerate this:
0 spam
1 spam
2 eggs
```

(下页继续)

(续上页)

```

3 spam
>>> fun(None)
Nothing.
>>> fun(1.23)
0.615

```

Where there is no registered implementation for a specific type, its method resolution order is used to find a more generic implementation. The original function decorated with `@singledispatch` is registered for the base `object` type, which means it is used if no better implementation is found.

To check which implementation will the generic function choose for a given type, use the `dispatch()` attribute:

```

>>> fun.dispatch(float)
<function fun_num at 0x1035a2840>
>>> fun.dispatch(dict)      # note: default implementation
<function fun at 0x103fe0000>

```

To access all registered implementations, use the read-only `registry` attribute:

```

>>> fun.registry.keys()
dict_keys([<class 'NoneType'>, <class 'int'>, <class 'object'>,
          <class 'decimal.Decimal'>, <class 'list'>,
          <class 'float'>])
>>> fun.registry[float]
<function fun_num at 0x1035a2840>
>>> fun.registry[object]
<function fun at 0x103fe0000>

```

3.4 新版功能.

`functools.update_wrapper(wrapper, wrapped, assigned=WRAPPER_ASSIGNMENTS, updated=WRAPPER_UPDATES)`

Update a *wrapper* function to look like the *wrapped* function. The optional arguments are tuples to specify which attributes of the original function are assigned directly to the matching attributes on the wrapper function and which attributes of the wrapper function are updated with the corresponding attributes from the original function. The default values for these arguments are the module level constants `WRAPPER_ASSIGNMENTS` (which assigns to the wrapper function's `__module__`, `__name__`, `__qualname__`, `__annotations__` and `__doc__`, the documentation string) and `WRAPPER_UPDATES` (which updates the wrapper function's `__dict__`, i.e. the instance dictionary).

To allow access to the original function for introspection and other purposes (e.g. bypassing a caching decorator such as `lru_cache()`), this function automatically adds a `__wrapped__` attribute to the wrapper that refers to the function being wrapped.

The main intended use for this function is in *decorator* functions which wrap the decorated function and return the wrapper. If the wrapper function is not updated, the metadata of the returned function will reflect the wrapper definition rather than the original function definition, which is typically less than helpful.

`update_wrapper()` may be used with callables other than functions. Any attributes named in *assigned* or *updated* that are missing from the object being wrapped are ignored (i.e. this function will not attempt to set them on the wrapper function). `AttributeError` is still raised if the wrapper function itself is missing any attributes named in *updated*.

3.2 新版功能: Automatic addition of the `__wrapped__` attribute.

3.2 新版功能: Copying of the `__annotations__` attribute by default.

在 3.2 版更改: Missing attributes no longer trigger an `AttributeError`.

在 3.4 版更改: The `__wrapped__` attribute now always refers to the wrapped function, even if that function defined a `__wrapped__` attribute. (see [bpo-17482](#))

`@functools.wraps(wrapped, assigned=WRAPPER_ASSIGNMENTS, updated=WRAPPER_UPDATES)`

This is a convenience function for invoking `update_wrapper()` as a function decorator when defining a wrapper function. It is equivalent to `partial(update_wrapper, wrapped=wrapped, assigned=assigned, updated=updated)`. For example:

```
>>> from functools import wraps
>>> def my_decorator(f):
...     @wraps(f)
...     def wrapper(*args, **kwds):
...         print('Calling decorated function')
...         return f(*args, **kwds)
...     return wrapper
...
>>> @my_decorator
... def example():
...     """Docstring"""
...     print('Called example function')
...
>>> example()
Calling decorated function
Called example function
>>> example.__name__
'example'
>>> example.__doc__
'Docstring'
```

Without the use of this decorator factory, the name of the example function would have been `'wrapper'`, and the docstring of the original `example()` would have been lost.

10.2.1 partial Objects

`partial` objects are callable objects created by `partial()`. They have three read-only attributes:

`partial.func`

A callable object or function. Calls to the `partial` object will be forwarded to `func` with new arguments and keywords.

`partial.args`

The leftmost positional arguments that will be prepended to the positional arguments provided to a `partial` object call.

`partial.keywords`

The keyword arguments that will be supplied when the `partial` object is called.

`partial` objects are like function objects in that they are callable, weak referencable, and can have attributes. There are some important differences. For instance, the `__name__` and `__doc__` attributes are not created automatically. Also, `partial` objects defined in classes behave like static methods and do not transform into bound methods during instance attribute look-up.

10.3 operator — 标准运算符替代函数

源代码: [Lib/operator.py](#)

`operator` 模块提供了一套与 Python 的内置运算符对应的高效率函数。例如, `operator.add(x, y)` 与表达式 `x+y` 相同。许多函数名与特殊方法名相同, 只是没有双下划线。为了向后兼容性, 也保留了许多包含双下划线的函数。为了表述清楚, 建议使用没有双下划线的函数。

函数包含的种类有: 对象的比较运算、逻辑运算、数学运算以及序列运算。

对象比较函数适用于所有的对象, 函数名根据它们对应的比较运算符命名。

```
operator.lt(a, b)
operator.le(a, b)
operator.eq(a, b)
operator.ne(a, b)
operator.ge(a, b)
operator.gt(a, b)
operator.__lt__(a, b)
operator.__le__(a, b)
operator.__eq__(a, b)
operator.__ne__(a, b)
operator.__ge__(a, b)
operator.__gt__(a, b)
```

在 `a` 和 `b` 之间进行全比较。具体的, `lt(a, b)` 与 `a < b` 相同, `le(a, b)` 与 `a <= b` 相同, `eq(a, b)` 与 `a == b` 相同, `ne(a, b)` 与 `a != b` 相同, `gt(a, b)` 与 `a > b` 相同, `ge(a, b)` 与 `a >= b` 相同。注意这些函数可以返回任何值, 无论它是否可当作布尔值。关于全比较的更多信息请参考 `comparisons`。

The logical operations are also generally applicable to all objects, and support truth tests, identity tests, and boolean operations:

```
operator.not_(obj)
operator.__not__(obj)
```

Return the outcome of `not obj`. (Note that there is no `__not__()` method for object instances; only the interpreter core defines this operation. The result is affected by the `__bool__()` and `__len__()` methods.)

```
operator.truth(obj)
```

Return `True` if `obj` is true, and `False` otherwise. This is equivalent to using the `bool` constructor.

```
operator.is_(a, b)
```

Return `a is b`. Tests object identity.

```
operator.is_not(a, b)
```

Return `a is not b`. Tests object identity.

The mathematical and bitwise operations are the most numerous:

```
operator.abs(obj)
operator.__abs__(obj)
```

Return the absolute value of `obj`.

```
operator.add(a, b)
operator.__add__(a, b)
```

Return `a + b`, for `a` and `b` numbers.

```
operator.and_(a, b)
operator.__and__(a, b)
```

Return the bitwise and of `a` and `b`.

```
operator.floordiv(a, b)
```

`operator.__floordiv__ (a, b)`

Return $a // b$.

`operator.index (a)`

`operator.__index__ (a)`

Return a converted to an integer. Equivalent to `a.__index__()`.

`operator.inv (obj)`

`operator.invert (obj)`

`operator.__inv__ (obj)`

`operator.__invert__ (obj)`

Return the bitwise inverse of the number *obj*. This is equivalent to $\sim obj$.

`operator.lshift (a, b)`

`operator.__lshift__ (a, b)`

Return a shifted left by b .

`operator.mod (a, b)`

`operator.__mod__ (a, b)`

Return $a \% b$.

`operator.mul (a, b)`

`operator.__mul__ (a, b)`

Return $a * b$, for a and b numbers.

`operator.matmul (a, b)`

`operator.__matmul__ (a, b)`

Return $a @ b$.

3.5 新版功能.

`operator.neg (obj)`

`operator.__neg__ (obj)`

Return *obj* negated ($-obj$).

`operator.or_ (a, b)`

`operator.__or__ (a, b)`

Return the bitwise or of a and b .

`operator.pos (obj)`

`operator.__pos__ (obj)`

Return *obj* positive ($+obj$).

`operator.pow (a, b)`

`operator.__pow__ (a, b)`

Return $a ** b$, for a and b numbers.

`operator.rshift (a, b)`

`operator.__rshift__ (a, b)`

Return a shifted right by b .

`operator.sub (a, b)`

`operator.__sub__ (a, b)`

Return $a - b$.

`operator.truediv (a, b)`

`operator.__truediv__ (a, b)`

Return a / b where $2/3$ is .66 rather than 0. This is also known as “true” division.

`operator.xor (a, b)`

`operator.__xor__(a, b)`

Return the bitwise exclusive or of *a* and *b*.

Operations which work with sequences (some of them with mappings too) include:

`operator.concat(a, b)`

`operator.__concat__(a, b)`

Return *a* + *b* for *a* and *b* sequences.

`operator.contains(a, b)`

`operator.__contains__(a, b)`

Return the outcome of the test *b* in *a*. Note the reversed operands.

`operator.countOf(a, b)`

Return the number of occurrences of *b* in *a*.

`operator.delitem(a, b)`

`operator.__delitem__(a, b)`

Remove the value of *a* at index *b*.

`operatorgetitem(a, b)`

`operator.__getitem__(a, b)`

Return the value of *a* at index *b*.

`operator.indexOf(a, b)`

Return the index of the first of occurrence of *b* in *a*.

`operator.setitem(a, b, c)`

`operator.__setitem__(a, b, c)`

Set the value of *a* at index *b* to *c*.

`operator.length_hint(obj, default=0)`

Return an estimated length for the object *o*. First try to return its actual length, then an estimate using `object.__length_hint__()`, and finally return the default value.

3.4 新版功能.

The `operator` module also defines tools for generalized attribute and item lookups. These are useful for making fast field extractors as arguments for `map()`, `sorted()`, `itertools.groupby()`, or other functions that expect a function argument.

`operator.attrgetter(attr)`

`operator.attrgetter(*attrs)`

Return a callable object that fetches *attr* from its operand. If more than one attribute is requested, returns a tuple of attributes. The attribute names can also contain dots. For example:

- After `f = attrgetter('name')`, the call `f(b)` returns `b.name`.
- After `f = attrgetter('name', 'date')`, the call `f(b)` returns `(b.name, b.date)`.
- After `f = attrgetter('name.first', 'name.last')`, the call `f(b)` returns `(b.name.first, b.name.last)`.

等价于:

```
def attrgetter(*items):
    if any(not isinstance(item, str) for item in items):
        raise TypeError('attribute name must be a string')
    if len(items) == 1:
        attr = items[0]
        def g(obj):
            return resolve_attr(obj, attr)
```

(下页继续)

(续上页)

```

else:
    def g(obj):
        return tuple(resolve_attr(obj, attr) for attr in items)
    return g

def resolve_attr(obj, attr):
    for name in attr.split("."):
        obj = getattr(obj, name)
    return obj

```

`operator.itemgetter(item)`

`operator.itemgetter(*items)`

Return a callable object that fetches *item* from its operand using the operand's `__getitem__()` method. If multiple items are specified, returns a tuple of lookup values. For example:

- After `f = itemgetter(2)`, the call `f(r)` returns `r[2]`.
- After `g = itemgetter(2, 5, 3)`, the call `g(r)` returns `(r[2], r[5], r[3])`.

等价于:

```

def itemgetter(*items):
    if len(items) == 1:
        item = items[0]
        def g(obj):
            return obj[item]
    else:
        def g(obj):
            return tuple(obj[item] for item in items)
    return g

```

The items can be any type accepted by the operand's `__getitem__()` method. Dictionaries accept any hashable value. Lists, tuples, and strings accept an index or a slice:

```

>>> itemgetter(1)('ABCDEFGH')
'B'
>>> itemgetter(1,3,5)('ABCDEFGH')
('B', 'D', 'F')
>>> itemgetter(slice(2,None))('ABCDEFGH')
'CDEFG'

```

Example of using `itemgetter()` to retrieve specific fields from a tuple record:

```

>>> inventory = [('apple', 3), ('banana', 2), ('pear', 5), ('orange', 1)]
>>> getcount = itemgetter(1)
>>> list(map(getcount, inventory))
[3, 2, 5, 1]
>>> sorted(inventory, key=getcount)
[('orange', 1), ('banana', 2), ('apple', 3), ('pear', 5)]

```

`operator.methodcaller(name[, args...])`

Return a callable object that calls the method *name* on its operand. If additional arguments and/or keyword arguments are given, they will be given to the method as well. For example:

- After `f = methodcaller('name')`, the call `f(b)` returns `b.name()`.
- After `f = methodcaller('name', 'foo', bar=1)`, the call `f(b)` returns `b.name('foo', bar=1)`.

等价于:

```
def methodcaller(name, *args, **kwargs):
    def caller(obj):
        return getattr(obj, name)(*args, **kwargs)
    return caller
```

10.3.1 Mapping Operators to Functions

This table shows how abstract operations correspond to operator symbols in the Python syntax and the functions in the `operator` module.

运算	Syntax	Function
Addition	<code>a + b</code>	<code>add(a, b)</code>
Concatenation	<code>seq1 + seq2</code>	<code>concat(seq1, seq2)</code>
Containment Test	<code>obj in seq</code>	<code>contains(seq, obj)</code>
Division	<code>a / b</code>	<code>truediv(a, b)</code>
Division	<code>a // b</code>	<code>floordiv(a, b)</code>
Bitwise And	<code>a & b</code>	<code>and_(a, b)</code>
Bitwise Exclusive Or	<code>a ^ b</code>	<code>xor(a, b)</code>
Bitwise Inversion	<code>~ a</code>	<code>invert(a)</code>
Bitwise Or	<code>a b</code>	<code>or_(a, b)</code>
Exponentiation	<code>a ** b</code>	<code>pow(a, b)</code>
Identity	<code>a 是 b</code>	<code>is_(a, b)</code>
Identity	<code>a 不是 b</code>	<code>is_not(a, b)</code>
Indexed Assignment	<code>obj[k] = v</code>	<code>setitem(obj, k, v)</code>
Indexed Deletion	<code>del obj[k]</code>	<code>delitem(obj, k)</code>
Indexing	<code>obj[k]</code>	<code>getitem(obj, k)</code>
Left Shift	<code>a << b</code>	<code>lshift(a, b)</code>
Modulo	<code>a % b</code>	<code>mod(a, b)</code>
Multiplication	<code>a * b</code>	<code>mul(a, b)</code>
Matrix Multiplication	<code>a @ b</code>	<code>matmul(a, b)</code>
Negation (Arithmetic)	<code>- a</code>	<code>neg(a)</code>
Negation (Logical)	<code>not a</code>	<code>not_(a)</code>
Positive	<code>+ a</code>	<code>pos(a)</code>
Right Shift	<code>a >> b</code>	<code>rshift(a, b)</code>
Slice Assignment	<code>seq[i:j] = values</code>	<code>setitem(seq, slice(i, j), values)</code>
Slice Deletion	<code>del seq[i:j]</code>	<code>delitem(seq, slice(i, j))</code>
Slicing	<code>seq[i:j]</code>	<code>getitem(seq, slice(i, j))</code>
String Formatting	<code>s % obj</code>	<code>mod(s, obj)</code>
Subtraction	<code>a - b</code>	<code>sub(a, b)</code>
Truth Test	<code>obj</code>	<code>truth(obj)</code>
Ordering	<code>a < b</code>	<code>lt(a, b)</code>
Ordering	<code>a <= b</code>	<code>le(a, b)</code>
Equality	<code>a == b</code>	<code>eq(a, b)</code>
Difference	<code>a != b</code>	<code>ne(a, b)</code>
Ordering	<code>a >= b</code>	<code>ge(a, b)</code>
Ordering	<code>a > b</code>	<code>gt(a, b)</code>

10.3.2 Inplace Operators

Many operations have an “in-place” version. Listed below are functions providing a more primitive access to in-place operators than the usual syntax does; for example, the *statement* `x += y` is equivalent to `x = operator.iadd(x, y)`. Another way to put it is to say that `z = operator.iadd(x, y)` is equivalent to the compound statement `z = x; z += y`.

In those examples, note that when an in-place method is called, the computation and assignment are performed in two separate steps. The in-place functions listed below only do the first step, calling the in-place method. The second step, assignment, is not handled.

For immutable targets such as strings, numbers, and tuples, the updated value is computed, but not assigned back to the input variable:

```
>>> a = 'hello'
>>> iadd(a, ' world')
'hello world'
>>> a
'hello'
```

For mutable targets such as lists and dictionaries, the inplace method will perform the update, so no subsequent assignment is necessary:

```
>>> s = ['h', 'e', 'l', 'l', 'o']
>>> iadd(s, [' ', 'w', 'o', 'r', 'l', 'd'])
['h', 'e', 'l', 'l', 'o', ' ', 'w', 'o', 'r', 'l', 'd']
>>> s
['h', 'e', 'l', 'l', 'o', ' ', 'w', 'o', 'r', 'l', 'd']
```

```
operator.iadd(a, b)
operator.__iadd__(a, b)
    a = iadd(a, b) is equivalent to a += b.
```

```
operator.iand(a, b)
operator.__iand__(a, b)
    a = iand(a, b) is equivalent to a &= b.
```

```
operator.iconcat(a, b)
operator.__iconcat__(a, b)
    a = iconcat(a, b) is equivalent to a += b for a and b sequences.
```

```
operator.ifloordiv(a, b)
operator.__ifloordiv__(a, b)
    a = ifloordiv(a, b) is equivalent to a //= b.
```

```
operator.ilshift(a, b)
operator.__ilshift__(a, b)
    a = ilshift(a, b) is equivalent to a <<= b.
```

```
operator.imod(a, b)
operator.__imod__(a, b)
    a = imod(a, b) is equivalent to a %= b.
```

```
operator.imul(a, b)
operator.__imul__(a, b)
    a = imul(a, b) is equivalent to a *= b.
```

```
operator.imatmul(a, b)
operator.__imatmul__(a, b)
    a = imatmul(a, b) is equivalent to a @= b.
```

3.5 新版功能.

```
operator.ior(a, b)
operator.__ior__(a, b)
    a = ior(a, b) is equivalent to a |= b.

operator.ipow(a, b)
operator.__ipow__(a, b)
    a = ipow(a, b) is equivalent to a **= b.

operator.irshift(a, b)
operator.__irshift__(a, b)
    a = irshift(a, b) is equivalent to a >>= b.

operator.isub(a, b)
operator.__isub__(a, b)
    a = isub(a, b) is equivalent to a -= b.

operator.itruediv(a, b)
operator.__itruediv__(a, b)
    a = itruediv(a, b) is equivalent to a /= b.

operator.ixor(a, b)
operator.__ixor__(a, b)
    a = ixor(a, b) is equivalent to a ^= b.
```

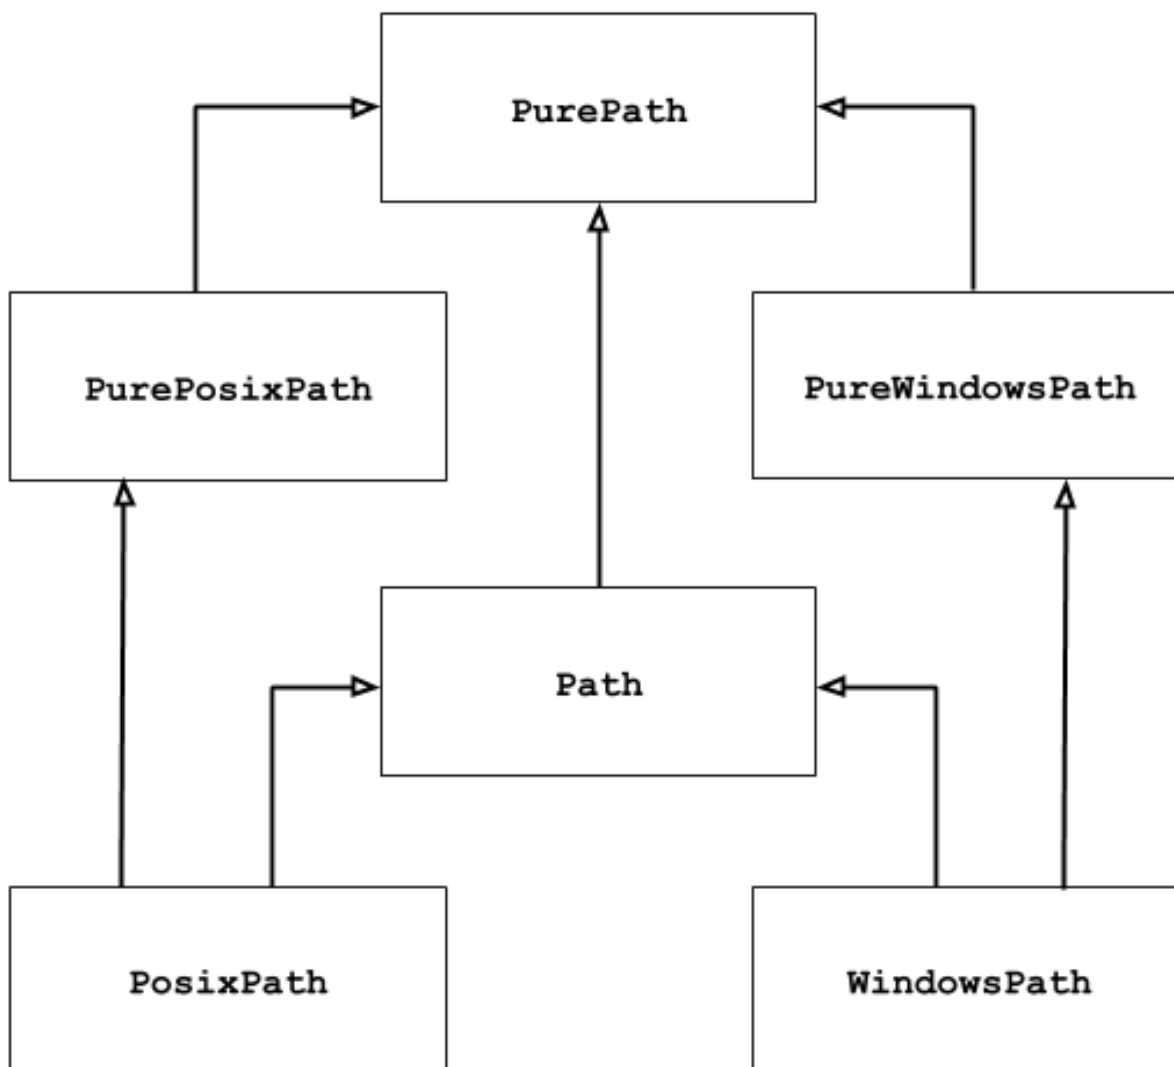
本章中描述的模块处理磁盘文件和目录。例如，有一些模块用于读取文件的属性，以可移植的方式操作路径以及创建临时文件。本章的完整模块列表如下：

11.1 `pathlib` — 面向对象的文件系统路径

3.4 新版功能.

源代码 [Lib/pathlib.py](#)

该模块提供表示文件系统路径的类，其语义适用于不同的操作系统。路径类被分为提供纯计算操作而没有 I/O 的 [纯路径](#)，以及从纯路径继承而来但提供 I/O 操作的 [具体路径](#)。



如果你以前从未使用过此模块或者不确定在项目中使用哪一个类是正确的，则`Path`总是你需要的。它在运行代码的平台上实例化为一个具体路径。

在一些用例中纯路径很有用，例如：

1. 如果你想要在 Unix 设备上操作 Windows 路径（或者相反）。你不应在 Unix 上实例化一个`WindowsPath`，但是你可以实例化`PureWindowsPath`。
2. 你只想操作路径但不想实际访问操作系统。在这种情况下，实例化一个纯路径是有用的，因为它们没有任何访问操作系统的操作。

参见：

PEP 428：pathlib 模块 – 面向对象的文件系统路径。

参见：

对于底层的路径字符串操作，你也可以使用`os.path`模块。

11.1.1 基础使用

导入主类:

```
>>> from pathlib import Path
```

列出子目录:

```
>>> p = Path('.')
>>> [x for x in p.iterdir() if x.is_dir()]
[PosixPath('.hg'), PosixPath('docs'), PosixPath('dist'),
 PosixPath('__pycache__'), PosixPath('build')]
```

列出当前目录树下的所有 Python 源代码文件:

```
>>> list(p.glob('**/*.py'))
[PosixPath('test_pathlib.py'), PosixPath('setup.py'),
 PosixPath('pathlib.py'), PosixPath('docs/conf.py'),
 PosixPath('build/lib/pathlib.py')]
```

在目录树中移动:

```
>>> p = Path('/etc')
>>> q = p / 'init.d' / 'reboot'
>>> q
PosixPath('/etc/init.d/reboot')
>>> q.resolve()
PosixPath('/etc/rc.d/init.d/halt')
```

查询路径的属性:

```
>>> q.exists()
True
>>> q.is_dir()
False
```

打开一个文件:

```
>>> with q.open() as f: f.readline()
...
'#!/bin/bash\n'
```

11.1.2 纯路径

纯路径对象提供了不实际访问文件系统的路径处理操作。有三种方式来访问这些类，也是不同的风格：

class `pathlib.PurePath` (**pathsegments*)

一个通用的类，代表当前系统的路径风格（实例化为 `PurePosixPath` 或者 `PureWindowsPath`）：

```
>>> PurePath('setup.py')           # Running on a Unix machine
PurePosixPath('setup.py')
```

每一个 *pathsegments* 的元素可能是一个代表路径片段的字符串，一个返回字符串的实现了 `os.PathLike` 接口的对象，或者另一个路径对象：

```
>>> PurePath('foo', 'some/path', 'bar')
PurePosixPath('foo/some/path/bar')
>>> PurePath(Path('foo'), Path('bar'))
PurePosixPath('foo/bar')
```

当 *pathsegments* 为空的时候, 假定为当前目录:

```
>>> PurePath()
PurePosixPath('.')
```

当给出一些绝对路径, 最后一位将被当作锚 (模仿 *os.path.join()* 的行为):

```
>>> PurePath('/etc', '/usr', 'lib64')
PurePosixPath('/usr/lib64')
>>> PureWindowsPath('c:/Windows', 'd:bar')
PureWindowsPath('d:bar')
```

但是, 在 Windows 路径中, 改变本地根目录并不会丢弃之前盘符的设置:

```
>>> PureWindowsPath('c:/Windows', '/Program Files')
PureWindowsPath('c:/Program Files')
```

假斜线和单独的点都会被消除, 但是双点 ('..') 不会, 以防改变符号链接的含义。

```
>>> PurePath('foo//bar')
PurePosixPath('foo/bar')
>>> PurePath('foo/./bar')
PurePosixPath('foo/bar')
>>> PurePath('foo/../bar')
PurePosixPath('foo/../bar')
```

(如果你想让 `PurePosixPath('foo/../bar')` 等同于 `PurePosixPath('bar')`, 那么 you are too young, too simple, sometimes naive! 如果 `foo` 是一个指向其他其他目录的符号链接, 那就出毛病啦。)

纯路径对象实现了 *os.PathLike* 接口, 允许它们在任何接受此接口的地方使用。

在 3.6 版更改: 添加了 *os.PathLike* 接口支持。

class `pathlib.PurePosixPath(*pathsegments)`
一个 *PurePath* 的子类, 路径风格不同于 Windows 文件系统:

```
>>> PurePosixPath('/etc')
PurePosixPath('/etc')
```

pathsegments 参数的指定和 *PurePath* 相同。

class `pathlib.PureWindowsPath(*pathsegments)`
PurePath 的一个子类, 路径风格为 Windows 文件系统路径:

```
>>> PureWindowsPath('c:/Program Files/')
PureWindowsPath('c:/Program Files')
```

pathsegments 参数的指定和 *PurePath* 相同。

无论你正运行什么系统, 你都可以实例化这些类, 因为它们提供的操作不做任何系统调用。

通用性质

路径是不可变并可哈希的。相同风格的路径可以排序与比较。这些性质尊重对应风格的大小写转换语义:

```
>>> PurePosixPath('foo') == PurePosixPath('FOO')
False
>>> PureWindowsPath('foo') == PureWindowsPath('FOO')
True
>>> PureWindowsPath('FOO') in { PureWindowsPath('foo') }
True
>>> PureWindowsPath('C:') < PureWindowsPath('d:')
True
```

不同风格的路径比较得到不等的结果并且无法被排序:

```
>>> PureWindowsPath('foo') == PurePosixPath('foo')
False
>>> PureWindowsPath('foo') < PurePosixPath('foo')
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
TypeError: '<' not supported between instances of 'PureWindowsPath' and 'PurePosixPath'
↪'
```

运算符

斜杠 / 操作符有助于创建子路径, 就像 `os.path.join()` 一样:

```
>>> p = PurePath('/etc')
>>> p
PurePosixPath('/etc')
>>> p / 'init.d' / 'apache2'
PurePosixPath('/etc/init.d/apache2')
>>> q = PurePath('bin')
>>> '/usr' / q
PurePosixPath('/usr/bin')
```

文件对象可用于任何接受 `os.PathLike` 接口实现的地方。

```
>>> import os
>>> p = PurePath('/etc')
>>> os.fspath(p)
'/etc'
```

路径的字符串表示法为它自己原始的文件系统路径 (以原生形式, 例如在 Windows 下使用反斜杠)。你可以传递给任何需要字符串形式路径的函数。

```
>>> p = PurePath('/etc')
>>> str(p)
'/etc'
>>> p = PureWindowsPath('c:/Program Files')
>>> str(p)
'c:\\Program Files'
```

类似地, 在路径上调用 `bytes` 将原始文件系统路径作为字节对象给出, 就像被 `os.fsencode()` 编码一样:

```
>>> bytes(p)
b'/etc'
```

注解：只推荐在 Unix 下调用 `bytes`。在 Windows，`unicode` 形式是文件系统路径的规范表示法。

访问个别部分

为了访问路径独立的部分（组件），使用以下特征属性：

`PurePath.parts`

一个元组，可以访问路径的多个组件：

```
>>> p = PurePath('/usr/bin/python3')
>>> p.parts
('/', 'usr', 'bin', 'python3')

>>> p = PureWindowsPath('c:/Program Files/PSF')
>>> p.parts
('c:\\', 'Program Files', 'PSF')
```

（注意盘符和本地根目录是如何重组的）

方法和特征属性

纯路径提供以下方法和特征属性：

`PurePath.drive`

一个表示驱动器盘符或命名的字符串，如果存在：

```
>>> PureWindowsPath('c:/Program Files/').drive
'c:'
>>> PureWindowsPath('/Program Files/').drive
''
>>> PurePosixPath('/etc').drive
''
```

UNC 分享也被认作驱动器：

```
>>> PureWindowsPath('//host/share/foo.txt').drive
'\\\\host\\share'
```

`PurePath.root`

一个表示（本地或全局）根的字符串，如果存在：

```
>>> PureWindowsPath('c:/Program Files/').root
'\\'
>>> PureWindowsPath('c:Program Files/').root
''
>>> PurePosixPath('/etc').root
'/'
```

UNC 分享一样拥有根：

```
>>> PureWindowsPath('//host/share').root
'\\'
```

`PurePath.anchor`

驱动器和根的联合：

```
>>> PureWindowsPath('c:/Program Files/').anchor
'c:\\'
>>> PureWindowsPath('c:Program Files/').anchor
'c:'
>>> PurePosixPath('/etc').anchor
 '/'
>>> PureWindowsPath('//host/share').anchor
 '\\\\host\\share\\'
```

PurePath.parents

An immutable sequence providing access to the logical ancestors of the path:

```
>>> p = PureWindowsPath('c:/foo/bar/setup.py')
>>> p.parents[0]
PureWindowsPath('c:/foo/bar')
>>> p.parents[1]
PureWindowsPath('c:/foo')
>>> p.parents[2]
PureWindowsPath('c:/')
```

PurePath.parent

此路径的逻辑父路径:

```
>>> p = PurePosixPath('/a/b/c/d')
>>> p.parent
PurePosixPath('/a/b/c')
```

你不能超过一个 anchor 或空路径:

```
>>> p = PurePosixPath('/')
>>> p.parent
PurePosixPath('/')
>>> p = PurePosixPath('.')
>>> p.parent
PurePosixPath('.')
```

注解: 这是一个单纯的词法操作, 因此有以下行为:

```
>>> p = PurePosixPath('foo/..')
>>> p.parent
PurePosixPath('foo')
```

如果你想要向上移动任意文件系统路径, 推荐先使用 `Path.resolve()` 来解析符号链接以及消除 `".."` 组件。

PurePath.name

一个表示最后路径组件的字符串, 排除了驱动器与根目录, 如果存在的话:

```
>>> PurePosixPath('my/library/setup.py').name
'setup.py'
```

UNC 驱动器名不被考虑:

```
>>> PureWindowsPath('//some/share/setup.py').name
'setup.py'
```

(下页继续)

(续上页)

```
>>> PureWindowsPath('//some/share').name
''
```

PurePath.suffix

最后一个组件的文件扩展名，如果存在：

```
>>> PurePosixPath('my/library/setup.py').suffix
'.py'
>>> PurePosixPath('my/library.tar.gz').suffix
'.gz'
>>> PurePosixPath('my/library').suffix
''
```

PurePath.suffixes

路径的文件扩展名列表：

```
>>> PurePosixPath('my/library.tar.gar').suffixes
['.tar', '.gar']
>>> PurePosixPath('my/library.tar.gz').suffixes
['.tar', '.gz']
>>> PurePosixPath('my/library').suffixes
[]
```

PurePath.stem

最后一个路径组件，除去后缀：

```
>>> PurePosixPath('my/library.tar.gz').stem
'library.tar'
>>> PurePosixPath('my/library.tar').stem
'library'
>>> PurePosixPath('my/library').stem
'library'
```

PurePath.as_posix()

返回使用正斜杠 (/) 的路径字符串：

```
>>> p = PureWindowsPath('c:\\windows')
>>> str(p)
'c:\\windows'
>>> p.as_posix()
'c:/windows'
```

PurePath.as_uri()

将路径表示为 file URL。如果并非绝对路径，抛出 *ValueError*。

```
>>> p = PurePosixPath('/etc/passwd')
>>> p.as_uri()
'file:///etc/passwd'
>>> p = PureWindowsPath('c:/Windows')
>>> p.as_uri()
'file:///c:/Windows'
```

PurePath.is_absolute()

返回此路径是否为绝对路径。如果路径同时拥有驱动器符与根路径（如果风格允许）则将被认作绝对路径。

```

>>> PurePosixPath('/a/b').is_absolute()
True
>>> PurePosixPath('a/b').is_absolute()
False

>>> PureWindowsPath('c:/a/b').is_absolute()
True
>>> PureWindowsPath('/a/b').is_absolute()
False
>>> PureWindowsPath('c:').is_absolute()
False
>>> PureWindowsPath('//some/share').is_absolute()
True

```

PurePath.is_reserved()

在 *PureWindowsPath*, 如果路径是被 Windows 保留的则返回 True, 否则 False。在 *PurePosixPath*, 总是返回 False。

```

>>> PureWindowsPath('nul').is_reserved()
True
>>> PurePosixPath('nul').is_reserved()
False

```

当保留路径上的文件系统被调用, 则可能出现玄学失败或者意料之外的效应。

PurePath.joinpath(*other)

调用此方法等同于将每个 *other* 参数中的项目连接在一起:

```

>>> PurePosixPath('/etc').joinpath('passwd')
PurePosixPath('/etc/passwd')
>>> PurePosixPath('/etc').joinpath(PurePosixPath('passwd'))
PurePosixPath('/etc/passwd')
>>> PurePosixPath('/etc').joinpath('init.d', 'apache2')
PurePosixPath('/etc/init.d/apache2')
>>> PureWindowsPath('c:').joinpath('/Program Files')
PureWindowsPath('c:/Program Files')

```

PurePath.match(pattern)

将此路径与提供的通配符风格的模式匹配。如果匹配成功则返回 True, 否则返回 False。

如果 *pattern* 是相对的, 则路径可以是相对路径或绝对路径, 并且匹配是从右侧完成的:

```

>>> PurePath('a/b.py').match('*.py')
True
>>> PurePath('/a/b/c.py').match('b/*.py')
True
>>> PurePath('/a/b/c.py').match('a/*.py')
False

```

如果 *pattern* 是绝对的, 则路径必须是绝对的, 并且路径必须完全匹配:

```

>>> PurePath('/a.py').match('/*.py')
True
>>> PurePath('a/b.py').match('/*.py')
False

```

与其他方法一样, 可以观察到大小写区分:

```
>>> PureWindowsPath('b.py').match('*.PY')
True
```

`PurePath.relative_to(*other)`

计算此路径相对 *other* 表示路径的版本。如果不可计算，则抛出 `ValueError`：

```
>>> p = PurePosixPath('/etc/passwd')
>>> p.relative_to('/')
PurePosixPath('etc/passwd')
>>> p.relative_to('/etc')
PurePosixPath('passwd')
>>> p.relative_to('/usr')
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
  File "pathlib.py", line 694, in relative_to
    .format(str(self), str(formatted)))
ValueError: '/etc/passwd' does not start with '/usr'
```

`PurePath.with_name(name)`

返回一个新的路径并修改 *name*。如果原本路径没有 *name*，`ValueError` 被抛出：

```
>>> p = PureWindowsPath('c:/Downloads/pathlib.tar.gz')
>>> p.with_name('setup.py')
PureWindowsPath('c:/Downloads/setup.py')
>>> p = PureWindowsPath('c:/')
>>> p.with_name('setup.py')
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
  File "/home/antoine/cpython/default/Lib/pathlib.py", line 751, in with_name
    raise ValueError("%r has an empty name" % (self,))
ValueError: PureWindowsPath('c:/') has an empty name
```

`PurePath.with_suffix(suffix)`

返回一个新的路径并修改 *suffix*。如果原本的路径没有后缀，新的 *suffix* 则被追加以代替。如果 *suffix* 是空字符串，则原本的后缀被移除：

```
>>> p = PureWindowsPath('c:/Downloads/pathlib.tar.gz')
>>> p.with_suffix('.bz2')
PureWindowsPath('c:/Downloads/pathlib.tar.bz2')
>>> p = PureWindowsPath('README')
>>> p.with_suffix('.txt')
PureWindowsPath('README.txt')
>>> p = PureWindowsPath('README.txt')
>>> p.with_suffix('')
PureWindowsPath('README')
```

11.1.3 具体路径

具体路径是纯路径的子类。除了后者提供的操作之外，它们还提供了对路径对象进行系统调用的方法。有三种方法可以实例化具体路径：

class `pathlib.Path(*pathsegments)`

一个 `PurePath` 的子类，此类以当前系统的路径风格表示路径（实例化为 `PosixPath` 或 `WindowsPath`）：

```
>>> Path('setup.py')
PosixPath('setup.py')
```

pathsegments 参数的指定和 *PurePath* 相同。

class `pathlib.PosixPath(*pathsegments)`

一个 *Path* 和 *PurePosixPath* 的子类，此类表示一个非 Windows 文件系统的具体路径：

```
>>> PosixPath('/etc')
PosixPath('/etc')
```

pathsegments 参数的指定和 *PurePath* 相同。

class `pathlib.WindowsPath(*pathsegments)`

Path 和 *PureWindowsPath* 的子类，从类表示一个 Windows 文件系统的具体路径：

```
>>> WindowsPath('c:/Program Files/')
WindowsPath('c:/Program Files')
```

pathsegments 参数的指定和 *PurePath* 相同。

你只能实例化与当前系统风格相同的类（允许系统调用作用于不兼容的路径风格可能在应用程序中导致缺陷或失败）：

```
>>> import os
>>> os.name
'posix'
>>> Path('setup.py')
PosixPath('setup.py')
>>> PosixPath('setup.py')
PosixPath('setup.py')
>>> WindowsPath('setup.py')
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
  File "pathlib.py", line 798, in __new__
    % (cls.__name__,))
NotImplementedError: cannot instantiate 'WindowsPath' on your system
```

方法

除纯路径方法外，具体路径还提供以下方法。如果系统调用失败（例如因为路径不存在），其中许多方法都会引发 *OSError*：

classmethod `Path.cwd()`

返回一个新的表示当前目录的路径对象（和 *os.getcwd()* 返回的相同）：

```
>>> Path.cwd()
PosixPath('/home/antoine/pathlib')
```

classmethod `Path.home()`

返回一个表示当前用户家目录的新路径对象（和 *os.path.expanduser()* 构造含 ~ 路径返回的相同）：

```
>>> Path.home()
PosixPath('/home/antoine')
```

3.5 新版功能.

`Path.stat()`

返回此路径的信息（类似于 `os.stat()`）。结果在每次调用此方法时都重新产生。

```
>>> p = Path('setup.py')
>>> p.stat().st_size
956
>>> p.stat().st_mtime
1327883547.852554
```

`Path.chmod(mode)`

改变文件的模式和权限，和 `os.chmod()` 一样：

```
>>> p = Path('setup.py')
>>> p.stat().st_mode
33277
>>> p.chmod(0o444)
>>> p.stat().st_mode
33060
```

`Path.exists()`

此路径是否指向一个已存在的文件或目录：

```
>>> Path('.').exists()
True
>>> Path('setup.py').exists()
True
>>> Path('/etc').exists()
True
>>> Path('nonexistentfile').exists()
False
```

注解： 如果路径指向一个符号链接，`exists()` 返回此符号链接是否指向存在的文件或目录。

`Path.expanduser()`

返回展开了包含 `~` 和 `~user` 的构造，就和 `os.path.expanduser()` 一样：

```
>>> p = PosixPath('~ /films/Monty Python')
>>> p.expanduser()
PosixPath('/home/eric/films/Monty Python')
```

3.5 新版功能.

`Path.glob(pattern)`

Glob the given *pattern* in the directory represented by this path, yielding all matching files (of any kind):

```
>>> sorted(Path('.').glob('*.py'))
[PosixPath('pathlib.py'), PosixPath('setup.py'), PosixPath('test_pathlib.py')]
>>> sorted(Path('.').glob('*/*.py'))
[PosixPath('docs/conf.py')]
```

“`***`” 模式表示 “此目录以及所有子目录，递归”。换句话说，它启用递归通配：

```
>>> sorted(Path('.').glob('**/*.py'))
[PosixPath('build/lib/pathlib.py'),
 PosixPath('docs/conf.py'),
 PosixPath('pathlib.py'),
```

(下页继续)

(续上页)

```
PosixPath('setup.py'),
PosixPath('test_pathlib.py')]
```

注解： 在一个较大的目录树中使用 “**” 模式可能会消耗非常多的时间。

Path.group()

返回拥有此文件的用户组。如果文件的 GID 无法在系统数据库中找到，将抛出 `KeyError`。

Path.is_dir()

如果路径指向一个目录（或者一个指向目录的符号链接）则返回 `True`，如果指向其他类型的文件则返回 `False`。

当路径不存在或者是一个破损的符号链接时也会返回 `False`；其他错误（例如权限错误）被传播。

Path.is_file()

如果路径指向一个正常的文件（或者一个指向正常文件的符号链接）则返回 `True`，如果指向其他类型的文件则返回 `False`。

当路径不存在或者是一个破损的符号链接时也会返回 `False`；其他错误（例如权限错误）被传播。

Path.is_symlink()

如果路径指向符号链接则返回 `True`，否则 `False`。

如果路径不存在也返回 `False`；其他错误（例如权限错误）被传播。

Path.is_socket()

如果路径指向一个 Unix socket 文件（或者指向 Unix socket 文件的符号链接）则返回 `True`，如果指向其他类型的文件则返回 `False`。

当路径不存在或者是一个破损的符号链接时也会返回 `False`；其他错误（例如权限错误）被传播。

Path.is_fifo()

如果路径指向一个先进先出存储（或者指向先进先出存储的符号链接）则返回 `True`，指向其他类型的文件则返回 `False`。

当路径不存在或者是一个破损的符号链接时也会返回 `False`；其他错误（例如权限错误）被传播。

Path.is_block_device()

如果文件指向一个块设备（或者指向块设备的符号链接）则返回 `True`，指向其他类型的文件则返回 `False`。

当路径不存在或者是一个破损的符号链接时也会返回 `False`；其他错误（例如权限错误）被传播。

Path.is_char_device()

如果路径指向一个字符设备（或指向字符设备的符号链接）则返回 `True`，指向其他类型的文件则返回 `False`。

当路径不存在或者是一个破损的符号链接时也会返回 `False`；其他错误（例如权限错误）被传播。

Path.iterdir()

当路径指向一个目录时，产生该路径下的对象的路径：

```
>>> p = Path('docs')
>>> for child in p.iterdir(): child
...
PosixPath('docs/conf.py')
PosixPath('docs/_templates')
PosixPath('docs/make.bat')
PosixPath('docs/index.rst')
```

(下页继续)

(续上页)

```
PosixPath('docs/_build')
PosixPath('docs/_static')
PosixPath('docs/Makefile')
```

Path.lchmod(mode)

就像 `Path.chmod()` 但是如果路径指向符号链接则是修改符号链接的模式，而不是修改符号链接的目标。

Path.lstat()

就和 `Path.stat()` 一样，但是如果路径指向符号链接，则是返回符号链接而不是目标的信息。

Path.mkdir(mode=0o777, parents=False, exist_ok=False)

新建给定路径的目录。如果给出了 `mode`，它将与当前进程的 `umask` 值合并来决定文件模式和访问标志。如果路径已经存在，则抛出 `FileExistsError`。

如果 `parents` 为 `true`，任何找不到的父目录都会伴随着此路径被创建；它们会以默认权限被创建，而不考虑 `mode` 设置（模仿 POSIX 的 `mkdir -p` 命令）。

如果 `parents` 为 `false`（默认），则找不到的父级目录会导致 `FileNotFoundError` 被抛出。

如果 `exist_ok` 为 `false`（默认），则在目标已存在的情况下抛出 `FileExistsError`。

如果 `exist_ok` 为 `true`，则 `FileExistsError` 异常将被忽略（和 POSIX `mkdir -p` 命令行为相同），但是只有在最后一个路径组件不是现存的非目录文件时才生效。

在 3.5 版更改: `exist_ok` 形参被加入。

Path.open(mode='r', buffering=-1, encoding=None, errors=None, newline=None)

打开路径指向的文件，就像内置的 `open()` 函数所做的一样：

```
>>> p = Path('setup.py')
>>> with p.open() as f:
...     f.readline()
...
'#!/usr/bin/env python3\n'
```

Path.owner()

返回拥有此文件的用户名。如果文件的 UID 无法在系统数据库中找到，则抛出 `KeyError`。

Path.read_bytes()

以字节对象的形式返回路径指向的文件的二进制内容：

```
>>> p = Path('my_binary_file')
>>> p.write_bytes(b'Binary file contents')
20
>>> p.read_bytes()
b'Binary file contents'
```

3.5 新版功能。

Path.read_text(encoding=None, errors=None)

以字符串形式返回路径指向的文件的解码后文本内容。

```
>>> p = Path('my_text_file')
>>> p.write_text('Text file contents')
18
>>> p.read_text()
'Text file contents'
```

文件先被打开然后关闭。有和`open()`一样的可选形参。

3.5 新版功能.

`Path.rename(target)`

使用给定的 *target* 将文件重命名。在 Unix 上, 如果 *target* 已经存在并且为文件, 则只要用户拥有权限, 其将被静默地被覆盖。*target* 可以是一个字符串或者另一个路径对象:

```
>>> p = Path('foo')
>>> p.open('w').write('some text')
9
>>> target = Path('bar')
>>> p.rename(target)
>>> target.open().read()
'some text'
```

`Path.replace(target)`

使用给定的 *target* 重命名文件或目录。如果 *target* 指向现存的文件或目录, 则将被无条件覆盖。

`Path.resolve(strict=False)`

将路径绝对化, 解析任何符号链接。返回新的路径对象:

```
>>> p = Path()
>>> p
PosixPath('.')
>>> p.resolve()
PosixPath('/home/antoine/pathlib')
```

“..” 组件也将被消除 (只有这一种方法这么做):

```
>>> p = Path('docs/../setup.py')
>>> p.resolve()
PosixPath('/home/antoine/pathlib/setup.py')
```

如果路径不存在并且 *strict* 设为 `True`, 则抛出 `FileNotFoundError`。如果 *strict* 为 `False`, 则路径将被尽可能地解析并且任何剩余部分都会被不检查是否存在地追加。如果在解析路径上发生无限循环, 则抛出 `RuntimeError`。

3.6 新版功能: The *strict* argument.

`Path.rglob(pattern)`

This is like calling `Path.glob()` with “**” added in front of the given *pattern*:

```
>>> sorted(Path().rglob("*.py"))
[PosixPath('build/lib/pathlib.py'),
 PosixPath('docs/conf.py'),
 PosixPath('pathlib.py'),
 PosixPath('setup.py'),
 PosixPath('test_pathlib.py')]
```

`Path.rmdir()`

移除此目录。此目录必须为空的。

`Path.samefile(other_path)`

返回此目录是否指向与可能是字符串或者另一个路径对象的 *other_path* 相同的文件。语义类似于 `os.path.samefile()` 与 `os.path.samestat()`。

如果两者都以同一原因无法访问, 则抛出 `OSError`。

```
>>> p = Path('spam')
>>> q = Path('eggs')
>>> p.samefile(q)
False
>>> p.samefile('spam')
True
```

3.5 新版功能.

`Path.symlink_to(target, target_is_directory=False)`

将此路径创建为指向 *target* 的符号链接。在 Windows 下，如果链接的目标是一个目录则 *target_is_directory* 必须为 `true`（默认为 `False`）。在 POSIX 下，*target_is_directory* 的值将被忽略。

```
>>> p = Path('mylink')
>>> p.symlink_to('setup.py')
>>> p.resolve()
PosixPath('/home/antoine/pathlib/setup.py')
>>> p.stat().st_size
956
>>> p.lstat().st_size
8
```

注解： 参数的顺序 (*link, target*) 和 `os.symlink()` 是相反的。

`Path.touch(mode=0o666, exist_ok=True)`

将给定的路径创建为文件。如果给出了 *mode* 它将与当前进程的 `umask` 值合并以确定文件的模式和访问标志。如果文件已经存在，则当 *exist_ok* 为 `true` 则函数仍会成功（并且将它的修改事件更新为当前事件），否则抛出 `FileExistsError`。

`Path.unlink()`

移除此文件或符号链接。如果路径指向目录，则用 `Path.rmdir()` 代替。

`Path.write_bytes(data)`

将文件以二进制模式打开，写入 *data* 并关闭：

```
>>> p = Path('my_binary_file')
>>> p.write_bytes(b'Binary file contents')
20
>>> p.read_bytes()
b'Binary file contents'
```

一个同名的现存文件将被覆盖。

3.5 新版功能.

`Path.write_text(data, encoding=None, errors=None)`

将文件以文本模式打开，写入 *data* 并关闭：

```
>>> p = Path('my_text_file')
>>> p.write_text('Text file contents')
18
>>> p.read_text()
'Text file contents'
```

3.5 新版功能.

11.2 `os.path` — 常见路径操作

源代码: `Lib/posixpath.py` (适用 POSIX), `Lib/ntpath.py` (适用 Windows NT), and `Lib/macpath.py` (适用 Macintosh)

This module implements some useful functions on pathnames. To read or write files see `open()`, and for accessing the filesystem see the `os` module. The path parameters can be passed as either strings, or bytes. Applications are encouraged to represent file names as (Unicode) character strings. Unfortunately, some file names may not be representable as strings on Unix, so applications that need to support arbitrary file names on Unix should use bytes objects to represent path names. Vice versa, using bytes objects cannot represent all file names on Windows (in the standard `mbs` encoding), hence Windows applications should use string objects to access all files.

Unlike a unix shell, Python does not do any *automatic* path expansions. Functions such as `expanduser()` and `expandvars()` can be invoked explicitly when an application desires shell-like path expansion. (See also the `glob` module.)

参见:

The `pathlib` module offers high-level path objects.

注解: All of these functions accept either only bytes or only string objects as their parameters. The result is an object of the same type, if a path or file name is returned.

注解: Since different operating systems have different path name conventions, there are several versions of this module in the standard library. The `os.path` module is always the path module suitable for the operating system Python is running on, and therefore usable for local paths. However, you can also import and use the individual modules if you want to manipulate a path that is *always* in one of the different formats. They all have the same interface:

- `posixpath` for UNIX-style paths
- `ntpath` for Windows paths
- `macpath` for old-style MacOS paths

`os.path.abspath(path)`

Return a normalized absolutized version of the pathname `path`. On most platforms, this is equivalent to calling the function `normpath()` as follows: `normpath(join(os.getcwd(), path))`.

在 3.6 版更改: 接受一个类路径对象。

`os.path.basename(path)`

Return the base name of pathname `path`. This is the second element of the pair returned by passing `path` to the function `split()`. Note that the result of this function is different from the Unix `basename` program; where `basename` for `'/foo/bar/'` returns `'bar'`, the `basename()` function returns an empty string `''`.

在 3.6 版更改: 接受一个类路径对象。

`os.path.commonpath(paths)`

Return the longest common sub-path of each pathname in the sequence `paths`. Raise `ValueError` if `paths` contains both absolute and relative pathnames, or if `paths` is empty. Unlike `commonprefix()`, this returns a valid path.

Availability: Unix, Windows

3.5 新版功能.

在 3.6 版更改: Accepts a sequence of *path-like objects*.

`os.path.commonprefix(list)`

Return the longest path prefix (taken character-by-character) that is a prefix of all paths in *list*. If *list* is empty, return the empty string ('').

注解: This function may return invalid paths because it works a character at a time. To obtain a valid path, see `commonpath()`.

```
>>> os.path.commonprefix(['/usr/lib', '/usr/local/lib'])
'/usr/l'

>>> os.path.commonpath(['/usr/lib', '/usr/local/lib'])
'/usr'
```

在 3.6 版更改: 接受一个类路径对象。

`os.path.dirname(path)`

Return the directory name of pathname *path*. This is the first element of the pair returned by passing *path* to the function `split()`.

在 3.6 版更改: 接受一个类路径对象。

`os.path.exists(path)`

Return True if *path* refers to an existing path or an open file descriptor. Returns False for broken symbolic links. On some platforms, this function may return False if permission is not granted to execute `os.stat()` on the requested file, even if the *path* physically exists.

在 3.3 版更改: *path* can now be an integer: True is returned if it is an open file descriptor, False otherwise.

在 3.6 版更改: 接受一个类路径对象。

`os.path.lexists(path)`

Return True if *path* refers to an existing path. Returns True for broken symbolic links. Equivalent to `exists()` on platforms lacking `os.lstat()`.

在 3.6 版更改: 接受一个类路径对象。

`os.path.expanduser(path)`

On Unix and Windows, return the argument with an initial component of ~ or ~user replaced by that *user*'s home directory.

On Unix, an initial ~ is replaced by the environment variable HOME if it is set; otherwise the current user's home directory is looked up in the password directory through the built-in module `pwd`. An initial ~user is looked up directly in the password directory.

On Windows, HOME and USERPROFILE will be used if set, otherwise a combination of HOMEPATH and HOMEDRIVE will be used. An initial ~user is handled by stripping the last directory component from the created user path derived above.

If the expansion fails or if the path does not begin with a tilde, the path is returned unchanged.

在 3.6 版更改: 接受一个类路径对象。

`os.path.expandvars(path)`

Return the argument with environment variables expanded. Substrings of the form \$name or \${name} are replaced by the value of environment variable *name*. Malformed variable names and references to non-existing variables are left unchanged.

On Windows, %name% expansions are supported in addition to \$name and \${name}.

在 3.6 版更改: 接受一个类路径对象。

`os.path.getatime(path)`

Return the time of last access of *path*. The return value is a number giving the number of seconds since the epoch (see the *time* module). Raise *OSError* if the file does not exist or is inaccessible.

If `os.stat_float_times()` returns *True*, the result is a floating point number.

`os.path.getmtime(path)`

Return the time of last modification of *path*. The return value is a number giving the number of seconds since the epoch (see the *time* module). Raise *OSError* if the file does not exist or is inaccessible.

If `os.stat_float_times()` returns *True*, the result is a floating point number.

在 3.6 版更改: 接受一个类路径对象。

`os.path.getctime(path)`

Return the system's ctime which, on some systems (like Unix) is the time of the last metadata change, and, on others (like Windows), is the creation time for *path*. The return value is a number giving the number of seconds since the epoch (see the *time* module). Raise *OSError* if the file does not exist or is inaccessible.

在 3.6 版更改: 接受一个类路径对象。

`os.path.getsize(path)`

Return the size, in bytes, of *path*. Raise *OSError* if the file does not exist or is inaccessible.

在 3.6 版更改: 接受一个类路径对象。

`os.path.isabs(path)`

Return *True* if *path* is an absolute pathname. On Unix, that means it begins with a slash, on Windows that it begins with a (back)slash after chopping off a potential drive letter.

在 3.6 版更改: 接受一个类路径对象。

`os.path.isfile(path)`

Return *True* if *path* is an *existing* regular file. This follows symbolic links, so both *islink()* and *isfile()* can be true for the same path.

在 3.6 版更改: 接受一个类路径对象。

`os.path.isdir(path)`

Return *True* if *path* is an *existing* directory. This follows symbolic links, so both *islink()* and *isdir()* can be true for the same path.

在 3.6 版更改: 接受一个类路径对象。

`os.path.islink(path)`

Return *True* if *path* refers to an *existing* directory entry that is a symbolic link. Always *False* if symbolic links are not supported by the Python runtime.

在 3.6 版更改: 接受一个类路径对象。

`os.path.ismount(path)`

Return *True* if pathname *path* is a *mount point*: a point in a file system where a different file system has been mounted. On POSIX, the function checks whether *path*'s parent, *path/..*, is on a different device than *path*, or whether *path/..* and *path* point to the same i-node on the same device —this should detect mount points for all Unix and POSIX variants. On Windows, a drive letter root and a share UNC are always mount points, and for any other path *GetVolumePathName* is called to see if it is different from the input path.

3.4 新版功能: Support for detecting non-root mount points on Windows.

在 3.6 版更改: 接受一个类路径对象。

`os.path.join(path, *paths)`

Join one or more path components intelligently. The return value is the concatenation of *path* and any members of **paths* with exactly one directory separator (`os.sep`) following each non-empty part except the last, meaning

that the result will only end in a separator if the last part is empty. If a component is an absolute path, all previous components are thrown away and joining continues from the absolute path component.

On Windows, the drive letter is not reset when an absolute path component (e.g., `r'\foo'`) is encountered. If a component contains a drive letter, all previous components are thrown away and the drive letter is reset. Note that since there is a current directory for each drive, `os.path.join("c:", "foo")` represents a path relative to the current directory on drive C: (`c:foo`), not `c:\foo`.

在 3.6 版更改: Accepts a *path-like object* for *path* and *paths*.

`os.path.normcase(path)`

Normalize the case of a pathname. On Unix and Mac OS X, this returns the path unchanged; on case-insensitive filesystems, it converts the path to lowercase. On Windows, it also converts forward slashes to backward slashes. Raise a `TypeError` if the type of *path* is not `str` or `bytes` (directly or indirectly through the `os.PathLike` interface).

在 3.6 版更改: 接受一个类路径对象。

`os.path.normpath(path)`

Normalize a pathname by collapsing redundant separators and up-level references so that `A//B`, `A/B/`, `A/./B` and `A/foo/./B` all become `A/B`. This string manipulation may change the meaning of a path that contains symbolic links. On Windows, it converts forward slashes to backward slashes. To normalize case, use `normcase()`.

在 3.6 版更改: 接受一个类路径对象。

`os.path.realpath(path)`

Return the canonical path of the specified filename, eliminating any symbolic links encountered in the path (if they are supported by the operating system).

在 3.6 版更改: 接受一个类路径对象。

`os.path.relpath(path, start=os.curdir)`

Return a relative filepath to *path* either from the current directory or from an optional *start* directory. This is a path computation: the filesystem is not accessed to confirm the existence or nature of *path* or *start*.

start defaults to `os.curdir`.

Availability: Unix, Windows.

在 3.6 版更改: 接受一个类路径对象。

`os.path.samefile(path1, path2)`

Return `True` if both pathname arguments refer to the same file or directory. This is determined by the device number and i-node number and raises an exception if an `os.stat()` call on either pathname fails.

Availability: Unix, Windows.

在 3.2 版更改: Added Windows support.

在 3.4 版更改: Windows now uses the same implementation as all other platforms.

在 3.6 版更改: 接受一个类路径对象。

`os.path.sameopenfile(fp1, fp2)`

Return `True` if the file descriptors *fp1* and *fp2* refer to the same file.

Availability: Unix, Windows.

在 3.2 版更改: Added Windows support.

在 3.6 版更改: 接受一个类路径对象。

`os.path.samestat(stat1, stat2)`

Return `True` if the stat tuples *stat1* and *stat2* refer to the same file. These structures may have been returned by

`os.fstat()`, `os.lstat()`, or `os.stat()`. This function implements the underlying comparison used by `samefile()` and `sameopenfile()`.

Availability: Unix, Windows.

在 3.4 版更改: Added Windows support.

在 3.6 版更改: 接受一个类路径对象。

`os.path.split(path)`

Split the pathname *path* into a pair, (*head*, *tail*) where *tail* is the last pathname component and *head* is everything leading up to that. The *tail* part will never contain a slash; if *path* ends in a slash, *tail* will be empty. If there is no slash in *path*, *head* will be empty. If *path* is empty, both *head* and *tail* are empty. Trailing slashes are stripped from *head* unless it is the root (one or more slashes only). In all cases, `join(head, tail)` returns a path to the same location as *path* (but the strings may differ). Also see the functions `dirname()` and `basename()`.

在 3.6 版更改: 接受一个类路径对象。

`os.path.splitdrive(path)`

Split the pathname *path* into a pair (*drive*, *tail*) where *drive* is either a mount point or the empty string. On systems which do not use drive specifications, *drive* will always be the empty string. In all cases, *drive* + *tail* will be the same as *path*.

On Windows, splits a pathname into drive/UNC sharepoint and relative path.

If the path contains a drive letter, *drive* will contain everything up to and including the colon. e.g. `splitdrive("c:/dir")` returns ("c:", "/dir")

If the path contains a UNC path, *drive* will contain the host name and share, up to but not including the fourth separator. e.g. `splitdrive("//host/computer/dir")` returns ("//host/computer", "/dir")

在 3.6 版更改: 接受一个类路径对象。

`os.path.splitext(path)`

Split the pathname *path* into a pair (*root*, *ext*) such that *root* + *ext* == *path*, and *ext* is empty or begins with a period and contains at most one period. Leading periods on the basename are ignored; `splitext('.cshrc')` returns ('.cshrc', '').

在 3.6 版更改: 接受一个类路径对象。

`os.path.splitunc(path)`

3.1 版后已移除: Use `splitdrive` instead.

Split the pathname *path* into a pair (*unc*, *rest*) so that *unc* is the UNC mount point (such as `r'\\host\\mount'`), if present, and *rest* the rest of the path (such as `r'\\path\\file.ext'`). For paths containing drive letters, *unc* will always be the empty string.

Availability: Windows.

`os.path.supports_unicode_filenames`

True if arbitrary Unicode strings can be used as file names (within limitations imposed by the file system).

11.3 fileinput —Iterate over lines from multiple input streams

Source code: [Lib/fileinput.py](#)

This module implements a helper class and functions to quickly write a loop over standard input or a list of files. If you just want to read or write one file see `open()`.

The typical use is:

```
import fileinput
for line in fileinput.input():
    process(line)
```

This iterates over the lines of all files listed in `sys.argv[1:]`, defaulting to `sys.stdin` if the list is empty. If a filename is `'-'`, it is also replaced by `sys.stdin`. To specify an alternative list of filenames, pass it as the first argument to `input()`. A single file name is also allowed.

All files are opened in text mode by default, but you can override this by specifying the `mode` parameter in the call to `input()` or `FileInput`. If an I/O error occurs during opening or reading a file, `IOError` is raised.

在 3.3 版更改: `IOError` used to be raised; it is now an alias of `IOError`.

If `sys.stdin` is used more than once, the second and further use will return no lines, except perhaps for interactive use, or if it has been explicitly reset (e.g. using `sys.stdin.seek(0)`).

Empty files are opened and immediately closed; the only time their presence in the list of filenames is noticeable at all is when the last file opened is empty.

Lines are returned with any newlines intact, which means that the last line in a file may not have one.

You can control how files are opened by providing an opening hook via the `openhook` parameter to `fileinput.input()` or `FileInput`. The hook must be a function that takes two arguments, `filename` and `mode`, and returns an accordingly opened file-like object. Two useful hooks are already provided by this module.

The following function is the primary interface of this module:

`fileinput.input(files=None, inplace=False, backup="", bufsize=0, mode='r', openhook=None)`

Create an instance of the `FileInput` class. The instance will be used as global state for the functions of this module, and is also returned to use during iteration. The parameters to this function will be passed along to the constructor of the `FileInput` class.

The `FileInput` instance can be used as a context manager in the `with` statement. In this example, `input` is closed after the `with` statement is exited, even if an exception occurs:

```
with fileinput.input(files=('spam.txt', 'eggs.txt')) as f:
    for line in f:
        process(line)
```

在 3.2 版更改: Can be used as a context manager.

Deprecated since version 3.6, will be removed in version 3.8: The `bufsize` parameter.

The following functions use the global state created by `fileinput.input()`; if there is no active state, `RuntimeError` is raised.

`fileinput.filename()`

Return the name of the file currently being read. Before the first line has been read, returns `None`.

`fileinput.fileno()`

Return the integer “file descriptor” for the current file. When no file is opened (before the first line and between files), returns `-1`.

`fileinput.lineno()`

Return the cumulative line number of the line that has just been read. Before the first line has been read, returns `0`. After the last line of the last file has been read, returns the line number of that line.

`fileinput.filelineno()`

Return the line number in the current file. Before the first line has been read, returns `0`. After the last line of the last file has been read, returns the line number of that line within the file.

`fileinput.isfirstline()`

Returns true if the line just read is the first line of its file, otherwise returns false.

`fileinput.isstdin()`

Returns true if the last line was read from `sys.stdin`, otherwise returns false.

`fileinput.nextfile()`

Close the current file so that the next iteration will read the first line from the next file (if any); lines not read from the file will not count towards the cumulative line count. The filename is not changed until after the first line of the next file has been read. Before the first line has been read, this function has no effect; it cannot be used to skip the first file. After the last line of the last file has been read, this function has no effect.

`fileinput.close()`

Close the sequence.

The class which implements the sequence behavior provided by the module is available for subclassing as well:

class `fileinput.FileInput` (*files=None*, *inplace=False*, *backup=""*, *bufsize=0*, *mode='r'*, *openhook=None*)

Class `FileInput` is the implementation; its methods `filename()`, `fileno()`, `lineno()`, `filelineno()`, `isfirstline()`, `isstdin()`, `nextfile()` and `close()` correspond to the functions of the same name in the module. In addition it has a `readline()` method which returns the next input line, and a `__getitem__()` method which implements the sequence behavior. The sequence must be accessed in strictly sequential order; random access and `readline()` cannot be mixed.

With *mode* you can specify which file mode will be passed to `open()`. It must be one of `'r'`, `'rU'`, `'U'` and `'rb'`.

The *openhook*, when given, must be a function that takes two arguments, *filename* and *mode*, and returns an accordingly opened file-like object. You cannot use *inplace* and *openhook* together.

A `FileInput` instance can be used as a context manager in the `with` statement. In this example, *input* is closed after the `with` statement is exited, even if an exception occurs:

```
with FileInput(files=('spam.txt', 'eggs.txt')) as input:
    process(input)
```

在 3.2 版更改: Can be used as a context manager.

3.4 版后已移除: The `'rU'` and `'U'` modes.

Deprecated since version 3.6, will be removed in version 3.8: The *bufsize* parameter.

Optional in-place filtering: if the keyword argument `inplace=True` is passed to `fileinput.input()` or to the `FileInput` constructor, the file is moved to a backup file and standard output is directed to the input file (if a file of the same name as the backup file already exists, it will be replaced silently). This makes it possible to write a filter that rewrites its input file in place. If the *backup* parameter is given (typically as `backup='.<some extension>'`), it specifies the extension for the backup file, and the backup file remains around; by default, the extension is `'.bak'` and it is deleted when the output file is closed. In-place filtering is disabled when standard input is read.

The two following opening hooks are provided by this module:

`fileinput.hook_compressed(filename, mode)`

Transparently opens files compressed with `gzip` and `bzip2` (recognized by the extensions `'.gz'` and `'.bz2'`) using the `gzip` and `bz2` modules. If the filename extension is not `'.gz'` or `'.bz2'`, the file is opened normally (ie, using `open()` without any decompression).

Usage example: `fi = fileinput.FileInput(openhook=fileinput.hook_compressed)`

`fileinput.hook_encoded(encoding, errors=None)`

Returns a hook which opens each file with `open()`, using the given *encoding* and *errors* to read the file.

Usage example: `fi = fileinput.FileInput(openhook=fileinput.hook_encoded("utf-8", "surrogateescape"))`

在 3.6 版更改: Added the optional *errors* parameter.

11.4 stat —Interpreting stat () results

Source code: [Lib/stat.py](#)

The *stat* module defines constants and functions for interpreting the results of *os.stat()*, *os.fstat()* and *os.lstat()* (if they exist). For complete details about the *stat()*, *fstat()* and *lstat()* calls, consult the documentation for your system.

在 3.4 版更改: The *stat* module is backed by a C implementation.

The *stat* module defines the following functions to test for specific file types:

stat.S_ISDIR(mode)
Return non-zero if the mode is from a directory.

stat.S_ISCHR(mode)
Return non-zero if the mode is from a character special device file.

stat.S_ISBLK(mode)
Return non-zero if the mode is from a block special device file.

stat.S_ISREG(mode)
Return non-zero if the mode is from a regular file.

stat.S_ISFIFO(mode)
Return non-zero if the mode is from a FIFO (named pipe).

stat.S_ISLNK(mode)
Return non-zero if the mode is from a symbolic link.

stat.S_ISSOCK(mode)
Return non-zero if the mode is from a socket.

stat.S_ISDOOR(mode)
Return non-zero if the mode is from a door.

3.4 新版功能.

stat.S_ISPORT(mode)
Return non-zero if the mode is from an event port.

3.4 新版功能.

stat.S_ISWHT(mode)
Return non-zero if the mode is from a whiteout.

3.4 新版功能.

Two additional functions are defined for more general manipulation of the file's mode:

stat.S_IMODE(mode)
Return the portion of the file's mode that can be set by *os.chmod()*—that is, the file's permission bits, plus the sticky bit, set-group-id, and set-user-id bits (on systems that support them).

stat.S_IFMT(mode)
Return the portion of the file's mode that describes the file type (used by the *S_IS**() functions above).

Normally, you would use the `os.path.is*()` functions for testing the type of a file; the functions here are useful when you are doing multiple tests of the same file and wish to avoid the overhead of the `stat()` system call for each test. These are also useful when checking for information about a file that isn't handled by `os.path`, like the tests for block and character devices.

示例:

```
import os, sys
from stat import *

def walktree(top, callback):
    '''recursively descend the directory tree rooted at top,
       calling the callback function for each regular file'''

    for f in os.listdir(top):
        pathname = os.path.join(top, f)
        mode = os.stat(pathname).st_mode
        if S_ISDIR(mode):
            # It's a directory, recurse into it
            walktree(pathname, callback)
        elif S_ISREG(mode):
            # It's a file, call the callback function
            callback(pathname)
        else:
            # Unknown file type, print a message
            print('Skipping %s' % pathname)

def visitfile(file):
    print('visiting', file)

if __name__ == '__main__':
    walktree(sys.argv[1], visitfile)
```

An additional utility function is provided to convert a file's mode in a human readable string:

`stat.filemode(mode)`
Convert a file's mode to a string of the form `'-rwxrwxrwx'`.

3.3 新版功能.

在 3.4 版更改: The function supports `S_IFDOOR`, `S_IFPORT` and `S_IFWHT`.

All the variables below are simply symbolic indexes into the 10-tuple returned by `os.stat()`, `os.fstat()` or `os.lstat()`.

`stat.ST_MODE`
Inode protection mode.

`stat.ST_INO`
Inode number.

`stat.ST_DEV`
Device inode resides on.

`stat.ST_NLINK`
Number of links to the inode.

`stat.ST_UID`
User id of the owner.

`stat.ST_GID`
Group id of the owner.

`stat.ST_SIZE`

Size in bytes of a plain file; amount of data waiting on some special files.

`stat.ST_ATIME`

Time of last access.

`stat.ST_MTIME`

Time of last modification.

`stat.ST_CTIME`

The “ctime” as reported by the operating system. On some systems (like Unix) is the time of the last metadata change, and, on others (like Windows), is the creation time (see platform documentation for details).

The interpretation of “file size” changes according to the file type. For plain files this is the size of the file in bytes. For FIFOs and sockets under most flavors of Unix (including Linux in particular), the “size” is the number of bytes waiting to be read at the time of the call to `os.stat()`, `os.fstat()`, or `os.lstat()`; this can sometimes be useful, especially for polling one of these special files after a non-blocking open. The meaning of the size field for other character and block devices varies more, depending on the implementation of the underlying system call.

The variables below define the flags used in the `ST_MODE` field.

Use of the functions above is more portable than use of the first set of flags:

`stat.S_IFSOCK`

Socket.

`stat.S_IFLNK`

Symbolic link.

`stat.S_IFREG`

Regular file.

`stat.S_IFBLK`

Block device.

`stat.S_IFDIR`

Directory.

`stat.S_IFCHR`

Character device.

`stat.S_IFIFO`

FIFO.

`stat.S_IFDOOR`

Door.

3.4 新版功能.

`stat.S_IFPORT`

Event port.

3.4 新版功能.

`stat.S_IFWHT`

Whiteout.

3.4 新版功能.

注解: `S_IFDOOR`, `S_IFPORT` or `S_IFWHT` are defined as 0 when the platform does not have support for the file types.

The following flags can also be used in the *mode* argument of `os.chmod()`:

`stat.S_ISUID`
Set UID bit.

`stat.S_ISGID`
Set-group-ID bit. This bit has several special uses. For a directory it indicates that BSD semantics is to be used for that directory: files created there inherit their group ID from the directory, not from the effective group ID of the creating process, and directories created there will also get the `S_ISGID` bit set. For a file that does not have the group execution bit (`S_IXGRP`) set, the set-group-ID bit indicates mandatory file/record locking (see also `S_ENFMT`).

`stat.S_ISVTX`
Sticky bit. When this bit is set on a directory it means that a file in that directory can be renamed or deleted only by the owner of the file, by the owner of the directory, or by a privileged process.

`stat.S_IRWXU`
Mask for file owner permissions.

`stat.S_IRUSR`
Owner has read permission.

`stat.S_IWUSR`
Owner has write permission.

`stat.S_IXUSR`
Owner has execute permission.

`stat.S_IRWXG`
Mask for group permissions.

`stat.S_IRGRP`
Group has read permission.

`stat.S_IWGRP`
Group has write permission.

`stat.S_IXGRP`
Group has execute permission.

`stat.S_IRWXO`
Mask for permissions for others (not in group).

`stat.S_IROTH`
Others have read permission.

`stat.S_IWOTH`
Others have write permission.

`stat.S_IXOTH`
Others have execute permission.

`stat.S_ENFMT`
System V file locking enforcement. This flag is shared with `S_ISGID`: file/record locking is enforced on files that do not have the group execution bit (`S_IXGRP`) set.

`stat.S_IREAD`
Unix V7 synonym for `S_IRUSR`.

`stat.S_IWRITE`
Unix V7 synonym for `S_IWUSR`.

`stat.S_IXEXEC`

Unix V7 synonym for `S_IXUSR`.

The following flags can be used in the *flags* argument of `os.chflags()`:

`stat.UF_NODUMP`

Do not dump the file.

`stat.UF_IMMUTABLE`

The file may not be changed.

`stat.UF_APPEND`

The file may only be appended to.

`stat.UF_OPAQUE`

The directory is opaque when viewed through a union stack.

`stat.UF_NOUNLINK`

The file may not be renamed or deleted.

`stat.UF_COMPRESSED`

The file is stored compressed (Mac OS X 10.6+).

`stat.UF_HIDDEN`

The file should not be displayed in a GUI (Mac OS X 10.5+).

`stat.SF_ARCHIVED`

The file may be archived.

`stat.SF_IMMUTABLE`

The file may not be changed.

`stat.SF_APPEND`

The file may only be appended to.

`stat.SF_NOUNLINK`

The file may not be renamed or deleted.

`stat.SF_SNAPSHOT`

The file is a snapshot file.

See the *BSD or Mac OS systems man page `chflags(2)` for more information.

On Windows, the following file attribute constants are available for use when testing bits in the `st_file_attributes` member returned by `os.stat()`. See the [Windows API documentation](#) for more detail on the meaning of these constants.

`stat.FILE_ATTRIBUTE_ARCHIVE`

`stat.FILE_ATTRIBUTE_COMPRESSED`

`stat.FILE_ATTRIBUTE_DEVICE`

`stat.FILE_ATTRIBUTE_DIRECTORY`

`stat.FILE_ATTRIBUTE_ENCRYPTED`

`stat.FILE_ATTRIBUTE_HIDDEN`

`stat.FILE_ATTRIBUTE_INTEGRITY_STREAM`

`stat.FILE_ATTRIBUTE_NORMAL`

`stat.FILE_ATTRIBUTE_NOT_CONTENT_INDEXED`

`stat.FILE_ATTRIBUTE_NO_SCRUB_DATA`

`stat.FILE_ATTRIBUTE_OFFLINE`

`stat.FILE_ATTRIBUTE_READONLY`

`stat.FILE_ATTRIBUTE_REPARSE_POINT`

`stat.FILE_ATTRIBUTE_SPARSE_FILE`

`stat.FILE_ATTRIBUTE_SYSTEM`

`stat.FILE_ATTRIBUTE_TEMPORARY`
`stat.FILE_ATTRIBUTE_VIRTUAL`
 3.5 新版功能.

11.5 filecmp —File and Directory Comparisons

Source code: [Lib/filecmp.py](#)

The *filecmp* module defines functions to compare files and directories, with various optional time/correctness trade-offs. For comparing files, see also the *difflib* module.

The *filecmp* module defines the following functions:

`filecmp.cmp(f1, f2, shallow=True)`

Compare the files named *f1* and *f2*, returning `True` if they seem equal, `False` otherwise.

If *shallow* is true, files with identical `os.stat()` signatures are taken to be equal. Otherwise, the contents of the files are compared.

Note that no external programs are called from this function, giving it portability and efficiency.

This function uses a cache for past comparisons and the results, with cache entries invalidated if the `os.stat()` information for the file changes. The entire cache may be cleared using `clear_cache()`.

`filecmp.cmpfiles(dir1, dir2, common, shallow=True)`

Compare the files in the two directories *dir1* and *dir2* whose names are given by *common*.

Returns three lists of file names: *match*, *mismatch*, *errors*. *match* contains the list of files that match, *mismatch* contains the names of those that don't, and *errors* lists the names of files which could not be compared. Files are listed in *errors* if they don't exist in one of the directories, the user lacks permission to read them or if the comparison could not be done for some other reason.

The *shallow* parameter has the same meaning and default value as for `filecmp.cmp()`.

For example, `cmpfiles('a', 'b', ['c', 'd/e'])` will compare *a/c* with *b/c* and *a/d/e* with *b/d/e*. 'c' and 'd/e' will each be in one of the three returned lists.

`filecmp.clear_cache()`

Clear the filecmp cache. This may be useful if a file is compared so quickly after it is modified that it is within the mtime resolution of the underlying filesystem.

3.4 新版功能.

11.5.1 The dircmp class

`class filecmp.dircmp(a, b, ignore=None, hide=None)`

Construct a new directory comparison object, to compare the directories *a* and *b*. *ignore* is a list of names to ignore, and defaults to `filecmp.DEFAULT_IGNORES`. *hide* is a list of names to hide, and defaults to `[os.curdir, os.pardir]`.

The *dircmp* class compares files by doing *shallow* comparisons as described for `filecmp.cmp()`.

The *dircmp* class provides the following methods:

`report()`

Print (to `sys.stdout`) a comparison between *a* and *b*.

report_partial_closure()

Print a comparison between *a* and *b* and common immediate subdirectories.

report_full_closure()

Print a comparison between *a* and *b* and common subdirectories (recursively).

The `dircmp` class offers a number of interesting attributes that may be used to get various bits of information about the directory trees being compared.

Note that via `__getattr__()` hooks, all attributes are computed lazily, so there is no speed penalty if only those attributes which are lightweight to compute are used.

left

The directory *a*.

right

The directory *b*.

left_list

Files and subdirectories in *a*, filtered by *hide* and *ignore*.

right_list

Files and subdirectories in *b*, filtered by *hide* and *ignore*.

common

Files and subdirectories in both *a* and *b*.

left_only

Files and subdirectories only in *a*.

right_only

Files and subdirectories only in *b*.

common_dirs

Subdirectories in both *a* and *b*.

common_files

Files in both *a* and *b*.

common_funny

Names in both *a* and *b*, such that the type differs between the directories, or names for which `os.stat()` reports an error.

same_files

Files which are identical in both *a* and *b*, using the class' s file comparison operator.

diff_files

Files which are in both *a* and *b*, whose contents differ according to the class' s file comparison operator.

funny_files

Files which are in both *a* and *b*, but could not be compared.

subdirs

A dictionary mapping names in `common_dirs` to `dircmp` objects.

`filecmp.DEFAULT_IGNORES`

3.4 新版功能.

List of directories ignored by `dircmp` by default.

Here is a simplified example of using the `subdirs` attribute to search recursively through two directories to show common different files:

```
>>> from filecmp import dircmp
>>> def print_diff_files(dcmp):
...     for name in dcmp.diff_files:
...         print("diff_file %s found in %s and %s" % (name, dcmp.left,
...             dcmp.right))
...     for sub_dcmp in dcmp.subdirs.values():
...         print_diff_files(sub_dcmp)
...
>>> dcmp = dircmp('dir1', 'dir2')
>>> print_diff_files(dcmp)
```

11.6 tempfile —Generate temporary files and directories

Source code: [Lib/tempfile.py](#)

This module creates temporary files and directories. It works on all supported platforms. *TemporaryFile*, *NamedTemporaryFile*, *TemporaryDirectory*, and *SpooledTemporaryFile* are high-level interfaces which provide automatic cleanup and can be used as context managers. *mkstemp()* and *mkdtemp()* are lower-level functions which require manual cleanup.

All the user-callable functions and constructors take additional arguments which allow direct control over the location and name of temporary files and directories. Files names used by this module include a string of random characters which allows those files to be securely created in shared temporary directories. To maintain backward compatibility, the argument order is somewhat odd; it is recommended to use keyword arguments for clarity.

The module defines the following user-callable items:

`tempfile.TemporaryFile(mode='w+b', buffering=None, encoding=None, newline=None, suffix=None, prefix=None, dir=None)`

Return a *file-like object* that can be used as a temporary storage area. The file is created securely, using the same rules as *mkstemp()*. It will be destroyed as soon as it is closed (including an implicit close when the object is garbage collected). Under Unix, the directory entry for the file is either not created at all or is removed immediately after the file is created. Other platforms do not support this; your code should not rely on a temporary file created using this function having or not having a visible name in the file system.

The resulting object can be used as a context manager (see 示例). On completion of the context or destruction of the file object the temporary file will be removed from the filesystem.

The *mode* parameter defaults to 'w+b' so that the file created can be read and written without being closed. Binary mode is used so that it behaves consistently on all platforms without regard for the data that is stored. *buffering*, *encoding* and *newline* are interpreted as for *open()*.

The *dir*, *prefix* and *suffix* parameters have the same meaning and defaults as with *mkstemp()*.

The returned object is a true file object on POSIX platforms. On other platforms, it is a file-like object whose *file* attribute is the underlying true file object.

The *os.O_TMPFILE* flag is used if it is available and works (Linux-specific, requires Linux kernel 3.11 or later).

在 3.5 版更改: The *os.O_TMPFILE* flag is now used if available.

`tempfile.NamedTemporaryFile(mode='w+b', buffering=None, encoding=None, newline=None, suffix=None, prefix=None, dir=None, delete=True)`

This function operates exactly as *TemporaryFile()* does, except that the file is guaranteed to have a visible name in the file system (on Unix, the directory entry is not unlinked). That name can be retrieved from the *name* attribute of the returned file-like object. Whether the name can be used to open the file a second time, while the

named temporary file is still open, varies across platforms (it can be so used on Unix; it cannot on Windows NT or later). If *delete* is true (the default), the file is deleted as soon as it is closed. The returned object is always a file-like object whose *file* attribute is the underlying true file object. This file-like object can be used in a *with* statement, just like a normal file.

`tempfile.SpooledTemporaryFile` (*max_size=0, mode='w+b', buffering=None, encoding=None, new-line=None, suffix=None, prefix=None, dir=None*)

This function operates exactly as `TemporaryFile()` does, except that data is spooled in memory until the file size exceeds *max_size*, or until the file's `fileno()` method is called, at which point the contents are written to disk and operation proceeds as with `TemporaryFile()`.

The resulting file has one additional method, `rollover()`, which causes the file to roll over to an on-disk file regardless of its size.

The returned object is a file-like object whose `_file` attribute is either an `io.BytesIO` or `io.StringIO` object (depending on whether binary or text *mode* was specified) or a true file object, depending on whether `rollover()` has been called. This file-like object can be used in a *with* statement, just like a normal file.

在 3.3 版更改: the truncate method now accepts a *size* argument.

`tempfile.TemporaryDirectory` (*suffix=None, prefix=None, dir=None*)

This function securely creates a temporary directory using the same rules as `mkdtemp()`. The resulting object can be used as a context manager (see 示例). On completion of the context or destruction of the temporary directory object the newly created temporary directory and all its contents are removed from the filesystem.

The directory name can be retrieved from the *name* attribute of the returned object. When the returned object is used as a context manager, the *name* will be assigned to the target of the *as* clause in the *with* statement, if there is one.

The directory can be explicitly cleaned up by calling the `cleanup()` method.

3.2 新版功能.

`tempfile.mkstemp` (*suffix=None, prefix=None, dir=None, text=False*)

Creates a temporary file in the most secure manner possible. There are no race conditions in the file's creation, assuming that the platform properly implements the `os.O_EXCL` flag for `os.open()`. The file is readable and writable only by the creating user ID. If the platform uses permission bits to indicate whether a file is executable, the file is executable by no one. The file descriptor is not inherited by child processes.

Unlike `TemporaryFile()`, the user of `mkstemp()` is responsible for deleting the temporary file when done with it.

If *suffix* is not *None*, the file name will end with that suffix, otherwise there will be no suffix. `mkstemp()` does not put a dot between the file name and the suffix; if you need one, put it at the beginning of *suffix*.

If *prefix* is not *None*, the file name will begin with that prefix; otherwise, a default prefix is used. The default is the return value of `gettempprefix()` or `gettempprefixb()`, as appropriate.

If *dir* is not *None*, the file will be created in that directory; otherwise, a default directory is used. The default directory is chosen from a platform-dependent list, but the user of the application can control the directory location by setting the `TMPDIR`, `TEMP` or `TMP` environment variables. There is thus no guarantee that the generated filename will have any nice properties, such as not requiring quoting when passed to external commands via `os.popen()`.

If any of *suffix*, *prefix*, and *dir* are not *None*, they must be the same type. If they are bytes, the returned name will be bytes instead of str. If you want to force a bytes return value with otherwise default behavior, pass `suffix=b''`.

If *text* is specified, it indicates whether to open the file in binary mode (the default) or text mode. On some platforms, this makes no difference.

`mkstemp()` returns a tuple containing an OS-level handle to an open file (as would be returned by `os.open()`) and the absolute pathname of that file, in that order.

在 3.5 版更改: *suffix*, *prefix*, and *dir* may now be supplied in bytes in order to obtain a bytes return value. Prior to this, only str was allowed. *suffix* and *prefix* now accept and default to `None` to cause an appropriate default value to be used.

`tempfile.mkdtemp(suffix=None, prefix=None, dir=None)`

Creates a temporary directory in the most secure manner possible. There are no race conditions in the directory's creation. The directory is readable, writable, and searchable only by the creating user ID.

The user of `mkdtemp()` is responsible for deleting the temporary directory and its contents when done with it.

The *prefix*, *suffix*, and *dir* arguments are the same as for `mkstemp()`.

`mkdtemp()` returns the absolute pathname of the new directory.

在 3.5 版更改: *suffix*, *prefix*, and *dir* may now be supplied in bytes in order to obtain a bytes return value. Prior to this, only str was allowed. *suffix* and *prefix* now accept and default to `None` to cause an appropriate default value to be used.

`tempfile.gettempdir()`

Return the name of the directory used for temporary files. This defines the default value for the *dir* argument to all functions in this module.

Python searches a standard list of directories to find one which the calling user can create files in. The list is:

1. The directory named by the `TMPDIR` environment variable.
2. The directory named by the `TEMP` environment variable.
3. The directory named by the `TMP` environment variable.
4. A platform-specific location:
 - On Windows, the directories `C:\TEMP`, `C:\TMP`, `\TEMP`, and `\TMP`, in that order.
 - On all other platforms, the directories `/tmp`, `/var/tmp`, and `/usr/tmp`, in that order.
5. As a last resort, the current working directory.

The result of this search is cached, see the description of `tempdir` below.

`tempfile.gettempdirb()`

Same as `gettempdir()` but the return value is in bytes.

3.5 新版功能.

`tempfile.gettempprefix()`

Return the filename prefix used to create temporary files. This does not contain the directory component.

`tempfile.gettempprefixb()`

Same as `gettempprefix()` but the return value is in bytes.

3.5 新版功能.

The module uses a global variable to store the name of the directory used for temporary files returned by `gettempdir()`. It can be set directly to override the selection process, but this is discouraged. All functions in this module take a *dir* argument which can be used to specify the directory and this is the recommended approach.

`tempfile.tempdir`

When set to a value other than `None`, this variable defines the default value for the *dir* argument to the functions defined in this module.

If `tempdir` is `None` (the default) at any call to any of the above functions except `gettempprefix()` it is initialized following the algorithm described in `gettempdir()`.

11.6.1 示例

Here are some examples of typical usage of the `tempfile` module:

```
>>> import tempfile

# create a temporary file and write some data to it
>>> fp = tempfile.TemporaryFile()
>>> fp.write(b'Hello world!')
# read data from file
>>> fp.seek(0)
>>> fp.read()
b'Hello world!'
# close the file, it will be removed
>>> fp.close()

# create a temporary file using a context manager
>>> with tempfile.TemporaryFile() as fp:
...     fp.write(b'Hello world!')
...     fp.seek(0)
...     fp.read()
b'Hello world!'
>>>
# file is now closed and removed

# create a temporary directory using the context manager
>>> with tempfile.TemporaryDirectory() as tmpdirname:
...     print('created temporary directory', tmpdirname)
>>>
# directory and contents have been removed
```

11.6.2 Deprecated functions and variables

A historical way to create temporary files was to first generate a file name with the `mktemp()` function and then create a file using this name. Unfortunately this is not secure, because a different process may create a file with this name in the time between the call to `mktemp()` and the subsequent attempt to create the file by the first process. The solution is to combine the two steps and create the file immediately. This approach is used by `mkstemp()` and the other functions described above.

`tempfile.mktemp(suffix='', prefix='tmp', dir=None)`

2.3 版后已移除: Use `mkstemp()` instead.

Return an absolute pathname of a file that did not exist at the time the call is made. The *prefix*, *suffix*, and *dir* arguments are similar to those of `mkstemp()`, except that bytes file names, `suffix=None` and `prefix=None` are not supported.

警告: Use of this function may introduce a security hole in your program. By the time you get around to doing anything with the file name it returns, someone else may have beaten you to the punch. `mktemp()` usage can be replaced easily with `NamedTemporaryFile()`, passing it the `delete=False` parameter:

```
>>> f = NamedTemporaryFile(delete=False)
>>> f.name
'/tmp/tmpjtujjt'
>>> f.write(b"Hello World!\n")
13
>>> f.close()
>>> os.unlink(f.name)
>>> os.path.exists(f.name)
False
```

11.7 glob —Unix style pathname pattern expansion

Source code: [Lib/glob.py](#)

The `glob` module finds all the pathnames matching a specified pattern according to the rules used by the Unix shell, although results are returned in arbitrary order. No tilde expansion is done, but `*`, `?`, and character ranges expressed with `[]` will be correctly matched. This is done by using the `os.scandir()` and `fnmatch.fnmatch()` functions in concert, and not by actually invoking a subshell. Note that unlike `fnmatch.fnmatch()`, `glob` treats filenames beginning with a dot (`.`) as special cases. (For tilde and shell variable expansion, use `os.path.expanduser()` and `os.path.expandvars()`.)

For a literal match, wrap the meta-characters in brackets. For example, `'[?]'` matches the character `'?'`.

参见:

The `pathlib` module offers high-level path objects.

`glob.glob(pathname, *, recursive=False)`

Return a possibly-empty list of path names that match `pathname`, which must be a string containing a path specification. `pathname` can be either absolute (like `/usr/src/Python-1.5/Makefile`) or relative (like `.././Tools/*/*.gif`), and can contain shell-style wildcards. Broken symlinks are included in the results (as in the shell).

If `recursive` is true, the pattern `"**"` will match any files and zero or more directories and subdirectories. If the pattern is followed by an `os.sep`, only directories and subdirectories match.

注解: 在一个较大的目录树中使用 `"**"` 模式可能会消耗非常多的时间。

在 3.5 版更改: Support for recursive globs using `"**"`.

`glob.iglob(pathname, *, recursive=False)`

Return an *iterator* which yields the same values as `glob()` without actually storing them all simultaneously.

`glob.escape(pathname)`

Escape all special characters (`'?'`, `'*'` and `'['`). This is useful if you want to match an arbitrary literal string that may have special characters in it. Special characters in drive/UNC sharepoints are not escaped, e.g. on Windows `escape('///?/c:/Quo vadis?.txt')` returns `'///?/c:/Quo vadis[?].txt'`.

3.4 新版功能.

For example, consider a directory containing the following files: `1.gif`, `2.txt`, `card.gif` and a subdirectory `sub` which contains only the file `3.txt`. `glob()` will produce the following results. Notice how any leading components of the path are preserved.

```
>>> import glob
>>> glob.glob('./[0-9].*')
['./1.gif', './2.txt']
>>> glob.glob('*.gif')
['1.gif', 'card.gif']
>>> glob.glob('?.gif')
['1.gif']
>>> glob.glob('**/*.txt', recursive=True)
```

(下页继续)

(续上页)

```
['2.txt', 'sub/3.txt']
>>> glob.glob('.**/', recursive=True)
['./', './sub/']
```

If the directory contains files starting with `.` they won't be matched by default. For example, consider a directory containing `card.gif` and `.card.gif`:

```
>>> import glob
>>> glob.glob('*.gif')
['card.gif']
>>> glob.glob('.*')
['.card.gif']
```

参见:

Module `fnmatch` Shell-style filename (not path) expansion

11.8 fnmatch — Unix filename pattern matching

Source code: [Lib/fnmatch.py](#)

This module provides support for Unix shell-style wildcards, which are *not* the same as regular expressions (which are documented in the `re` module). The special characters used in shell-style wildcards are:

模式	意义
<code>*</code>	匹配所有
<code>?</code>	匹配任何单个字符
<code>[seq]</code>	匹配 <i>seq</i> 中的任何字符
<code>[!seq]</code>	匹配任何不在 <i>seq</i> 中的字符

For a literal match, wrap the meta-characters in brackets. For example, `'[?]'` matches the character `'?'`.

Note that the filename separator (`'/'` on Unix) is *not* special to this module. See module `glob` for pathname expansion (`glob` uses `filter()` to match pathname segments). Similarly, filenames starting with a period are not special for this module, and are matched by the `*` and `?` patterns.

`fnmatch.fnmatch(filename, pattern)`

Test whether the *filename* string matches the *pattern* string, returning `True` or `False`. Both parameters are case-normalized using `os.path.normcase()`. `fnmatchcase()` can be used to perform a case-sensitive comparison, regardless of whether that's standard for the operating system.

This example will print all file names in the current directory with the extension `.txt`:

```
import fnmatch
import os

for file in os.listdir('.'):
    if fnmatch.fnmatch(file, '*.txt'):
        print(file)
```

`fnmatch.fnmatchcase(filename, pattern)`

Test whether *filename* matches *pattern*, returning `True` or `False`; the comparison is case-sensitive and does not apply `os.path.normcase()`.

`fnmatch.filter(names, pattern)`

Return the subset of the list of *names* that match *pattern*. It is the same as `[n for n in names if fnmatch(n, pattern)]`, but implemented more efficiently.

`fnmatch.translate(pattern)`

Return the shell-style *pattern* converted to a regular expression for using with `re.match()`.

示例:

```
>>> import fnmatch, re
>>>
>>> regex = fnmatch.translate('*.txt')
>>> regex
'(?s:.*\\.txt)\\Z'
>>> reobj = re.compile(regex)
>>> reobj.match('foobar.txt')
<_sre.SRE_Match object; span=(0, 10), match='foobar.txt'>
```

参见:

Module `glob` Unix shell-style path expansion.

11.9 linecache — Random access to text lines

Source code: [Lib/linecache.py](#)

The `linecache` module allows one to get any line from a Python source file, while attempting to optimize internally, using a cache, the common case where many lines are read from a single file. This is used by the `traceback` module to retrieve source lines for inclusion in the formatted traceback.

The `tokenize.open()` function is used to open files. This function uses `tokenize.detect_encoding()` to get the encoding of the file; in the absence of an encoding token, the file encoding defaults to UTF-8.

The `linecache` module defines the following functions:

`linecache.getline(filename, lineno, module_globals=None)`

Get line *lineno* from file named *filename*. This function will never raise an exception—it will return `' '` on errors (the terminating newline character will be included for lines that are found).

If a file named *filename* is not found, the function will look for it in the module search path, `sys.path`, after first checking for a [PEP 302](#) `__loader__` in *module_globals*, in case the module was imported from a zipfile or other non-filesystem import source.

`linecache.clearcache()`

Clear the cache. Use this function if you no longer need lines from files previously read using `getline()`.

`linecache.checkcache(filename=None)`

Check the cache for validity. Use this function if files in the cache may have changed on disk, and you require the updated version. If *filename* is omitted, it will check all the entries in the cache.

`linecache.lazycache(filename, module_globals)`

Capture enough detail about a non-file-based module to permit getting its lines later via `getline()` even if *module_globals* is `None` in the later call. This avoids doing I/O until a line is actually needed, without having to carry the module globals around indefinitely.

3.5 新版功能.

示例:

```
>>> import linecache
>>> linecache.getline(linecache.__file__, 8)
'import sys\n'
```

11.10 `shutil` — High-level file operations

Source code: [Lib/shutil.py](#)

The `shutil` module offers a number of high-level operations on files and collections of files. In particular, functions are provided which support file copying and removal. For operations on individual files, see also the `os` module.

警告: Even the higher-level file copying functions (`shutil.copy()`, `shutil.copy2()`) cannot copy all file metadata.

On POSIX platforms, this means that file owner and group are lost as well as ACLs. On Mac OS, the resource fork and other metadata are not used. This means that resources will be lost and file type and creator codes will not be correct. On Windows, file owners, ACLs and alternate data streams are not copied.

11.10.1 Directory and files operations

`shutil.copyfileobj(fsrc, fdst[, length])`

Copy the contents of the file-like object *fsrc* to the file-like object *fdst*. The integer *length*, if given, is the buffer size. In particular, a negative *length* value means to copy the data without looping over the source data in chunks; by default the data is read in chunks to avoid uncontrolled memory consumption. Note that if the current file position of the *fsrc* object is not 0, only the contents from the current file position to the end of the file will be copied.

`shutil.copyfile(src, dst, *, follow_symlinks=True)`

Copy the contents (no metadata) of the file named *src* to a file named *dst* and return *dst*. *src* and *dst* are path names given as strings. *dst* must be the complete target file name; look at `shutil.copy()` for a copy that accepts a target directory path. If *src* and *dst* specify the same file, `SameFileError` is raised.

The destination location must be writable; otherwise, an `OSError` exception will be raised. If *dst* already exists, it will be replaced. Special files such as character or block devices and pipes cannot be copied with this function.

If *follow_symlinks* is false and *src* is a symbolic link, a new symbolic link will be created instead of copying the file *src* points to.

在 3.3 版更改: `IOError` used to be raised instead of `OSError`. Added *follow_symlinks* argument. Now returns *dst*.

在 3.4 版更改: Raise `SameFileError` instead of `Error`. Since the former is a subclass of the latter, this change is backward compatible.

exception `shutil.SameFileError`

This exception is raised if source and destination in `copyfile()` are the same file.

3.4 新版功能.

`shutil.copymode(src, dst, *, follow_symlinks=True)`

Copy the permission bits from *src* to *dst*. The file contents, owner, and group are unaffected. *src* and *dst* are path names given as strings. If *follow_symlinks* is false, and both *src* and *dst* are symbolic links, `copymode()` will attempt to modify the mode of *dst* itself (rather than the file it points to). This functionality is not available on every

platform; please see `copystat()` for more information. If `copymode()` cannot modify symbolic links on the local platform, and it is asked to do so, it will do nothing and return.

在 3.3 版更改: Added `follow_symlinks` argument.

`shutil.copystat(src, dst, *, follow_symlinks=True)`

Copy the permission bits, last access time, last modification time, and flags from `src` to `dst`. On Linux, `copystat()` also copies the “extended attributes” where possible. The file contents, owner, and group are unaffected. `src` and `dst` are path names given as strings.

If `follow_symlinks` is false, and `src` and `dst` both refer to symbolic links, `copystat()` will operate on the symbolic links themselves rather than the files the symbolic links refer to—reading the information from the `src` symbolic link, and writing the information to the `dst` symbolic link.

注解: Not all platforms provide the ability to examine and modify symbolic links. Python itself can tell you what functionality is locally available.

- If `os.chmod` in `os.supports_follow_symlinks` is True, `copystat()` can modify the permission bits of a symbolic link.
- If `os.utime` in `os.supports_follow_symlinks` is True, `copystat()` can modify the last access and modification times of a symbolic link.
- If `os.chflags` in `os.supports_follow_symlinks` is True, `copystat()` can modify the flags of a symbolic link. (`os.chflags` is not available on all platforms.)

On platforms where some or all of this functionality is unavailable, when asked to modify a symbolic link, `copystat()` will copy everything it can. `copystat()` never returns failure.

Please see `os.supports_follow_symlinks` for more information.

在 3.3 版更改: Added `follow_symlinks` argument and support for Linux extended attributes.

`shutil.copy(src, dst, *, follow_symlinks=True)`

Copies the file `src` to the file or directory `dst`. `src` and `dst` should be strings. If `dst` specifies a directory, the file will be copied into `dst` using the base filename from `src`. Returns the path to the newly created file.

If `follow_symlinks` is false, and `src` is a symbolic link, `dst` will be created as a symbolic link. If `follow_symlinks` is true and `src` is a symbolic link, `dst` will be a copy of the file `src` refers to.

`copy()` copies the file data and the file’s permission mode (see `os.chmod()`). Other metadata, like the file’s creation and modification times, is not preserved. To preserve all file metadata from the original, use `copy2()` instead.

在 3.3 版更改: Added `follow_symlinks` argument. Now returns path to the newly created file.

`shutil.copy2(src, dst, *, follow_symlinks=True)`

Identical to `copy()` except that `copy2()` also attempts to preserve file metadata.

When `follow_symlinks` is false, and `src` is a symbolic link, `copy2()` attempts to copy all metadata from the `src` symbolic link to the newly-created `dst` symbolic link. However, this functionality is not available on all platforms. On platforms where some or all of this functionality is unavailable, `copy2()` will preserve all the metadata it can; `copy2()` never returns failure.

`copy2()` uses `copystat()` to copy the file metadata. Please see `copystat()` for more information about platform support for modifying symbolic link metadata.

在 3.3 版更改: Added `follow_symlinks` argument, try to copy extended file system attributes too (currently Linux only). Now returns path to the newly created file.

`shutil.ignore_patterns(*patterns)`

This factory function creates a function that can be used as a callable for `copytree()`'s `ignore` argument, ignoring files and directories that match one of the glob-style *patterns* provided. See the example below.

`shutil.copytree(src, dst, symlinks=False, ignore=None, copy_function=copy2, ignore_dangling_symlinks=False)`

Recursively copy an entire directory tree rooted at *src*, returning the destination directory. The destination directory, named by *dst*, must not already exist; it will be created as well as missing parent directories. Permissions and times of directories are copied with `copystat()`, individual files are copied using `shutil.copy2()`.

If *symlinks* is true, symbolic links in the source tree are represented as symbolic links in the new tree and the metadata of the original links will be copied as far as the platform allows; if false or omitted, the contents and metadata of the linked files are copied to the new tree.

When *symlinks* is false, if the file pointed by the symlink doesn't exist, an exception will be added in the list of errors raised in an `Error` exception at the end of the copy process. You can set the optional *ignore_dangling_symlinks* flag to true if you want to silence this exception. Notice that this option has no effect on platforms that don't support `os.symlink()`.

If *ignore* is given, it must be a callable that will receive as its arguments the directory being visited by `copytree()`, and a list of its contents, as returned by `os.listdir()`. Since `copytree()` is called recursively, the *ignore* callable will be called once for each directory that is copied. The callable must return a sequence of directory and file names relative to the current directory (i.e. a subset of the items in its second argument); these names will then be ignored in the copy process. `ignore_patterns()` can be used to create such a callable that ignores names based on glob-style patterns.

If exception(s) occur, an `Error` is raised with a list of reasons.

If *copy_function* is given, it must be a callable that will be used to copy each file. It will be called with the source path and the destination path as arguments. By default, `shutil.copy2()` is used, but any function that supports the same signature (like `shutil.copy()`) can be used.

在 3.3 版更改: Copy metadata when *symlinks* is false. Now returns *dst*.

在 3.2 版更改: Added the *copy_function* argument to be able to provide a custom copy function. Added the *ignore_dangling_symlinks* argument to silent dangling symlinks errors when *symlinks* is false.

`shutil.rmtree(path, ignore_errors=False, onerror=None)`

Delete an entire directory tree; *path* must point to a directory (but not a symbolic link to a directory). If *ignore_errors* is true, errors resulting from failed removals will be ignored; if false or omitted, such errors are handled by calling a handler specified by *onerror* or, if that is omitted, they raise an exception.

注解: On platforms that support the necessary fd-based functions a symlink attack resistant version of `rmtree()` is used by default. On other platforms, the `rmtree()` implementation is susceptible to a symlink attack: given proper timing and circumstances, attackers can manipulate symlinks on the filesystem to delete files they wouldn't be able to access otherwise. Applications can use the `rmtree.avoids_symlink_attacks` function attribute to determine which case applies.

If *onerror* is provided, it must be a callable that accepts three parameters: *function*, *path*, and *excinfo*.

The first parameter, *function*, is the function which raised the exception; it depends on the platform and implementation. The second parameter, *path*, will be the path name passed to *function*. The third parameter, *excinfo*, will be the exception information returned by `sys.exc_info()`. Exceptions raised by *onerror* will not be caught.

在 3.3 版更改: Added a symlink attack resistant version that is used automatically if platform supports fd-based functions.

`rmtree.avoids_symlink_attacks`

Indicates whether the current platform and implementation provides a symlink attack resistant version of

`rmtree()`. Currently this is only true for platforms supporting fd-based directory access functions.

3.3 新版功能.

`shutil.move(src, dst, copy_function=copy2)`

Recursively move a file or directory (*src*) to another location (*dst*) and return the destination.

If the destination is an existing directory, then *src* is moved inside that directory. If the destination already exists but is not a directory, it may be overwritten depending on `os.rename()` semantics.

If the destination is on the current filesystem, then `os.rename()` is used. Otherwise, *src* is copied to *dst* using `copy_function` and then removed. In case of symlinks, a new symlink pointing to the target of *src* will be created in or as *dst* and *src* will be removed.

If `copy_function` is given, it must be a callable that takes two arguments *src* and *dst*, and will be used to copy *src* to *dst* if `os.rename()` cannot be used. If the source is a directory, `copytree()` is called, passing it the `copy_function()`. The default `copy_function` is `copy2()`. Using `copy()` as the `copy_function` allows the move to succeed when it is not possible to also copy the metadata, at the expense of not copying any of the metadata.

在 3.3 版更改: Added explicit symlink handling for foreign filesystems, thus adapting it to the behavior of GNU's `mv`. Now returns *dst*.

在 3.5 版更改: Added the `copy_function` keyword argument.

`shutil.disk_usage(path)`

Return disk usage statistics about the given path as a *named tuple* with the attributes *total*, *used* and *free*, which are the amount of total, used and free space, in bytes. On Windows, *path* must be a directory; on Unix, it can be a file or directory.

3.3 新版功能.

Availability: Unix, Windows.

`shutil.chown(path, user=None, group=None)`

Change owner *user* and/or *group* of the given *path*.

user can be a system user name or a uid; the same applies to *group*. At least one argument is required.

See also `os.chown()`, the underlying function.

Availability: Unix.

3.3 新版功能.

`shutil.which(cmd, mode=os.F_OK | os.X_OK, path=None)`

Return the path to an executable which would be run if the given *cmd* was called. If no *cmd* would be called, return *None*.

mode is a permission mask passed to `os.access()`, by default determining if the file exists and executable.

When no *path* is specified, the results of `os.environ()` are used, returning either the “PATH” value or a fallback of `os.defpath`.

On Windows, the current directory is always prepended to the *path* whether or not you use the default or provide your own, which is the behavior the command shell uses when finding executables. Additionally, when finding the *cmd* in the *path*, the `PATHEXT` environment variable is checked. For example, if you call `shutil.which("python")`, `which()` will search `PATHEXT` to know that it should look for `python.exe` within the *path* directories. For example, on Windows:

```
>>> shutil.which("python")
'C:\\Python33\\python.EXE'
```

3.3 新版功能.

exception `shutil.Error`

This exception collects exceptions that are raised during a multi-file operation. For `copytree()`, the exception argument is a list of 3-tuples (`srcname`, `dstname`, `exception`).

copytree example

This example is the implementation of the `copytree()` function, described above, with the docstring omitted. It demonstrates many of the other functions provided by this module.

```
def copytree(src, dst, symlinks=False):
    names = os.listdir(src)
    os.makedirs(dst)
    errors = []
    for name in names:
        srcname = os.path.join(src, name)
        dstname = os.path.join(dst, name)
        try:
            if symlinks and os.path.islink(srcname):
                linkto = os.readlink(srcname)
                os.symlink(linkto, dstname)
            elif os.path.isdir(srcname):
                copytree(srcname, dstname, symlinks)
            else:
                copy2(srcname, dstname)
                # XXX What about devices, sockets etc.?
        except OSError as why:
            errors.append((srcname, dstname, str(why)))
        # catch the Error from the recursive copytree so that we can
        # continue with other files
        except Error as err:
            errors.extend(err.args[0])
    try:
        copystat(src, dst)
    except OSError as why:
        # can't copy file access times on Windows
        if why.winerror is None:
            errors.extend((src, dst, str(why)))
    if errors:
        raise Error(errors)
```

Another example that uses the `ignore_patterns()` helper:

```
from shutil import copytree, ignore_patterns

copytree(source, destination, ignore=ignore_patterns('*.pyc', 'tmp*'))
```

This will copy everything except `.pyc` files and files or directories whose name starts with `tmp`.

Another example that uses the `ignore` argument to add a logging call:

```
from shutil import copytree
import logging

def _logpath(path, names):
    logging.info('Working in %s', path)
    return [] # nothing will be ignored
```

(下页继续)

(续上页)

```
copytree(source, destination, ignore=_logpath)
```

rmtree example

This example shows how to remove a directory tree on Windows where some of the files have their read-only bit set. It uses the `onerror` callback to clear the readonly bit and reattempt the remove. Any subsequent failure will propagate.

```
import os, stat
import shutil

def remove_readonly(func, path, _):
    "Clear the readonly bit and reattempt the removal"
    os.chmod(path, stat.S_IWRITE)
    func(path)

shutil.rmtree(directory, onerror=remove_readonly)
```

11.10.2 Archiving operations

3.2 新版功能.

在 3.5 版更改: Added support for the *xz*tar format.

High-level utilities to create and read compressed and archived files are also provided. They rely on the *zipfile* and *tarfile* modules.

```
shutil.make_archive(base_name, format[, root_dir[, base_dir[, verbose[, dry_run[, owner[, group[,
                        logger]]]]]])
```

Create an archive file (such as zip or tar) and return its name.

base_name is the name of the file to create, including the path, minus any format-specific extension. *format* is the archive format: one of “zip” (if the *zlib* module is available), “tar”, “gztar” (if the *zlib* module is available), “bztar” (if the *bz2* module is available), or “xztar” (if the *lzma* module is available).

root_dir is a directory that will be the root directory of the archive; for example, we typically `chdir` into *root_dir* before creating the archive.

base_dir is the directory where we start archiving from; i.e. *base_dir* will be the common prefix of all files and directories in the archive.

root_dir and *base_dir* both default to the current directory.

If *dry_run* is true, no archive is created, but the operations that would be executed are logged to *logger*.

owner and *group* are used when creating a tar archive. By default, uses the current owner and group.

logger must be an object compatible with **PEP 282**, usually an instance of *logging.Logger*.

The *verbose* argument is unused and deprecated.

```
shutil.get_archive_formats()
```

Return a list of supported formats for archiving. Each element of the returned sequence is a tuple (name, description).

By default *shutil* provides these formats:

- *zip*: ZIP file (if the *zlib* module is available).

- *tar*: uncompressed tar file.
- *gztar*: gzip' ed tar-file (if the *zlib* module is available).
- *bztar*: bzip2' ed tar-file (if the *bz2* module is available).
- *xztar*: xz' ed tar-file (if the *lzma* module is available).

You can register new formats or provide your own archiver for any existing formats, by using `register_archive_format()`.

`shutil.register_archive_format(name, function[, extra_args[, description]])`

Register an archiver for the format *name*.

function is the callable that will be used to unpack archives. The callable will receive the *base_name* of the file to create, followed by the *base_dir* (which defaults to `os.curdir`) to start archiving from. Further arguments are passed as keyword arguments: *owner*, *group*, *dry_run* and *logger* (as passed in `make_archive()`).

If given, *extra_args* is a sequence of (name, value) pairs that will be used as extra keywords arguments when the archiver callable is used.

description is used by `get_archive_formats()` which returns the list of archivers. Defaults to an empty string.

`shutil.unregister_archive_format(name)`

Remove the archive format *name* from the list of supported formats.

`shutil.unpack_archive(filename[, extract_dir[, format]])`

Unpack an archive. *filename* is the full path of the archive.

extract_dir is the name of the target directory where the archive is unpacked. If not provided, the current working directory is used.

format is the archive format: one of “zip”, “tar”, “gztar”, “bztar”, or “xztar”. Or any other format registered with `register_unpack_format()`. If not provided, `unpack_archive()` will use the archive file name extension and see if an unpacker was registered for that extension. In case none is found, a `ValueError` is raised.

`shutil.register_unpack_format(name, extensions, function[, extra_args[, description]])`

Registers an unpack format. *name* is the name of the format and *extensions* is a list of extensions corresponding to the format, like `.zip` for Zip files.

function is the callable that will be used to unpack archives. The callable will receive the path of the archive, followed by the directory the archive must be extracted to.

When provided, *extra_args* is a sequence of (name, value) tuples that will be passed as keywords arguments to the callable.

description can be provided to describe the format, and will be returned by the `get_unpack_formats()` function.

`shutil.unregister_unpack_format(name)`

Unregister an unpack format. *name* is the name of the format.

`shutil.get_unpack_formats()`

Return a list of all registered formats for unpacking. Each element of the returned sequence is a tuple (name, extensions, description).

By default *shutil* provides these formats:

- *zip*: ZIP file (unpacking compressed files works only if the corresponding module is available).
- *tar*: uncompressed tar file.
- *gztar*: gzip' ed tar-file (if the *zlib* module is available).

- *bztar*: bzip2' ed tar-file (if the *bz2* module is available).
- *xztar*: xz' ed tar-file (if the *lzma* module is available).

You can register new formats or provide your own unpacker for any existing formats, by using `register_unpack_format()`.

Archiving example

In this example, we create a gzip' ed tar-file archive containing all files found in the `.ssh` directory of the user:

```
>>> from shutil import make_archive
>>> import os
>>> archive_name = os.path.expanduser(os.path.join('~', 'myarchive'))
>>> root_dir = os.path.expanduser(os.path.join('~', '.ssh'))
>>> make_archive(archive_name, 'gztar', root_dir)
'/Users/tarek/myarchive.tar.gz'
```

The resulting archive contains:

```
$ tar -tzvf /Users/tarek/myarchive.tar.gz
drwx----- tarek/staff      0 2010-02-01 16:23:40 ./
-rw-r--r-- tarek/staff     609 2008-06-09 13:26:54 ./authorized_keys
-rwxr-xr-x tarek/staff      65 2008-06-09 13:26:54 ./config
-rwx----- tarek/staff     668 2008-06-09 13:26:54 ./id_dsa
-rwxr-xr-x tarek/staff     609 2008-06-09 13:26:54 ./id_dsa.pub
-rw----- tarek/staff    1675 2008-06-09 13:26:54 ./id_rsa
-rw-r--r-- tarek/staff     397 2008-06-09 13:26:54 ./id_rsa.pub
-rw-r--r-- tarek/staff   37192 2010-02-06 18:23:10 ./known_hosts
```

11.10.3 Querying the size of the output terminal

`shutil.get_terminal_size(fallback=(columns, lines))`

Get the size of the terminal window.

For each of the two dimensions, the environment variable, `COLUMNS` and `LINES` respectively, is checked. If the variable is defined and the value is a positive integer, it is used.

When `COLUMNS` or `LINES` is not defined, which is the common case, the terminal connected to `sys.__stdout__` is queried by invoking `os.get_terminal_size()`.

If the terminal size cannot be successfully queried, either because the system doesn't support querying, or because we are not connected to a terminal, the value given in `fallback` parameter is used. `fallback` defaults to (80, 24) which is the default size used by many terminal emulators.

The value returned is a named tuple of type `os.terminal_size`.

See also: The Single UNIX Specification, Version 2, [Other Environment Variables](#).

3.3 新版功能.

11.11 macpath —Mac OS 9 路径操作函数

源码: [Lib/macpath.py](#)

该模块是 `os.path` 模块的 Mac OS 9（及更早版本）实现。它可用于在 Mac OS X（或任何其他平台）上操作旧式 Macintosh 路径名。

该模块中提供了以下函数：`normcase()`、`normpath()`、`isabs()`、`join()`、`split()`、`isdir()`、`isfile()`、`walk()`、`exists()`。对于其他可用的函数 `os.path` 虚拟对应可用。

参见：

模块 `os` 操作系统接口，包括处理比 Python 文件对象 更低级别文件的功能。

模块 `io` Python 的内置 I/O 库，包括抽象类和一些具体的类，如文件 I/O。

内置函数 `open()` 使用 Python 打开文件进行读写的标准方法。

本章中描述的模块支持在磁盘上以持久形式存储 Python 数据。`pickle` 和 `marshal` 模块可以将许多 Python 数据类型转换为字节流，然后从字节中重新创建对象。各种与 DBM 相关的模块支持一系列基于散列的文件格式，这些格式存储字符串到其他字符串的映射。

本章中描述的模块列表是：

12.1 `pickle` —— Python 对象序列化

源代码：Lib/pickle.py

模块 `pickle` 实现了对一个 Python 对象结构的二进制序列化和反序列化。“Pickling”是将 Python 对象和所拥有的层次结构被转化为一个字节流的过程，而 “unpickling”是相反的操作，会将（来自一个 *binary file* 或者 *bytes-like object* 的）字节流转化回一个对象层次结构。Pickling（和 unpickling）也被称为“序列化”，“编组”¹ 或者“平面化”。而为了避免混乱，此处采用术语 “pickling” 和 “unpickling”。

警告： `pickle` 模块在接受被错误地构造或者被恶意地构造的数据时不安全。永远不要 unpickle 来自于不受信任的或者未经验证的来源的数据。

12.1.1 与其他 Python 模块间的关系

与 `marshal` 间的关系

Python 有一个更原始的序列化模块称为 `marshal`，但一般地 `pickle` 应该是序列化 Python 对象时的首选。`marshal` 存在主要是为了支持 Python 的 `.pyc` 文件。

`pickle` 模块与 `marshal` 在如下几方面显著地不同：

¹ Don't confuse this with the `marshal` module

- `pickle` 模块会跟踪已被序列化的对象，所以该对象之后再次被引用时不会再次被序列化。`marshal` 不会这么做。

这隐含了递归对象和共享对象。递归对象指包含对自己的引用的对象。这种对象并不会被 `marshal` 接受，并且实际上尝试 `marshal` 递归对象会让你的 Python 解释器崩溃。对象共享发生在对象层级中存在多处引用同一对象时。`pickle` 只会存储这些对象一次，并确保其他的引用指向同一个主副本。共享对象将保持共享，这可能对可变对象非常重要。

- `marshal` 不能被用于序列化用户定义类及其实例。`pickle` 能够透明地存储并保存类实例，然而此时类定义必须能够从与被存储时相同的模块被引入。
- The `marshal` serialization format is not guaranteed to be portable across Python versions. Because its primary job in life is to support `.pyc` files, the Python implementers reserve the right to change the serialization format in non-backwards compatible ways should the need arise. The `pickle` serialization format is guaranteed to be backwards compatible across Python releases.

与 json 模块的比较

Pickle 协议和 JSON (JavaScript Object Notation) 间有着本质的不同：

- JSON 是一个文本序列化格式（它输出 unicode 文本，尽管在大多数时候它会接着以 utf-8 编码），而 pickle 是一个二进制序列化格式；
- JSON is human-readable, while pickle is not; JSON 是我们很容易理解的，而 pickle 不是；
- JSON 是可互操作的，在 Python 系统之外广泛使用，而 pickle 则是 Python 专用的；
- JSON, by default, can only represent a subset of the Python built-in types, and no custom classes; pickle can represent an extremely large number of Python types (many of them automatically, by clever usage of Python's introspection facilities; complex cases can be tackled by implementing *specific object APIs*).

参见：

`json` 模块：一个允许 JSON 序列化和反序列化的标准库模块

12.1.2 Data stream format

The data format used by `pickle` is Python-specific. This has the advantage that there are no restrictions imposed by external standards such as JSON or XDR (which can't represent pointer sharing); however it means that non-Python programs may not be able to reconstruct pickled Python objects.

By default, the `pickle` data format uses a relatively compact binary representation. If you need optimal size characteristics, you can efficiently *compress* pickled data.

The module `pickletools` contains tools for analyzing data streams generated by `pickle`. `pickletools` source code has extensive comments about opcodes used by pickle protocols.

There are currently 5 different protocols which can be used for pickling. The higher the protocol used, the more recent the version of Python needed to read the pickle produced.

- Protocol version 0 is the original “human-readable” protocol and is backwards compatible with earlier versions of Python.
- Protocol version 1 is an old binary format which is also compatible with earlier versions of Python.
- Protocol version 2 was introduced in Python 2.3. It provides much more efficient pickling of *new-style classes*. Refer to **PEP 307** for information about improvements brought by protocol 2.

- Protocol version 3 was added in Python 3.0. It has explicit support for *bytes* objects and cannot be unpickled by Python 2.x. This is the default protocol, and the recommended protocol when compatibility with other Python 3 versions is required.
- Protocol version 4 was added in Python 3.4. It adds support for very large objects, pickling more kinds of objects, and some data format optimizations. Refer to [PEP 3154](#) for information about improvements brought by protocol 4.

注解: Serialization is a more primitive notion than persistence; although *pickle* reads and writes file objects, it does not handle the issue of naming persistent objects, nor the (even more complicated) issue of concurrent access to persistent objects. The *pickle* module can transform a complex object into a byte stream and it can transform the byte stream into an object with the same internal structure. Perhaps the most obvious thing to do with these byte streams is to write them onto a file, but it is also conceivable to send them across a network or store them in a database. The *shelve* module provides a simple interface to pickle and unpickle objects on DBM-style database files.

12.1.3 Module Interface

To serialize an object hierarchy, you simply call the *dumps()* function. Similarly, to de-serialize a data stream, you call the *loads()* function. However, if you want more control over serialization and de-serialization, you can create a *Pickler* or an *Unpickler* object, respectively.

The *pickle* module provides the following constants:

`pickle.HIGHEST_PROTOCOL`

An integer, the highest *protocol version* available. This value can be passed as a *protocol* value to functions *dump()* and *dumps()* as well as the *Pickler* constructor.

`pickle.DEFAULT_PROTOCOL`

An integer, the default *protocol version* used for pickling. May be less than *HIGHEST_PROTOCOL*. Currently the default protocol is 3, a new protocol designed for Python 3.

The *pickle* module provides the following functions to make the pickling process more convenient:

`pickle.dump(obj, file, protocol=None, *, fix_imports=True)`

Write a pickled representation of *obj* to the open *file object file*. This is equivalent to `Pickler(file, protocol).dump(obj)`.

The optional *protocol* argument, an integer, tells the pickler to use the given protocol; supported protocols are 0 to *HIGHEST_PROTOCOL*. If not specified, the default is *DEFAULT_PROTOCOL*. If a negative number is specified, *HIGHEST_PROTOCOL* is selected.

The *file* argument must have a `write()` method that accepts a single bytes argument. It can thus be an on-disk file opened for binary writing, an *io.BytesIO* instance, or any other custom object that meets this interface.

If *fix_imports* is true and *protocol* is less than 3, pickle will try to map the new Python 3 names to the old module names used in Python 2, so that the pickle data stream is readable with Python 2.

`pickle.dumps(obj, protocol=None, *, fix_imports=True)`

Return the pickled representation of the object as a *bytes* object, instead of writing it to a file.

Arguments *protocol* and *fix_imports* have the same meaning as in *dump()*.

`pickle.load(file, *, fix_imports=True, encoding="ASCII", errors="strict")`

Read a pickled object representation from the open *file object file* and return the reconstituted object hierarchy specified therein. This is equivalent to `Unpickler(file).load()`.

The protocol version of the pickle is detected automatically, so no protocol argument is needed. Bytes past the pickled object's representation are ignored.

The argument *file* must have two methods, a `read()` method that takes an integer argument, and a `readline()` method that requires no arguments. Both methods should return bytes. Thus *file* can be an on-disk file opened for binary reading, an `io.BytesIO` object, or any other custom object that meets this interface.

Optional keyword arguments are *fix_imports*, *encoding* and *errors*, which are used to control compatibility support for pickle stream generated by Python 2. If *fix_imports* is true, pickle will try to map the old Python 2 names to the new names used in Python 3. The *encoding* and *errors* tell pickle how to decode 8-bit string instances pickled by Python 2; these default to 'ASCII' and 'strict', respectively. The *encoding* can be 'bytes' to read these 8-bit string instances as bytes objects. Using *encoding*='latin1' is required for unpickling NumPy arrays and instances of *datetime*, *date* and *time* pickled by Python 2.

`pickle.loads(bytes_object, *, fix_imports=True, encoding="ASCII", errors="strict")`

Read a pickled object hierarchy from a *bytes* object and return the reconstituted object hierarchy specified therein.

The protocol version of the pickle is detected automatically, so no protocol argument is needed. Bytes past the pickled object's representation are ignored.

Optional keyword arguments are *fix_imports*, *encoding* and *errors*, which are used to control compatibility support for pickle stream generated by Python 2. If *fix_imports* is true, pickle will try to map the old Python 2 names to the new names used in Python 3. The *encoding* and *errors* tell pickle how to decode 8-bit string instances pickled by Python 2; these default to 'ASCII' and 'strict', respectively. The *encoding* can be 'bytes' to read these 8-bit string instances as bytes objects. Using *encoding*='latin1' is required for unpickling NumPy arrays and instances of *datetime*, *date* and *time* pickled by Python 2.

The *pickle* module defines three exceptions:

exception `pickle.PickleError`

Common base class for the other pickling exceptions. It inherits *Exception*.

exception `pickle.PicklingError`

Error raised when an unpicklable object is encountered by *Pickler*. It inherits *PickleError*.

Refer to *What can be pickled and unpickled?* to learn what kinds of objects can be pickled.

exception `pickle.UnpicklingError`

Error raised when there is a problem unpickling an object, such as a data corruption or a security violation. It inherits *PickleError*.

Note that other exceptions may also be raised during unpickling, including (but not necessarily limited to) *AttributeError*, *EOFError*, *ImportError*, and *IndexError*.

The *pickle* module exports two classes, *Pickler* and *Unpickler*:

class `pickle.Pickler(file, protocol=None, *, fix_imports=True)`

This takes a binary file for writing a pickle data stream.

The optional *protocol* argument, an integer, tells the pickler to use the given protocol; supported protocols are 0 to *HIGHEST_PROTOCOL*. If not specified, the default is *DEFAULT_PROTOCOL*. If a negative number is specified, *HIGHEST_PROTOCOL* is selected.

The *file* argument must have a `write()` method that accepts a single bytes argument. It can thus be an on-disk file opened for binary writing, an `io.BytesIO` instance, or any other custom object that meets this interface.

If *fix_imports* is true and *protocol* is less than 3, pickle will try to map the new Python 3 names to the old module names used in Python 2, so that the pickle data stream is readable with Python 2.

dump (*obj*)

Write a pickled representation of *obj* to the open file object given in the constructor.

persistent_id (*obj*)

Do nothing by default. This exists so a subclass can override it.

If `persistent_id()` returns `None`, `obj` is pickled as usual. Any other value causes `Pickler` to emit the returned value as a persistent ID for `obj`. The meaning of this persistent ID should be defined by `Unpickler.persistent_load()`. Note that the value returned by `persistent_id()` cannot itself have a persistent ID.

See *Persistence of External Objects* for details and examples of uses.

dispatch_table

A pickler object's dispatch table is a registry of *reduction functions* of the kind which can be declared using `copyreg.pickle()`. It is a mapping whose keys are classes and whose values are reduction functions. A reduction function takes a single argument of the associated class and should conform to the same interface as a `__reduce__()` method.

By default, a pickler object will not have a `dispatch_table` attribute, and it will instead use the global dispatch table managed by the `copyreg` module. However, to customize the pickling for a specific pickler object one can set the `dispatch_table` attribute to a dict-like object. Alternatively, if a subclass of `Pickler` has a `dispatch_table` attribute then this will be used as the default dispatch table for instances of that class.

See *Dispatch Tables* for usage examples.

3.3 新版功能.

fast

Deprecated. Enable fast mode if set to a true value. The fast mode disables the usage of memo, therefore speeding the pickling process by not generating superfluous PUT opcodes. It should not be used with self-referential objects, doing otherwise will cause `Pickler` to recurse infinitely.

Use `pickletools.optimize()` if you need more compact pickles.

class `pickle.Unpickler` (*file*, *, *fix_imports*=`True`, *encoding*=`"ASCII"`, *errors*=`"strict"`)

This takes a binary file for reading a pickle data stream.

The protocol version of the pickle is detected automatically, so no protocol argument is needed.

The argument *file* must have two methods, a `read()` method that takes an integer argument, and a `readline()` method that requires no arguments. Both methods should return bytes. Thus *file* can be an on-disk file object opened for binary reading, an `io.BytesIO` object, or any other custom object that meets this interface.

Optional keyword arguments are *fix_imports*, *encoding* and *errors*, which are used to control compatibility support for pickle stream generated by Python 2. If *fix_imports* is true, pickle will try to map the old Python 2 names to the new names used in Python 3. The *encoding* and *errors* tell pickle how to decode 8-bit string instances pickled by Python 2; these default to `'ASCII'` and `'strict'`, respectively. The *encoding* can be `'bytes'` to read these 8-bit string instances as bytes objects.

load()

Read a pickled object representation from the open file object given in the constructor, and return the reconstituted object hierarchy specified therein. Bytes past the pickled object's representation are ignored.

persistent_load(pid)

Raise an `UnpicklingError` by default.

If defined, `persistent_load()` should return the object specified by the persistent ID *pid*. If an invalid persistent ID is encountered, an `UnpicklingError` should be raised.

See *Persistence of External Objects* for details and examples of uses.

find_class (*module*, *name*)

Import *module* if necessary and return the object called *name* from it, where the *module* and *name* arguments are `str` objects. Note, unlike its name suggests, `find_class()` is also used for finding functions.

Subclasses may override this to gain control over what type of objects and how they can be loaded, potentially reducing security risks. Refer to *Restricting Globals* for details.

12.1.4 What can be pickled and unpickled?

The following types can be pickled:

- `None`, `True`, and `False`
- integers, floating point numbers, complex numbers
- strings, bytes, bytearrays
- tuples, lists, sets, and dictionaries containing only picklable objects
- functions defined at the top level of a module (using `def`, not `lambda`)
- built-in functions defined at the top level of a module
- classes that are defined at the top level of a module
- instances of such classes whose `__dict__` or the result of calling `__getstate__()` is picklable (see section *Pickling Class Instances* for details).

Attempts to pickle unpicklable objects will raise the *PicklingError* exception; when this happens, an unspecified number of bytes may have already been written to the underlying file. Trying to pickle a highly recursive data structure may exceed the maximum recursion depth, a *RecursionError* will be raised in this case. You can carefully raise this limit with `sys.setrecursionlimit()`.

Note that functions (built-in and user-defined) are pickled by “fully qualified” name reference, not by value.² This means that only the function name is pickled, along with the name of the module the function is defined in. Neither the function’s code, nor any of its function attributes are pickled. Thus the defining module must be importable in the unpickling environment, and the module must contain the named object, otherwise an exception will be raised.³

Similarly, classes are pickled by named reference, so the same restrictions in the unpickling environment apply. Note that none of the class’s code or data is pickled, so in the following example the class attribute `attr` is not restored in the unpickling environment:

```
class Foo:
    attr = 'A class attribute'

picklestring = pickle.dumps(Foo)
```

These restrictions are why picklable functions and classes must be defined in the top level of a module.

Similarly, when class instances are pickled, their class’s code and data are not pickled along with them. Only the instance data are pickled. This is done on purpose, so you can fix bugs in a class or add methods to the class and still load objects that were created with an earlier version of the class. If you plan to have long-lived objects that will see many versions of a class, it may be worthwhile to put a version number in the objects so that suitable conversions can be made by the class’s `__setstate__()` method.

12.1.5 Pickling Class Instances

In this section, we describe the general mechanisms available to you to define, customize, and control how class instances are pickled and unpickled.

² This is why `lambda` functions cannot be pickled: all `lambda` functions share the same name: `<lambda>`.

³ The exception raised will likely be an *ImportError* or an *AttributeError* but it could be something else.

In most cases, no additional code is needed to make instances picklable. By default, pickle will retrieve the class and the attributes of an instance via introspection. When a class instance is unpickled, its `__init__()` method is usually *not* invoked. The default behaviour first creates an uninitialized instance and then restores the saved attributes. The following code shows an implementation of this behaviour:

```
def save(obj):
    return (obj.__class__, obj.__dict__)

def load(cls, attributes):
    obj = cls.__new__(cls)
    obj.__dict__.update(attributes)
    return obj
```

Classes can alter the default behaviour by providing one or several special methods:

`object.__getnewargs_ex__()`

In protocols 2 and newer, classes that implements the `__getnewargs_ex__()` method can dictate the values passed to the `__new__()` method upon unpickling. The method must return a pair (args, kwargs) where *args* is a tuple of positional arguments and *kwargs* a dictionary of named arguments for constructing the object. Those will be passed to the `__new__()` method upon unpickling.

You should implement this method if the `__new__()` method of your class requires keyword-only arguments. Otherwise, it is recommended for compatibility to implement `__getnewargs__()`.

在 3.6 版更改: `__getnewargs_ex__()` is now used in protocols 2 and 3.

`object.__getnewargs__()`

This method serves a similar purpose as `__getnewargs_ex__()`, but supports only positional arguments. It must return a tuple of arguments *args* which will be passed to the `__new__()` method upon unpickling.

`__getnewargs__()` will not be called if `__getnewargs_ex__()` is defined.

在 3.6 版更改: Before Python 3.6, `__getnewargs__()` was called instead of `__getnewargs_ex__()` in protocols 2 and 3.

`object.__getstate__()`

Classes can further influence how their instances are pickled; if the class defines the method `__getstate__()`, it is called and the returned object is pickled as the contents for the instance, instead of the contents of the instance's dictionary. If the `__getstate__()` method is absent, the instance's `__dict__` is pickled as usual.

`object.__setstate__(state)`

Upon unpickling, if the class defines `__setstate__()`, it is called with the unpickled state. In that case, there is no requirement for the state object to be a dictionary. Otherwise, the pickled state must be a dictionary and its items are assigned to the new instance's dictionary.

注解: If `__getstate__()` returns a false value, the `__setstate__()` method will not be called upon unpickling.

Refer to the section *Handling Stateful Objects* for more information about how to use the methods `__getstate__()` and `__setstate__()`.

注解: At unpickling time, some methods like `__getattr__()`, `__getattribute__()`, or `__setattr__()` may be called upon the instance. In case those methods rely on some internal invariant being true, the type should implement `__getnewargs__()` or `__getnewargs_ex__()` to establish such an invariant; otherwise, neither `__new__()` nor `__init__()` will be called.

As we shall see, pickle does not use directly the methods described above. In fact, these methods are part of the copy protocol which implements the `__reduce__()` special method. The copy protocol provides a unified interface for retrieving the data necessary for pickling and copying objects.⁴

Although powerful, implementing `__reduce__()` directly in your classes is error prone. For this reason, class designers should use the high-level interface (i.e., `__getnewargs_ex__()`, `__getstate__()` and `__setstate__()`) whenever possible. We will show, however, cases where using `__reduce__()` is the only option or leads to more efficient pickling or both.

`object.__reduce__()`

The interface is currently defined as follows. The `__reduce__()` method takes no argument and shall return either a string or preferably a tuple (the returned object is often referred to as the “reduce value”).

If a string is returned, the string should be interpreted as the name of a global variable. It should be the object’s local name relative to its module; the pickle module searches the module namespace to determine the object’s module. This behaviour is typically useful for singletons.

When a tuple is returned, it must be between two and five items long. Optional items can either be omitted, or `None` can be provided as their value. The semantics of each item are in order:

- A callable object that will be called to create the initial version of the object.
- A tuple of arguments for the callable object. An empty tuple must be given if the callable does not accept any argument.
- Optionally, the object’s state, which will be passed to the object’s `__setstate__()` method as previously described. If the object has no such method then, the value must be a dictionary and it will be added to the object’s `__dict__` attribute.
- Optionally, an iterator (and not a sequence) yielding successive items. These items will be appended to the object either using `obj.append(item)` or, in batch, using `obj.extend(list_of_items)`. This is primarily used for list subclasses, but may be used by other classes as long as they have `append()` and `extend()` methods with the appropriate signature. (Whether `append()` or `extend()` is used depends on which pickle protocol version is used as well as the number of items to append, so both must be supported.)
- Optionally, an iterator (not a sequence) yielding successive key-value pairs. These items will be stored to the object using `obj[key] = value`. This is primarily used for dictionary subclasses, but may be used by other classes as long as they implement `__setitem__()`.

`object.__reduce_ex__(protocol)`

Alternatively, a `__reduce_ex__()` method may be defined. The only difference is this method should take a single integer argument, the protocol version. When defined, pickle will prefer it over the `__reduce__()` method. In addition, `__reduce__()` automatically becomes a synonym for the extended version. The main use for this method is to provide backwards-compatible reduce values for older Python releases.

Persistence of External Objects

For the benefit of object persistence, the `pickle` module supports the notion of a reference to an object outside the pickled data stream. Such objects are referenced by a persistent ID, which should be either a string of alphanumeric characters (for protocol 0)⁵ or just an arbitrary object (for any newer protocol).

The resolution of such persistent IDs is not defined by the `pickle` module; it will delegate this resolution to the user defined methods on the pickler and unpickler, `persistent_id()` and `persistent_load()` respectively.

To pickle objects that have an external persistent id, the pickler must have a custom `persistent_id()` method that takes an object as an argument and returns either `None` or the persistent id for that object. When `None` is returned, the

⁴ The `copy` module uses this protocol for shallow and deep copying operations.

⁵ The limitation on alphanumeric characters is due to the fact the persistent IDs, in protocol 0, are delimited by the newline character. Therefore if any kind of newline characters occurs in persistent IDs, the resulting pickle will become unreadable.

pickler simply pickles the object as normal. When a persistent ID string is returned, the pickler will pickle that object, along with a marker so that the unpickler will recognize it as a persistent ID.

To unpickle external objects, the unpickler must have a custom `persistent_load()` method that takes a persistent ID object and returns the referenced object.

Here is a comprehensive example presenting how persistent ID can be used to pickle external objects by reference.

```
# Simple example presenting how persistent ID can be used to pickle
# external objects by reference.

import pickle
import sqlite3
from collections import namedtuple

# Simple class representing a record in our database.
MemoRecord = namedtuple("MemoRecord", "key, task")

class DBPickler(pickle.Pickler):

    def persistent_id(self, obj):
        # Instead of pickling MemoRecord as a regular class instance, we emit a
        # persistent ID.
        if isinstance(obj, MemoRecord):
            # Here, our persistent ID is simply a tuple, containing a tag and a
            # key, which refers to a specific record in the database.
            return ("MemoRecord", obj.key)
        else:
            # If obj does not have a persistent ID, return None. This means obj
            # needs to be pickled as usual.
            return None

class DBUnpickler(pickle.Unpickler):

    def __init__(self, file, connection):
        super().__init__(file)
        self.connection = connection

    def persistent_load(self, pid):
        # This method is invoked whenever a persistent ID is encountered.
        # Here, pid is the tuple returned by DBPickler.
        cursor = self.connection.cursor()
        type_tag, key_id = pid
        if type_tag == "MemoRecord":
            # Fetch the referenced record from the database and return it.
            cursor.execute("SELECT * FROM memos WHERE key=?", (str(key_id),))
            key, task = cursor.fetchone()
            return MemoRecord(key, task)
        else:
            # Always raises an error if you cannot return the correct object.
            # Otherwise, the unpickler will think None is the object referenced
            # by the persistent ID.
            raise pickle.UnpicklingError("unsupported persistent object")

def main():
    import io
    import pprint
```

(下页继续)

(续上页)

```

# Initialize and populate our database.
conn = sqlite3.connect(":memory:")
cursor = conn.cursor()
cursor.execute("CREATE TABLE memos(key INTEGER PRIMARY KEY, task TEXT)")
tasks = (
    'give food to fish',
    'prepare group meeting',
    'fight with a zebra',
)
for task in tasks:
    cursor.execute("INSERT INTO memos VALUES(NULL, ?)", (task,))

# Fetch the records to be pickled.
cursor.execute("SELECT * FROM memos")
memos = [MemoRecord(key, task) for key, task in cursor]
# Save the records using our custom DBPickler.
file = io.BytesIO()
DBPickler(file).dump(memos)

print("Pickled records:")
pprint.pprint(memos)

# Update a record, just for good measure.
cursor.execute("UPDATE memos SET task='learn italian' WHERE key=1")

# Load the records from the pickle data stream.
file.seek(0)
memos = DBUnpickler(file, conn).load()

print("Unpickled records:")
pprint.pprint(memos)

if __name__ == '__main__':
    main()

```

Dispatch Tables

If one wants to customize pickling of some classes without disturbing any other code which depends on pickling, then one can create a pickler with a private dispatch table.

The global dispatch table managed by the `copyreg` module is available as `copyreg.dispatch_table`. Therefore, one may choose to use a modified copy of `copyreg.dispatch_table` as a private dispatch table.

For example

```

f = io.BytesIO()
p = pickle.Pickler(f)
p.dispatch_table = copyreg.dispatch_table.copy()
p.dispatch_table[SomeClass] = reduce_SomeClass

```

creates an instance of `pickle.Pickler` with a private dispatch table which handles the `SomeClass` class specially. Alternatively, the code

```
class MyPickler(pickle.Pickler):
    dispatch_table = copyreg.dispatch_table.copy()
    dispatch_table[SomeClass] = reduce_SomeClass
f = io.BytesIO()
p = MyPickler(f)
```

does the same, but all instances of `MyPickler` will by default share the same dispatch table. The equivalent code using the `copyreg` module is

```
copyreg.pickle(SomeClass, reduce_SomeClass)
f = io.BytesIO()
p = pickle.Pickler(f)
```

Handling Stateful Objects

Here's an example that shows how to modify pickling behavior for a class. The `TextReader` class opens a text file, and returns the line number and line contents each time its `readline()` method is called. If a `TextReader` instance is pickled, all attributes *except* the file object member are saved. When the instance is unpickled, the file is reopened, and reading resumes from the last location. The `__setstate__()` and `__getstate__()` methods are used to implement this behavior.

```
class TextReader:
    """Print and number lines in a text file."""

    def __init__(self, filename):
        self.filename = filename
        self.file = open(filename)
        self.lineno = 0

    def readline(self):
        self.lineno += 1
        line = self.file.readline()
        if not line:
            return None
        if line.endswith('\n'):
            line = line[:-1]
        return "%i: %s" % (self.lineno, line)

    def __getstate__(self):
        # Copy the object's state from self.__dict__ which contains
        # all our instance attributes. Always use the dict.copy()
        # method to avoid modifying the original state.
        state = self.__dict__.copy()
        # Remove the unpicklable entries.
        del state['file']
        return state

    def __setstate__(self, state):
        # Restore instance attributes (i.e., filename and lineno).
        self.__dict__.update(state)
        # Restore the previously opened file's state. To do so, we need to
        # reopen it and read from it until the line count is restored.
        file = open(self.filename)
        for _ in range(self.lineno):
            file.readline()
```

(下页继续)

(续上页)

```
# Finally, save the file.
self.file = file
```

A sample usage might be something like this:

```
>>> reader = TextReader("hello.txt")
>>> reader.readline()
'1: Hello world!'
>>> reader.readline()
'2: I am line number two.'
>>> new_reader = pickle.loads(pickle.dumps(reader))
>>> new_reader.readline()
'3: Goodbye!'
```

12.1.6 Restricting Globals

By default, unpickling will import any class or function that it finds in the pickle data. For many applications, this behaviour is unacceptable as it permits the unpickler to import and invoke arbitrary code. Just consider what this hand-crafted pickle data stream does when loaded:

```
>>> import pickle
>>> pickle.loads(b"cos\nsystem\n(S'echo hello world'\ntr.")
hello world
0
```

In this example, the unpickler imports the `os.system()` function and then apply the string argument “echo hello world”. Although this example is inoffensive, it is not difficult to imagine one that could damage your system.

For this reason, you may want to control what gets unpickled by customizing `Unpickler.find_class()`. Unlike its name suggests, `Unpickler.find_class()` is called whenever a global (i.e., a class or a function) is requested. Thus it is possible to either completely forbid globals or restrict them to a safe subset.

Here is an example of an unpickler allowing only few safe classes from the `builtins` module to be loaded:

```
import builtins
import io
import pickle

safe_builtins = {
    'range',
    'complex',
    'set',
    'frozenset',
    'slice',
}

class RestrictedUnpickler(pickle.Unpickler):

    def find_class(self, module, name):
        # Only allow safe classes from builtins.
        if module == "builtins" and name in safe_builtins:
            return getattr(builtins, name)
        # Forbid everything else.
        raise pickle.UnpicklingError("global '%s.%s' is forbidden" %
                                      (module, name))
```

(下页继续)

(续上页)

```
def restricted_loads(s):
    """Helper function analogous to pickle.loads()."""
    return RestrictedUnpickler(io.BytesIO(s)).load()
```

A sample usage of our unpickler working has intended:

```
>>> restricted_loads(pickle.dumps([1, 2, range(15)]))
[1, 2, range(0, 15)]
>>> restricted_loads(b"cos\nsystem\n(S'echo hello world'\nR.")
Traceback (most recent call last):
...
pickle.UnpicklingError: global 'os.system' is forbidden
>>> restricted_loads(b'cbuiltins\neval\n'
...                  b'(S\'getattr(__import__("os"), "system")'
...                  b'("echo hello world")\'\nR.')
Traceback (most recent call last):
...
pickle.UnpicklingError: global 'builtins.eval' is forbidden
```

As our examples shows, you have to be careful with what you allow to be unpickled. Therefore if security is a concern, you may want to consider alternatives such as the marshalling API in *xmlrpc.client* or third-party solutions.

12.1.7 性能

Recent versions of the pickle protocol (from protocol 2 and upwards) feature efficient binary encodings for several common features and built-in types. Also, the *pickle* module has a transparent optimizer written in C.

12.1.8 示例

For the simplest code, use the *dump()* and *load()* functions.

```
import pickle

# An arbitrary collection of objects supported by pickle.
data = {
    'a': [1, 2.0, 3, 4+6j],
    'b': ("character string", b"byte string"),
    'c': {None, True, False}
}

with open('data.pickle', 'wb') as f:
    # Pickle the 'data' dictionary using the highest protocol available.
    pickle.dump(data, f, pickle.HIGHEST_PROTOCOL)
```

The following example reads the resulting pickled data.

```
import pickle

with open('data.pickle', 'rb') as f:
    # The protocol version used is detected automatically, so we do not
    # have to specify it.
    data = pickle.load(f)
```

参见:

Module `copyreg` Pickle interface constructor registration for extension types.

Module `pickletools` Tools for working with and analyzing pickled data.

模块 `shelve` Indexed databases of objects; uses `pickle`.

Module `copy` Shallow and deep object copying.

Module `marshal` High-performance serialization of built-in types.

12.2 copyreg — Register pickle support functions

Source code: `Lib/copyreg.py`

The `copyreg` module offers a way to define functions used while pickling specific objects. The `pickle` and `copy` modules use those functions when pickling/copying those objects. The module provides configuration information about object constructors which are not classes. Such constructors may be factory functions or class instances.

`copyreg.constructor(object)`

Declares *object* to be a valid constructor. If *object* is not callable (and hence not valid as a constructor), raises `TypeError`.

`copyreg.pickle(type, function, constructor=None)`

Declares that *function* should be used as a “reduction” function for objects of type *type*. *function* should return either a string or a tuple containing two or three elements.

The optional *constructor* parameter, if provided, is a callable object which can be used to reconstruct the object when called with the tuple of arguments returned by *function* at pickling time. `TypeError` will be raised if *object* is a class or *constructor* is not callable.

See the `pickle` module for more details on the interface expected of *function* and *constructor*. Note that the `dispatch_table` attribute of a pickler object or subclass of `pickle.Pickler` can also be used for declaring reduction functions.

12.2.1 示例

The example below would like to show how to register a pickle function and how it will be used:

```
>>> import copyreg, copy, pickle
>>> class C(object):
...     def __init__(self, a):
...         self.a = a
...
>>> def pickle_c(c):
...     print("pickling a C instance...")
...     return C, (c.a,)
...
>>> copyreg.pickle(C, pickle_c)
>>> c = C(1)
>>> d = copy.copy(c)
pickling a C instance...
>>> p = pickle.dumps(c)
pickling a C instance...
```

12.3 `shelve` —Python object persistence

Source code: [Lib/shelve.py](#)

A “shelf” is a persistent, dictionary-like object. The difference with “dbm” databases is that the values (not the keys!) in a shelf can be essentially arbitrary Python objects —anything that the `pickle` module can handle. This includes most class instances, recursive data types, and objects containing lots of shared sub-objects. The keys are ordinary strings.

`shelve.open(filename, flag='c', protocol=None, writeback=False)`

Open a persistent dictionary. The filename specified is the base filename for the underlying database. As a side-effect, an extension may be added to the filename and more than one file may be created. By default, the underlying database file is opened for reading and writing. The optional `flag` parameter has the same interpretation as the `flag` parameter of `dbm.open()`.

By default, version 3 pickles are used to serialize values. The version of the pickle protocol can be specified with the `protocol` parameter.

Because of Python semantics, a shelf cannot know when a mutable persistent-dictionary entry is modified. By default modified objects are written *only* when assigned to the shelf (see 示例). If the optional `writeback` parameter is set to `True`, all entries accessed are also cached in memory, and written back on `sync()` and `close()`; this can make it handier to mutate mutable entries in the persistent dictionary, but, if many entries are accessed, it can consume vast amounts of memory for the cache, and it can make the close operation very slow since all accessed entries are written back (there is no way to determine which accessed entries are mutable, nor which ones were actually mutated).

注解: Do not rely on the shelf being closed automatically; always call `close()` explicitly when you don't need it any more, or use `shelve.open()` as a context manager:

```
with shelve.open('spam') as db:
    db['eggs'] = 'eggs'
```

警告: Because the `shelve` module is backed by `pickle`, it is insecure to load a shelf from an untrusted source. Like with pickle, loading a shelf can execute arbitrary code.

Shelf objects support all methods supported by dictionaries. This eases the transition from dictionary based scripts to those requiring persistent storage.

Two additional methods are supported:

`Shelf.sync()`

Write back all entries in the cache if the shelf was opened with `writeback` set to `True`. Also empty the cache and synchronize the persistent dictionary on disk, if feasible. This is called automatically when the shelf is closed with `close()`.

`Shelf.close()`

Synchronize and close the persistent `dict` object. Operations on a closed shelf will fail with a `ValueError`.

参见:

[Persistent dictionary recipe](#) with widely supported storage formats and having the speed of native dictionaries.

12.3.1 Restrictions

- The choice of which database package will be used (such as `dbm.ndbm` or `dbm.gnu`) depends on which interface is available. Therefore it is not safe to open the database directly using `dbm`. The database is also (unfortunately) subject to the limitations of `dbm`, if it is used —this means that (the pickled representation of) the objects stored in the database should be fairly small, and in rare cases key collisions may cause the database to refuse updates.
- The `shelve` module does not support *concurrent* read/write access to shelved objects. (Multiple simultaneous read accesses are safe.) When a program has a shelf open for writing, no other program should have it open for reading or writing. Unix file locking can be used to solve this, but this differs across Unix versions and requires knowledge about the database implementation used.

class `shelve.Shelf` (*dict*, *protocol=None*, *writeback=False*, *keyencoding='utf-8'*)

A subclass of `collections.abc.MutableMapping` which stores pickled values in the *dict* object.

By default, version 3 pickles are used to serialize values. The version of the pickle protocol can be specified with the *protocol* parameter. See the `pickle` documentation for a discussion of the pickle protocols.

If the *writeback* parameter is `True`, the object will hold a cache of all entries accessed and write them back to the *dict* at sync and close times. This allows natural operations on mutable entries, but can consume much more memory and make sync and close take a long time.

The *keyencoding* parameter is the encoding used to encode keys before they are used with the underlying dict.

A `Shelf` object can also be used as a context manager, in which case it will be automatically closed when the `with` block ends.

在 3.2 版更改: Added the *keyencoding* parameter; previously, keys were always encoded in UTF-8.

在 3.4 版更改: 添加了上下文管理器支持

class `shelve.BsddbShelf` (*dict*, *protocol=None*, *writeback=False*, *keyencoding='utf-8'*)

A subclass of `Shelf` which exposes `first()`, `next()`, `previous()`, `last()` and `set_location()` which are available in the third-party `bsddb` module from `pybsddb` but not in other database modules. The *dict* object passed to the constructor must support those methods. This is generally accomplished by calling one of `bsddb.hashopen()`, `bsddb.btopen()` or `bsddb.rnopen()`. The optional *protocol*, *writeback*, and *keyencoding* parameters have the same interpretation as for the `Shelf` class.

class `shelve.DbfilenameShelf` (*filename*, *flag='c'*, *protocol=None*, *writeback=False*)

A subclass of `Shelf` which accepts a *filename* instead of a dict-like object. The underlying file will be opened using `dbm.open()`. By default, the file will be created and opened for both read and write. The optional *flag* parameter has the same interpretation as for the `open()` function. The optional *protocol* and *writeback* parameters have the same interpretation as for the `Shelf` class.

12.3.2 示例

To summarize the interface (key is a string, data is an arbitrary object):

```
import shelve

d = shelve.open(filename) # open -- file may get suffix added by low-level
                           # library

d[key] = data              # store data at key (overwrites old data if
                           # using an existing key)
data = d[key]              # retrieve a COPY of data at key (raise KeyError
```

(下页继续)

(续上页)

```

                                # if no such key)
del d[key]                      # delete data stored at key (raises KeyError)
                                # if no such key)

flag = key in d                  # true if the key exists
klist = list(d.keys())          # a list of all existing keys (slow!)

# as d was opened WITHOUT writeback=True, beware:
d['xx'] = [0, 1, 2]             # this works as expected, but...
d['xx'].append(3)               # *this doesn't!* -- d['xx'] is STILL [0, 1, 2]!

# having opened d without writeback=True, you need to code carefully:
temp = d['xx']                  # extracts the copy
temp.append(5)                  # mutates the copy
d['xx'] = temp                  # stores the copy right back, to persist it

# or, d=shelve.open(filename,writeback=True) would let you just code
# d['xx'].append(5) and have it work as expected, BUT it would also
# consume more memory and make the d.close() operation slower.

d.close()                      # close it

```

参见:**Module `dbm`** Generic interface to dbm-style databases.**模块 `pickle`** Object serialization used by `shelve`.

12.4 `marshal` —Internal Python object serialization

This module contains functions that can read and write Python values in a binary format. The format is specific to Python, but independent of machine architecture issues (e.g., you can write a Python value to a file on a PC, transport the file to a Sun, and read it back there). Details of the format are undocumented on purpose; it may change between Python versions (although it rarely does).¹

This is not a general “persistence” module. For general persistence and transfer of Python objects through RPC calls, see the modules `pickle` and `shelve`. The `marshal` module exists mainly to support reading and writing the “pseudo-compiled” code for Python modules of `.pyc` files. Therefore, the Python maintainers reserve the right to modify the marshal format in backward incompatible ways should the need arise. If you’re serializing and de-serializing Python objects, use the `pickle` module instead – the performance is comparable, version independence is guaranteed, and pickle supports a substantially wider range of objects than marshal.

警告: The `marshal` module is not intended to be secure against erroneous or maliciously constructed data. Never unmarshal data received from an untrusted or unauthenticated source.

Not all Python object types are supported; in general, only objects whose value is independent from a particular invocation of Python can be written and read by this module. The following types are supported: booleans, integers, floating point numbers, complex numbers, strings, bytes, bytearray, tuples, lists, sets, frozensets, dictionaries, and code objects, where

¹ The name of this module stems from a bit of terminology used by the designers of Modula-3 (amongst others), who use the term “marshalling” for shipping of data around in a self-contained form. Strictly speaking, “to marshal” means to convert some data from internal to external form (in an RPC buffer for instance) and “unmarshalling” for the reverse process.

it should be understood that tuples, lists, sets, frozensets and dictionaries are only supported as long as the values contained therein are themselves supported. The singletons *None*, *Ellipsis* and *StopIteration* can also be marshalled and unmarshalled. For format *version* lower than 3, recursive lists, sets and dictionaries cannot be written (see below).

There are functions that read/write files as well as functions operating on bytes-like objects.

The module defines these functions:

`marshal.dump(value, file[, version])`

Write the value on the open file. The value must be a supported type. The file must be a writeable *binary file*.

If the value has (or contains an object that has) an unsupported type, a *ValueError* exception is raised—but garbage data will also be written to the file. The object will not be properly read back by *load()*.

The *version* argument indicates the data format that *dump* should use (see below).

`marshal.load(file)`

Read one value from the open file and return it. If no valid value is read (e.g. because the data has a different Python version's incompatible marshal format), raise *EOFError*, *ValueError* or *TypeError*. The file must be a readable *binary file*.

注解: If an object containing an unsupported type was marshalled with *dump()*, *load()* will substitute *None* for the unmarshallable type.

`marshal.dumps(value[, version])`

Return the bytes object that would be written to a file by *dump(value, file)*. The value must be a supported type. Raise a *ValueError* exception if value has (or contains an object that has) an unsupported type.

The *version* argument indicates the data format that *dumps* should use (see below).

`marshal.loads(bytes)`

Convert the *bytes-like object* to a value. If no valid value is found, raise *EOFError*, *ValueError* or *TypeError*. Extra bytes in the input are ignored.

In addition, the following constants are defined:

`marshal.version`

Indicates the format that the module uses. Version 0 is the historical format, version 1 shares interned strings and version 2 uses a binary format for floating point numbers. Version 3 adds support for object instancing and recursion. The current version is 4.

12.5 dbm —Interfaces to Unix “databases”

源代码: [Lib/dbm/__init__.py](#)

dbm is a generic interface to variants of the DBM database—*dbm.gnu* or *dbm.ndbm*. If none of these modules is installed, the slow-but-simple implementation in module *dbm.dumb* will be used. There is a [third party interface](#) to the Oracle Berkeley DB.

exception `dbm.error`

A tuple containing the exceptions that can be raised by each of the supported modules, with a unique exception also named *dbm.error* as the first item—the latter is used when *dbm.error* is raised.

`dbm.whichdb(filename)`

This function attempts to guess which of the several simple database modules available—*dbm.gnu*, *dbm.ndbm* or *dbm.dumb*—should be used to open a given file.

Returns one of the following values: `None` if the file can't be opened because it's unreadable or doesn't exist; the empty string (`' '`) if the file's format can't be guessed; or a string containing the required module name, such as `'dbm.ndbm'` or `'dbm.gnu'`.

`dbm.open(file, flag='r', mode=0o666)`

Open the database file *file* and return a corresponding object.

If the database file already exists, the `whichdb()` function is used to determine its type and the appropriate module is used; if it does not exist, the first module listed above that can be imported is used.

可选的 *flag* 参数可以是：

值	意义
'r'	以只读方式打开现有数据库（默认）
'w'	以读写方式打开现有数据库
'c'	以读写方式打开数据库，如果不存在则创建它
'n'	始终创建一个新的空数据库，以读写方式打开

The optional *mode* argument is the Unix mode of the file, used only when the database has to be created. It defaults to octal `0o666` (and will be modified by the prevailing `umask`).

The object returned by `open()` supports the same basic functionality as dictionaries; keys and their corresponding values can be stored, retrieved, and deleted, and the `in` operator and the `keys()` method are available, as well as `get()` and `setdefault()`.

在 3.2 版更改: `get()` and `setdefault()` are now available in all database modules.

Key and values are always stored as bytes. This means that when strings are used they are implicitly converted to the default encoding before being stored.

These objects also support being used in a `with` statement, which will automatically close them when done.

在 3.4 版更改: Added native support for the context management protocol to the objects returned by `open()`.

The following example records some hostnames and a corresponding title, and then prints out the contents of the database:

```
import dbm

# Open database, creating it if necessary.
with dbm.open('cache', 'c') as db:

    # Record some values
    db[b'hello'] = b'there'
    db['www.python.org'] = 'Python Website'
    db['www.cnn.com'] = 'Cable News Network'

    # Note that the keys are considered bytes now.
    assert db[b'www.python.org'] == b'Python Website'
    # Notice how the value is now in bytes.
    assert db['www.cnn.com'] == b'Cable News Network'

    # Often-used methods of the dict interface work too.
    print(db.get('python.org', b'not present'))

    # Storing a non-string key or value will raise an exception (most
    # likely a TypeError).
    db['www.yahoo.com'] = 4

# db is automatically closed when leaving the with statement.
```

参见:

模块 [shelve](#) Persistence module which stores non-string data.

The individual submodules are described in the following sections.

12.5.1 `dbm.gnu` —GNU' s reinterpretation of `dbm`

源代码: [Lib/dbm/gnu.py](#)

This module is quite similar to the `dbm` module, but uses the GNU library `gdbm` instead to provide some additional functionality. Please note that the file formats created by `dbm.gnu` and `dbm.ndbm` are incompatible.

The `dbm.gnu` module provides an interface to the GNU DBM library. `dbm.gnu.gdbm` objects behave like mappings (dictionaries), except that keys and values are always converted to bytes before storing. Printing a `gdbm` object doesn't print the keys and values, and the `items()` and `values()` methods are not supported.

exception `dbm.gnu.error`

Raised on `dbm.gnu`-specific errors, such as I/O errors. `KeyError` is raised for general mapping errors like specifying an incorrect key.

`dbm.gnu.open(filename[, flag[, mode]])`

Open a `gdbm` database and return a `gdbm` object. The `filename` argument is the name of the database file.

可选的 `flag` 参数可以是:

值	意义
'r'	以只读方式打开现有数据库 (默认)
'w'	以读写方式打开现有数据库
'c'	以读写方式打开数据库, 如果不存在则创建它
'n'	始终创建一个新的空数据库, 以读写方式打开

The following additional characters may be appended to the flag to control how the database is opened:

值	意义
'f'	以快速模式打开数据库。写入数据库将不会同步。
's'	同步模式。这将导致数据库的更改立即写入文件。
'u'	不要锁定数据库。

Not all flags are valid for all versions of `gdbm`. The module constant `open_flags` is a string of supported flag characters. The exception `error` is raised if an invalid flag is specified.

The optional `mode` argument is the Unix mode of the file, used only when the database has to be created. It defaults to octal `0o666`.

In addition to the dictionary-like methods, `gdbm` objects have the following methods:

`gdbm.firstkey()`

It's possible to loop over every key in the database using this method and the `nextkey()` method. The traversal is ordered by `gdbm`'s internal hash values, and won't be sorted by the key values. This method returns the starting key.

`gdbm.nextkey(key)`

Returns the key that follows `key` in the traversal. The following code prints every key in the database `db`, without having to create a list in memory that contains them all:

```
k = db.firstkey()
while k != None:
    print(k)
    k = db.nextkey(k)
```

`gdbm.reorganize()`

If you have carried out a lot of deletions and would like to shrink the space used by the `gdbm` file, this routine will reorganize the database. `gdbm` objects will not shorten the length of a database file except by using this reorganization; otherwise, deleted file space will be kept and reused as new (key, value) pairs are added.

`gdbm.sync()`

When the database has been opened in fast mode, this method forces any unwritten data to be written to the disk.

`gdbm.close()`

Close the `gdbm` database.

12.5.2 `dbm.ndbm` —Interface based on `ndbm`

源代码: [Lib/dbm/ndbm.py](#)

The `dbm.ndbm` module provides an interface to the Unix “(n)dbm” library. Dbm objects behave like mappings (dictionaries), except that keys and values are always stored as bytes. Printing a dbm object doesn’t print the keys and values, and the `items()` and `values()` methods are not supported.

This module can be used with the “classic” `ndbm` interface or the GNU GDBM compatibility interface. On Unix, the **configure** script will attempt to locate the appropriate header file to simplify building this module.

exception `dbm.ndbm.error`

Raised on `dbm.ndbm`-specific errors, such as I/O errors. `KeyError` is raised for general mapping errors like specifying an incorrect key.

`dbm.ndbm.library`

Name of the `ndbm` implementation library used.

`dbm.ndbm.open(filename[, flag[, mode]])`

Open a dbm database and return a `ndbm` object. The *filename* argument is the name of the database file (without the `.dir` or `.pag` extensions).

The optional *flag* argument must be one of these values:

值	意义
'r'	以只读方式打开现有数据库（默认）
'w'	以读写方式打开现有数据库
'c'	以读写方式打开数据库，如果不存在则创建它
'n'	始终创建一个新的空数据库，以读写方式打开

The optional *mode* argument is the Unix mode of the file, used only when the database has to be created. It defaults to octal 0o666 (and will be modified by the prevailing `umask`).

In addition to the dictionary-like methods, `ndbm` objects provide the following method:

`ndbm.close()`

Close the `ndbm` database.

12.5.3 dbm.dumb —Portable DBM implementation

源代码: [Lib/dbm/dumb.py](#)

注解: The `dbm.dumb` module is intended as a last resort fallback for the `dbm` module when a more robust module is not available. The `dbm.dumb` module is not written for speed and is not nearly as heavily used as the other database modules.

The `dbm.dumb` module provides a persistent dictionary-like interface which is written entirely in Python. Unlike other modules such as `dbm.gnu` no external library is required. As with other persistent mappings, the keys and values are always stored as bytes.

The module defines the following:

exception `dbm.dumb.error`

Raised on `dbm.dumb`-specific errors, such as I/O errors. `KeyError` is raised for general mapping errors like specifying an incorrect key.

`dbm.dumb.open(filename[, flag[, mode]])`

Open a dumbdbm database and return a dumbdbm object. The `filename` argument is the basename of the database file (without any specific extensions). When a dumbdbm database is created, files with `.dat` and `.dir` extensions are created.

The optional `flag` argument supports only the semantics of `'c'` and `'n'` values. Other values will default to database being always opened for update, and will be created if it does not exist.

The optional `mode` argument is the Unix mode of the file, used only when the database has to be created. It defaults to octal `0o666` (and will be modified by the prevailing `umask`).

在 3.5 版更改: `open()` always creates a new database when the flag has the value `'n'`.

Deprecated since version 3.6, will be removed in version 3.8: Creating database in `'r'` and `'w'` modes. Modifying database in `'r'` mode.

In addition to the methods provided by the `collections.abc.MutableMapping` class, dumbdbm objects provide the following methods:

`dumbdbm.sync()`

Synchronize the on-disk directory and data files. This method is called by the `Shelve.sync()` method.

`dumbdbm.close()`

Close the dumbdbm database.

12.6 sqlite3 —SQLite 数据库 DB-API 2.0 接口模块

源代码: [Lib/sqlite3/](#)

SQLite 是一个 C 语言库, 它可以提供一种轻量级的基于磁盘的数据库, 这种数据库不需要独立的服务器进程, 也允许需要使用一种非标准的 SQL 查询语言来访问它。一些应用程序可以使用 SQLite 作为内部数据存储。可以用它来创建一个应用程序原型, 然后再迁移到更大的数据库, 比如 PostgreSQL 或 Oracle。

sqlite3 模块是由 Gerhard Häring 编写。它提供了符合 DB-API 2.0 规范的接口, 这个规范是 [PEP 249](#)。

要使用这个模块, 必须先创建一个 `Connection` 对象, 它代表数据库。下面例子中, 数据将存储在 `example.db` 文件中:

```
import sqlite3
conn = sqlite3.connect('example.db')
```

你也可以使用 `:memory:` 来创建一个内存中的数据库

当有了 `Connection` 对象后，你可以创建一个 `Cursor` 游标对象，然后调用它的 `execute()` 方法来执行 SQL 语句：

```
c = conn.cursor()

# Create table
c.execute('CREATE TABLE stocks
          (date text, trans text, symbol text, qty real, price real)')

# Insert a row of data
c.execute("INSERT INTO stocks VALUES ('2006-01-05','BUY','RHAT',100,35.14)")

# Save (commit) the changes
conn.commit()

# We can also close the connection if we are done with it.
# Just be sure any changes have been committed or they will be lost.
conn.close()
```

这些数据被持久化保存了，而且可以在之后的会话中使用它们：

```
import sqlite3
conn = sqlite3.connect('example.db')
c = conn.cursor()
```

通常你的 SQL 操作需要使用一些 Python 变量的值。你不应该使用 Python 的字符串操作来创建你的查询语句，因为那样做不安全；它会使你的程序容易受到 SQL 注入攻击（在 <https://xkcd.com/327/> 上有一个搞笑的例子，看看有什么后果）

推荐另外一种方法：使用 DB-API 的参数替换。在你的 SQL 语句中，使用 `?` 占位符来代替值，然后把对应的值组成的元组做为 `execute()` 方法的第二个参数。（其他数据库可能会使用不同的占位符，比如 `%s` 或者 `:1`）例如：

```
# Never do this -- insecure!
symbol = 'RHAT'
c.execute("SELECT * FROM stocks WHERE symbol = '%s'" % symbol)

# Do this instead
t = ('RHAT',)
c.execute('SELECT * FROM stocks WHERE symbol=?', t)
print(c.fetchone())

# Larger example that inserts many records at a time
purchases = [('2006-03-28', 'BUY', 'IBM', 1000, 45.00),
              ('2006-04-05', 'BUY', 'MSFT', 1000, 72.00),
              ('2006-04-06', 'SELL', 'IBM', 500, 53.00),
              ]
c.executemany('INSERT INTO stocks VALUES (?, ?, ?, ?, ?)', purchases)
```

要在执行 `SELECT` 语句后获取数据，你可以把游标作为 `iterator`，然后调用它的 `fetchone()` 方法来获取一条匹配的行，也可以调用 `fetchall()` 来得到包含多个匹配行的列表。

下面是一个使用迭代器形式的例子：

```
>>> for row in c.execute('SELECT * FROM stocks ORDER BY price'):
    print(row)

('2006-01-05', 'BUY', 'RHAT', 100, 35.14)
('2006-03-28', 'BUY', 'IBM', 1000, 45.0)
('2006-04-06', 'SELL', 'IBM', 500, 53.0)
('2006-04-05', 'BUY', 'MSFT', 1000, 72.0)
```

参见:

<https://github.com/ghaering/pysqlite> pysqlite 的主页 – sqlite3 在外部使用 “pysqlite” 名字进行开发。

<https://www.sqlite.org> SQLite 的主页；它的文档详细描述了它所支持的 SQL 方言的语法和可用的数据类型。

<http://www.w3schools.com/sql/> 学习 SQL 语法的教程、参考和例子。

PEP 249 - DB-API 2.0 规范 Marc-André Lemburg 写的 PEP。

12.6.1 模块函数和常量

sqlite3.version

这个模块的版本号，是一个字符串。不是 SQLite 库的版本号。

sqlite3.version_info

这个模块的版本号，是一个由整数组成的元组。不是 SQLite 库的版本号。

sqlite3.sqlite_version

使用中的 SQLite 库的版本号，是一个字符串。

sqlite3.sqlite_version_info

使用中的 SQLite 库的版本号，是一个整数组成的元组。

sqlite3.PARSE_DECLTYPES

这个常量可以作为 `connect()` 函数的 `detect_types` 参数。

设置这个参数后，`sqlite3` 模块将解析它返回的每一列声明的类型。它会声明的类型的第一个单词，比如 “integer primary key”，它会解析出 “integer”，再比如 “number(10)”，它会解析出 “number”。然后，它会在转换器字典里查找那个类型注册的转换器函数，并调用它。

sqlite3.PARSE_COLNAMES

这个常量可以作为 `connect()` 函数的 `detect_types` 参数。

设置这个参数后，SQLite 接口将解析它返回的每一列的列名。它会在其中查找 [mytype] 这个形式的字符串，然后用 ‘mytype’ 来决定那个列的类型。它会尝试在转换器字典中查找 ‘mytype’ 键对应的转换器函数，然后用这个转换器函数返回的值来做为列的类型。在 `Cursor.description` 中找到的列名仅仅是列名的第一个单词，比如你在 SQL 中使用 ‘as “x [datetime]”’，然后它会解析出第一个空白字符前的所有字符来作为列名：列名就是 “x”。

sqlite3.connect (*database* [, *timeout*, *detect_types*, *isolation_level*, *check_same_thread*, *factory*, *cached_statements*, *uri*])

Opens a connection to the SQLite database file *database*. You can use “:memory:” to open a database connection to a database that resides in RAM instead of on disk.

当一个数据库被多个连接访问的时候，如果其中一个进程修改这个数据库，在这个事务提交之前，这个 SQLite 数据库将会被一直锁定。`timeout` 参数指定了这个连接等待锁释放的超时时间，超时之后会引发一个异常。这个超时时间默认是 5.0（5 秒）。

`isolation_level` 参数，请查看 `Connection` 对象的 `isolation_level` 属性。

SQLite 原生只支持 5 种类型：TEXT，INTEGER，REAL，BLOB 和 NULL。如果你想用其它类型，你必须自己添加相应的支持。使用 `detect_types` 参数和模块级别的 `register_converter()` 函数注册 ** 转换器 ** 可以简单的实现。

`detect_types` 默认为 0 (即关闭, 没有类型检测)。你也可以组合 `PARSE_DECLTYPES` 和 `PARSE_COLNAMES` 来开启类型检测。

默认情况下, `check_same_thread` 为 `True`, 只有当前的线程可以使用该连接。如果设置为 `False`, 则多个线程可以共享返回的连接。当多个线程使用同一个连接的时候, 用户应该把写操作进行序列化, 以避免数据损坏。

默认情况下, 当调用 `connect` 方法的时候, `sqlite3` 模块使用了它的 `Connection` 类。当然, 你也可以创建 `Connection` 类的子类, 然后创建提供了 `factory` 参数的 `connect()` 方法。

详情请查阅当前手册的 [SQLite 与 Python 类型](#) 部分。

`sqlite3` 模块在内部使用语句缓存来避免 SQL 解析开销。如果要显式设置当前连接可以缓存的语句数, 可以设置 `cached_statements` 参数。当前实现的默认值是缓存 100 条语句。

如果 `uri` 为真, 则 `database` 被解释为 URI。它允许您指定选项。例如, 以只读模式打开数据库:

```
db = sqlite3.connect('file:path/to/database?mode=ro', uri=True)
```

有关此功能的更多信息, 包括已知选项的列表, 可以在 ‘SQLite URI 文档 <<https://www.sqlite.org/uri.html>>’ 中找到。

在 3.4 版更改: 增加了 `uri` 参数。

`sqlite3.register_converter(typename, callable)`

注册一个回调对象 `callable`, 用来转换数据库中的字节串为自定的 Python 类型。所有类型为 `typename` 的数据库的值在转换时, 都会调用这个回调对象。通过指定 `connect()` 函数的 `detect-types` 参数来设置类型检测的方式。注意, `typename` 与查询语句中的类型名进行匹配时不区分大小写。

`sqlite3.register_adapter(type, callable)`

注册一个回调对象 `callable`, 用来转换自定义 Python 类型为一个 SQLite 支持的类型。这个回调对象 `callable` 仅接受一个 Python 值作为参数, 而且必须返回以下某个类型的值: `int`, `float`, `str` 或 `bytes`。

`sqlite3.complete_statement(sql)`

如果字符串 `sql` 包含一个或多个完整的 SQL 语句 (以分号结束) 则返回 `True`。它不会验证 SQL 语法是否正确, 仅会验证字符串字面上是否完整, 以及是否以分号结束。

它可以用来构建一个 SQLite shell, 下面是一个例子:

```
# A minimal SQLite shell for experiments

import sqlite3

con = sqlite3.connect(":memory:")
con.isolation_level = None
cur = con.cursor()

buffer = ""

print("Enter your SQL commands to execute in sqlite3.")
print("Enter a blank line to exit.")

while True:
    line = input()
    if line == "":
        break
    buffer += line
```

(下页继续)

(续上页)

```

if sqlite3.complete_statement(buffer):
    try:
        buffer = buffer.strip()
        cur.execute(buffer)

        if buffer.lstrip().upper().startswith("SELECT"):
            print(cur.fetchall())
    except sqlite3.Error as e:
        print("An error occurred:", e.args[0])
    buffer = ""

con.close()

```

`sqlite3.enable_callback_tracebacks(flag)`

默认情况下，您不会获得任何用户定义函数中的回溯消息，比如聚合，转换器，授权器回调等。如果要调试它们，可以设置 *flag* 参数为 `True` 并调用此函数。之后，回调中的回溯信息将会输出到 `sys.stderr`。再次使用 *False* 来禁用该功能。

12.6.2 连接对象 (Connection)

class `sqlite3.Connection`

SQLite 数据库连接对象有如下的属性和方法：

isolation_level

获取或设置当前默认的隔离级别。表示自动提交模式的 *None* 以及 “DEFERRED”，“IMMEDIATE” 或 “EXCLUSIVE” 其中之一。详细描述请参阅 *Controlling Transactions*。

in_transaction

如果是在活动事务中（还没有提交改变），返回 *True*，否则，返回 *False*。它是一个只读属性。

3.2 新版功能。

cursor (*factory=Cursor*)

这个方法接受一个可选参数 *factory*，如果要指定这个参数，它必须是一个可调用对象，而且必须返回 *Cursor* 类的一个实例或者子类。

commit()

这个方法提交当前事务。如果没有调用这个方法，那么从上一次提交 `commit()` 以来所有的变化在其他数据库连接上都是不可见的。如果你往数据库里写了数据，但是又查询不到，请检查是否忘记了调用这个方法。

rollback()

这个方法回滚从上一次调用 `commit()` 以来所有数据库的改变。

close()

关闭数据库连接。注意，它不会自动调用 `commit()` 方法。如果在关闭数据库连接之前没有调用 `commit()`，那么你的修改将会丢失！

execute (*sql[, parameters]*)

这是一个非标准的快捷方法，它会调用 `cursor()` 方法来创建一个游标对象，并使用给定的 *parameters* 参数来调用游标对象的 `execute()` 方法，最后返回这个游标对象。

executemany (*sql[, parameters]*)

这是一个非标准的快捷方法，它会调用 `cursor()` 方法来创建一个游标对象，并使用给定的 *parameters* 参数来调用游标对象的 `executemany()` 方法，最后返回这个游标对象。

executescript (*sql_script*)

这是一个非标准的快捷方法，它会调用 `cursor()` 方法来创建一个游标对象，并使用给定的 `sql_script` 参数来调用游标对象的 `executescript()` 方法，最后返回这个游标对象。

create_function (*name, num_params, func*)

创建一个可以在 SQL 语句中使用的自定义函数，其中参数 *name* 为 SQL 语句中使用的函数名，*num_params* 是这个函数接受的参数个数（如果 *num_params* 为 -1，那这个函数可以接受任意数量的参数），最后一个参数 *func* 是作为 SQL 函数调用的一个 Python 可调用对象。

此函数可返回任何 SQLite 所支持的类型：bytes, str, int, float 和 None。

示例：

```
import sqlite3
import hashlib

def md5sum(t):
    return hashlib.md5(t).hexdigest()

con = sqlite3.connect(":memory:")
con.create_function("md5", 1, md5sum)
cur = con.cursor()
cur.execute("select md5(?)", (b"foo",))
print(cur.fetchone()[0])
```

create_aggregate (*name, num_params, aggregate_class*)

创建一个自定义的聚合函数。

参数中 *aggregate_class* 类必须实现两个方法：step 和 finalize。step 方法接受 *num_params* 个参数（如果 *num_params* 为 -1，那么这个函数可以接受任意数量的参数）；finalize 方法返回最终的聚合结果。

finalize 方法可以返回任何 SQLite 支持的类型：bytes, str, int, float 和 None。

示例：

```
import sqlite3

class MySum:
    def __init__(self):
        self.count = 0

    def step(self, value):
        self.count += value

    def finalize(self):
        return self.count

con = sqlite3.connect(":memory:")
con.create_aggregate("mysum", 1, MySum)
cur = con.cursor()
cur.execute("create table test(i)")
cur.execute("insert into test(i) values (1)")
cur.execute("insert into test(i) values (2)")
cur.execute("select mysum(i) from test")
print(cur.fetchone()[0])
```

create_collation (*name, callable*)

使用 *name* 和 *callable* 创建排序规则。这个 *callable* 接受两个字符串对象，如果第一个小于第二个

则返回 -1，如果两个相等则返回 0，如果第一个大于第二个则返回 1。注意，这是用来控制排序的 (SQL 中的 ORDER BY)，所以它不会影响其它的 SQL 操作。

注意，这个 *callable* 可调用对象会把它的参数作为 Python 字节串，通常会以 UTF-8 编码格式对它进行编码。

以下示例显示了使用“错误方式”进行排序的自定义排序规则：

```
import sqlite3

def collate_reverse(string1, string2):
    if string1 == string2:
        return 0
    elif string1 < string2:
        return 1
    else:
        return -1

con = sqlite3.connect(":memory:")
con.create_collation("reverse", collate_reverse)

cur = con.cursor()
cur.execute("create table test(x)")
cur.executemany("insert into test(x) values (?)", [("a",), ("b",)])
cur.execute("select x from test order by x collate reverse")
for row in cur:
    print(row)
con.close()
```

要移除一个排序规则，需要调用 `create_collation` 并设置 `callable` 参数为 `None`。

```
con.create_collation("reverse", None)
```

`interrupt()`

可以从不同的线程调用这个方法来终止所有查询操作，这些查询操作可能正在连接上执行。此方法调用之后，查询将会终止，而且查询的调用者会获得一个异常。

`set_authorizer(authorizer_callback)`

此方法注册一个授权回调对象。每次在访问数据库中某个表的某一列的时候，这个回调对象将会被调用。如果要允许访问，则返回 `SQLITE_OK`，如果要终止整个 SQL 语句，则返回 `SQLITE_DENY`，如果这一列需要当做 `NULL` 值处理，则返回 `SQLITE_IGNORE`。这些常量可以在 `sqlite3` 模块中找到。

回调的第一个参数表示要授权的操作类型。第二个和第三个参数将是参数或 `None`，具体取决于第一个参数的值。第 4 个参数是数据库的名称（“main”，“temp”等），如果需要的话。第 5 个参数是负责访问尝试的最内层触发器或视图的名称，或者如果此访问尝试直接来自输入 SQL 代码，则为 `None`。

请参阅 SQLite 文档，了解第一个参数的可能值以及第二个和第三个参数的含义，具体取决于第一个参数。所有必需的常量都可以在 `sqlite3` 模块中找到。

`set_progress_handler(handler, n)`

此例程注册回调。对 SQLite 虚拟机的每个多指令调用回调。如果要在长时间运行的操作期间从 SQLite 调用（例如更新用户界面），这非常有用。

如果要清除以前安装的任何进度处理程序，调用该方法时请将 `handler` 参数设置为 `None`。

从处理函数返回非零值将终止当前正在执行的查询并导致它引发 `OperationalError` 异常。

`set_trace_callback(trace_callback)`

为每个 SQLite 后端实际执行的 SQL 语句注册要调用的 `trace_callback`。

传递给回调的唯一参数是正在执行的语句（作为字符串）。回调的返回值将被忽略。请注意，后端不仅运行传递给 `Cursor.execute()` 方法的语句。其他来源包括 Python 模块的事务管理和当前数据库中定义的触发器的执行。

将传入的 `trace_callback` 设为 `None` 将禁用跟踪回调。

3.3 新版功能.

`enable_load_extension(enabled)`

此例程允许/禁止 SQLite 引擎从共享库加载 SQLite 扩展。SQLite 扩展可以定义新功能，聚合或全新的虚拟表实现。一个众所周知的扩展是与 SQLite 一起分发的全文搜索扩展。

默认情况下禁用可加载扩展。见¹。

3.2 新版功能.

```
import sqlite3

con = sqlite3.connect(":memory:")

# enable extension loading
con.enable_load_extension(True)

# Load the fulltext search extension
con.execute("select load_extension('./fts3.so')")

# alternatively you can load the extension using an API call:
# con.load_extension("./fts3.so")

# disable extension loading again
con.enable_load_extension(False)

# example from SQLite wiki
con.execute("create virtual table recipe using fts3(name, ingredients)")
con.executescript("""
    insert into recipe (name, ingredients) values ('broccoli stew', 'broccoli
↪peppers cheese tomatoes');
    insert into recipe (name, ingredients) values ('pumpkin stew', 'pumpkin
↪onions garlic celery');
    insert into recipe (name, ingredients) values ('broccoli pie', 'broccoli
↪cheese onions flour');
    insert into recipe (name, ingredients) values ('pumpkin pie', 'pumpkin
↪sugar flour butter');
""")
for row in con.execute("select rowid, name, ingredients from recipe where
↪name match 'pie'"):
    print(row)
```

`load_extension(path)`

此例程从共享库加载 SQLite 扩展。在使用此例程之前，必须使用 `enable_load_extension()` 启用扩展加载。

默认情况下禁用可加载扩展。见¹。

3.2 新版功能.

`row_factory`

您可以将此属性更改为可接受游标和原始行作为元组的可调用对象，并将返回实际结果行。这样，您可以实现更高级的返回结果的方法，例如返回一个可以按名称访问列的对象。

¹ `sqlite3` 模块默认没有构建可加载扩展支持，因为有一些平台带有不支持这个特性的 SQLite 库（特别是 Mac OS X）。要获得可加载扩展的支持，那么在编译配置的时候必须指定 `-enable-loadable-sqlite-extensions` 选项。

示例:

```
import sqlite3

def dict_factory(cursor, row):
    d = {}
    for idx, col in enumerate(cursor.description):
        d[col[0]] = row[idx]
    return d

con = sqlite3.connect(":memory:")
con.row_factory = dict_factory
cur = con.cursor()
cur.execute("select 1 as a")
print(cur.fetchone()["a"])
```

如果返回一个元组是不够的, 并且你想要对列进行基于名称的访问, 你应该考虑将 `row_factory` 设置为高度优化的 `sqlite3.Row` 类型。 `Row` 提供基于索引和不区分大小写的基于名称的访问, 几乎没有内存开销。它可能比您自己的基于字典的自定义方法甚至基于 `db_row` 的解决方案更好。

text_factory

使用此属性可以控制为 TEXT 数据类型返回的对象。默认情况下, 此属性设置为 `str` 和 `sqlite3` 模块将返回 TEXT 的 Unicode 对象。如果要返回字节串, 可以将其设置为 `bytes`。

您还可以将其设置为接受单个 `bytestring` 参数的任何其他可调用对象, 并返回结果对象。

请参阅以下示例代码以进行说明:

```
import sqlite3

con = sqlite3.connect(":memory:")
cur = con.cursor()

AUSTRIA = "\xd6sterreich"

# by default, rows are returned as Unicode
cur.execute("select ?", (AUSTRIA,))
row = cur.fetchone()
assert row[0] == AUSTRIA

# but we can make sqlite3 always return bytestrings ...
con.text_factory = bytes
cur.execute("select ?", (AUSTRIA,))
row = cur.fetchone()
assert type(row[0]) is bytes
# the bytestrings will be encoded in UTF-8, unless you stored garbage in the
# database ...
assert row[0] == AUSTRIA.encode("utf-8")

# we can also implement a custom text_factory ...
# here we implement one that appends "foo" to all strings
con.text_factory = lambda x: x.decode("utf-8") + "foo"
cur.execute("select ?", ("bar",))
row = cur.fetchone()
assert row[0] == "barfoo"
```

total_changes

返回自打开数据库连接以来已修改, 插入或删除的数据库行的总数。

iterdump()

返回以 SQL 文本格式转储数据库的迭代器。保存内存数据库以便以后恢复时很有用。此函数提供与 `sqlite3` shell 中的 `.dump` 命令相同的功能。

示例:

```
# Convert file existing_db.db to SQL dump file dump.sql
import sqlite3

con = sqlite3.connect('existing_db.db')
with open('dump.sql', 'w') as f:
    for line in con.iterdump():
        f.write('%s\n' % line)
```

12.6.3 Cursor 对象

class `sqlite3.Cursor`

Cursor 游标实例具有以下属性和方法。

execute (*sql* [, *parameters*])

执行 SQL 语句。可以是参数化 SQL 语句（即，在 SQL 语句中使用占位符）。`sqlite3` 模块支持两种占位符：问号（qmark 风格）和命名占位符（命名风格）。

以下是两种风格的示例：

```
import sqlite3

con = sqlite3.connect(":memory:")
cur = con.cursor()
cur.execute("create table people (name_last, age)")

who = "Yeltsin"
age = 72

# This is the qmark style:
cur.execute("insert into people values (?, ?)", (who, age))

# And this is the named style:
cur.execute("select * from people where name_last=:who and age=:age", {"who": who, "age": age})

print(cur.fetchone())
```

`execute()` will only execute a single SQL statement. If you try to execute more than one statement with it, it will raise a *Warning*. Use `executescript()` if you want to execute multiple SQL statements with one call.

executemany (*sql*, *seq_of_parameters*)

Executes an SQL command against all parameter sequences or mappings found in the sequence *seq_of_parameters*. The `sqlite3` module also allows using an *iterator* yielding parameters instead of a sequence.

```
import sqlite3

class IterChars:
    def __init__(self):
        self.count = ord('a')
```

(下页继续)

(续上页)

```

def __iter__(self):
    return self

def __next__(self):
    if self.count > ord('z'):
        raise StopIteration
    self.count += 1
    return (chr(self.count - 1),) # this is a 1-tuple

con = sqlite3.connect(":memory:")
cur = con.cursor()
cur.execute("create table characters(c)")

theIter = IterChars()
cur.executemany("insert into characters(c) values (?)", theIter)

cur.execute("select c from characters")
print(cur.fetchall())

```

这是一个使用生成器generator的简短示例：

```

import sqlite3
import string

def char_generator():
    for c in string.ascii_lowercase:
        yield (c,)

con = sqlite3.connect(":memory:")
cur = con.cursor()
cur.execute("create table characters(c)")

cur.executemany("insert into characters(c) values (?)", char_generator())

cur.execute("select c from characters")
print(cur.fetchall())

```

executescript (sql_script)

This is a nonstandard convenience method for executing multiple SQL statements at once. It issues a COMMIT statement first, then executes the SQL script it gets as a parameter.

sql_script can be an instance of *str*.

示例：

```

import sqlite3

con = sqlite3.connect(":memory:")
cur = con.cursor()
cur.executescript("""
    create table person(
        firstname,
        lastname,
        age
    );

    create table book(

```

(下页继续)

(续上页)

```

        title,
        author,
        published
    );

    insert into book(title, author, published)
    values (
        'Dirk Gently's Holistic Detective Agency',
        'Douglas Adams',
        1987
    );
    """

```

fetchone()

Fetches the next row of a query result set, returning a single sequence, or *None* when no more data is available.

fetchmany() (*size=cursor.arraysize*)

Fetches the next set of rows of a query result, returning a list. An empty list is returned when no more rows are available.

The number of rows to fetch per call is specified by the *size* parameter. If it is not given, the cursor's *arraysize* determines the number of rows to be fetched. The method should try to fetch as many rows as indicated by the *size* parameter. If this is not possible due to the specified number of rows not being available, fewer rows may be returned.

Note there are performance considerations involved with the *size* parameter. For optimal performance, it is usually best to use the *arraysize* attribute. If the *size* parameter is used, then it is best for it to retain the same value from one *fetchmany()* call to the next.

fetchall()

Fetches all (remaining) rows of a query result, returning a list. Note that the cursor's *arraysize* attribute can affect the performance of this operation. An empty list is returned when no rows are available.

close()

Close the cursor now (rather than whenever `__del__` is called).

The cursor will be unusable from this point forward; a *ProgrammingError* exception will be raised if any operation is attempted with the cursor.

rowcount

Although the *Cursor* class of the *sqlite3* module implements this attribute, the database engine's own support for the determination of “rows affected” / “rows selected” is quirky.

For *executemany()* statements, the number of modifications are summed up into *rowcount*.

As required by the Python DB API Spec, the *rowcount* attribute “is -1 in case no *executeXX()* has been performed on the cursor or the rowcount of the last operation is not determinable by the interface”. This includes *SELECT* statements because we cannot determine the number of rows a query produced until all rows were fetched.

With SQLite versions before 3.6.5, *rowcount* is set to 0 if you make a *DELETE FROM table* without any condition.

lastrowid

This read-only attribute provides the *rowid* of the last modified row. It is only set if you issued an *INSERT* or a *REPLACE* statement using the *execute()* method. For operations other than *INSERT* or *REPLACE* or when *executemany()* is called, *lastrowid* is set to *None*.

If the *INSERT* or *REPLACE* statement failed to insert the previous successful *rowid* is returned.

在 3.6 版更改: 增加了 REPLACE 语句的支持。

arraysize

Read/write attribute that controls the number of rows returned by `fetchmany()`. The default value is 1 which means a single row would be fetched per call.

description

This read-only attribute provides the column names of the last query. To remain compatible with the Python DB API, it returns a 7-tuple for each column where the last six items of each tuple are `None`.

It is set for SELECT statements without any matching rows as well.

connection

This read-only attribute provides the SQLite database `Connection` used by the `Cursor` object. A `Cursor` object created by calling `con.cursor()` will have a `connection` attribute that refers to `con`:

```
>>> con = sqlite3.connect(":memory:")
>>> cur = con.cursor()
>>> cur.connection == con
True
```

12.6.4 行对象 *Row*

class `sqlite3.Row`

A `Row` instance serves as a highly optimized `row_factory` for `Connection` objects. It tries to mimic a tuple in most of its features.

It supports mapping access by column name and index, iteration, representation, equality testing and `len()`.

If two `Row` objects have exactly the same columns and their members are equal, they compare equal.

keys()

This method returns a list of column names. Immediately after a query, it is the first member of each tuple in `Cursor.description`.

在 3.5 版更改: Added support of slicing.

Let's assume we initialize a table as in the example given above:

```
conn = sqlite3.connect(":memory:")
c = conn.cursor()
c.execute('create table stocks
(date text, trans text, symbol text,
qty real, price real)')
c.execute("""insert into stocks
values ('2006-01-05', 'BUY', 'RHAT', 100, 35.14)""")
conn.commit()
c.close()
```

Now we plug `Row` in:

```
>>> conn.row_factory = sqlite3.Row
>>> c = conn.cursor()
>>> c.execute('select * from stocks')
<sqlite3.Cursor object at 0x7f4e7dd8fa80>
>>> r = c.fetchone()
>>> type(r)
<class 'sqlite3.Row'>
```

(下页继续)

(续上页)

```

>>> tuple(r)
('2006-01-05', 'BUY', 'RHAT', 100.0, 35.14)
>>> len(r)
5
>>> r[2]
'RHAT'
>>> r.keys()
['date', 'trans', 'symbol', 'qty', 'price']
>>> r['qty']
100.0
>>> for member in r:
...     print(member)
...
2006-01-05
BUY
RHAT
100.0
35.14

```

12.6.5 异常

exception `sqlite3.Warning`

Exception 的一个子类。

exception `sqlite3.Error`

此模块中其他异常的基类。它是 *Exception* 的一个子类。

exception `sqlite3.DatabaseError`

Exception raised for errors that are related to the database.

exception `sqlite3.IntegrityError`

Exception raised when the relational integrity of the database is affected, e.g. a foreign key check fails. It is a subclass of *DatabaseError*.

exception `sqlite3.ProgrammingError`

Exception raised for programming errors, e.g. table not found or already exists, syntax error in the SQL statement, wrong number of parameters specified, etc. It is a subclass of *DatabaseError*.

exception `sqlite3.OperationalError`

Exception raised for errors that are related to the database's operation and not necessarily under the control of the programmer, e.g. an unexpected disconnect occurs, the data source name is not found, a transaction could not be processed, etc. It is a subclass of *DatabaseError*.

exception `sqlite3.NotSupportedError`

Exception raised in case a method or database API was used which is not supported by the database, e.g. calling the *rollback()* method on a connection that does not support transaction or has transactions turned off. It is a subclass of *DatabaseError*.

12.6.6 SQLite 与 Python 类型

概述

SQLite 原生支持如下的类型：NULL，INTEGER，REAL，TEXT，BLOB。

因此可以将以下 Python 类型发送到 SQLite 而不会出现任何问题：

Python 类型	SQLite 类型
<code>None</code>	NULL
<code>int</code>	INTEGER
<code>float</code>	REAL
<code>str</code>	TEXT
<code>bytes</code>	BLOB

这是 SQLite 类型默认转换为 Python 类型的方式:

SQLite 类型	Python 类型
NULL	<code>None</code>
INTEGER	<code>int</code>
REAL	<code>float</code>
TEXT	取决于 <code>text_factory</code> , 默认为 <code>str</code>
BLOB	<code>bytes</code>

The type system of the `sqlite3` module is extensible in two ways: you can store additional Python types in a SQLite database via object adaptation, and you can let the `sqlite3` module convert SQLite types to different Python types via converters.

Using adapters to store additional Python types in SQLite databases

As described before, SQLite supports only a limited set of types natively. To use other Python types with SQLite, you must **adapt** them to one of the `sqlite3` module's supported types for SQLite: one of `NoneType`, `int`, `float`, `str`, `bytes`.

There are two ways to enable the `sqlite3` module to adapt a custom Python type to one of the supported ones.

Letting your object adapt itself

This is a good approach if you write the class yourself. Let's suppose you have a class like this:

```
class Point:
    def __init__(self, x, y):
        self.x, self.y = x, y
```

Now you want to store the point in a single SQLite column. First you'll have to choose one of the supported types first to be used for representing the point. Let's just use `str` and separate the coordinates using a semicolon. Then you need to give your class a method `__conform__(self, protocol)` which must return the converted value. The parameter `protocol` will be `PrepareProtocol`.

```
import sqlite3

class Point:
    def __init__(self, x, y):
        self.x, self.y = x, y

    def __conform__(self, protocol):
        if protocol is sqlite3.PrepareProtocol:
            return "%f;%f" % (self.x, self.y)

con = sqlite3.connect(":memory:")
cur = con.cursor()
```

(下页继续)

(续上页)

```
p = Point(4.0, -3.2)
cur.execute("select ?", (p,))
print(cur.fetchone()[0])
```

Registering an adapter callable

The other possibility is to create a function that converts the type to the string representation and register the function with `register_adapter()`.

```
import sqlite3

class Point:
    def __init__(self, x, y):
        self.x, self.y = x, y

def adapt_point(point):
    return "%f;%f" % (point.x, point.y)

sqlite3.register_adapter(Point, adapt_point)

con = sqlite3.connect(":memory:")
cur = con.cursor()

p = Point(4.0, -3.2)
cur.execute("select ?", (p,))
print(cur.fetchone()[0])
```

The `sqlite3` module has two default adapters for Python's built-in `datetime.date` and `datetime.datetime` types. Now let's suppose we want to store `datetime.datetime` objects not in ISO representation, but as a Unix timestamp.

```
import sqlite3
import datetime
import time

def adapt_datetime(ts):
    return time.mktime(ts.timetuple())

sqlite3.register_adapter(datetime.datetime, adapt_datetime)

con = sqlite3.connect(":memory:")
cur = con.cursor()

now = datetime.datetime.now()
cur.execute("select ?", (now,))
print(cur.fetchone()[0])
```

Converting SQLite values to custom Python types

Writing an adapter lets you send custom Python types to SQLite. But to make it really useful we need to make the Python to SQLite to Python roundtrip work.

Enter converters.

Let's go back to the `Point` class. We stored the `x` and `y` coordinates separated via semicolons as strings in SQLite. First, we'll define a converter function that accepts the string as a parameter and constructs a `Point` object from it.

注解: Converter functions **always** get called with a `bytes` object, no matter under which data type you sent the value to SQLite.

```
def convert_point(s):
    x, y = map(float, s.split(b";"))
    return Point(x, y)
```

Now you need to make the `sqlite3` module know that what you select from the database is actually a point. There are two ways of doing this:

- Implicitly via the declared type
- Explicitly via the column name

Both ways are described in section 模块函数和常量, in the entries for the constants `PARSE_DECLTYPES` and `PARSE_COLNAMES`.

The following example illustrates both approaches.

```
import sqlite3

class Point:
    def __init__(self, x, y):
        self.x, self.y = x, y

    def __repr__(self):
        return "(%f;%f)" % (self.x, self.y)

def adapt_point(point):
    return "(%f;%f)" % (point.x, point.y).encode('ascii')

def convert_point(s):
    x, y = list(map(float, s.split(b";")))
    return Point(x, y)

# Register the adapter
sqlite3.register_adapter(Point, adapt_point)

# Register the converter
sqlite3.register_converter("point", convert_point)

p = Point(4.0, -3.2)

#####
# 1) Using declared types
con = sqlite3.connect(":memory:", detect_types=sqlite3.PARSE_DECLTYPES)
cur = con.cursor()
cur.execute("create table test(p point)")

cur.execute("insert into test(p) values (?)", (p,))
cur.execute("select p from test")
print("with declared types:", cur.fetchone()[0])
cur.close()
con.close()
```

(下页继续)

(续上页)

```
#####
# 1) Using column names
con = sqlite3.connect(":memory:", detect_types=sqlite3.PARSE_COLNAMES)
cur = con.cursor()
cur.execute("create table test(p)")

cur.execute("insert into test(p) values (?)", (p,))
cur.execute('select p as "p [point]" from test')
print("with column names:", cur.fetchone()[0])
cur.close()
con.close()
```

Default adapters and converters

There are default adapters for the date and datetime types in the datetime module. They will be sent as ISO dates/ISO timestamps to SQLite.

The default converters are registered under the name “date” for *datetime.date* and under the name “timestamp” for *datetime.datetime*.

This way, you can use date/timestamps from Python without any additional fiddling in most cases. The format of the adapters is also compatible with the experimental SQLite date/time functions.

The following example demonstrates this.

```
import sqlite3
import datetime

con = sqlite3.connect(":memory:", detect_types=sqlite3.PARSE_DECLTYPES|sqlite3.PARSE_
↳COLNAMES)
cur = con.cursor()
cur.execute("create table test(d date, ts timestamp)")

today = datetime.date.today()
now = datetime.datetime.now()

cur.execute("insert into test(d, ts) values (?, ?)", (today, now))
cur.execute("select d, ts from test")
row = cur.fetchone()
print(today, "=>", row[0], type(row[0]))
print(now, "=>", row[1], type(row[1]))

cur.execute('select current_date as "d [date]", current_timestamp as "ts [timestamp]"
↳')
row = cur.fetchone()
print("current_date", row[0], type(row[0]))
print("current_timestamp", row[1], type(row[1]))
```

If a timestamp stored in SQLite has a fractional part longer than 6 numbers, its value will be truncated to microsecond precision by the timestamp converter.

12.6.7 Controlling Transactions

The underlying `sqlite3` library operates in `autocommit` mode by default, but the Python `sqlite3` module by default does not.

`autocommit` mode means that statements that modify the database take effect immediately. A `BEGIN` or `SAVEPOINT` statement disables `autocommit` mode, and a `COMMIT`, a `ROLLBACK`, or a `RELEASE` that ends the outermost transaction, turns `autocommit` mode back on.

The Python `sqlite3` module by default issues a `BEGIN` statement implicitly before a Data Modification Language (DML) statement (i.e. `INSERT/UPDATE/DELETE/REPLACE`).

You can control which kind of `BEGIN` statements `sqlite3` implicitly executes via the `isolation_level` parameter to the `connect()` call, or via the `isolation_level` property of connections. If you specify no `isolation_level`, a plain `BEGIN` is used, which is equivalent to specifying `DEFERRED`. Other possible values are `IMMEDIATE` and `EXCLUSIVE`.

You can disable the `sqlite3` module's implicit transaction management by setting `isolation_level` to `None`. This will leave the underlying `sqlite3` library operating in `autocommit` mode. You can then completely control the transaction state by explicitly issuing `BEGIN`, `ROLLBACK`, `SAVEPOINT`, and `RELEASE` statements in your code.

在 3.6 版更改: `sqlite3` used to implicitly commit an open transaction before DDL statements. This is no longer the case.

12.6.8 Using `sqlite3` efficiently

Using shortcut methods

Using the nonstandard `execute()`, `executemany()` and `executescript()` methods of the `Connection` object, your code can be written more concisely because you don't have to create the (often superfluous) `Cursor` objects explicitly. Instead, the `Cursor` objects are created implicitly and these shortcut methods return the cursor objects. This way, you can execute a `SELECT` statement and iterate over it directly using only a single call on the `Connection` object.

```
import sqlite3

persons = [
    ("Hugo", "Boss"),
    ("Calvin", "Klein")
]

con = sqlite3.connect(":memory:")

# Create the table
con.execute("create table person(firstname, lastname)")

# Fill the table
con.executemany("insert into person(firstname, lastname) values (?, ?)", persons)

# Print the table contents
for row in con.execute("select firstname, lastname from person"):
    print(row)

print("I just deleted", con.execute("delete from person").rowcount, "rows")
```

Accessing columns by name instead of by index

One useful feature of the `sqlite3` module is the built-in `sqlite3.Row` class designed to be used as a row factory.

Rows wrapped with this class can be accessed both by index (like tuples) and case-insensitively by name:

```
import sqlite3

con = sqlite3.connect(":memory:")
con.row_factory = sqlite3.Row

cur = con.cursor()
cur.execute("select 'John' as name, 42 as age")
for row in cur:
    assert row[0] == row["name"]
    assert row["name"] == row["nAmE"]
    assert row[1] == row["age"]
    assert row[1] == row["AgE"]
```

使用连接作为上下文管理器

连接对象可以用来作为上下文管理器，它可以自动提交或者回滚事务。如果出现异常，事务会被回滚；否则，事务会被提交。

```
import sqlite3

con = sqlite3.connect(":memory:")
con.execute("create table person (id integer primary key, firstname varchar unique)")

# Successful, con.commit() is called automatically afterwards
with con:
    con.execute("insert into person(firstname) values (?)", ("Joe",))

# con.rollback() is called after the with block finishes with an exception, the
# exception is still raised and must be caught
try:
    with con:
        con.execute("insert into person(firstname) values (?)", ("Joe",))
except sqlite3.IntegrityError:
    print("couldn't add Joe twice")
```

12.6.9 常见问题

多线程

Older SQLite versions had issues with sharing connections between threads. That's why the Python module disallows sharing connections and cursors between threads. If you still try to do so, you will get an exception at runtime.

The only exception is calling the `interrupt()` method, which only makes sense to call from a different thread.

数据压缩和存档

本章中描述的模块支持 `zlib`、`gzip`、`bzip2` 和 `lzma` 数据压缩算法，以及创建 ZIP 和 tar 格式的归档文件。参见由 `shutil` 模块提供的 *Archiving operations*。

13.1 `zlib` 一与 `gzip` 兼容的压缩

此模块为需要数据压缩的程序提供了一系列函数，用于压缩和解压缩。这些函数使用了 `zlib` 库。`zlib` 库的项目主页是 <http://www.zlib.net>。版本低于 1.1.3 的 `zlib` 与此 Python 模块之间存在已知的不兼容。1.1.3 版本的 `zlib` 存在一个安全漏洞，我们推荐使用 1.1.4 或更新的版本。

`zlib` 的函数有很多选项，一般需要按特定顺序使用。本文档没有覆盖全部的用法。更多详细信息请于 <http://www.zlib.net/manual.html> 参阅官方手册。

要读写 `.gz` 格式的文件，请参考 `gzip` 模块。

此模块中可用的异常和函数如下：

exception `zlib.error`

在压缩或解压缩过程中发生错误时的异常。

`zlib.adler32(data[, value])`

计算 `data` 的 Adler-32 校验值。(Adler-32 校验的可靠性与 CRC32 基本相当，但比计算 CRC32 更高效。) 计算的结果是一个 32 位的整数。参数 `value` 是校验时的起始值，其默认值为 1。借助参数 `value` 可为分段的输入计算校验值。此算法没有加密强度，不应用于身份验证和数字签名。此算法的目的仅为验证数据的正确性，不适合作为通用散列算法。

在 3.0 版更改：返回值永远是无符号数。要在所有的 Python 版本和平台上获得相同的值，请使用 `adler32(data) & 0xffffffff`。

`zlib.compress(data, level=-1)`

压缩 `data` 中的字节，返回含有已压缩内容的 `bytes` 对象。参数 `level` 为整数，可取值为 0 到 9 或 -1，用于指定压缩等级。1 (`Z_BEST_SPEED`) 表示最快速度和最低压缩率，9 (`Z_BEST_COMPRESSION`) 表示最慢速度和最高压缩率。0 (`Z_NO_COMPRESSION`) 表示不压缩。参数默认值为 -1

(Z_DEFAULT_COMPRESSION)。Z_DEFAULT_COMPRESSION 是速度和压缩率之间的平衡 (一般相当于设压缩等级为 6)。函数发生错误时抛出 `error` 异常。

在 3.6 版更改: 现在, `level` 可作为关键字参数。

```
zlib.compressobj (level=-1, method=DEFLATED, wbits=MAX_WBITS, memLevel=DEF_MEM_LEVEL,
                  strategy=Z_DEFAULT_STRATEGY[, zdict ])
```

返回一个压缩对象, 用来压缩内存中难以容下的数据流。

参数 `level` 为压缩等级, 是整数, 可取值为 0 到 9 或 -1。1 (Z_BEST_SPEED) 表示最快速度和最低压缩率, 9 (Z_BEST_COMPRESSION) 表示最慢速度和最高压缩率。0 (Z_NO_COMPRESSION) 表示不压缩。参数默认值为 -1 (Z_DEFAULT_COMPRESSION)。Z_DEFAULT_COMPRESSION 是速度和压缩率之间的平衡 (一般相当于设压缩等级为 6)。

`method` 表示压缩算法。现在只支持 DEFLATED 这个算法。

参数 `wbits` 指定压缩数据时所使用的历史缓冲区的大小 (窗口大小), 并指定压缩输出是否包含头部或尾部。参数的默认值是 15 (MAX_WBITS)。参数的值分为几个范围:

- +9 到 +15: 窗口大小以 2 为底的对数。即这些值对应着 512 到 32768 的窗口大小。更大的值会提供更好的压缩, 同时内存开销也会更大。压缩输出会包含 zlib 特定格式的头部和尾部。
- -9 到 -15: 绝对值为窗口大小以 2 为底的对数。压缩输出仅包含压缩数据, 没有头部和尾部。
- +25 到 +31 = 16 + (9 到 15): 后 4 个比特位为窗口大小以 2 为底的对数。压缩输出包含一个基本的 **gzip** 头部, 并以校验和为尾部。

参数 `memLevel` 指定内部压缩操作时所占用内存大小。参数取 1 到 9。更大的值占用更多的内存, 同时速度也更快输出也更小。

参数 `strategy` 用于调节压缩算法。可取值为 Z_DEFAULT_STRATEGY、Z_FILTERED、Z_HUFFMAN_ONLY、Z_RLE (zlib 1.2.0.1) 或 Z_FIXED (zlib 1.2.2.2)。

参数 `zdict` 指定预定义的压缩字典。它是一个字节序列 (如 `bytes` 对象), 其中包含用户认为要压缩的数据中可能频繁出现的子序列。频率高的子序列应当放在字典的尾部。

在 3.3 版更改: 添加关键字参数 `zdict`。

```
zlib.crc32 (data[, value ])
```

计算 `data` 的 CRC (循环冗余校验) 值。计算的结果是一个 32 位的整数。参数 `value` 是校验时的起始值, 其默认值为 0。借助参数 `value` 可为分段的输入计算校验值。此算法没有加密强度, 不应用于身份验证和数字签名。此算法的目的仅为验证数据的正确性, 不适合作为通用散列算法。

在 3.0 版更改: 返回值永远是无符号数。要在所有的 Python 版本和平台上获得相同的值, 请使用 `crc32(data) & 0xffffffff`。

```
zlib.decompress (data, wbits=MAX_WBITS, bufsize=DEF_BUF_SIZE)
```

解压 `data` 中的字节, 返回含有已解压内容的 `bytes` 对象。参数 `wbits` 取决于 `data` 的格式, 具体参见下边的说明。`bufsize` 为输出缓冲区的起始大小。函数发生错误时抛出 `error` 异常。

The `wbits` parameter controls the size of the history buffer (or “window size”), and what header and trailer format is expected. It is similar to the parameter for `compressobj()`, but accepts more ranges of values:

- +8 to +15: The base-two logarithm of the window size. The input must include a zlib header and trailer.
- 0: Automatically determine the window size from the zlib header. Only supported since zlib 1.2.3.5.
- -8 to -15: Uses the absolute value of `wbits` as the window size logarithm. The input must be a raw stream with no header or trailer.
- +24 to +31 = 16 + (8 to 15): Uses the low 4 bits of the value as the window size logarithm. The input must include a gzip header and trailer.

- +40 to +47 = 32 + (8 to 15): Uses the low 4 bits of the value as the window size logarithm, and automatically accepts either the zlib or gzip format.

When decompressing a stream, the window size must not be smaller than the size originally used to compress the stream; using a too-small value may result in an *error* exception. The default *wbits* value corresponds to the largest window size and requires a zlib header and trailer to be included.

bufsize is the initial size of the buffer used to hold decompressed data. If more space is required, the buffer size will be increased as needed, so you don't have to get this value exactly right; tuning it will only save a few calls to `malloc()`.

在 3.6 版更改: *wbits* and *bufsize* can be used as keyword arguments.

`zlib.decompressobj(wbits=MAX_WBITS[, zdict])`

Returns a decompression object, to be used for decompressing data streams that won't fit into memory at once.

The *wbits* parameter controls the size of the history buffer (or the “window size”), and what header and trailer format is expected. It has the same meaning as *described for decompress()*.

The *zdict* parameter specifies a predefined compression dictionary. If provided, this must be the same dictionary as was used by the compressor that produced the data that is to be decompressed.

注解: If *zdict* is a mutable object (such as a *bytearray*), you must not modify its contents between the call to *decompressobj()* and the first call to the decompressor's *decompress()* method.

在 3.3 版更改: Added the *zdict* parameter.

压缩对象支持以下方法:

`Compress.compress(data)`

压缩 *data* 并返回 *bytes* 对象, 这个对象含有 *data* 的部分或全部内容的已压缩数据。所得的对象必须拼接在上一次调用 *compress()* 方法所得数据的后面。缓冲区中可能留存部分输入以供下一次调用。

`Compress.flush([mode])`

压缩所有缓冲区的数据并返回已压缩的数据。参数 *mode* 可以传入的常量为: `Z_NO_FLUSH`、`Z_PARTIAL_FLUSH`、`Z_SYNC_FLUSH`、`Z_FULL_FLUSH`、`Z_BLOCK` (zlib 1.2.3.4) 或 `Z_FINISH`。默认值为 `Z_FINISH`。`Z_FINISH` 关闭已压缩数据流并不允许再压缩其他数据, `Z_FINISH` 以外的值皆允许这个对象继续压缩数据。调用 *flush()* 方法并将 *mode* 设为 `Z_FINISH` 后会无法再次调用 *compress()*, 此时只能删除这个对象。

`Compress.copy()`

返回此压缩对象的一个拷贝。它可以用来高效压缩一系列拥有相同前缀的数据。

解压缩对象支持以下方法:

`Decompress.unused_data`

A *bytes* object which contains any bytes past the end of the compressed data. That is, this remains `b""` until the last byte that contains compression data is available. If the whole bytestring turned out to contain compressed data, this is `b""`, an empty *bytes* object.

`Decompress.unconsumed_tail`

A *bytes* object that contains any data that was not consumed by the last *decompress()* call because it exceeded the limit for the uncompressed data buffer. This data has not yet been seen by the zlib machinery, so you must feed it (possibly with further data concatenated to it) back to a subsequent *decompress()* method call in order to get correct output.

`Decompress.eof`

A boolean indicating whether the end of the compressed data stream has been reached.

This makes it possible to distinguish between a properly-formed compressed stream, and an incomplete or truncated one.

3.3 新版功能.

`Decompress.decompress (data, max_length=0)`

Decompress *data*, returning a bytes object containing the uncompressed data corresponding to at least part of the data in *string*. This data should be concatenated to the output produced by any preceding calls to the `decompress()` method. Some of the input data may be preserved in internal buffers for later processing.

If the optional parameter *max_length* is non-zero then the return value will be no longer than *max_length*. This may mean that not all of the compressed input can be processed; and unconsumed data will be stored in the attribute `unconsumed_tail`. This bytestring must be passed to a subsequent call to `decompress()` if decompression is to continue. If *max_length* is zero then the whole input is decompressed, and `unconsumed_tail` is empty.

在 3.6 版更改: *max_length* can be used as a keyword argument.

`Decompress.flush ([length])`

All pending input is processed, and a bytes object containing the remaining uncompressed output is returned. After calling `flush()`, the `decompress()` method cannot be called again; the only realistic action is to delete the object.

The optional parameter *length* sets the initial size of the output buffer.

`Decompress.copy ()`

Returns a copy of the decompression object. This can be used to save the state of the decompressor midway through the data stream in order to speed up random seeks into the stream at a future point.

通过下列常量可获取模块所使用的 zlib 库的版本信息:

`zlib.ZLIB_VERSION`

构建此模块时所用的 zlib 库的版本字符串。它的值可能与运行时所加载的 zlib 不同。运行时加载的 zlib 库的版本字符串为 `ZLIB_RUNTIME_VERSION`。

`zlib.ZLIB_RUNTIME_VERSION`

解释器所加载的 zlib 库的版本字符串。

3.3 新版功能.

参见:

模块 `gzip` 读写 `gzip` 格式的文件。

<http://www.zlib.net> zlib 库项目主页。

<http://www.zlib.net/manual.html> zlib 库用户手册。提供了库的许多功能的解释和用法。

13.2 gzip — 对 gzip 格式的支持

源代码: `Lib/gzip.py`

此模块提供的简单接口帮助用户压缩和解压缩文件, 功能类似于 GNU 应用程序 `gzip` 和 `gunzip`。

数据压缩由 `zlib` 模块提供。

`gzip` 模块提供 `GzipFile` 类和 `open()`、`compress()`、`decompress()` 几个便利的函数。`GzipFile` 类可以读写 `gzip` 格式的文件, 还能自动压缩和解压缩数据, 这让操作压缩文件如同操作普通的 *file object* 一样方便。

注意, 此模块不支持部分可以被 `gzip` 和 `gunzip` 解压的格式, 如利用 `compress` 或 `pack` 压缩所得的文件。

这个模块定义了以下内容：

`gzip.open(filename, mode='rb', compresslevel=9, encoding=None, errors=None, newline=None)`

以二进制方式或者文本方式打开一个 `gzip` 格式的压缩文件，返回一个 *file object*。

filename 参数可以是一个实际的文件名（一个 *a str* 对象或者 *bytes* 对象），或者是一个用来读写的已存在的文件对象。

mode 参数可以是二进制模式：'r', 'rb', 'a', 'ab', 'w', 'wb', 'x' or 'xb'，或者是文本模式 'rt', 'at', 'wt', or 'xt'。默认值是 'rb'。

The *compresslevel* argument is an integer from 0 to 9, as for the *GzipFile* constructor.

对于二进制模式，这个函数等价于 *GzipFile* 构造器：`GzipFile(filename, mode, compresslevel)`。在这个例子中，*encoding*, *errors* 和 *newline* 三个参数一定不要设置。

对于文本模式，将会创建一个 *GzipFile* 对象，并将它封装到一个 *io.TextIOWrapper* 实例中，这个实例默认了指定编码，错误捕获行为和行。

在 3.3 版更改：支持 *filename* 为一个文件对象，支持文本模式和 *encoding*, *errors* 和 *newline* 参数。

在 3.4 版更改：支持 'x', 'xb' 和 "xt" 三种模式。

在 3.6 版更改：接受一个 *path-like object*。

class `gzip.GzipFile(filename=None, mode=None, compresslevel=9, fileobj=None, mtime=None)`

Constructor for the *GzipFile* class, which simulates most of the methods of a *file object*, with the exception of the `truncate()` method. At least one of *fileobj* and *filename* must be given a non-trivial value.

新的实例基于 *fileobj*，它可以是一个普通文件，一个 *io.BytesIO* 对象，或者任何一个与文件相似的对象。当 *filename* 是一个文件对象时，它的默认值是 `None`。

当 *fileobj* 为 `None` 时，*filename* 参数只用于 `gzip` 文件头中，这个文件有可能包含未压缩文件的源文件名。如果文件可以被识别，默认 *fileobj* 的文件名；否则默认为空字符串，在这种情况下文件头将不包含源文件名。

The *mode* argument can be any of 'r', 'rb', 'a', 'ab', 'w', 'wb', 'x', or 'xb', depending on whether the file will be read or written. The default is the mode of *fileobj* if discernible; otherwise, the default is 'rb'.

需要注意的是，文件默认使用二进制模式打开。如果要以文本模式打开文件一个压缩文件，请使用 `open()` 方法（或者使用 *io.TextIOWrapper* 包装 *GzipFile*）。

compresslevel 参数是一个从 0 to 9 的整数，用于控制压缩等级；1 最快但压缩比例最小，9 最慢但压缩比例最大。0 不压缩。默认为 9。

The *mtime* argument is an optional numeric timestamp to be written to the last modification time field in the stream when compressing. It should only be provided in compression mode. If omitted or `None`, the current time is used. See the *mtime* attribute for more details.

调用 *GzipFile* 的 `close()` 方法不会关闭 *fileobj*，因为你可以希望增加其它内容到已经压缩的数中。你可以将一个 *io.BytesIO* 对象作为 *fileobj*，也可以使用 *io.BytesIO* 的 `getvalue()` 方法从内存缓存中恢复数据。

GzipFile 支持 *io.BufferedIOBase* 类的接口，包括迭代和 `with` 语句。只有 `truncate()` 方法没有实现。

GzipFile 还提供了以下的方法和属性：

peek(n)

在不移动文件指针的情况下读取 *n* 个未压缩字节。最多只有一个单独的读取流来服务这个方法调用。返回的字节数不一定刚好等于要求的数量。

注解: 调用`peek()` 并没有改变`GzipFile` 的文件指针, 它可能改变潜在文件对象 (例如: `GzipFile` 使用 `fileobj` 参数进行初始化)。

3.2 新版功能.

mtime

在解压的过程中, 最后修改时间字段的值可能来自于这个属性, 以整数的形式出现。在读取任何文件头信息前, 初始值为 `None`。

所有 **gzip** 东方压缩流中必须包含时间戳这个字段。以便于像 **gunzip** 这样的程序可以使用时间戳。格式与`time.time()` 的返回值和`os.stat()` 对象的`st_mtime` 属性值一样。

在 3.1 版更改: 支持 `with` 语句, 构造器参数 `mtime` 和`mtime` 属性。

在 3.2 版更改: 添加了对零填充和不可搜索文件的支持。

在 3.3 版更改: 实现`io.BufferedIOBase.read1()` 方法。

在 3.4 版更改: 支持 `'x'` and `'xb'` 两种模式。

在 3.5 版更改: 支持写入任意`bytes-like objects`。 `read()` 方法可以接受 `"None"` 为参数。

在 3.6 版更改: 接受一个`path-like object`。

`gzip.compress(data, compresslevel=9)`

压缩 `data` 返回一个包含压缩数据的`bytes` 对象。`compresslevel` 的意义与上面提到的`GzipFile` 构造器一至。

3.2 新版功能.

`gzip.decompress(data)`

解压数据, 返回一个`bytes` 包含未解压数据的对象。

3.2 新版功能.

13.2.1 用法示例

读取压缩文件示例:

```
import gzip
with gzip.open('/home/joe/file.txt.gz', 'rb') as f:
    file_content = f.read()
```

创建 GZIP 文件示例:

```
import gzip
content = b"Lots of content here"
with gzip.open('/home/joe/file.txt.gz', 'wb') as f:
    f.write(content)
```

使用 GZIP 压缩已有的文件示例:

```
import gzip
import shutil
with open('/home/joe/file.txt', 'rb') as f_in:
    with gzip.open('/home/joe/file.txt.gz', 'wb') as f_out:
        shutil.copyfileobj(f_in, f_out)
```

使用 GZIP 压缩二进制字符串示例:

```
import gzip
s_in = b"Lots of content here"
s_out = gzip.compress(s_in)
```

参见:

模块 `zlib` 支持 `gzip` 格式所需要的基本压缩模块。

13.3 bz2 — 对 bzip2 压缩算法的支持

源代码: `Lib/bz2.py`

此模块提供了使用 `bzip2` 压缩算法压缩和解压数据的一套完整的接口。

`bz2` 模块包含:

- 用于读写压缩文件的 `open()` 函数和 `BZ2File` 类。
- 用于增量压缩和解压的 `BZ2Compressor` 和 `BZ2Decompressor` 类。
- 用于一次性压缩和解压的 `compress()` 和 `decompress()` 函数。

此模块中的所有类都能安全地从多个线程访问。

13.3.1 文件压缩和解压

`bz2.open(filename, mode='r', compresslevel=9, encoding=None, errors=None, newline=None)`

以二进制或文本模式打开 `bzip2` 压缩文件, 返回一个 *file object*。

和 `BZ2File` 的构造函数类似, `filename` 参数可以是一个实际的文件名 (*str* 或 *bytes* 对象), 或是已有的可供读取或写入的文件对象。

`mode` 参数可设为二进制模式的 `'r'`、`'rb'`、`'w'`、`'wb'`、`'x'`、`'xb'`、`'a'` 或 `'ab'`, 或者文本模式的 `'rt'`、`'wt'`、`'xt'` 或 `'at'`。默认是 `'rb'`。

`compresslevel` 参数是 1 到 9 的整数, 和 `BZ2File` 的构造函数一样。

For binary mode, this function is equivalent to the `BZ2File` constructor: `BZ2File(filename, mode, compresslevel=compresslevel)`. In this case, the `encoding`, `errors` and `newline` arguments must not be provided.

For text mode, a `BZ2File` object is created, and wrapped in an `io.TextIOWrapper` instance with the specified encoding, error handling behavior, and line ending(s).

3.3 新版功能.

在 3.4 版更改: 添加了 `'x'` (仅创建) 模式。

在 3.6 版更改: 接受一个类路径对象。

class `bz2.BZ2File(filename, mode='r', buffering=None, compresslevel=9)`

用二进制模式打开 `bzip2` 压缩文件。

If `filename` is a *str* or *bytes* object, open the named file directly. Otherwise, `filename` should be a *file object*, which will be used to read or write the compressed data.

The `mode` argument can be either `'r'` for reading (default), `'w'` for overwriting, `'x'` for exclusive creation, or `'a'` for appending. These can equivalently be given as `'rb'`, `'wb'`, `'xb'` and `'ab'` respectively.

If *filename* is a file object (rather than an actual file name), a mode of 'w' does not truncate the file, and is instead equivalent to 'a'.

The *buffering* argument is ignored. Its use is deprecated.

If *mode* is 'w' or 'a', *compresslevel* can be a number between 1 and 9 specifying the level of compression: 1 produces the least compression, and 9 (default) produces the most compression.

If *mode* is 'r', the input file may be the concatenation of multiple compressed streams.

BZ2File provides all of the members specified by the *io.BufferedIOBase*, except for *detach()* and *truncate()*. Iteration and the *with* statement are supported.

BZ2File also provides the following method:

peek (*[n]*)

Return buffered data without advancing the file position. At least one byte of data will be returned (unless at EOF). The exact number of bytes returned is unspecified.

注解: While calling *peek()* does not change the file position of the *BZ2File*, it may change the position of the underlying file object (e.g. if the *BZ2File* was constructed by passing a file object for *filename*).

3.3 新版功能.

在 3.1 版更改: 支持了 *with* 语句。

在 3.3 版更改: The *fileno()*, *readable()*, *seekable()*, *writable()*, *read1()* and *readinto()* methods were added.

在 3.3 版更改: Support was added for *filename* being a *file object* instead of an actual filename.

在 3.3 版更改: The 'a' (append) mode was added, along with support for reading multi-stream files.

在 3.4 版更改: 添加了 'x' (仅创建) 模式。

在 3.5 版更改: The *read()* method now accepts an argument of *None*.

在 3.6 版更改: 接受一个类路径对象。

13.3.2 增量压缩和解压

class *bz2.BZ2Compressor* (*compresslevel=9*)

Create a new compressor object. This object may be used to compress data incrementally. For one-shot compression, use the *compress()* function instead.

compresslevel, if given, must be a number between 1 and 9. The default is 9.

compress (*data*)

Provide data to the compressor object. Returns a chunk of compressed data if possible, or an empty byte string otherwise.

When you have finished providing data to the compressor, call the *flush()* method to finish the compression process.

flush ()

结束压缩进程, 返回内部缓冲中剩余的压缩完成的数据。

The compressor object may not be used after this method has been called.

class bz2.BZ2Decompressor

Create a new decompressor object. This object may be used to decompress data incrementally. For one-shot compression, use the `decompress()` function instead.

注解: This class does not transparently handle inputs containing multiple compressed streams, unlike `decompress()` and `BZ2File`. If you need to decompress a multi-stream input with `BZ2Decompressor`, you must use a new decompressor for each stream.

decompress (*data*, *max_length*=-1)

Decompress *data* (a *bytes-like object*), returning uncompressed data as bytes. Some of *data* may be buffered internally, for use in later calls to `decompress()`. The returned data should be concatenated with the output of any previous calls to `decompress()`.

If *max_length* is nonnegative, returns at most *max_length* bytes of decompressed data. If this limit is reached and further output can be produced, the `needs_input` attribute will be set to `False`. In this case, the next call to `decompress()` may provide *data* as `b''` to obtain more of the output.

If all of the input data was decompressed and returned (either because this was less than *max_length* bytes, or because *max_length* was negative), the `needs_input` attribute will be set to `True`.

Attempting to decompress data after the end of stream is reached raises an `EOFError`. Any data found after the end of the stream is ignored and saved in the `unused_data` attribute.

在 3.5 版更改: Added the *max_length* parameter.

eof

若达到了数据流末尾标识符则为 `True`。

3.3 新版功能。

unused_data

压缩数据流的末尾还有数据。

If this attribute is accessed before the end of the stream has been reached, its value will be `b''`.

needs_input

`False` if the `decompress()` method can provide more decompressed data before requiring new uncompressed input.

3.5 新版功能。

13.3.3 一次性压缩或解压

bz2.compress (*data*, *compresslevel*=9)

Compress *data*.

compresslevel, if given, must be a number between 1 and 9. The default is 9.

For incremental compression, use a `BZ2Compressor` instead.

bz2.decompress (*data*)

Decompress *data*.

If *data* is the concatenation of multiple compressed streams, decompress all of the streams.

For incremental decompression, use a `BZ2Decompressor` instead.

在 3.3 版更改: 支持了多数据流的输入。

13.4 lzma — 用 LZMA 算法压缩

3.3 新版功能.

源代码: [Lib/lzma.py](#)

This module provides classes and convenience functions for compressing and decompressing data using the LZMA compression algorithm. Also included is a file interface supporting the `.xz` and legacy `.lzma` file formats used by the `xz` utility, as well as raw compressed streams.

The interface provided by this module is very similar to that of the `bz2` module. However, note that `LZMAFile` is *not* thread-safe, unlike `bz2.BZ2File`, so if you need to use a single `LZMAFile` instance from multiple threads, it is necessary to protect it with a lock.

exception `lzma.LZMAError`

This exception is raised when an error occurs during compression or decompression, or while initializing the compressor/decompressor state.

13.4.1 Reading and writing compressed files

`lzma.open(filename, mode="rb", *, format=None, check=-1, preset=None, filters=None, encoding=None, errors=None, newline=None)`

Open an LZMA-compressed file in binary or text mode, returning a *file object*.

The *filename* argument can be either an actual file name (given as a *str*, *bytes* or *path-like* object), in which case the named file is opened, or it can be an existing file object to read from or write to.

The *mode* argument can be any of "r", "rb", "w", "wb", "x", "xb", "a" or "ab" for binary mode, or "rt", "wt", "xt", or "at" for text mode. The default is "rb".

When opening a file for reading, the *format* and *filters* arguments have the same meanings as for `LZMADecompressor`. In this case, the *check* and *preset* arguments should not be used.

When opening a file for writing, the *format*, *check*, *preset* and *filters* arguments have the same meanings as for `LZMACompressor`.

For binary mode, this function is equivalent to the `LZMAFile` constructor: `LZMAFile(filename, mode, ...)`. In this case, the *encoding*, *errors* and *newline* arguments must not be provided.

For text mode, a `LZMAFile` object is created, and wrapped in an `io.TextIOWrapper` instance with the specified encoding, error handling behavior, and line ending(s).

在 3.4 版更改: Added support for the "x", "xb" and "xt" modes.

在 3.6 版更改: 接受一个类路径对象。

class `lzma.LZMAFile(filename=None, mode="r", *, format=None, check=-1, preset=None, filters=None)`

Open an LZMA-compressed file in binary mode.

An `LZMAFile` can wrap an already-open *file object*, or operate directly on a named file. The *filename* argument specifies either the file object to wrap, or the name of the file to open (as a *str*, *bytes* or *path-like* object). When wrapping an existing file object, the wrapped file will not be closed when the `LZMAFile` is closed.

The *mode* argument can be either "r" for reading (default), "w" for overwriting, "x" for exclusive creation, or "a" for appending. These can equivalently be given as "rb", "wb", "xb" and "ab" respectively.

If *filename* is a file object (rather than an actual file name), a mode of "w" does not truncate the file, and is instead equivalent to "a".

When opening a file for reading, the input file may be the concatenation of multiple separate compressed streams. These are transparently decoded as a single logical stream.

When opening a file for reading, the *format* and *filters* arguments have the same meanings as for *LZMADecompressor*. In this case, the *check* and *preset* arguments should not be used.

When opening a file for writing, the *format*, *check*, *preset* and *filters* arguments have the same meanings as for *LZMACompressor*.

LZMAFile supports all the members specified by *io.BufferedIOBase*, except for *detach()* and *truncate()*. Iteration and the *with* statement are supported.

The following method is also provided:

peek (*size=-1*)

Return buffered data without advancing the file position. At least one byte of data will be returned, unless EOF has been reached. The exact number of bytes returned is unspecified (the *size* argument is ignored).

注解: While calling *peek()* does not change the file position of the *LZMAFile*, it may change the position of the underlying file object (e.g. if the *LZMAFile* was constructed by passing a file object for *filename*).

在 3.4 版更改: Added support for the "x" and "xb" modes.

在 3.5 版更改: The *read()* method now accepts an argument of *None*.

在 3.6 版更改: 接受一个类路径对象。

13.4.2 Compressing and decompressing data in memory

class *lzma.LZMACompressor* (*format=FORMAT_XZ, check=-1, preset=None, filters=None*)

Create a compressor object, which can be used to compress data incrementally.

For a more convenient way of compressing a single chunk of data, see *compress()*.

The *format* argument specifies what container format should be used. Possible values are:

- **FORMAT_XZ: The .xz container format.** This is the default format.
- **FORMAT_ALONE: The legacy .lzma container format.** This format is more limited than .xz – it does not support integrity checks or multiple filters.
- **FORMAT_RAW: A raw data stream, not using any container format.** This format specifier does not support integrity checks, and requires that you always specify a custom filter chain (for both compression and decompression). Additionally, data compressed in this manner cannot be decompressed using *FORMAT_AUTO* (see *LZMADecompressor*).

The *check* argument specifies the type of integrity check to include in the compressed data. This check is used when decompressing, to ensure that the data has not been corrupted. Possible values are:

- **CHECK_NONE: No integrity check.** This is the default (and the only acceptable value) for *FORMAT_ALONE* and *FORMAT_RAW*.
- **CHECK_CRC32: 32-bit Cyclic Redundancy Check.**
- **CHECK_CRC64: 64-bit Cyclic Redundancy Check.** This is the default for *FORMAT_XZ*.
- **CHECK_SHA256: 256-bit Secure Hash Algorithm.**

If the specified check is not supported, an *LZMAError* is raised.

The compression settings can be specified either as a preset compression level (with the *preset* argument), or in detail as a custom filter chain (with the *filters* argument).

The *preset* argument (if provided) should be an integer between 0 and 9 (inclusive), optionally OR-ed with the constant `PRESET_EXTREME`. If neither *preset* nor *filters* are given, the default behavior is to use `PRESET_DEFAULT` (preset level 6). Higher presets produce smaller output, but make the compression process slower.

注解: In addition to being more CPU-intensive, compression with higher presets also requires much more memory (and produces output that needs more memory to decompress). With preset 9 for example, the overhead for an `LZMACompressor` object can be as high as 800 MiB. For this reason, it is generally best to stick with the default preset.

The *filters* argument (if provided) should be a filter chain specifier. See *Specifying custom filter chains* for details.

compress (*data*)

Compress *data* (a *bytes* object), returning a *bytes* object containing compressed data for at least part of the input. Some of *data* may be buffered internally, for use in later calls to `compress()` and `flush()`. The returned data should be concatenated with the output of any previous calls to `compress()`.

flush ()

Finish the compression process, returning a *bytes* object containing any data stored in the compressor's internal buffers.

The compressor cannot be used after this method has been called.

class `lzma.LZMADecompressor` (*format*=`FORMAT_AUTO`, *memlimit*=`None`, *filters*=`None`)

Create a decompressor object, which can be used to decompress data incrementally.

For a more convenient way of decompressing an entire compressed stream at once, see `decompress()`.

The *format* argument specifies the container format that should be used. The default is `FORMAT_AUTO`, which can decompress both `.xz` and `.lzma` files. Other possible values are `FORMAT_XZ`, `FORMAT_ALONE`, and `FORMAT_RAW`.

The *memlimit* argument specifies a limit (in bytes) on the amount of memory that the decompressor can use. When this argument is used, decompression will fail with an `LZMAError` if it is not possible to decompress the input within the given memory limit.

The *filters* argument specifies the filter chain that was used to create the stream being decompressed. This argument is required if *format* is `FORMAT_RAW`, but should not be used for other formats. See *Specifying custom filter chains* for more information about filter chains.

注解: This class does not transparently handle inputs containing multiple compressed streams, unlike `decompress()` and `LZMAFile`. To decompress a multi-stream input with `LZMADecompressor`, you must create a new decompressor for each stream.

decompress (*data*, *max_length*=-1)

Decompress *data* (a *bytes-like object*), returning uncompressed data as bytes. Some of *data* may be buffered internally, for use in later calls to `decompress()`. The returned data should be concatenated with the output of any previous calls to `decompress()`.

If *max_length* is nonnegative, returns at most *max_length* bytes of decompressed data. If this limit is reached and further output can be produced, the *needs_input* attribute will be set to `False`. In this case, the next call to `decompress()` may provide *data* as `b''` to obtain more of the output.

If all of the input data was decompressed and returned (either because this was less than *max_length* bytes, or because *max_length* was negative), the *needs_input* attribute will be set to `True`.

Attempting to decompress data after the end of stream is reached raises an *EOFError*. Any data found after the end of the stream is ignored and saved in the *unused_data* attribute.

在 3.5 版更改: Added the *max_length* parameter.

check

The ID of the integrity check used by the input stream. This may be *CHECK_UNKNOWN* until enough of the input has been decoded to determine what integrity check it uses.

eof

若达到了数据流末尾标识符则为 *True*。

unused_data

压缩数据流的末尾还有数据。

Before the end of the stream is reached, this will be *b""*.

needs_input

False if the *decompress()* method can provide more decompressed data before requiring new uncompressed input.

3.5 新版功能.

`lzma.compress(data, format=FORMAT_XZ, check=-1, preset=None, filters=None)`

Compress *data* (a *bytes* object), returning the compressed data as a *bytes* object.

See *LZMACompressor* above for a description of the *format*, *check*, *preset* and *filters* arguments.

`lzma.decompress(data, format=FORMAT_AUTO, memlimit=None, filters=None)`

Decompress *data* (a *bytes* object), returning the uncompressed data as a *bytes* object.

If *data* is the concatenation of multiple distinct compressed streams, decompress all of these streams, and return the concatenation of the results.

See *LZMADecompressor* above for a description of the *format*, *memlimit* and *filters* arguments.

13.4.3 杂项

`lzma.is_check_supported(check)`

Returns true if the given integrity check is supported on this system.

CHECK_NONE and *CHECK_CRC32* are always supported. *CHECK_CRC64* and *CHECK_SHA256* may be unavailable if you are using a version of **liblzma** that was compiled with a limited feature set.

13.4.4 Specifying custom filter chains

A filter chain specifier is a sequence of dictionaries, where each dictionary contains the ID and options for a single filter. Each dictionary must contain the key *"id"*, and may contain additional keys to specify filter-dependent options. Valid filter IDs are as follows:

- **Compression filters:**
 - *FILTER_LZMA1* (for use with *FORMAT_ALONE*)
 - *FILTER_LZMA2* (for use with *FORMAT_XZ* and *FORMAT_RAW*)
- **Delta filter:**
 - *FILTER_DELTA*
- **Branch-Call-Jump (BCJ) filters:**

- FILTER_X86
- FILTER_IA64
- FILTER_ARM
- FILTER_ARMTHUMB
- FILTER_POWERPC
- FILTER_SPARC

A filter chain can consist of up to 4 filters, and cannot be empty. The last filter in the chain must be a compression filter, and any other filters must be delta or BCJ filters.

Compression filters support the following options (specified as additional entries in the dictionary representing the filter):

- `preset`: A compression preset to use as a source of default values for options that are not specified explicitly.
- `dict_size`: Dictionary size in bytes. This should be between 4 KiB and 1.5 GiB (inclusive).
- `lc`: Number of literal context bits.
- `lp`: Number of literal position bits. The sum `lc + lp` must be at most 4.
- `pb`: Number of position bits; must be at most 4.
- `mode`: `MODE_FAST` or `MODE_NORMAL`.
- `nice_len`: What should be considered a “nice length” for a match. This should be 273 or less.
- `mf`: What match finder to use – `MF_HC3`, `MF_HC4`, `MF_BT2`, `MF_BT3`, or `MF_BT4`.
- `depth`: Maximum search depth used by match finder. 0 (default) means to select automatically based on other filter options.

The delta filter stores the differences between bytes, producing more repetitive input for the compressor in certain circumstances. It supports one option, `dist`. This indicates the distance between bytes to be subtracted. The default is 1, i.e. take the differences between adjacent bytes.

The BCJ filters are intended to be applied to machine code. They convert relative branches, calls and jumps in the code to use absolute addressing, with the aim of increasing the redundancy that can be exploited by the compressor. These filters support one option, `start_offset`. This specifies the address that should be mapped to the beginning of the input data. The default is 0.

13.4.5 示例

Reading in a compressed file:

```
import lzma
with lzma.open("file.xz") as f:
    file_content = f.read()
```

Creating a compressed file:

```
import lzma
data = b"Insert Data Here"
with lzma.open("file.xz", "w") as f:
    f.write(data)
```

Compressing data in memory:

```
import lzma
data_in = b"Insert Data Here"
data_out = lzma.compress(data_in)
```

Incremental compression:

```
import lzma
lzc = lzma.LZMACompressor()
out1 = lzc.compress(b"Some data\n")
out2 = lzc.compress(b"Another piece of data\n")
out3 = lzc.compress(b"Even more data\n")
out4 = lzc.flush()
# Concatenate all the partial results:
result = b"".join([out1, out2, out3, out4])
```

Writing compressed data to an already-open file:

```
import lzma
with open("file.xz", "wb") as f:
    f.write(b"This data will not be compressed\n")
    with lzma.open(f, "w") as lzf:
        lzf.write(b"This *will* be compressed\n")
    f.write(b"Not compressed\n")
```

Creating a compressed file using a custom filter chain:

```
import lzma
my_filters = [
    {"id": lzma.FILTER_DELTA, "dist": 5},
    {"id": lzma.FILTER_LZMA2, "preset": 7 | lzma.PRESET_EXTREME},
]
with lzma.open("file.xz", "w", filters=my_filters) as f:
    f.write(b"blah blah blah")
```

13.5 zipfile — 使用 ZIP 存档

源代码: [Lib/zipfile.py](#)

ZIP 文件格式是一个常用的归档与压缩标准。这个模块提供了创建、读取、写入、添加及列出 ZIP 文件的工具。任何对此模块的进阶使用都将需要理解此格式，其定义参见 [PKZIP 应用程序笔记](#)。

此模块目前不能处理分卷 ZIP 文件。它可以处理使用 ZIP64 扩展（超过 4 GB 的 ZIP 文件）的 ZIP 文件。它支持解密 ZIP 归档中的加密文件，但是目前不能创建一个加密的文件。解密非常慢，因为它是使用原生 Python 而不是 C 实现的。

这个模块定义了以下内容：

exception `zipfile.BadZipFile`

为损坏的 ZIP 文件抛出的错误。

3.2 新版功能。

exception `zipfile.BadZipfile`

`BadZipFile` 的别名，与旧版本 Python 保持兼容性。

3.2 版后已移除。

exception zipfile.LargeZipFile

当 ZIP 文件取药 ZIP64 功能但是未启用时抛出此错误。

class zipfile.ZipFile

用于读写 ZIP 文件的类。欲了解构造函数的描述，参阅段落 [ZipFile 对象](#)。

class zipfile.PyZipFile

用于创建包含 Python 库的 ZIP 归档的类。

class zipfile.ZipInfo (filename='NoName', date_time=(1980, 1, 1, 0, 0, 0))

用于表示档案内一个成员信息的类。此类的实例会由 [ZipFile](#) 对象的 [getinfo\(\)](#) 和 [infolist\(\)](#) 方法返回。大多数 [zipfile](#) 模块的用户都不必创建它们，只需使用此模块所创建的实例。*filename* 应当是档案成员的全名，*date_time* 应当是包含六个字段的描述最近修改时间的元组；这些字段的描述请参阅 [ZipInfo Objects](#)。

zipfile.is_zipfile (filename)

根据文件的 Magic Number，如果 *filename* 是一个有效的 ZIP 文件则返回 True，否则返回 False。*filename* 也可能是一个文件或类文件对象。

在 3.1 版更改：支持文件或类文件对象。

zipfile.ZIP_STORED

未被压缩的归档成员的数字常数。

zipfile.ZIP_DEFLATED

常用的 ZIP 压缩方法的数字常数。需要 [zlib](#) 模块。

zipfile.ZIP_BZIP2

BZIP2 压缩方法的数字常数。需要 [bz2](#) 模块。

3.3 新版功能。

zipfile.ZIP_LZMA

LZMA 压缩方法的数字常数。需要 [lzma](#) 模块。

3.3 新版功能。

注解： ZIP 文件格式规范包括自 2001 年以来对 bzip2 压缩的支持，以及自 2006 年以来对 LZMA 压缩的支持。但是，一些工具（包括较旧的 Python 版本）不支持这些压缩方法，并且可能拒绝完全处理 ZIP 文件，或者无法提取单个文件。

参见：

PKZIP 应用程序笔记 Phil Katz 编写的 ZIP 文件格式文档，此格式和使用的算法的创建者。

Info-ZIP 主页 有关 Info-ZIP 项目的 ZIP 存档程序和开发库的信息。

13.5.1 ZipFile 对象

class zipfile.ZipFile (file, mode='r', compression=ZIP_STORED, allowZip64=True)

Open a ZIP file, where *file* can be a path to a file (a string), a file-like object or a [path-like object](#). The *mode* parameter should be 'r' to read an existing file, 'w' to truncate and write a new file, 'a' to append to an existing file, or 'x' to exclusively create and write a new file. If *mode* is 'x' and *file* refers to an existing file, a [FileExistsError](#) will be raised. If *mode* is 'a' and *file* refers to an existing ZIP file, then additional files are added to it. If *file* does not refer to a ZIP file, then a new ZIP archive is appended to the file. This is meant for adding a ZIP archive to another file (such as `python.exe`). If *mode* is 'a' and the file does not exist at all, it is created. If *mode* is 'r' or 'a', the file should be seekable. *compression* is the ZIP compression method to use when writing the archive, and should be [ZIP_STORED](#), [ZIP_DEFLATED](#), [ZIP_BZIP2](#) or [ZIP_LZMA](#);

unrecognized values will cause `NotImplementedError` to be raised. If `ZIP_DEFLATED`, `ZIP_BZIP2` or `ZIP_LZMA` is specified but the corresponding module (`zlib`, `bz2` or `lzma`) is not available, `RuntimeError` is raised. The default is `ZIP_STORED`. If `allowZip64` is `True` (the default) zipfile will create ZIP files that use the ZIP64 extensions when the zipfile is larger than 4 GiB. If it is false `zipfile` will raise an exception when the ZIP file would require ZIP64 extensions.

If the file is created with mode `'w'`, `'x'` or `'a'` and then `closed` without adding any files to the archive, the appropriate ZIP structures for an empty archive will be written to the file.

`ZipFile` is also a context manager and therefore supports the `with` statement. In the example, `myzip` is closed after the `with` statement's suite is finished—even if an exception occurs:

```
with ZipFile('spam.zip', 'w') as myzip:
    myzip.write('eggs.txt')
```

3.2 新版功能: Added the ability to use `ZipFile` as a context manager.

在 3.3 版更改: Added support for `bzip2` and `lzma` compression.

在 3.4 版更改: ZIP64 extensions are enabled by default.

在 3.5 版更改: Added support for writing to unseekable streams. Added support for the `'x'` mode.

在 3.6 版更改: Previously, a plain `RuntimeError` was raised for unrecognized compression values.

在 3.6.2 版更改: The `file` parameter accepts a *path-like object*.

`ZipFile.close()`

Close the archive file. You must call `close()` before exiting your program or essential records will not be written.

`ZipFile.getinfo(name)`

Return a `ZipInfo` object with information about the archive member `name`. Calling `getinfo()` for a name not currently contained in the archive will raise a `KeyError`.

`ZipFile.infolist()`

Return a list containing a `ZipInfo` object for each member of the archive. The objects are in the same order as their entries in the actual ZIP file on disk if an existing archive was opened.

`ZipFile.namelist()`

Return a list of archive members by name.

`ZipFile.open(name, mode='r', pwd=None, *, force_zip64=False)`

Access a member of the archive as a binary file-like object. `name` can be either the name of a file within the archive or a `ZipInfo` object. The `mode` parameter, if included, must be `'r'` (the default) or `'w'`. `pwd` is the password used to decrypt encrypted ZIP files.

`open()` is also a context manager and therefore supports the `with` statement:

```
with ZipFile('spam.zip') as myzip:
    with myzip.open('eggs.txt') as myfile:
        print(myfile.read())
```

With mode `'r'` the file-like object (`ZipExtFile`) is read-only and provides the following methods: `read()`, `readline()`, `readlines()`, `__iter__()`, `__next__()`. These objects can operate independently of the `ZipFile`.

With mode `'w'`, a writable file handle is returned, which supports the `write()` method. While a writable file handle is open, attempting to read or write other files in the ZIP file will raise a `ValueError`.

When writing a file, if the file size is not known in advance but may exceed 2 GiB, pass `force_zip64=True` to ensure that the header format is capable of supporting large files. If the file size is known in advance, construct a `ZipInfo` object with `file_size` set, and use that as the `name` parameter.

注解: The `open()`, `read()` and `extract()` methods can take a filename or a `ZipInfo` object. You will appreciate this when trying to read a ZIP file that contains members with duplicate names.

在 3.6 版更改: Removed support of `mode='U'`. Use `io.TextIOWrapper` for reading compressed text files in *universal newlines* mode.

在 3.6 版更改: `open()` can now be used to write files into the archive with the `mode='w'` option.

在 3.6 版更改: Calling `open()` on a closed `ZipFile` will raise a `ValueError`. Previously, a `RuntimeError` was raised.

`ZipFile.extract(member, path=None, pwd=None)`

Extract a member from the archive to the current working directory; *member* must be its full name or a `ZipInfo` object. Its file information is extracted as accurately as possible. *path* specifies a different directory to extract to. *member* can be a filename or a `ZipInfo` object. *pwd* is the password used for encrypted files.

Returns the normalized path created (a directory or new file).

注解: If a member filename is an absolute path, a drive/UNC sharepoint and leading (back)slashes will be stripped, e.g.: `///foo/bar` becomes `foo/bar` on Unix, and `C:\foo\bar` becomes `foo\bar` on Windows. And all `".."` components in a member filename will be removed, e.g.: `../..foo../..ba..r` becomes `foo../ba..r`. On Windows illegal characters (`:`, `<`, `>`, `|`, `"`, `?`, and `*`) replaced by underscore (`_`).

在 3.6 版更改: Calling `extract()` on a closed `ZipFile` will raise a `ValueError`. Previously, a `RuntimeError` was raised.

在 3.6.2 版更改: The *path* parameter accepts a *path-like object*.

`ZipFile.extractall(path=None, members=None, pwd=None)`

Extract all members from the archive to the current working directory. *path* specifies a different directory to extract to. *members* is optional and must be a subset of the list returned by `namelist()`. *pwd* is the password used for encrypted files.

警告: Never extract archives from untrusted sources without prior inspection. It is possible that files are created outside of *path*, e.g. members that have absolute filenames starting with `"/"` or filenames with two dots `".."`. This module attempts to prevent that. See `extract()` note.

在 3.6 版更改: Calling `extractall()` on a closed `ZipFile` will raise a `ValueError`. Previously, a `RuntimeError` was raised.

在 3.6.2 版更改: The *path* parameter accepts a *path-like object*.

`ZipFile.printdir()`

Print a table of contents for the archive to `sys.stdout`.

`ZipFile.setpassword(pwd)`

Set *pwd* as default password to extract encrypted files.

`ZipFile.read(name, pwd=None)`

Return the bytes of the file *name* in the archive. *name* is the name of the file in the archive, or a `ZipInfo` object. The archive must be open for read or append. *pwd* is the password used for encrypted files and, if specified, it will override the default password set with `setpassword()`. Calling `read()` on a `ZipFile` that uses a compression method other than `ZIP_STORED`, `ZIP_DEFLATED`, `ZIP_BZIP2` or `ZIP_LZMA` will raise a `NotImplementedError`. An error will also be raised if the corresponding compression module is not available.

在 3.6 版更改: Calling `read()` on a closed `ZipFile` will raise a `ValueError`. Previously, a `RuntimeError` was raised.

`ZipFile.testzip()`

Read all the files in the archive and check their CRC's and file headers. Return the name of the first bad file, or else return `None`.

在 3.6 版更改: Calling `testzip()` on a closed `ZipFile` will raise a `ValueError`. Previously, a `RuntimeError` was raised.

`ZipFile.write(filename, arcname=None, compress_type=None)`

Write the file named `filename` to the archive, giving it the archive name `arcname` (by default, this will be the same as `filename`, but without a drive letter and with leading path separators removed). If given, `compress_type` overrides the value given for the `compression` parameter to the constructor for the new entry. The archive must be open with mode `'w'`, `'x'` or `'a'`.

注解: Archive names should be relative to the archive root, that is, they should not start with a path separator.

注解: If `arcname` (or `filename`, if `arcname` is not given) contains a null byte, the name of the file in the archive will be truncated at the null byte.

在 3.6 版更改: Calling `write()` on a `ZipFile` created with mode `'r'` or a closed `ZipFile` will raise a `ValueError`. Previously, a `RuntimeError` was raised.

`ZipFile.writestr(zinfo_or_arcname, data[, compress_type])`

Write a file into the archive. The contents is `data`, which may be either a `str` or a `bytes` instance; if it is a `str`, it is encoded as UTF-8 first. `zinfo_or_arcname` is either the file name it will be given in the archive, or a `ZipInfo` instance. If it's an instance, at least the filename, date, and time must be given. If it's a name, the date and time is set to the current date and time. The archive must be opened with mode `'w'`, `'x'` or `'a'`.

If given, `compress_type` overrides the value given for the `compression` parameter to the constructor for the new entry, or in the `zinfo_or_arcname` (if that is a `ZipInfo` instance).

注解: When passing a `ZipInfo` instance as the `zinfo_or_arcname` parameter, the compression method used will be that specified in the `compress_type` member of the given `ZipInfo` instance. By default, the `ZipInfo` constructor sets this member to `ZIP_STORED`.

在 3.2 版更改: The `compress_type` argument.

在 3.6 版更改: Calling `writestr()` on a `ZipFile` created with mode `'r'` or a closed `ZipFile` will raise a `ValueError`. Previously, a `RuntimeError` was raised.

The following data attributes are also available:

`ZipFile.filename`

Name of the ZIP file.

`ZipFile.debug`

The level of debug output to use. This may be set from 0 (the default, no output) to 3 (the most output). Debugging information is written to `sys.stdout`.

`ZipFile.comment`

The comment associated with the ZIP file as a `bytes` object. If assigning a comment to a `ZipFile` instance created with mode `'w'`, `'x'` or `'a'`, it should be no longer than 65535 bytes. Comments longer than this will be truncated.

13.5.2 PyZipFile Objects

The `PyZipFile` constructor takes the same parameters as the `ZipFile` constructor, and one additional parameter, `optimize`.

class `zipfile.PyZipFile` (`file`, `mode='r'`, `compression=ZIP_STORED`, `allowZip64=True`, `optimize=-1`)

3.2 新版功能: The `optimize` parameter.

在 3.4 版更改: ZIP64 extensions are enabled by default.

Instances have one method in addition to those of `ZipFile` objects:

writepy (`pathname`, `basename=""`, `filterfunc=None`)

Search for files `*.py` and add the corresponding file to the archive.

If the `optimize` parameter to `PyZipFile` was not given or `-1`, the corresponding file is a `*.pyc` file, compiling if necessary.

If the `optimize` parameter to `PyZipFile` was `0`, `1` or `2`, only files with that optimization level (see `compile()`) are added to the archive, compiling if necessary.

If `pathname` is a file, the filename must end with `.py`, and just the (corresponding `*.pyc`) file is added at the top level (no path information). If `pathname` is a file that does not end with `.py`, a `RuntimeError` will be raised. If it is a directory, and the directory is not a package directory, then all the `*.pyc` are added at the top level. If the directory is a package directory, then all `*.pyc` are added under the package name as a file path, and if any subdirectories are package directories, all of these are added recursively.

`basename` is intended for internal use only.

`filterfunc`, if given, must be a function taking a single string argument. It will be passed each path (including each individual full file path) before it is added to the archive. If `filterfunc` returns a false value, the path will not be added, and if it is a directory its contents will be ignored. For example, if our test files are all either in test directories or start with the string `test_`, we can use a `filterfunc` to exclude them:

```
>>> zf = PyZipFile('myprog.zip')
>>> def notests(s):
...     fn = os.path.basename(s)
...     return (not (fn == 'test' or fn.startswith('test_')))
>>> zf.writepy('myprog', filterfunc=notests)
```

The `writepy()` method makes archives with file names like this:

```
string.pyc           # Top level name
test/__init__.pyc    # Package directory
test/testall.pyc     # Module test.testall
test/bogus/__init__.pyc # Subpackage directory
test/bogus/myfile.pyc # Submodule test.bogus.myfile
```

3.4 新版功能: The `filterfunc` parameter.

在 3.6.2 版更改: The `pathname` parameter accepts a *path-like object*.

13.5.3 ZipInfo Objects

Instances of the `ZipInfo` class are returned by the `getinfo()` and `infolist()` methods of `ZipFile` objects. Each object stores information about a single member of the ZIP archive.

There is one classmethod to make a `ZipInfo` instance for a filesystem file:

classmethod `ZipInfo.from_file(filename, arcname=None)`

Construct a `ZipInfo` instance for a file on the filesystem, in preparation for adding it to a zip file.

filename should be the path to a file or directory on the filesystem.

If *arcname* is specified, it is used as the name within the archive. If *arcname* is not specified, the name will be the same as *filename*, but with any drive letter and leading path separators removed.

3.6 新版功能.

在 3.6.2 版更改: The *filename* parameter accepts a *path-like object*.

Instances have the following methods and attributes:

`ZipInfo.is_dir()`

Return `True` if this archive member is a directory.

This uses the entry's name: directories should always end with `/`.

3.6 新版功能.

`ZipInfo.filename`

Name of the file in the archive.

`ZipInfo.date_time`

上次修改存档成员的时间和日期。这是六个值的元组:

索引	值
0	Year (≥ 1980)
1	月 (1 为基数)
2	月份中的日期 (1 为基数)
3	小时 (0 为基数)
4	分钟 (0 为基数)
5	秒 (0 为基数)

注解: ZIP 文件格式不支持 1980 年以前的时间戳。

`ZipInfo.compress_type`

Type of compression for the archive member.

`ZipInfo.comment`

Comment for the individual archive member as a *bytes* object.

`ZipInfo.extra`

Expansion field data. The [PKZIP Application Note](#) contains some comments on the internal structure of the data contained in this *bytes* object.

`ZipInfo.create_system`

System which created ZIP archive.

`ZipInfo.create_version`

PKZIP version which created ZIP archive.

`ZipInfo.extract_version`

PKZIP version needed to extract archive.

`ZipInfo.reserved`

必须为零。

`ZipInfo.flag_bits`
ZIP 标志位。

`ZipInfo.volume`
Volume number of file header.

`ZipInfo.internal_attr`
Internal attributes.

`ZipInfo.external_attr`
External file attributes.

`ZipInfo.header_offset`
Byte offset to the file header.

`ZipInfo.CRC`
CRC-32 of the uncompressed file.

`ZipInfo.compress_size`
Size of the compressed data.

`ZipInfo.file_size`
Size of the uncompressed file.

13.5.4 命令行界面

The `zipfile` module provides a simple command-line interface to interact with ZIP archives.

If you want to create a new ZIP archive, specify its name after the `-c` option and then list the filename(s) that should be included:

```
$ python -m zipfile -c monty.zip spam.txt eggs.txt
```

Passing a directory is also acceptable:

```
$ python -m zipfile -c monty.zip life-of-brian_1979/
```

If you want to extract a ZIP archive into the specified directory, use the `-e` option:

```
$ python -m zipfile -e monty.zip target-dir/
```

For a list of the files in a ZIP archive, use the `-l` option:

```
$ python -m zipfile -l monty.zip
```

命令行选项

`-l <zipfile>`
List files in a zipfile.

`-c <zipfile> <source1> ... <sourceN>`
Create zipfile from source files.

`-e <zipfile> <output_dir>`
Extract zipfile into target directory.

`-t <zipfile>`
Test whether the zipfile is valid or not.

13.6 tarfile —读写 tar 归档文件

源代码: [Lib/tarfile.py](#)

The `tarfile` module makes it possible to read and write tar archives, including those using `gzip`, `bz2` and `lzma` compression. Use the `zipfile` module to read or write `.zip` files, or the higher-level functions in `shutil`.

Some facts and figures:

- reads and writes `gzip`, `bz2` and `lzma` compressed archives if the respective modules are available.
- read/write support for the POSIX.1-1988 (ustar) format.
- read/write support for the GNU tar format including `longname` and `longlink` extensions, read-only support for all variants of the `sparse` extension including restoration of sparse files.
- read/write support for the POSIX.1-2001 (pax) format.
- handles directories, regular files, hardlinks, symbolic links, fifos, character devices and block devices and is able to acquire and restore file information like timestamp, access permissions and owner.

在 3.3 版更改: Added support for `lzma` compression.

`tarfile.open` (`name=None`, `mode='r'`, `fileobj=None`, `bufsize=10240`, `**kwargs`)

Return a `TarFile` object for the pathname `name`. For detailed information on `TarFile` objects and the keyword arguments that are allowed, see [TarFile Objects](#).

`mode` has to be a string of the form `'filemode[:compression]'`, it defaults to `'r'`. Here is a full list of mode combinations:

模式	动作
'r' or 'r:*	打开和读取使用透明压缩（推荐）。
'r:'	打开和读取不使用压缩。
'r:gz'	打开和读取使用 <code>gzip</code> 压缩。
'r:bz2'	打开和读取使用 <code>bzip2</code> 压缩。
'r:xz'	打开和读取使用 <code>lzma</code> 压缩。
'x' or 'x:'	创建 <code>tarfile</code> 不进行压缩。如果文件已经存在，则抛出 <code>FileExistsError</code> 异常。
'x:gz'	使用 <code>gzip</code> 压缩创建 <code>tarfile</code> 。如果文件已经存在，则抛出 <code>FileExistsError</code> 异常。
'x:bz2'	使用 <code>bzip2</code> 压缩创建 <code>tarfile</code> 。如果文件已经存在，则抛出 <code>FileExistsError</code> 异常。
'x:xz'	使用 <code>lzma</code> 压缩创建 <code>tarfile</code> 。如果文件已经存在，则抛出 <code>FileExistsError</code> 异常。
'a' or 'a:'	打开以便在没有压缩的情况下追加。如果文件不存在，则创建该文件。
'w' or 'w:'	Open for uncompressed writing.
'w:gz'	Open for <code>gzip</code> compressed writing.
'w:bz2'	Open for <code>bzip2</code> compressed writing.
'w:xz'	Open for <code>lzma</code> compressed writing.

Note that `'a:gz'`, `'a:bz2'` or `'a:xz'` is not possible. If `mode` is not suitable to open a certain (compressed) file for reading, `ReadError` is raised. Use `mode 'r'` to avoid this. If a compression method is not supported, `CompressionError` is raised.

If `fileobj` is specified, it is used as an alternative to a `file object` opened in binary mode for `name`. It is supposed to be at position 0.

For modes `'w:gz'`, `'r:gz'`, `'w:bz2'`, `'r:bz2'`, `'x:gz'`, `'x:bz2'`, `tarfile.open()` accepts the keyword argument `compresslevel` (default 9) to specify the compression level of the file.

For special purposes, there is a second format for `mode`: `'filemode|[compression]'`. `tarfile.open()` will return a `TarFile` object that processes its data as a stream of blocks. No random seeking will be done on the file. If given, `fileobj` may be any object that has a `read()` or `write()` method (depending on the `mode`). `bufsize` specifies the blocksize and defaults to `20 * 512` bytes. Use this variant in combination with e.g. `sys.stdin`, a socket *file object* or a tape device. However, such a `TarFile` object is limited in that it does not allow random access, see 示例. The currently possible modes:

模式	动作
<code>'r *'</code>	Open a <i>stream</i> of tar blocks for reading with transparent compression.
<code>'r '</code>	Open a <i>stream</i> of uncompressed tar blocks for reading.
<code>'r gz'</code>	Open a gzip compressed <i>stream</i> for reading.
<code>'r bz2'</code>	Open a bzip2 compressed <i>stream</i> for reading.
<code>'r xz'</code>	Open an lzma compressed <i>stream</i> for reading.
<code>'w '</code>	Open an uncompressed <i>stream</i> for writing.
<code>'w gz'</code>	Open a gzip compressed <i>stream</i> for writing.
<code>'w bz2'</code>	Open a bzip2 compressed <i>stream</i> for writing.
<code>'w xz'</code>	Open an lzma compressed <i>stream</i> for writing.

在 3.5 版更改: 添加了 `'x'` (仅创建) 模式。

在 3.6 版更改: The `name` parameter accepts a *path-like object*.

class `tarfile.TarFile`

Class for reading and writing tar archives. Do not use this class directly: use `tarfile.open()` instead. See *TarFile Objects*.

`tarfile.is_tarfile(name)`

Return `True` if `name` is a tar archive file, that the `tarfile` module can read.

The `tarfile` module defines the following exceptions:

exception `tarfile.TarError`

Base class for all `tarfile` exceptions.

exception `tarfile.ReadError`

Is raised when a tar archive is opened, that either cannot be handled by the `tarfile` module or is somehow invalid.

exception `tarfile.CompressionError`

Is raised when a compression method is not supported or when the data cannot be decoded properly.

exception `tarfile.StreamError`

Is raised for the limitations that are typical for stream-like `TarFile` objects.

exception `tarfile.ExtractError`

Is raised for *non-fatal* errors when using `TarFile.extract()`, but only if `TarFile.errorlevel== 2`.

exception `tarfile.HeaderError`

Is raised by `TarInfo.frombuf()` if the buffer it gets is invalid.

The following constants are available at the module level:

`tarfile.ENCODING`

The default character encoding: `'utf-8'` on Windows, the value returned by `sys.getfilesystemencoding()` otherwise.

Each of the following constants defines a tar archive format that the `tarfile` module is able to create. See section *Supported tar formats* for details.

`tarfile.USTAR_FORMAT`
POSIX.1-1988 (ustar) format.

`tarfile.GNU_FORMAT`
GNU tar format.

`tarfile.PAX_FORMAT`
POSIX.1-2001 (pax) format.

`tarfile.DEFAULT_FORMAT`
The default format for creating archives. This is currently `GNU_FORMAT`.

参见:

Module `zipfile` Documentation of the `zipfile` standard module.

Archiving operations Documentation of the higher-level archiving facilities provided by the standard `shutil` module.

GNU tar manual, Basic Tar Format Documentation for tar archive files, including GNU tar extensions.

13.6.1 TarFile Objects

The `TarFile` object provides an interface to a tar archive. A tar archive is a sequence of blocks. An archive member (a stored file) is made up of a header block followed by data blocks. It is possible to store a file in a tar archive several times. Each archive member is represented by a `TarInfo` object, see *TarInfo Objects* for details.

A `TarFile` object can be used as a context manager in a `with` statement. It will automatically be closed when the block is completed. Please note that in the event of an exception an archive opened for writing will not be finalized; only the internally used file object will be closed. See the 示例 section for a use case.

3.2 新版功能: Added support for the context management protocol.

```
class tarfile.TarFile(name=None, mode='r', fileobj=None, format=DEFAULT_FORMAT, tar-
                      info=TarInfo, dereference=False, ignore_zeros=False, encoding=ENCODING,
                      errors='surrogateescape', pax_headers=None, debug=0, errorlevel=0)
```

All following arguments are optional and can be accessed as instance attributes as well.

`name` is the pathname of the archive. `name` may be a *path-like object*. It can be omitted if `fileobj` is given. In this case, the file object's `name` attribute is used if it exists.

`mode` is either `'r'` to read from an existing archive, `'a'` to append data to an existing file, `'w'` to create a new file overwriting an existing one, or `'x'` to create a new file only if it does not already exist.

If `fileobj` is given, it is used for reading or writing data. If it can be determined, `mode` is overridden by `fileobj`'s `mode`. `fileobj` will be used from position 0.

注解: `fileobj` is not closed, when `TarFile` is closed.

`format` controls the archive format. It must be one of the constants `USTAR_FORMAT`, `GNU_FORMAT` or `PAX_FORMAT` that are defined at module level.

The `tarinfo` argument can be used to replace the default `TarInfo` class with a different one.

If `dereference` is `False`, add symbolic and hard links to the archive. If it is `True`, add the content of the target files to the archive. This has no effect on systems that do not support symbolic links.

If `ignore_zeros` is `False`, treat an empty block as the end of the archive. If it is `True`, skip empty (and invalid) blocks and try to get as many members as possible. This is only useful for reading concatenated or damaged archives.

`debug` can be set from 0 (no debug messages) up to 3 (all debug messages). The messages are written to `sys.stderr`.

If `errorlevel` is 0, all errors are ignored when using `TarFile.extract()`. Nevertheless, they appear as error messages in the debug output, when debugging is enabled. If 1, all *fatal* errors are raised as `OSError` exceptions. If 2, all *non-fatal* errors are raised as `TarError` exceptions as well.

The `encoding` and `errors` arguments define the character encoding to be used for reading or writing the archive and how conversion errors are going to be handled. The default settings will work for most users. See section [Unicode issues](#) for in-depth information.

The `pax_headers` argument is an optional dictionary of strings which will be added as a pax global header if `format` is `PAX_FORMAT`.

在 3.2 版更改: Use 'surrogateescape' as the default for the `errors` argument.

在 3.5 版更改: 添加了 'x' (仅创建) 模式。

在 3.6 版更改: The `name` parameter accepts a *path-like object*.

classmethod `TarFile.open(...)`

Alternative constructor. The `tarfile.open()` function is actually a shortcut to this classmethod.

`TarFile.getmember(name)`

Return a `TarInfo` object for member `name`. If `name` can not be found in the archive, `KeyError` is raised.

注解: If a member occurs more than once in the archive, its last occurrence is assumed to be the most up-to-date version.

`TarFile.getmembers()`

Return the members of the archive as a list of `TarInfo` objects. The list has the same order as the members in the archive.

`TarFile.getnames()`

Return the members as a list of their names. It has the same order as the list returned by `getmembers()`.

`TarFile.list(verbose=True, *, members=None)`

Print a table of contents to `sys.stdout`. If `verbose` is `False`, only the names of the members are printed. If it is `True`, output similar to that of `ls -l` is produced. If optional `members` is given, it must be a subset of the list returned by `getmembers()`.

在 3.5 版更改: Added the `members` parameter.

`TarFile.next()`

Return the next member of the archive as a `TarInfo` object, when `TarFile` is opened for reading. Return `None` if there is no more available.

`TarFile.extractall(path=".", members=None, *, numeric_owner=False)`

Extract all members from the archive to the current working directory or directory `path`. If optional `members` is given, it must be a subset of the list returned by `getmembers()`. Directory information like owner, modification time and permissions are set after all members have been extracted. This is done to work around two problems: A directory's modification time is reset each time a file is created in it. And, if a directory's permissions do not allow writing, extracting files to it will fail.

If `numeric_owner` is `True`, the uid and gid numbers from the tarfile are used to set the owner/group for the extracted files. Otherwise, the named values from the tarfile are used.

警告: Never extract archives from untrusted sources without prior inspection. It is possible that files are created outside of *path*, e.g. members that have absolute filenames starting with "/" or filenames with two dots "...".

在 3.5 版更改: Added the *numeric_owner* parameter.

在 3.6 版更改: The *path* parameter accepts a *path-like object*.

`TarFile.extract(member, path="", set_attrs=True, *, numeric_owner=False)`

Extract a member from the archive to the current working directory, using its full name. Its file information is extracted as accurately as possible. *member* may be a filename or a *TarInfo* object. You can specify a different directory using *path*. *path* may be a *path-like object*. File attributes (owner, mtime, mode) are set unless *set_attrs* is false.

If *numeric_owner* is *True*, the uid and gid numbers from the tarfile are used to set the owner/group for the extracted files. Otherwise, the named values from the tarfile are used.

注解: The *extract()* method does not take care of several extraction issues. In most cases you should consider using the *extractall()* method.

警告: See the warning for *extractall()*.

在 3.2 版更改: Added the *set_attrs* parameter.

在 3.5 版更改: Added the *numeric_owner* parameter.

在 3.6 版更改: The *path* parameter accepts a *path-like object*.

`TarFile.extractfile(member)`

Extract a member from the archive as a file object. *member* may be a filename or a *TarInfo* object. If *member* is a regular file or a link, an *io.BufferedReader* object is returned. Otherwise, *None* is returned.

在 3.3 版更改: Return an *io.BufferedReader* object.

`TarFile.add(name, arcname=None, recursive=True, exclude=None, *, filter=None)`

Add the file *name* to the archive. *name* may be any type of file (directory, fifo, symbolic link, etc.). If given, *arcname* specifies an alternative name for the file in the archive. Directories are added recursively by default. This can be avoided by setting *recursive* to *False*. If *exclude* is given, it must be a function that takes one filename argument and returns a boolean value. Depending on this value the respective file is either excluded (*True*) or added (*False*). If *filter* is specified it must be a keyword argument. It should be a function that takes a *TarInfo* object argument and returns the changed *TarInfo* object. If it instead returns *None* the *TarInfo* object will be excluded from the archive. See 示例 for an example.

在 3.2 版更改: Added the *filter* parameter.

3.2 版后已移除: The *exclude* parameter is deprecated, please use the *filter* parameter instead.

`TarFile.addfile(tarinfo, fileobj=None)`

Add the *TarInfo* object *tarinfo* to the archive. If *fileobj* is given, it should be a *binary file*, and *tarinfo.size* bytes are read from it and added to the archive. You can create *TarInfo* objects directly, or by using *gettinfo()*.

`TarFile.gettarinfo(name=None, arcname=None, fileobj=None)`

Create a *TarInfo* object from the result of *os.stat()* or equivalent on an existing file. The file is either named by *name*, or specified as a *file object* *fileobj* with a file descriptor. *name* may be a *path-like object*. If given, *arcname*

specifies an alternative name for the file in the archive, otherwise, the name is taken from *fileobj*'s *name* attribute, or the *name* argument. The name should be a text string.

You can modify some of the *TarInfo*'s attributes before you add it using *addfile()*. If the file object is not an ordinary file object positioned at the beginning of the file, attributes such as *size* may need modifying. This is the case for objects such as *GzipFile*. The *name* may also be modified, in which case *arcname* could be a dummy string.

在 3.6 版更改: The *name* parameter accepts a *path-like object*.

TarFile.close()

Close the *TarFile*. In write mode, two finishing zero blocks are appended to the archive.

TarFile.pax_headers

A dictionary containing key-value pairs of pax global headers.

13.6.2 TarInfo Objects

A *TarInfo* object represents one member in a *TarFile*. Aside from storing all required attributes of a file (like file type, size, time, permissions, owner etc.), it provides some useful methods to determine its type. It does *not* contain the file's data itself.

TarInfo objects are returned by *TarFile*'s methods *getmember()*, *getmembers()* and *gettarinfo()*.

class tarfile.TarInfo (name='')

Create a *TarInfo* object.

classmethod TarInfo.frombuf (buf, encoding, errors)

Create and return a *TarInfo* object from string buffer *buf*.

Raises *HeaderError* if the buffer is invalid.

classmethod TarInfo.fromtarfile (tarfile)

Read the next member from the *TarFile* object *tarfile* and return it as a *TarInfo* object.

TarInfo.tobuf (format=DEFAULT_FORMAT, encoding=ENCODING, errors='surrogateescape')

Create a string buffer from a *TarInfo* object. For information on the arguments see the constructor of the *TarFile* class.

在 3.2 版更改: Use 'surrogateescape' as the default for the *errors* argument.

A *TarInfo* object has the following public data attributes:

TarInfo.name

Name of the archive member.

TarInfo.size

Size in bytes.

TarInfo.mtime

Time of last modification.

TarInfo.mode

Permission bits.

TarInfo.type

File type. *type* is usually one of these constants: REGTYPE, AREGTYPE, LNKTYPE, SYMTYPE, DIRTYPE, FIFOTYPE, CONTTYPE, CHRTYPE, BLKTYPE, GNUTYPE_SPARSE. To determine the type of a *TarInfo* object more conveniently, use the *is*()* methods below.

TarInfo.linkname

Name of the target file name, which is only present in *TarInfo* objects of type LNKTYPE and SYMTYPE.

`TarInfo.uid`
User ID of the user who originally stored this member.

`TarInfo.gid`
Group ID of the user who originally stored this member.

`TarInfo.uname`
User name.

`TarInfo.gname`
Group name.

`TarInfo.pax_headers`
A dictionary containing key-value pairs of an associated pax extended header.

A `TarInfo` object also provides some convenient query methods:

`TarInfo.isfile()`
Return `True` if the `TarInfo` object is a regular file.

`TarInfo.isreg()`
Same as `isfile()`.

`TarInfo.isdir()`
Return `True` if it is a directory.

`TarInfo.issym()`
Return `True` if it is a symbolic link.

`TarInfo.islnk()`
Return `True` if it is a hard link.

`TarInfo.ischr()`
Return `True` if it is a character device.

`TarInfo.isblk()`
Return `True` if it is a block device.

`TarInfo.isfifo()`
Return `True` if it is a FIFO.

`TarInfo.isdev()`
Return `True` if it is one of character device, block device or FIFO.

13.6.3 命令行界面

3.4 新版功能.

The `tarfile` module provides a simple command-line interface to interact with tar archives.

If you want to create a new tar archive, specify its name after the `-c` option and then list the filename(s) that should be included:

```
$ python -m tarfile -c monty.tar spam.txt eggs.txt
```

Passing a directory is also acceptable:

```
$ python -m tarfile -c monty.tar life-of-brian_1979/
```

If you want to extract a tar archive into the current directory, use the `-e` option:

```
$ python -m tarfile -e monty.tar
```

You can also extract a tar archive into a different directory by passing the directory's name:

```
$ python -m tarfile -e monty.tar other-dir/
```

For a list of the files in a tar archive, use the `-l` option:

```
$ python -m tarfile -l monty.tar
```

命令行选项

```
-l <tarfile>
--list <tarfile>
    List files in a tarfile.

-c <tarfile> <source1> ... <sourceN>
--create <tarfile> <source1> ... <sourceN>
    Create tarfile from source files.

-e <tarfile> [<output_dir>]
--extract <tarfile> [<output_dir>]
    Extract tarfile into the current directory if output_dir is not specified.

-t <tarfile>
--test <tarfile>
    Test whether the tarfile is valid or not.

-v, --verbose
    Verbose output.
```

13.6.4 示例

How to extract an entire tar archive to the current working directory:

```
import tarfile
tar = tarfile.open("sample.tar.gz")
tar.extractall()
tar.close()
```

How to extract a subset of a tar archive with `TarFile.extractall()` using a generator function instead of a list:

```
import os
import tarfile

def py_files(members):
    for tarinfo in members:
        if os.path.splitext(tarinfo.name)[1] == ".py":
            yield tarinfo

tar = tarfile.open("sample.tar.gz")
tar.extractall(members=py_files(tar))
tar.close()
```

How to create an uncompressed tar archive from a list of filenames:

```
import tarfile
tar = tarfile.open("sample.tar", "w")
for name in ["foo", "bar", "quux"]:
    tar.add(name)
tar.close()
```

The same example using the with statement:

```
import tarfile
with tarfile.open("sample.tar", "w") as tar:
    for name in ["foo", "bar", "quux"]:
        tar.add(name)
```

How to read a gzip compressed tar archive and display some member information:

```
import tarfile
tar = tarfile.open("sample.tar.gz", "r:gz")
for tarinfo in tar:
    print(tarinfo.name, "is", tarinfo.size, "bytes in size and is", end="")
    if tarinfo.isreg():
        print("a regular file.")
    elif tarinfo.isdir():
        print("a directory.")
    else:
        print("something else.")
tar.close()
```

How to create an archive and reset the user information using the *filter* parameter in *TarFile.add()*:

```
import tarfile
def reset(tarinfo):
    tarinfo.uid = tarinfo.gid = 0
    tarinfo.uname = tarinfo.gname = "root"
    return tarinfo
tar = tarfile.open("sample.tar.gz", "w:gz")
tar.add("foo", filter=reset)
tar.close()
```

13.6.5 Supported tar formats

There are three tar formats that can be created with the *tarfile* module:

- The POSIX.1-1988 ustar format (*USTAR_FORMAT*). It supports filenames up to a length of at best 256 characters and linknames up to 100 characters. The maximum file size is 8 GiB. This is an old and limited but widely supported format.
- The GNU tar format (*GNU_FORMAT*). It supports long filenames and linknames, files bigger than 8 GiB and sparse files. It is the de facto standard on GNU/Linux systems. *tarfile* fully supports the GNU tar extensions for long names, sparse file support is read-only.
- The POSIX.1-2001 pax format (*PAX_FORMAT*). It is the most flexible format with virtually no limits. It supports long filenames and linknames, large files and stores pathnames in a portable way. However, not all tar implementations today are able to handle pax archives properly.

The *pax* format is an extension to the existing *ustar* format. It uses extra headers for information that cannot be stored otherwise. There are two flavours of pax headers: Extended headers only affect the subsequent file header,

global headers are valid for the complete archive and affect all following files. All the data in a pax header is encoded in *UTF-8* for portability reasons.

There are some more variants of the tar format which can be read, but not created:

- The ancient V7 format. This is the first tar format from Unix Seventh Edition, storing only regular files and directories. Names must not be longer than 100 characters, there is no user/group name information. Some archives have miscalculated header checksums in case of fields with non-ASCII characters.
- The SunOS tar extended format. This format is a variant of the POSIX.1-2001 pax format, but is not compatible.

13.6.6 Unicode issues

The tar format was originally conceived to make backups on tape drives with the main focus on preserving file system information. Nowadays tar archives are commonly used for file distribution and exchanging archives over networks. One problem of the original format (which is the basis of all other formats) is that there is no concept of supporting different character encodings. For example, an ordinary tar archive created on a *UTF-8* system cannot be read correctly on a *Latin-1* system if it contains non-ASCII characters. Textual metadata (like filenames, linknames, user/group names) will appear damaged. Unfortunately, there is no way to autodetect the encoding of an archive. The pax format was designed to solve this problem. It stores non-ASCII metadata using the universal character encoding *UTF-8*.

The details of character conversion in *tarfile* are controlled by the *encoding* and *errors* keyword arguments of the *TarFile* class.

encoding defines the character encoding to use for the metadata in the archive. The default value is *sys.getfilesystemencoding()* or 'ascii' as a fallback. Depending on whether the archive is read or written, the metadata must be either decoded or encoded. If *encoding* is not set appropriately, this conversion may fail.

The *errors* argument defines how characters are treated that cannot be converted. Possible values are listed in section *Error Handlers*. The default scheme is 'surrogateescape' which Python also uses for its file system calls, see 文件名, 命令行参数, 以及环境变量。.

In case of *PAX_FORMAT* archives, *encoding* is generally not needed because all the metadata is stored using *UTF-8*. *encoding* is only used in the rare cases when binary pax headers are decoded or when strings with surrogate characters are stored.

本章中描述的模块解析各种不是标记语言且与电子邮件无关的杂项文件格式。

14.1 csv — CSV 文件读写

源代码： [Lib/csv.py](#)

CSV (Comma Separated Vaules) 格式是电子表格和数据库中最常见的输入、输出文件格式。在 [RFC 4180](#) 规范推出的很多年前，CSV 格式就已经被开始使用了，由于当时并没有合理的标准，不同应用程序读写的数据会存在细微的差别。这种差别让处理多个来源的 CSV 文件变得困难。但尽管分隔符会变化，此类文件的大致格式是相似的，所以编写一个单独的模块以高效处理此类数据，将程序员从读写数据的繁琐细节中解放出来是有可能的。

The `csv` module implements classes to read and write tabular data in CSV format. It allows programmers to say, “write this data in the format preferred by Excel,” or “read data from this file which was generated by Excel,” without knowing the precise details of the CSV format used by Excel. Programmers can also describe the CSV formats understood by other applications or define their own special-purpose CSV formats.

`csv` 模块中的 `reader` 类和 `writer` 类可用于读写序列化的数据。也可使用 `DictReader` 类和 `DictWriter` 类以字典的形式读写数据。

参见：

该实现在“Python 增强提议” - [PEP 305 \(CSV 文件 API\)](#) 中被提出 The Python Enhancement Proposal which proposed this addition to Python.

14.1.1 模块内容

`csv` 模块定义了以下函数：

`csv.reader(csvfile, dialect='excel', **fmtparams)`

Return a reader object which will iterate over lines in the given *csvfile*. *csvfile* can be any object which supports the *iterator* protocol and returns a string each time its `__next__()` method is called — *file objects* and list objects are both suitable. If *csvfile* is a file object, it should be opened with `newline=''`.¹ An optional *dialect* parameter can be given which is used to define a set of parameters specific to a particular CSV dialect. It may be an instance of a subclass of the *Dialect* class or one of the strings returned by the `list_dialects()` function. The other optional *fmtparams* keyword arguments can be given to override individual formatting parameters in the current dialect. For full details about the dialect and formatting parameters, see section *Dialects and Formatting Parameters*.

Each row read from the csv file is returned as a list of strings. No automatic data type conversion is performed unless the `QUOTE_NONNUMERIC` format option is specified (in which case unquoted fields are transformed into floats).

一个简短的用法示例:

```
>>> import csv
>>> with open('eggs.csv', newline='') as csvfile:
...     spamreader = csv.reader(csvfile, delimiter=' ', quotechar='|')
...     for row in spamreader:
...         print(', '.join(row))
Spam, Spam, Spam, Spam, Baked Beans
Spam, Lovely Spam, Wonderful Spam
```

`csv.writer(csvfile, dialect='excel', **fmtparams)`

Return a writer object responsible for converting the user's data into delimited strings on the given file-like object. *csvfile* can be any object with a `write()` method. If *csvfile* is a file object, it should be opened with `newline=''`. An optional *dialect* parameter can be given which is used to define a set of parameters specific to a particular CSV dialect. It may be an instance of a subclass of the *Dialect* class or one of the strings returned by the `list_dialects()` function. The other optional *fmtparams* keyword arguments can be given to override individual formatting parameters in the current dialect. For full details about the dialect and formatting parameters, see section *Dialects and Formatting Parameters*. To make it as easy as possible to interface with modules which implement the DB API, the value *None* is written as the empty string. While this isn't a reversible transformation, it makes it easier to dump SQL NULL data values to CSV files without preprocessing the data returned from a `cursor.fetch*` call. All other non-string data are stringified with `str()` before being written.

一个简短的用法示例:

```
import csv
with open('eggs.csv', 'w', newline='') as csvfile:
    spamwriter = csv.writer(csvfile, delimiter=' ',
                            quotechar='|', quoting=csv.QUOTE_MINIMAL)
    spamwriter.writerow(['Spam'] * 5 + ['Baked Beans'])
    spamwriter.writerow(['Spam', 'Lovely Spam', 'Wonderful Spam'])
```

`csv.register_dialect(name[, dialect[, **fmtparams]])`

Associate *dialect* with *name*. *name* must be a string. The dialect can be specified either by passing a sub-class of *Dialect*, or by *fmtparams* keyword arguments, or both, with keyword arguments overriding parameters of the dialect. For full details about the dialect and formatting parameters, see section *Dialects and Formatting Parameters*.

`csv.unregister_dialect(name)`

Delete the dialect associated with *name* from the dialect registry. An *Error* is raised if *name* is not a registered dialect name.

`csv.get_dialect(name)`

Return the dialect associated with *name*. An *Error* is raised if *name* is not a registered dialect name. This function

¹ If `newline=''` is not specified, newlines embedded inside quoted fields will not be interpreted correctly, and on platforms that use `\r\n` line endings on write an extra `\r` will be added. It should always be safe to specify `newline=''`, since the csv module does its own (*universal*) newline handling.

returns an immutable *Dialect*.

`csv.list_dialects()`
Return the names of all registered dialects.

`csv.field_size_limit([new_limit])`
Returns the current maximum field size allowed by the parser. If *new_limit* is given, this becomes the new limit.

The `csv` module defines the following classes:

class `csv.DictReader` (*f*, *fieldnames*=None, *restkey*=None, *restval*=None, *dialect*='excel', *args, **kwargs)
Create an object that operates like a regular reader but maps the information in each row to an *OrderedDict* whose keys are given by the optional *fieldnames* parameter.

The *fieldnames* parameter is a *sequence*. If *fieldnames* is omitted, the values in the first row of file *f* will be used as the fieldnames. Regardless of how the fieldnames are determined, the ordered dictionary preserves their original ordering.

If a row has more fields than fieldnames, the remaining data is put in a list and stored with the fieldname specified by *restkey* (which defaults to None). If a non-blank row has fewer fields than fieldnames, the missing values are filled-in with None.

All other optional or keyword arguments are passed to the underlying *reader* instance.

在 3.6 版更改: Returned rows are now of type *OrderedDict*.

一个简短的用法示例:

```
>>> import csv
>>> with open('names.csv', newline='') as csvfile:
...     reader = csv.DictReader(csvfile)
...     for row in reader:
...         print(row['first_name'], row['last_name'])
...
Eric Idle
John Cleese

>>> print(row)
OrderedDict([('first_name', 'John'), ('last_name', 'Cleese')])
```

class `csv.DictWriter` (*f*, *fieldnames*, *restval*=", *extrasaction*='raise', *dialect*='excel', *args, **kwargs)
Create an object which operates like a regular writer but maps dictionaries onto output rows. The *fieldnames* parameter is a *sequence* of keys that identify the order in which values in the dictionary passed to the *writerow()* method are written to file *f*. The optional *restval* parameter specifies the value to be written if the dictionary is missing a key in *fieldnames*. If the dictionary passed to the *writerow()* method contains a key not found in *fieldnames*, the optional *extrasaction* parameter indicates what action to take. If it is set to 'raise', the default value, a *ValueError* is raised. If it is set to 'ignore', extra values in the dictionary are ignored. Any other optional or keyword arguments are passed to the underlying *writer* instance.

Note that unlike the *DictReader* class, the *fieldnames* parameter of the *DictWriter* is not optional. Since Python's *dict* objects are not ordered, there is not enough information available to deduce the order in which the row should be written to file *f*.

一个简短的用法示例:

```
import csv

with open('names.csv', 'w', newline='') as csvfile:
    fieldnames = ['first_name', 'last_name']
    writer = csv.DictWriter(csvfile, fieldnames=fieldnames)
```

(下页继续)

(续上页)

```
writer.writeheader()
writer.writerow({'first_name': 'Baked', 'last_name': 'Beans'})
writer.writerow({'first_name': 'Lovely', 'last_name': 'Spam'})
writer.writerow({'first_name': 'Wonderful', 'last_name': 'Spam'})
```

class csv.Dialect

The *Dialect* class is a container class relied on primarily for its attributes, which are used to define the parameters for a specific *reader* or *writer* instance.

class csv.excel

The *excel* class defines the usual properties of an Excel-generated CSV file. It is registered with the dialect name 'excel'.

class csv.excel_tab

The *excel_tab* class defines the usual properties of an Excel-generated TAB-delimited file. It is registered with the dialect name 'excel-tab'.

class csv.unix_dialect

The *unix_dialect* class defines the usual properties of a CSV file generated on UNIX systems, i.e. using '\n' as line terminator and quoting all fields. It is registered with the dialect name 'unix'.

3.2 新版功能.

class csv.Sniffer

The *Sniffer* class is used to deduce the format of a CSV file.

Sniffer 类提供了两个方法:

sniff (*sample*, *delimiters=None*)

Analyze the given *sample* and return a *Dialect* subclass reflecting the parameters found. If the optional *delimiters* parameter is given, it is interpreted as a string containing possible valid delimiter characters.

has_header (*sample*)

Analyze the sample text (presumed to be in CSV format) and return *True* if the first row appears to be a series of column headers.

使用 *Sniffer* 的示例:

```
with open('example.csv', newline='') as csvfile:
    dialect = csv.Sniffer().sniff(csvfile.read(1024))
    csvfile.seek(0)
    reader = csv.reader(csvfile, dialect)
    # ... process CSV file contents here ...
```

The *csv* module defines the following constants:

csv.QUOTE_ALL

Instructs *writer* objects to quote all fields.

csv.QUOTE_MINIMAL

Instructs *writer* objects to only quote those fields which contain special characters such as *delimiter*, *quotechar* or any of the characters in *lineterminator*.

csv.QUOTE_NONNUMERIC

Instructs *writer* objects to quote all non-numeric fields.

Instructs the reader to convert all non-quoted fields to type *float*.

csv.QUOTE_NONE

Instructs *writer* objects to never quote fields. When the current *delimiter* occurs in output data it is preceded by

the current *escapechar* character. If *escapechar* is not set, the writer will raise *Error* if any characters that require escaping are encountered.

Instructs *reader* to perform no special processing of quote characters.

The *csv* module defines the following exception:

exception *csv.Error*

Raised by any of the functions when an error is detected.

14.1.2 Dialects and Formatting Parameters

To make it easier to specify the format of input and output records, specific formatting parameters are grouped together into dialects. A dialect is a subclass of the *Dialect* class having a set of specific methods and a single *validate()* method. When creating *reader* or *writer* objects, the programmer can specify a string or a subclass of the *Dialect* class as the dialect parameter. In addition to, or instead of, the *dialect* parameter, the programmer can also specify individual formatting parameters, which have the same names as the attributes defined below for the *Dialect* class.

Dialects support the following attributes:

Dialect.delimiter

A one-character string used to separate fields. It defaults to *','*.

Dialect.doublequote

Controls how instances of *quotechar* appearing inside a field should themselves be quoted. When *True*, the character is doubled. When *False*, the *escapechar* is used as a prefix to the *quotechar*. It defaults to *True*.

On output, if *doublequote* is *False* and no *escapechar* is set, *Error* is raised if a *quotechar* is found in a field.

Dialect.escapechar

A one-character string used by the writer to escape the *delimiter* if *quoting* is set to *QUOTE_NONE* and the *quotechar* if *doublequote* is *False*. On reading, the *escapechar* removes any special meaning from the following character. It defaults to *None*, which disables escaping.

Dialect.lineterminator

The string used to terminate lines produced by the *writer*. It defaults to *'\\r\\n'*.

注解: The *reader* is hard-coded to recognise either *'\\r'* or *'\\n'* as end-of-line, and ignores *lineterminator*. This behavior may change in the future.

Dialect.quotechar

A one-character string used to quote fields containing special characters, such as the *delimiter* or *quotechar*, or which contain new-line characters. It defaults to *'\"'*.

Dialect.quoting

Controls when quotes should be generated by the writer and recognised by the reader. It can take on any of the *QUOTE_** constants (see section 模块内容) and defaults to *QUOTE_MINIMAL*.

Dialect.skipinitialspace

When *True*, whitespace immediately following the *delimiter* is ignored. The default is *False*.

Dialect.strict

When *True*, raise exception *Error* on bad CSV input. The default is *False*.

14.1.3 Reader Objects

Reader objects (*DictReader* instances and objects returned by the *reader()* function) have the following public methods:

`csvreader.__next__()`

Return the next row of the reader's iterable object as a list (if the object was returned from *reader()*) or a dict (if it is a *DictReader* instance), parsed according to the current dialect. Usually you should call this as `next(reader)`.

Reader objects have the following public attributes:

`csvreader.dialect`

A read-only description of the dialect in use by the parser.

`csvreader.line_num`

The number of lines read from the source iterator. This is not the same as the number of records returned, as records can span multiple lines.

DictReader objects have the following public attribute:

`csvreader.fieldnames`

If not passed as a parameter when creating the object, this attribute is initialized upon first access or when the first record is read from the file.

14.1.4 Writer Objects

Writer objects (*DictWriter* instances and objects returned by the *writer()* function) have the following public methods. A *row* must be an iterable of strings or numbers for *Writer* objects and a dictionary mapping fieldnames to strings or numbers (by passing them through *str()* first) for *DictWriter* objects. Note that complex numbers are written out surrounded by parens. This may cause some problems for other programs which read CSV files (assuming they support complex numbers at all).

`csvwriter.writerow(row)`

Write the *row* parameter to the writer's file object, formatted according to the current dialect.

在 3.5 版更改: Added support of arbitrary iterables.

`csvwriter.writerows(rows)`

Write all elements in *rows* (an iterable of *row* objects as described above) to the writer's file object, formatted according to the current dialect.

Writer objects have the following public attribute:

`csvwriter.dialect`

A read-only description of the dialect in use by the writer.

DictWriter objects have the following public method:

`DictWriter.writeheader()`

Write a row with the field names (as specified in the constructor).

3.2 新版功能.

14.1.5 示例

读取 CSV 文件最简单的一个例子:

```
import csv
with open('some.csv', newline='') as f:
    reader = csv.reader(f)
    for row in reader:
        print(row)
```

Reading a file with an alternate format:

```
import csv
with open('passwd', newline='') as f:
    reader = csv.reader(f, delimiter=':', quoting=csv.QUOTE_NONE)
    for row in reader:
        print(row)
```

The corresponding simplest possible writing example is:

```
import csv
with open('some.csv', 'w', newline='') as f:
    writer = csv.writer(f)
    writer.writerows(someiterable)
```

Since `open()` is used to open a CSV file for reading, the file will by default be decoded into unicode using the system default encoding (see `locale.getpreferredencoding()`). To decode a file using a different encoding, use the encoding argument of `open`:

```
import csv
with open('some.csv', newline='', encoding='utf-8') as f:
    reader = csv.reader(f)
    for row in reader:
        print(row)
```

The same applies to writing in something other than the system default encoding: specify the encoding argument when opening the output file.

Registering a new dialect:

```
import csv
csv.register_dialect('unixpwd', delimiter=':', quoting=csv.QUOTE_NONE)
with open('passwd', newline='') as f:
    reader = csv.reader(f, 'unixpwd')
```

A slightly more advanced use of the reader — catching and reporting errors:

```
import csv, sys
filename = 'some.csv'
with open(filename, newline='') as f:
    reader = csv.reader(f)
    try:
        for row in reader:
            print(row)
    except csv.Error as e:
        sys.exit('file {}, line {}: {}'.format(filename, reader.line_num, e))
```

And while the module doesn't directly support parsing strings, it can easily be done:

```
import csv
for row in csv.reader(['one,two,three']):
    print(row)
```

14.2 configparser — Configuration file parser

Source code: [Lib/configparser.py](#)

This module provides the `ConfigParser` class which implements a basic configuration language which provides a structure similar to what's found in Microsoft Windows INI files. You can use this to write Python programs which can be customized by end users easily.

注解: This library does *not* interpret or write the value-type prefixes used in the Windows Registry extended version of INI syntax.

参见:

模块 `shlex` Support for creating Unix shell-like mini-languages which can be used as an alternate format for application configuration files.

模块 `json` The json module implements a subset of JavaScript syntax which can also be used for this purpose.

14.2.1 Quick Start

Let's take a very basic configuration file that looks like this:

```
[DEFAULT]
ServerAliveInterval = 45
Compression = yes
CompressionLevel = 9
ForwardX11 = yes

[bitbucket.org]
User = hg

[topsecret.server.com]
Port = 50022
ForwardX11 = no
```

The structure of INI files is described *in the following section*. Essentially, the file consists of sections, each of which contains keys with values. `configparser` classes can read and write such files. Let's start by creating the above configuration file programmatically.

```
>>> import configparser
>>> config = configparser.ConfigParser()
>>> config['DEFAULT'] = {'ServerAliveInterval': '45',
...                     'Compression': 'yes',
...                     'CompressionLevel': '9'}
>>> config['bitbucket.org'] = {}
>>> config['bitbucket.org']['User'] = 'hg'
>>> config['topsecret.server.com'] = {}
>>> topsecret = config['topsecret.server.com']
>>> topsecret['Port'] = '50022'      # mutates the parser
>>> topsecret['ForwardX11'] = 'no'  # same here
>>> config['DEFAULT']['ForwardX11'] = 'yes'
>>> with open('example.ini', 'w') as configfile:
...     config.write(configfile)
... 
```

As you can see, we can treat a config parser much like a dictionary. There are differences, *outlined later*, but the behavior is very close to what you would expect from a dictionary.

Now that we have created and saved a configuration file, let's read it back and explore the data it holds.

```
>>> import configparser
>>> config = configparser.ConfigParser()
>>> config.sections()
[]
>>> config.read('example.ini')
['example.ini']
>>> config.sections()
['bitbucket.org', 'topsecret.server.com']
>>> 'bitbucket.org' in config
True
>>> 'bytebong.com' in config
False
>>> config['bitbucket.org']['User']
'hg'
>>> config['DEFAULT']['Compression']
'yes'
>>> topsecret = config['topsecret.server.com']
>>> topsecret['ForwardX11']
'no'
>>> topsecret['Port']
'50022'
>>> for key in config['bitbucket.org']: print(key)
...
user
compressionlevel
serveraliveinterval
compression
forwardx11
>>> config['bitbucket.org']['ForwardX11']
'yes'
```

As we can see above, the API is pretty straightforward. The only bit of magic involves the `DEFAULT` section which provides default values for all other sections¹. Note also that keys in sections are case-insensitive and stored in lowercase¹.

14.2.2 Supported Datatypes

Config parsers do not guess datatypes of values in configuration files, always storing them internally as strings. This means that if you need other datatypes, you should convert on your own:

```
>>> int(topsecret['Port'])
50022
>>> float(topsecret['CompressionLevel'])
9.0
```

Since this task is so common, config parsers provide a range of handy getter methods to handle integers, floats and booleans. The last one is the most interesting because simply passing the value to `bool()` would do no good since `bool('False')` is still `True`. This is why config parsers also provide `getboolean()`. This method is case-insensitive and recognizes Boolean values from 'yes'/'no', 'on'/'off', 'true'/'false' and '1'/'0'¹. For example:

¹ Config parsers allow for heavy customization. If you are interested in changing the behaviour outlined by the footnote reference, consult the *Customizing Parser Behaviour* section.

```
>>> topsecret.getboolean('ForwardX11')
False
>>> config['bitbucket.org'].getboolean('ForwardX11')
True
>>> config.getboolean('bitbucket.org', 'Compression')
True
```

Apart from *getboolean()*, config parsers also provide equivalent *getint()* and *getfloat()* methods. You can register your own converters and customize the provided ones.¹

14.2.3 Fallback Values

As with a dictionary, you can use a section's *get()* method to provide fallback values:

```
>>> topsecret.get('Port')
'50022'
>>> topsecret.get('CompressionLevel')
'9'
>>> topsecret.get('Cipher')
>>> topsecret.get('Cipher', '3des-cbc')
'3des-cbc'
```

Please note that default values have precedence over fallback values. For instance, in our example the 'CompressionLevel' key was specified only in the 'DEFAULT' section. If we try to get it from the section 'topsecret.server.com', we will always get the default, even if we specify a fallback:

```
>>> topsecret.get('CompressionLevel', '3')
'9'
```

One more thing to be aware of is that the parser-level *get()* method provides a custom, more complex interface, maintained for backwards compatibility. When using this method, a fallback value can be provided via the *fallback* keyword-only argument:

```
>>> config.get('bitbucket.org', 'monster',
...           fallback='No such things as monsters')
'No such things as monsters'
```

The same fallback argument can be used with the *getint()*, *getfloat()* and *getboolean()* methods, for example:

```
>>> 'BatchMode' in topsecret
False
>>> topsecret.getboolean('BatchMode', fallback=True)
True
>>> config['DEFAULT']['BatchMode'] = 'no'
>>> topsecret.getboolean('BatchMode', fallback=True)
False
```

14.2.4 Supported INI File Structure

A configuration file consists of sections, each led by a [section] header, followed by key/value entries separated by a specific string (= or : by default¹). By default, section names are case sensitive but keys are not¹. Leading and trailing whitespace is removed from keys and values. Values can be omitted, in which case the key/value delimiter may also be

left out. Values can also span multiple lines, as long as they are indented deeper than the first line of the value. Depending on the parser's mode, blank lines may be treated as parts of multiline values or ignored.

Configuration files may include comments, prefixed by specific characters (`#` and `;` by default¹). Comments may appear on their own on an otherwise empty line, possibly indented.¹

例如:

```
[Simple Values]
key=value
spaces in keys=allowed
spaces in values=allowed as well
spaces around the delimiter = obviously
you can also use : to delimit keys from values

[All Values Are Strings]
values like this: 1000000
or this: 3.14159265359
are they treated as numbers? : no
integers, floats and booleans are held as: strings
can use the API to get converted values directly: true

[Multiline Values]
chorus: I'm a lumberjack, and I'm okay
       I sleep all night and I work all day

[No Values]
key_without_value
empty string value here =

[You can use comments]
# like this
; or this

# By default only in an empty line.
# Inline comments can be harmful because they prevent users
# from using the delimiting characters as parts of values.
# That being said, this can be customized.

    [Sections Can Be Indented]
        can_values_be_as_well = True
        does_that_mean_anything_special = False
        purpose = formatting for readability
        multiline_values = are
            handled just fine as
            long as they are indented
            deeper than the first line
            of a value
        # Did I mention we can indent comments, too?
```

14.2.5 Interpolation of values

On top of the core functionality, `ConfigParser` supports interpolation. This means values can be preprocessed before returning them from `get()` calls.

class `configparser.BasicInterpolation`

The default implementation used by `ConfigParser`. It enables values to contain format strings which refer to

other values in the same section, or values in the special default section¹. Additional default values can be provided on initialization.

例如:

```
[Paths]
home_dir: /Users
my_dir: %(home_dir)s/lumberjack
my_pictures: %(my_dir)s/Pictures
```

In the example above, *ConfigParser* with *interpolation* set to *BasicInterpolation()* would resolve `%(home_dir)s` to the value of `home_dir` (`/Users` in this case). `%(my_dir)s` in effect would resolve to `/Users/lumberjack`. All interpolations are done on demand so keys used in the chain of references do not have to be specified in any specific order in the configuration file.

With *interpolation* set to `None`, the parser would simply return `%(my_dir)s/Pictures` as the value of `my_pictures` and `%(home_dir)s/lumberjack` as the value of `my_dir`.

class `configparser.ExtendedInterpolation`

An alternative handler for interpolation which implements a more advanced syntax, used for instance in `zc.buildout`. Extended interpolation is using `${section:option}` to denote a value from a foreign section. Interpolation can span multiple levels. For convenience, if the `section:` part is omitted, interpolation defaults to the current section (and possibly the default values from the special section).

For example, the configuration specified above with basic interpolation, would look like this with extended interpolation:

```
[Paths]
home_dir: /Users
my_dir: ${home_dir}/lumberjack
my_pictures: ${my_dir}/Pictures
```

Values from other sections can be fetched as well:

```
[Common]
home_dir: /Users
library_dir: /Library
system_dir: /System
macports_dir: /opt/local

[Frameworks]
Python: 3.2
path: ${Common:system_dir}/Library/Frameworks/

[Arthur]
nickname: Two Sheds
last_name: Jackson
my_dir: ${Common:home_dir}/twosheds
my_pictures: ${my_dir}/Pictures
python_dir: ${Frameworks:path}/Python/Versions/${Frameworks:Python}
```

14.2.6 Mapping Protocol Access

3.2 新版功能.

Mapping protocol access is a generic name for functionality that enables using custom objects as if they were dictionaries. In case of *configparser*, the mapping interface implementation is using the `parser['section']['option']` notation.

`parser['section']` in particular returns a proxy for the section's data in the parser. This means that the values are not copied but they are taken from the original parser on demand. What's even more important is that when values are changed on a section proxy, they are actually mutated in the original parser.

`configparser` objects behave as close to actual dictionaries as possible. The mapping interface is complete and adheres to the *MutableMapping* ABC. However, there are a few differences that should be taken into account:

- By default, all keys in sections are accessible in a case-insensitive manner¹. E.g. for option in `parser["section"]` yields only optionxform'ed option key names. This means lowercased keys by default. At the same time, for a section that holds the key 'a', both expressions return `True`:

```
"a" in parser["section"]
"A" in parser["section"]
```

- All sections include `DEFAULTSECT` values as well which means that `.clear()` on a section may not leave the section visibly empty. This is because default values cannot be deleted from the section (because technically they are not there). If they are overridden in the section, deleting causes the default value to be visible again. Trying to delete a default value causes a `KeyError`.
- `DEFAULTSECT` cannot be removed from the parser:
 - trying to delete it raises `ValueError`,
 - `parser.clear()` leaves it intact,
 - `parser.popitem()` never returns it.
- `parser.get(section, option, **kwargs)` - the second argument is **not** a fallback value. Note however that the section-level `get()` methods are compatible both with the mapping protocol and the classic config-parser API.
- `parser.items()` is compatible with the mapping protocol (returns a list of *section_name, section_proxy* pairs including the `DEFAULTSECT`). However, this method can also be invoked with arguments: `parser.items(section, raw, vars)`. The latter call returns a list of *option, value* pairs for a specified section, with all interpolations expanded (unless `raw=True` is provided).

The mapping protocol is implemented on top of the existing legacy API so that subclasses overriding the original interface still should have mappings working as expected.

14.2.7 Customizing Parser Behaviour

There are nearly as many INI format variants as there are applications using it. `configparser` goes a long way to provide support for the largest sensible set of INI styles available. The default functionality is mainly dictated by historical background and it's very likely that you will want to customize some of the features.

The most common way to change the way a specific config parser works is to use the `__init__()` options:

- *defaults*, default value: `None`

This option accepts a dictionary of key-value pairs which will be initially put in the `DEFAULT` section. This makes for an elegant way to support concise configuration files that don't specify values which are the same as the documented default.

Hint: if you want to specify default values for a specific section, use `read_dict()` before you read the actual file.

- *dict_type*, default value: `collections.OrderedDict`

This option has a major impact on how the mapping protocol will behave and how the written configuration files look. With the default ordered dictionary, every section is stored in the order they were added to the parser. Same goes for options within sections.

An alternative dictionary type can be used for example to sort sections and options on write-back. You can also use a regular dictionary for performance reasons.

Please note: there are ways to add a set of key-value pairs in a single operation. When you use a regular dictionary in those operations, the order of the keys may be random. For example:

```
>>> parser = configparser.ConfigParser()
>>> parser.read_dict({'section1': {'key1': 'value1',
...                               'key2': 'value2',
...                               'key3': 'value3'},
...                  'section2': {'keyA': 'valueA',
...                               'keyB': 'valueB',
...                               'keyC': 'valueC'},
...                  'section3': {'foo': 'x',
...                               'bar': 'y',
...                               'baz': 'z'}})
>>> parser.sections()
['section3', 'section2', 'section1']
>>> [option for option in parser['section3']]
['baz', 'foo', 'bar']
```

In these operations you need to use an ordered dictionary as well:

```
>>> from collections import OrderedDict
>>> parser = configparser.ConfigParser()
>>> parser.read_dict(
...     OrderedDict((
...         ('s1',
...          OrderedDict((
...              ('1', '2'),
...              ('3', '4'),
...              ('5', '6'),
...          ))),
...         ('s2',
...          OrderedDict((
...              ('a', 'b'),
...              ('c', 'd'),
...              ('e', 'f'),
...          ))),
...     ))
>>> parser.sections()
['s1', 's2']
>>> [option for option in parser['s1']]
['1', '3', '5']
>>> [option for option in parser['s2'].values()]
['b', 'd', 'f']
```

- *allow_no_value*, default value: False

Some configuration files are known to include settings without values, but which otherwise conform to the syntax supported by *configparser*. The *allow_no_value* parameter to the constructor can be used to indicate that such values should be accepted:

```
>>> import configparser
```

(下页继续)

(续上页)

```

>>> sample_config = """
... [mysqld]
...   user = mysql
...   pid-file = /var/run/mysqld/mysqld.pid
...   skip-external-locking
...   old_passwords = 1
...   skip-bdb
...   # we don't need ACID today
...   skip-innodb
... """
>>> config = configparser.ConfigParser(allow_no_value=True)
>>> config.read_string(sample_config)

>>> # Settings with values are treated as before:
>>> config["mysqld"]["user"]
'mysql'

>>> # Settings without values provide None:
>>> config["mysqld"]["skip-bdb"]

>>> # Settings which aren't specified still raise an error:
>>> config["mysqld"]["does-not-exist"]
Traceback (most recent call last):
...
KeyError: 'does-not-exist'

```

- *delimiters*, default value: ('=', ':')

Delimiters are substrings that delimit keys from values within a section. The first occurrence of a delimiting substring on a line is considered a delimiter. This means values (but not keys) can contain the delimiters.

See also the *space_around_delimiters* argument to *ConfigParser.write()*.

- *comment_prefixes*, default value: ('# ', '; ')
- *inline_comment_prefixes*, default value: None

Comment prefixes are strings that indicate the start of a valid comment within a config file. *comment_prefixes* are used only on otherwise empty lines (optionally indented) whereas *inline_comment_prefixes* can be used after every valid value (e.g. section names, options and empty lines as well). By default inline comments are disabled and '# ' and ';' are used as prefixes for whole line comments.

在 3.2 版更改: In previous versions of *configparser* behaviour matched *comment_prefixes*=('# ', '; ') and *inline_comment_prefixes*=('; ',).

Please note that config parsers don't support escaping of comment prefixes so using *inline_comment_prefixes* may prevent users from specifying option values with characters used as comment prefixes. When in doubt, avoid setting *inline_comment_prefixes*. In any circumstances, the only way of storing comment prefix characters at the beginning of a line in multiline values is to interpolate the prefix, for example:

```

>>> from configparser import ConfigParser, ExtendedInterpolation
>>> parser = ConfigParser(interpolation=ExtendedInterpolation())
>>> # the default BasicInterpolation could be used as well
>>> parser.read_string("""
... [DEFAULT]
... hash = #
...

```

(下页继续)

(续上页)

```

... [hashes]
... shebang =
...     ${hash}!/usr/bin/env python
...     ${hash} -*- coding: utf-8 -*-
...
... extensions =
...     enabled_extension
...     another_extension
...     #disabled_by_comment
...     yet_another_extension
...
... interpolation not necessary = if # is not at line start
... even in multiline values = line #1
...     line #2
...     line #3
... """
>>> print(parser['hashes']['shebang'])

#!/usr/bin/env python
# -*- coding: utf-8 -*-
>>> print(parser['hashes']['extensions'])

enabled_extension
another_extension
yet_another_extension
>>> print(parser['hashes']['interpolation not necessary'])
if # is not at line start
>>> print(parser['hashes']['even in multiline values'])
line #1
line #2
line #3

```

- *strict*, default value: True

When set to True, the parser will not allow for any section or option duplicates while reading from a single source (using `read_file()`, `read_string()` or `read_dict()`). It is recommended to use strict parsers in new applications.

在 3.2 版更改: In previous versions of *configparser* behaviour matched `strict=False`.

- *empty_lines_in_values*, default value: True

In config parsers, values can span multiple lines as long as they are indented more than the key that holds them. By default parsers also let empty lines to be parts of values. At the same time, keys can be arbitrarily indented themselves to improve readability. In consequence, when configuration files get big and complex, it is easy for the user to lose track of the file structure. Take for instance:

```

[Section]
key = multiline
    value with a gotcha

    this = is still a part of the multiline value of 'key'

```

This can be especially problematic for the user to see if she's using a proportional font to edit the file. That is why when your application does not need values with empty lines, you should consider disallowing them. This will make empty lines split keys every time. In the example above, it would produce two keys, `key` and `this`.

- *default_section*, default value: `configparser.DEFAULTSECT` (that is: "DEFAULT")

The convention of allowing a special section of default values for other sections or interpolation purposes is a powerful concept of this library, letting users create complex declarative configurations. This section is normally called "DEFAULT" but this can be customized to point to any other valid section name. Some typical values include: "general" or "common". The name provided is used for recognizing default sections when reading from any source and is used when writing configuration back to a file. Its current value can be retrieved using the `parser_instance.default_section` attribute and may be modified at runtime (i.e. to convert files from one format to another).

- *interpolation*, default value: `configparser.BasicInterpolation`

Interpolation behaviour may be customized by providing a custom handler through the *interpolation* argument. `None` can be used to turn off interpolation completely, `ExtendedInterpolation()` provides a more advanced variant inspired by `zc.buildout`. More on the subject in the [dedicated documentation section](#). `RawConfigParser` has a default value of `None`.

- *converters*, default value: not set

Config parsers provide option value getters that perform type conversion. By default `getint()`, `getfloat()`, and `getboolean()` are implemented. Should other getters be desirable, users may define them in a subclass or pass a dictionary where each key is a name of the converter and each value is a callable implementing said conversion. For instance, passing `{'decimal': decimal.Decimal}` would add `getdecimal()` on both the parser object and all section proxies. In other words, it will be possible to write both `parser_instance.getdecimal('section', 'key', fallback=0)` and `parser_instance['section'].getdecimal('key', 0)`.

If the converter needs to access the state of the parser, it can be implemented as a method on a config parser subclass. If the name of this method starts with `get`, it will be available on all section proxies, in the dict-compatible form (see the `getdecimal()` example above).

More advanced customization may be achieved by overriding default values of these parser attributes. The defaults are defined on the classes, so they may be overridden by subclasses or by attribute assignment.

`ConfigParser.BOOLEAN_STATES`

By default when using `getboolean()`, config parsers consider the following values `True`: '1', 'yes', 'true', 'on' and the following values `False`: '0', 'no', 'false', 'off'. You can override this by specifying a custom dictionary of strings and their Boolean outcomes. For example:

```
>>> custom = configparser.ConfigParser()
>>> custom['section1'] = {'funky': 'nope'}
>>> custom['section1'].getboolean('funky')
Traceback (most recent call last):
...
ValueError: Not a boolean: nope
>>> custom.BOOLEAN_STATES = {'sure': True, 'nope': False}
>>> custom['section1'].getboolean('funky')
False
```

Other typical Boolean pairs include `accept/reject` or `enabled/disabled`.

`ConfigParser.optionxform(option)`

This method transforms option names on every read, get, or set operation. The default converts the name to lowercase. This also means that when a configuration file gets written, all keys will be lowercase. Override this method if that's unsuitable. For example:

```
>>> config = """
... [Section1]
... Key = Value
...
... [Section2]
```

(下页继续)

(续上页)

```

... AnotherKey = Value
... """
>>> typical = configparser.ConfigParser()
>>> typical.read_string(config)
>>> list(typical['Section1'].keys())
['key']
>>> list(typical['Section2'].keys())
['anotherkey']
>>> custom = configparser.RawConfigParser()
>>> custom.optionxform = lambda option: option
>>> custom.read_string(config)
>>> list(custom['Section1'].keys())
['Key']
>>> list(custom['Section2'].keys())
['AnotherKey']

```

ConfigParser.SECTCRE

A compiled regular expression used to parse section headers. The default matches `[section]` to the name "section". Whitespace is considered part of the section name, thus `[larch]` will be read as a section of name " larch ". Override this attribute if that's unsuitable. For example:

```

>>> config = """
... [Section 1]
... option = value
...
... [ Section 2 ]
... another = val
... """
>>> typical = ConfigParser()
>>> typical.read_string(config)
>>> typical.sections()
['Section 1', ' Section 2 ']
>>> custom = ConfigParser()
>>> custom.SECTCRE = re.compile(r"\[ *(?P<header>[^\]]+?) *\]")
>>> custom.read_string(config)
>>> custom.sections()
['Section 1', 'Section 2']

```

注解: While ConfigParser objects also use an OPTCRE attribute for recognizing option lines, it's not recommended to override it because that would interfere with constructor options *allow_no_value* and *delimiters*.

14.2.8 Legacy API Examples

Mainly because of backwards compatibility concerns, `configparser` provides also a legacy API with explicit get/set methods. While there are valid use cases for the methods outlined below, mapping protocol access is preferred for new projects. The legacy API is at times more advanced, low-level and downright counterintuitive.

An example of writing to a configuration file:

```

import configparser

config = configparser.RawConfigParser()

```

(下页继续)

(续上页)

```
# Please note that using RawConfigParser's set functions, you can assign
# non-string values to keys internally, but will receive an error when
# attempting to write to a file or when you get it in non-row mode. Setting
# values using the mapping protocol or ConfigParser's set() does not allow
# such assignments to take place.
config.add_section('Section1')
config.set('Section1', 'an_int', '15')
config.set('Section1', 'a_bool', 'true')
config.set('Section1', 'a_float', '3.1415')
config.set('Section1', 'baz', 'fun')
config.set('Section1', 'bar', 'Python')
config.set('Section1', 'foo', '%(bar)s is %(baz)s!')

# Writing our configuration file to 'example.cfg'
with open('example.cfg', 'w') as configfile:
    config.write(configfile)
```

An example of reading the configuration file again:

```
import configparser

config = configparser.RawConfigParser()
config.read('example.cfg')

# getfloat() raises an exception if the value is not a float
# getint() and getboolean() also do this for their respective types
a_float = config.getfloat('Section1', 'a_float')
an_int = config.getint('Section1', 'an_int')
print(a_float + an_int)

# Notice that the next output does not interpolate '%(bar)s' or '%(baz)s'.
# This is because we are using a RawConfigParser().
if config.getboolean('Section1', 'a_bool'):
    print(config.get('Section1', 'foo'))
```

To get interpolation, use *ConfigParser*:

```
import configparser

cfg = configparser.ConfigParser()
cfg.read('example.cfg')

# Set the optional *raw* argument of get() to True if you wish to disable
# interpolation in a single get operation.
print(cfg.get('Section1', 'foo', raw=False)) # -> "Python is fun!"
print(cfg.get('Section1', 'foo', raw=True))  # -> "%(bar)s is %(baz)s!"

# The optional *vars* argument is a dict with members that will take
# precedence in interpolation.
print(cfg.get('Section1', 'foo', vars={'bar': 'Documentation',
                                       'baz': 'evil'}))

# The optional *fallback* argument can be used to provide a fallback value
print(cfg.get('Section1', 'foo'))
# -> "Python is fun!"
```

(下页继续)

(续上页)

```
print(cfg.get('Section1', 'foo', fallback='Monty is not.'))
# -> "Python is fun!"

print(cfg.get('Section1', 'monster', fallback='No such things as monsters.'))
# -> "No such things as monsters."

# A bare print(cfg.get('Section1', 'monster')) would raise NoOptionError
# but we can also use:

print(cfg.get('Section1', 'monster', fallback=None))
# -> None
```

Default values are available in both types of ConfigParsers. They are used in interpolation if an option used is not defined elsewhere.

```
import configparser

# New instance with 'bar' and 'baz' defaulting to 'Life' and 'hard' each
config = configparser.ConfigParser({'bar': 'Life', 'baz': 'hard'})
config.read('example.cfg')

print(config.get('Section1', 'foo'))      # -> "Python is fun!"
config.remove_option('Section1', 'bar')
config.remove_option('Section1', 'baz')
print(config.get('Section1', 'foo'))      # -> "Life is hard!"
```

14.2.9 ConfigParser Objects

```
class configparser.ConfigParser (defaults=None, dict_type=collections.OrderedDict,
                                allow_no_value=False, delimiters=('=', ':'), comment_prefixes=(';', '#'), inline_comment_prefixes=None,
                                strict=True, empty_lines_in_values=True, default_section=configparser.DEFAULTSECT, interpolation=BasicInterpolation(), converters={})
```

The main configuration parser. When *defaults* is given, it is initialized into the dictionary of intrinsic defaults. When *dict_type* is given, it will be used to create the dictionary objects for the list of sections, for the options within a section, and for the default values.

When *delimiters* is given, it is used as the set of substrings that divide keys from values. When *comment_prefixes* is given, it will be used as the set of substrings that prefix comments in otherwise empty lines. Comments can be indented. When *inline_comment_prefixes* is given, it will be used as the set of substrings that prefix comments in non-empty lines.

When *strict* is *True* (the default), the parser won't allow for any section or option duplicates while reading from a single source (file, string or dictionary), raising *DuplicateSectionError* or *DuplicateOptionError*. When *empty_lines_in_values* is *False* (default: *True*), each empty line marks the end of an option. Otherwise, internal empty lines of a multiline option are kept as part of the value. When *allow_no_value* is *True* (default: *False*), options without values are accepted; the value held for these is *None* and they are serialized without the trailing delimiter.

When *default_section* is given, it specifies the name for the special section holding default values for other sections and interpolation purposes (normally named "DEFAULT"). This value can be retrieved and changed on runtime using the *default_section* instance attribute.

Interpolation behaviour may be customized by providing a custom handler through the *interpolation* argument.

`None` can be used to turn off interpolation completely, `ExtendedInterpolation()` provides a more advanced variant inspired by `zc.buildout`. More on the subject in the [dedicated documentation section](#).

All option names used in interpolation will be passed through the `optionxform()` method just like any other option name reference. For example, using the default implementation of `optionxform()` (which converts option names to lower case), the values `foo %(bar)s` and `foo %(BAR)s` are equivalent.

When `converters` is given, it should be a dictionary where each key represents the name of a type converter and each value is a callable implementing the conversion from string to the desired datatype. Every converter gets its own corresponding `get*()` method on the parser object and section proxies.

在 3.1 版更改: The default `dict_type` is `collections.OrderedDict`.

在 3.2 版更改: `allow_no_value`, `delimiters`, `comment_prefixes`, `strict`, `empty_lines_in_values`, `default_section` and `interpolation` were added.

在 3.5 版更改: The `converters` argument was added.

defaults()

Return a dictionary containing the instance-wide defaults.

sections()

Return a list of the sections available; the *default section* is not included in the list.

add_section(section)

Add a section named *section* to the instance. If a section by the given name already exists, `DuplicateSectionError` is raised. If the *default section* name is passed, `ValueError` is raised. The name of the section must be a string; if not, `TypeError` is raised.

在 3.2 版更改: Non-string section names raise `TypeError`.

has_section(section)

Indicates whether the named *section* is present in the configuration. The *default section* is not acknowledged.

options(section)

Return a list of options available in the specified *section*.

has_option(section, option)

If the given *section* exists, and contains the given *option*, return `True`; otherwise return `False`. If the specified *section* is `None` or an empty string, `DEFAULT` is assumed.

read(filename, encoding=None)

Attempt to read and parse an iterable of filenames, returning a list of filenames which were successfully parsed.

If *filenames* is a string or *path-like object*, it is treated as a single filename. If a file named in *filenames* cannot be opened, that file will be ignored. This is designed so that you can specify an iterable of potential configuration file locations (for example, the current directory, the user's home directory, and some system-wide directory), and all existing configuration files in the iterable will be read.

If none of the named files exist, the `ConfigParser` instance will contain an empty dataset. An application which requires initial values to be loaded from a file should load the required file or files using `read_file()` before calling `read()` for any optional files:

```
import configparser, os

config = configparser.ConfigParser()
config.read_file(open('defaults.cfg'))
config.read(['site.cfg', os.path.expanduser('~/.myapp.cfg')],
            encoding='cp1250')
```

3.2 新版功能: The *encoding* parameter. Previously, all files were read using the default encoding for `open()`.

3.6.1 新版功能: The *filenames* parameter accepts a *path-like object*.

read_file (*f*, *source=None*)

Read and parse configuration data from *f* which must be an iterable yielding Unicode strings (for example files opened in text mode).

Optional argument *source* specifies the name of the file being read. If not given and *f* has a *name* attribute, that is used for *source*; the default is '`<???`'.

3.2 新版功能: Replaces `readfp()`.

read_string (*string*, *source='<string>'*)

Parse configuration data from a string.

Optional argument *source* specifies a context-specific name of the string passed. If not given, '`<string>`' is used. This should commonly be a filesystem path or a URL.

3.2 新版功能.

read_dict (*dictionary*, *source='<dict>'*)

Load configuration from any object that provides a dict-like `items()` method. Keys are section names, values are dictionaries with keys and values that should be present in the section. If the used dictionary type preserves order, sections and their keys will be added in order. Values are automatically converted to strings.

Optional argument *source* specifies a context-specific name of the dictionary passed. If not given, `<dict>` is used.

This method can be used to copy state between parsers.

3.2 新版功能.

get (*section*, *option*, *, *raw=False*, *vars=None*[, *fallback*])

Get an *option* value for the named *section*. If *vars* is provided, it must be a dictionary. The *option* is looked up in *vars* (if provided), *section*, and in `DEFAULTSECT` in that order. If the key is not found and *fallback* is provided, it is used as a fallback value. `None` can be provided as a *fallback* value.

All the '`%`' interpolations are expanded in the return values, unless the *raw* argument is true. Values for interpolation keys are looked up in the same manner as the *option*.

在 3.2 版更改: Arguments *raw*, *vars* and *fallback* are keyword only to protect users from trying to use the third argument as the *fallback* fallback (especially when using the mapping protocol).

getint (*section*, *option*, *, *raw=False*, *vars=None*[, *fallback*])

A convenience method which coerces the *option* in the specified *section* to an integer. See `get()` for explanation of *raw*, *vars* and *fallback*.

getfloat (*section*, *option*, *, *raw=False*, *vars=None*[, *fallback*])

A convenience method which coerces the *option* in the specified *section* to a floating point number. See `get()` for explanation of *raw*, *vars* and *fallback*.

getboolean (*section*, *option*, *, *raw=False*, *vars=None*[, *fallback*])

A convenience method which coerces the *option* in the specified *section* to a Boolean value. Note that the accepted values for the *option* are '`1`', '`yes`', '`true`', and '`on`', which cause this method to return `True`, and '`0`', '`no`', '`false`', and '`off`', which cause it to return `False`. These string values are checked in a case-insensitive manner. Any other value will cause it to raise `ValueError`. See `get()` for explanation of *raw*, *vars* and *fallback*.

items (*raw=False*, *vars=None*)

items (*section*, *raw=False*, *vars=None*)

When *section* is not given, return a list of *section_name*, *section_proxy* pairs, including `DEFAULTSECT`.

Otherwise, return a list of *name*, *value* pairs for the options in the given *section*. Optional arguments have the same meaning as for the `get()` method.

set (*section*, *option*, *value*)

If the given section exists, set the given option to the specified value; otherwise raise `NoSectionError`. *option* and *value* must be strings; if not, `TypeError` is raised.

write (*fileobject*, *space_around_delimiters=True*)

Write a representation of the configuration to the specified *file object*, which must be opened in text mode (accepting strings). This representation can be parsed by a future `read()` call. If *space_around_delimiters* is true, delimiters between keys and values are surrounded by spaces.

remove_option (*section*, *option*)

Remove the specified *option* from the specified *section*. If the section does not exist, raise `NoSectionError`. If the option existed to be removed, return `True`; otherwise return `False`.

remove_section (*section*)

Remove the specified *section* from the configuration. If the section in fact existed, return `True`. Otherwise return `False`.

optionxform (*option*)

Transforms the option name *option* as found in an input file or as passed in by client code to the form that should be used in the internal structures. The default implementation returns a lower-case version of *option*; subclasses may override this or client code can set an attribute of this name on instances to affect this behavior.

You don't need to subclass the parser to use this method, you can also set it on an instance, to a function that takes a string argument and returns a string. Setting it to `str`, for example, would make option names case sensitive:

```
cfgparser = ConfigParser()
cfgparser.optionxform = str
```

Note that when reading configuration files, whitespace around the option names is stripped before `optionxform()` is called.

readfp (*fp*, *filename=None*)

3.2 版后已移除: Use `read_file()` instead.

在 3.2 版更改: `readfp()` now iterates on *fp* instead of calling `fp.readline()`.

For existing code calling `readfp()` with arguments which don't support iteration, the following generator may be used as a wrapper around the file-like object:

```
def readline_generator(fp):
    line = fp.readline()
    while line:
        yield line
        line = fp.readline()
```

Instead of `parser.readfp(fp)` use `parser.read_file(readline_generator(fp))`.

`configparser.MAX_INTERPOLATION_DEPTH`

The maximum depth for recursive interpolation for `get()` when the *raw* parameter is false. This is relevant only when the default *interpolation* is used.

14.2.10 RawConfigParser Objects

```
class configparser.RawConfigParser (defaults=None, dict_type=collections.OrderedDict,
                                     allow_no_value=False, *, delimiters=('=', ':'),
                                     comment_prefixes=(';', '#'), inline_comment_prefixes=None,
                                     strict=True, empty_lines_in_values=True,
                                     default_section=configparser.DEFAULTSECT[, interpolation
                                     ])
```

Legacy variant of the `ConfigParser` with interpolation disabled by default and unsafe `add_section` and `set` methods.

注解: Consider using `ConfigParser` instead which checks types of the values to be stored internally. If you don't want interpolation, you can use `ConfigParser(interpolation=None)`.

add_section (*section*)

Add a section named *section* to the instance. If a section by the given name already exists, `DuplicateSectionError` is raised. If the *default section* name is passed, `ValueError` is raised.

Type of *section* is not checked which lets users create non-string named sections. This behaviour is unsupported and may cause internal errors.

set (*section, option, value*)

If the given section exists, set the given option to the specified value; otherwise raise `NoSectionError`. While it is possible to use `RawConfigParser` (or `ConfigParser` with *raw* parameters set to true) for *internal* storage of non-string values, full functionality (including interpolation and output to files) can only be achieved using string values.

This method lets users assign non-string values to keys internally. This behaviour is unsupported and will cause errors when attempting to write to a file or get it in non-raw mode. **Use the mapping protocol API** which does not allow such assignments to take place.

14.2.11 异常

exception configparser.Error

Base class for all other `configparser` exceptions.

exception configparser.NoSectionError

Exception raised when a specified section is not found.

exception configparser.DuplicateSectionError

Exception raised if `add_section()` is called with the name of a section that is already present or in strict parsers when a section is found more than once in a single input file, string or dictionary.

3.2 新版功能: Optional `source` and `lineno` attributes and arguments to `__init__()` were added.

exception configparser.DuplicateOptionError

Exception raised by strict parsers if a single option appears twice during reading from a single file, string or dictionary. This catches misspellings and case sensitivity-related errors, e.g. a dictionary may have two keys representing the same case-insensitive configuration key.

exception configparser.NoOptionError

Exception raised when a specified option is not found in the specified section.

exception configparser.InterpolationError

Base class for exceptions raised when problems occur performing string interpolation.

exception `configparser.InterpolationDepthError`

Exception raised when string interpolation cannot be completed because the number of iterations exceeds `MAX_INTERPOLATION_DEPTH`. Subclass of `InterpolationError`.

exception `configparser.InterpolationMissingOptionError`

Exception raised when an option referenced from a value does not exist. Subclass of `InterpolationError`.

exception `configparser.InterpolationSyntaxError`

Exception raised when the source text into which substitutions are made does not conform to the required syntax. Subclass of `InterpolationError`.

exception `configparser.MissingSectionHeaderError`

Exception raised when attempting to parse a file which has no section headers.

exception `configparser.ParsingError`

Exception raised when errors occur attempting to parse a file.

在 3.2 版更改: The `filename` attribute and `__init__()` argument were renamed to `source` for consistency.

14.3 netrc —netrc file processing

Source code: [Lib/netrc.py](#)

The `netrc` class parses and encapsulates the netrc file format used by the Unix `ftp` program and other FTP clients.

class `netrc.netrc([file])`

A `netrc` instance or subclass instance encapsulates data from a netrc file. The initialization argument, if present, specifies the file to parse. If no argument is given, the file `.netrc` in the user's home directory will be read. Parse errors will raise `NetrcParseError` with diagnostic information including the file name, line number, and terminating token. If no argument is specified on a POSIX system, the presence of passwords in the `.netrc` file will raise a `NetrcParseError` if the file ownership or permissions are insecure (owned by a user other than the user running the process, or accessible for read or write by any other user). This implements security behavior equivalent to that of `ftp` and other programs that use `.netrc`.

在 3.4 版更改: Added the POSIX permission check.

exception `netrc.NetrcParseError`

Exception raised by the `netrc` class when syntactical errors are encountered in source text. Instances of this exception provide three interesting attributes: `msg` is a textual explanation of the error, `filename` is the name of the source file, and `lineno` gives the line number on which the error was found.

14.3.1 netrc Objects

A `netrc` instance has the following methods:

`netrc.authenticators(host)`

Return a 3-tuple (`login`, `account`, `password`) of authenticators for `host`. If the netrc file did not contain an entry for the given host, return the tuple associated with the 'default' entry. If neither matching host nor default entry is available, return `None`.

`netrc.__repr__()`

Dump the class data as a string in the format of a netrc file. (This discards comments and may reorder the entries.)

Instances of `netrc` have public instance variables:

netrc.hosts

Dictionary mapping host names to (login, account, password) tuples. The ‘default’ entry, if any, is represented as a pseudo-host by that name.

netrc.macros

Dictionary mapping macro names to string lists.

注解: Passwords are limited to a subset of the ASCII character set. All ASCII punctuation is allowed in passwords, however, note that whitespace and non-printable characters are not allowed in passwords. This is a limitation of the way the .netrc file is parsed and may be removed in the future.

14.4 xdrlib —Encode and decode XDR data

Source code: [Lib/xdrlib.py](#)

The *xdrlib* module supports the External Data Representation Standard as described in **RFC 1014**, written by Sun Microsystems, Inc. June 1987. It supports most of the data types described in the RFC.

The *xdrlib* module defines two classes, one for packing variables into XDR representation, and another for unpacking from XDR representation. There are also two exception classes.

class xdrlib.Packer

Packer is the class for packing data into XDR representation. The *Packer* class is instantiated with no arguments.

class xdrlib.Unpacker (data)

Unpacker is the complementary class which unpacks XDR data values from a string buffer. The input buffer is given as *data*.

参见:

RFC 1014 - XDR: External Data Representation Standard This RFC defined the encoding of data which was XDR at the time this module was originally written. It has apparently been obsoleted by **RFC 1832**.

RFC 1832 - XDR: External Data Representation Standard Newer RFC that provides a revised definition of XDR.

14.4.1 Packer Objects

Packer instances have the following methods:

Packer.get_buffer ()

Returns the current pack buffer as a string.

Packer.reset ()

Resets the pack buffer to the empty string.

In general, you can pack any of the most common XDR data types by calling the appropriate *pack_type()* method. Each method takes a single argument, the value to pack. The following simple data type packing methods are supported: *pack_uint()*, *pack_int()*, *pack_enum()*, *pack_bool()*, *pack_uhyper()*, and *pack_hyper()*.

Packer.pack_float (value)

Packs the single-precision floating point number *value*.

Packer.pack_double (value)

Packs the double-precision floating point number *value*.

The following methods support packing strings, bytes, and opaque data:

`Packer.pack_fstring(n, s)`

Packs a fixed length string, *s*. *n* is the length of the string but it is *not* packed into the data buffer. The string is padded with null bytes if necessary to guaranteed 4 byte alignment.

`Packer.pack_fopaque(n, data)`

Packs a fixed length opaque data stream, similarly to `pack_fstring()`.

`Packer.pack_string(s)`

Packs a variable length string, *s*. The length of the string is first packed as an unsigned integer, then the string data is packed with `pack_fstring()`.

`Packer.pack_opaque(data)`

Packs a variable length opaque data string, similarly to `pack_string()`.

`Packer.pack_bytes(bytes)`

Packs a variable length byte stream, similarly to `pack_string()`.

The following methods support packing arrays and lists:

`Packer.pack_list(list, pack_item)`

Packs a *list* of homogeneous items. This method is useful for lists with an indeterminate size; i.e. the size is not available until the entire list has been walked. For each item in the list, an unsigned integer 1 is packed first, followed by the data value from the list. *pack_item* is the function that is called to pack the individual item. At the end of the list, an unsigned integer 0 is packed.

For example, to pack a list of integers, the code might appear like this:

```
import xdrlib
p = xdrlib.Packer()
p.pack_list([1, 2, 3], p.pack_int)
```

`Packer.pack_farray(n, array, pack_item)`

Packs a fixed length list (*array*) of homogeneous items. *n* is the length of the list; it is *not* packed into the buffer, but a `ValueError` exception is raised if `len(array)` is not equal to *n*. As above, *pack_item* is the function used to pack each element.

`Packer.pack_array(list, pack_item)`

Packs a variable length *list* of homogeneous items. First, the length of the list is packed as an unsigned integer, then each element is packed as in `pack_farray()` above.

14.4.2 Unpacker Objects

The `Unpacker` class offers the following methods:

`Unpacker.reset(data)`

Resets the string buffer with the given *data*.

`Unpacker.get_position()`

Returns the current unpack position in the data buffer.

`Unpacker.set_position(position)`

Sets the data buffer unpack position to *position*. You should be careful about using `get_position()` and `set_position()`.

`Unpacker.get_buffer()`

Returns the current unpack data buffer as a string.

`Unpacker.done()`

Indicates unpack completion. Raises an *Error* exception if all of the data has not been unpacked.

In addition, every data type that can be packed with a *Packer*, can be unpacked with an *Unpacker*. Unpacking methods are of the form `unpack_type()`, and take no arguments. They return the unpacked object.

`Unpacker.unpack_float()`

Unpacks a single-precision floating point number.

`Unpacker.unpack_double()`

Unpacks a double-precision floating point number, similarly to `unpack_float()`.

In addition, the following methods unpack strings, bytes, and opaque data:

`Unpacker.unpack_fstring(n)`

Unpacks and returns a fixed length string. *n* is the number of characters expected. Padding with null bytes to guaranteed 4 byte alignment is assumed.

`Unpacker.unpack_fopaque(n)`

Unpacks and returns a fixed length opaque data stream, similarly to `unpack_fstring()`.

`Unpacker.unpack_string()`

Unpacks and returns a variable length string. The length of the string is first unpacked as an unsigned integer, then the string data is unpacked with `unpack_fstring()`.

`Unpacker.unpack_opaque()`

Unpacks and returns a variable length opaque data string, similarly to `unpack_string()`.

`Unpacker.unpack_bytes()`

Unpacks and returns a variable length byte stream, similarly to `unpack_string()`.

The following methods support unpacking arrays and lists:

`Unpacker.unpack_list(unpack_item)`

Unpacks and returns a list of homogeneous items. The list is unpacked one element at a time by first unpacking an unsigned integer flag. If the flag is 1, then the item is unpacked and appended to the list. A flag of 0 indicates the end of the list. `unpack_item` is the function that is called to unpack the items.

`Unpacker.unpack_farray(n, unpack_item)`

Unpacks and returns (as a list) a fixed length array of homogeneous items. *n* is number of list elements to expect in the buffer. As above, `unpack_item` is the function used to unpack each element.

`Unpacker.unpack_array(unpack_item)`

Unpacks and returns a variable length *list* of homogeneous items. First, the length of the list is unpacked as an unsigned integer, then each element is unpacked as in `unpack_farray()` above.

14.4.3 异常

Exceptions in this module are coded as class instances:

exception `xdrlib.Error`

The base exception class. *Error* has a single public attribute `msg` containing the description of the error.

exception `xdrlib.ConversionError`

Class derived from *Error*. Contains no additional instance variables.

Here is an example of how you would catch one of these exceptions:

```
import xdrlib
p = xdrlib.Packer()
try:
```

(下页继续)

(续上页)

```
p.pack_double(8.01)
except xdrlib.ConversionError as instance:
    print('packing the double failed:', instance.msg)
```

14.5 plistlib —Generate and parse Mac OS X .plist files

Source code: [Lib/plistlib.py](#)

This module provides an interface for reading and writing the “property list” files used mainly by Mac OS X and supports both binary and XML plist files.

The property list (.plist) file format is a simple serialization supporting basic object types, like dictionaries, lists, numbers and strings. Usually the top level object is a dictionary.

To write out and to parse a plist file, use the `dump()` and `load()` functions.

To work with plist data in bytes objects, use `dumps()` and `loads()`.

Values can be strings, integers, floats, booleans, tuples, lists, dictionaries (but only with string keys), `Data`, `bytes`, `bytesarray` or `datetime.datetime` objects.

在 3.4 版更改: New API, old API deprecated. Support for binary format plists added.

参见:

PList manual page Apple’s documentation of the file format.

这个模块定义了以下函数:

`plistlib.load(fp, *, fmt=None, use_builtintypes=True, dict_type=dict)`

Read a plist file. `fp` should be a readable and binary file object. Return the unpacked root object (which usually is a dictionary).

The `fmt` is the format of the file and the following values are valid:

- `None`: Autodetect the file format
- `FMT_XML`: XML file format
- `FMT_BINARY`: Binary plist format

If `use_builtintypes` is true (the default) binary data will be returned as instances of `bytes`, otherwise it is returned as instances of `Data`.

The `dict_type` is the type used for dictionaries that are read from the plist file. The exact structure of the plist can be recovered by using `collections.OrderedDict` (although the order of keys shouldn’t be important in plist files).

XML data for the `FMT_XML` format is parsed using the Expat parser from `xml.parsers.expat` – see its documentation for possible exceptions on ill-formed XML. Unknown elements will simply be ignored by the plist parser.

The parser for the binary format raises `InvalidFileException` when the file cannot be parsed.

3.4 新版功能.

`plistlib.loads(data, *, fmt=None, use_builtintypes=True, dict_type=dict)`

Load a plist from a bytes object. See `load()` for an explanation of the keyword arguments.

3.4 新版功能.

`plistlib.dump(value, fp, *, fmt=FMT_XML, sort_keys=True, skipkeys=False)`

Write *value* to a plist file. *Fp* should be a writable, binary file object.

The *fmt* argument specifies the format of the plist file and can be one of the following values:

- `FMT_XML`: XML formatted plist file
- `FMT_BINARY`: Binary formatted plist file

When *sort_keys* is true (the default) the keys for dictionaries will be written to the plist in sorted order, otherwise they will be written in the iteration order of the dictionary.

When *skipkeys* is false (the default) the function raises `TypeError` when a key of a dictionary is not a string, otherwise such keys are skipped.

A `TypeError` will be raised if the object is of an unsupported type or a container that contains objects of unsupported types.

An `OverflowError` will be raised for integer values that cannot be represented in (binary) plist files.

3.4 新版功能.

`plistlib.dumps(value, *, fmt=FMT_XML, sort_keys=True, skipkeys=False)`

Return *value* as a plist-formatted bytes object. See the documentation for `dump()` for an explanation of the keyword arguments of this function.

3.4 新版功能.

The following functions are deprecated:

`plistlib.readPlist(pathOrFile)`

Read a plist file. *pathOrFile* may be either a file name or a (readable and binary) file object. Returns the unpacked root object (which usually is a dictionary).

This function calls `load()` to do the actual work, see the documentation of *that function* for an explanation of the keyword arguments.

注解: Dict values in the result have a `__getattr__` method that defers to `__getitem__`. This means that you can use attribute access to access items of these dictionaries.

3.4 版后已移除: Use `load()` instead.

`plistlib.writePlist(rootObject, pathOrFile)`

Write *rootObject* to an XML plist file. *pathOrFile* may be either a file name or a (writable and binary) file object

3.4 版后已移除: Use `dump()` instead.

`plistlib.readPlistFromBytes(data)`

Read a plist data from a bytes object. Return the root object.

See `load()` for a description of the keyword arguments.

注解: Dict values in the result have a `__getattr__` method that defers to `__getitem__`. This means that you can use attribute access to access items of these dictionaries.

3.4 版后已移除: Use `loads()` instead.

`plistlib.writePlistToBytes(rootObject)`

Return *rootObject* as an XML plist-formatted bytes object.

3.4 版后已移除: Use `dumps()` instead.

The following classes are available:

Dict ([dict]) :

Return an extended mapping object with the same value as dictionary *dict*.

This class is a subclass of *dict* where attribute access can be used to access items. That is, `aDict.key` is the same as `aDict['key']` for getting, setting and deleting items in the mapping.

3.0 版后已移除.

class plistlib.Data (data)

Return a “data” wrapper object around the bytes object *data*. This is used in functions converting from/to plists to represent the `<data>` type available in plists.

It has one attribute, `data`, that can be used to retrieve the Python bytes object stored in it.

3.4 版后已移除: Use a *bytes* object instead.

The following constants are available:

plistlib.FMT_XML

The XML format for plist files.

3.4 新版功能.

plistlib.FMT_BINARY

The binary format for plist files

3.4 新版功能.

14.5.1 示例

Generating a plist:

```
pl = dict(
    aString = "Doodah",
    aList = ["A", "B", 12, 32.1, [1, 2, 3]],
    aFloat = 0.1,
    anInt = 728,
    aDict = dict(
        anotherString = "<hello & hi there!>",
        aThirdString = "M\ue4ssig, Ma\xdf",
        aTrueValue = True,
        aFalseValue = False,
    ),
    someData = b"<binary gunk>",
    someMoreData = b"<lots of binary gunk>" * 10,
    aDate = datetime.datetime.fromtimestamp(time.mktime(time.gmtime()))),
)
with open(fileName, 'wb') as fp:
    dump(pl, fp)
```

Parsing a plist:

```
with open(fileName, 'rb') as fp:
    pl = load(fp)
    print(pl["aKey"])
```


本章中描述的模块实现了加密性质的各种算法。它们可由安装人员自行决定。在 Unix 系统上，`crypt` 模块也可以使用。这是一个概述：

15.1 hashlib — 安全哈希与消息摘要

Source code: [Lib/hashlib.py](#)

This module implements a common interface to many different secure hash and message digest algorithms. Included are the FIPS secure hash algorithms SHA1, SHA224, SHA256, SHA384, and SHA512 (defined in FIPS 180-2) as well as RSA's MD5 algorithm (defined in Internet [RFC 1321](#)). The terms “secure hash” and “message digest” are interchangeable. Older algorithms were called message digests. The modern term is secure hash.

注解：如果你想找到 `adler32` 或 `crc32` 哈希函数，它们在 `zlib` 模块中。

警告：有些算法已知存在哈希碰撞弱点，请参考最后的“另请参阅”段。

15.1.1 哈希算法

There is one constructor method named for each type of *hash*. All return a hash object with the same simple interface. For example: use `sha256()` to create a SHA-256 hash object. You can now feed this object with *bytes-like objects* (normally *bytes*) using the `update()` method. At any point you can ask it for the *digest* of the concatenation of the data fed to it so far using the `digest()` or `hexdigest()` methods.

注解: For better multithreading performance, the Python *GIL* is released for data larger than 2047 bytes at object creation or on update.

注解: Feeding string objects into `update()` is not supported, as hashes work on bytes, not on characters.

Constructors for hash algorithms that are always present in this module are `sha1()`, `sha224()`, `sha256()`, `sha384()`, `sha512()`, `blake2b()`, and `blake2s()`. `md5()` is normally available as well, though it may be missing if you are using a rare “FIPS compliant” build of Python. Additional algorithms may also be available depending upon the OpenSSL library that Python uses on your platform. On most platforms the `sha3_224()`, `sha3_256()`, `sha3_384()`, `sha3_512()`, `shake_128()`, `shake_256()` are also available.

3.6 新版功能: SHA3 (Keccak) and SHAKE constructors `sha3_224()`, `sha3_256()`, `sha3_384()`, `sha3_512()`, `shake_128()`, `shake_256()`.

3.6 新版功能: 添加了 `blake2b()` 和 `blake2s()`。

For example, to obtain the digest of the byte string `b'Nobody inspects the spammish repetition'`:

```
>>> import hashlib
>>> m = hashlib.sha256()
>>> m.update(b"Nobody inspects")
>>> m.update(b" the spammish repetition")
>>> m.digest()
b'\x03\x1e\xdd
↪Ae\x15\x93\xc5\xfe\\x00o\xa5u+7\xfd\xdf\xf7\xbcN\x84:\xa6\xaf\x0c\x95\x0fK\x94\x06'
>>> m.digest_size
32
>>> m.block_size
64
```

More condensed:

```
>>> hashlib.sha224(b"Nobody inspects the spammish repetition").hexdigest()
'a4337bc45a8fc544c03f52dc550cd6e1e87021bc896588bd79e901e2'
```

`hashlib.new(name[, data])`

Is a generic constructor that takes the string *name* of the desired algorithm as its first parameter. It also exists to allow access to the above listed hashes as well as any other algorithms that your OpenSSL library may offer. The named constructors are much faster than `new()` and should be preferred.

Using `new()` with an algorithm provided by OpenSSL:

```
>>> h = hashlib.new('ripemd160')
>>> h.update(b"Nobody inspects the spammish repetition")
>>> h.hexdigest()
'cc4a5ce1b3df48aec5d22d1f16b894a0b894eccc'
```

Hashlib 提供下列常量属性:

`hashlib.algorithms_guaranteed`

A set containing the names of the hash algorithms guaranteed to be supported by this module on all platforms. Note that ‘md5’ is in this list despite some upstream vendors offering an odd “FIPS compliant” Python build that excludes it.

3.2 新版功能.

hashlib.algorithms_available

A set containing the names of the hash algorithms that are available in the running Python interpreter. These names will be recognized when passed to `new()`. `algorithms_guaranteed` will always be a subset. The same algorithm may appear multiple times in this set under different names (thanks to OpenSSL).

3.2 新版功能.

The following values are provided as constant attributes of the hash objects returned by the constructors:

hash.digest_size

The size of the resulting hash in bytes.

hash.block_size

The internal block size of the hash algorithm in bytes.

A hash object has the following attributes:

hash.name

The canonical name of this hash, always lowercase and always suitable as a parameter to `new()` to create another hash of this type.

在 3.4 版更改: The name attribute has been present in CPython since its inception, but until Python 3.4 was not formally specified, so may not exist on some platforms.

A hash object has the following methods:

hash.update(data)

Update the hash object with the *bytes-like object*. Repeated calls are equivalent to a single call with the concatenation of all the arguments: `m.update(a); m.update(b)` is equivalent to `m.update(a+b)`.

在 3.1 版更改: The Python GIL is released to allow other threads to run while hash updates on data larger than 2047 bytes is taking place when using hash algorithms supplied by OpenSSL.

hash.digest()

Return the digest of the data passed to the `update()` method so far. This is a bytes object of size `digest_size` which may contain bytes in the whole range from 0 to 255.

hash.hexdigest()

Like `digest()` except the digest is returned as a string object of double length, containing only hexadecimal digits. This may be used to exchange the value safely in email or other non-binary environments.

hash.copy()

Return a copy (“clone”) of the hash object. This can be used to efficiently compute the digests of data sharing a common initial substring.

15.1.2 SHAKE variable length digests

The `shake_128()` and `shake_256()` algorithms provide variable length digests with `length_in_bits//2` up to 128 or 256 bits of security. As such, their digest methods require a length. Maximum length is not limited by the SHAKE algorithm.

shake.digest(length)

Return the digest of the data passed to the `update()` method so far. This is a bytes object of size `length` which may contain bytes in the whole range from 0 to 255.

shake.hexdigest(length)

Like `digest()` except the digest is returned as a string object of double length, containing only hexadecimal digits. This may be used to exchange the value safely in email or other non-binary environments.

15.1.3 Key derivation

Key derivation and key stretching algorithms are designed for secure password hashing. Naive algorithms such as `sha1(password)` are not resistant against brute-force attacks. A good password hashing function must be tunable, slow, and include a *salt*.

`hashlib.pbkdf2_hmac` (*hash_name*, *password*, *salt*, *iterations*, *dklen=None*)

The function provides PKCS#5 password-based key derivation function 2. It uses HMAC as pseudorandom function.

The string *hash_name* is the desired name of the hash digest algorithm for HMAC, e.g. ‘sha1’ or ‘sha256’. *password* and *salt* are interpreted as buffers of bytes. Applications and libraries should limit *password* to a sensible length (e.g. 1024). *salt* should be about 16 or more bytes from a proper source, e.g. `os.urandom()`.

The number of *iterations* should be chosen based on the hash algorithm and computing power. As of 2013, at least 100,000 iterations of SHA-256 are suggested.

dklen is the length of the derived key. If *dklen* is `None` then the digest size of the hash algorithm *hash_name* is used, e.g. 64 for SHA-512.

```
>>> import hashlib, binascii
>>> dk = hashlib.pbkdf2_hmac('sha256', b'password', b'salt', 100000)
>>> binascii.hexlify(dk)
b'0394a2ede332c9a13eb82e9b24631604c31df978b4e2f0fbd2c549944f9d79a5'
```

3.4 新版功能.

注解: A fast implementation of *pbkdf2_hmac* is available with OpenSSL. The Python implementation uses an inline version of *hmac*. It is about three times slower and doesn't release the GIL.

`hashlib.scrypt` (*password*, *, *salt*, *n*, *r*, *p*, *maxmem=0*, *dklen=64*)

The function provides scrypt password-based key derivation function as defined in [RFC 7914](#).

password and *salt* must be *bytes-like objects*. Applications and libraries should limit *password* to a sensible length (e.g. 1024). *salt* should be about 16 or more bytes from a proper source, e.g. `os.urandom()`.

n is the CPU/Memory cost factor, *r* the block size, *p* parallelization factor and *maxmem* limits memory (OpenSSL 1.1.0 defaults to 32 MB). *dklen* is the length of the derived key.

Availability: OpenSSL 1.1+

3.6 新版功能.

15.1.4 BLAKE2

BLAKE2 is a cryptographic hash function defined in [RFC 7693](#) that comes in two flavors:

- **BLAKE2b**, optimized for 64-bit platforms and produces digests of any size between 1 and 64 bytes,
- **BLAKE2s**, optimized for 8- to 32-bit platforms and produces digests of any size between 1 and 32 bytes.

BLAKE2 supports **keyed mode** (a faster and simpler replacement for **HMAC**), **salted hashing**, **personalization**, and **tree hashing**.

Hash objects from this module follow the API of standard library's *hashlib* objects.

Creating hash objects

New hash objects are created by calling constructor functions:

```
hashlib.blake2b (data=b", *, digest_size=64, key=b", salt=b", person=b", fanout=1, depth=1, leaf_size=0,
                  node_offset=0, node_depth=0, inner_size=0, last_node=False)
```

```
hashlib.blake2s (data=b", *, digest_size=32, key=b", salt=b", person=b", fanout=1, depth=1, leaf_size=0,
                  node_offset=0, node_depth=0, inner_size=0, last_node=False)
```

These functions return the corresponding hash objects for calculating BLAKE2b or BLAKE2s. They optionally take these general parameters:

- *data*: initial chunk of data to hash, which must be *bytes-like object*. It can be passed only as positional argument.
- *digest_size*: size of output digest in bytes.
- *key*: key for keyed hashing (up to 64 bytes for BLAKE2b, up to 32 bytes for BLAKE2s).
- *salt*: salt for randomized hashing (up to 16 bytes for BLAKE2b, up to 8 bytes for BLAKE2s).
- *person*: personalization string (up to 16 bytes for BLAKE2b, up to 8 bytes for BLAKE2s).

The following table shows limits for general parameters (in bytes):

Hash	digest_size	len(key)	len(salt)	len(person)
BLAKE2b	64	64	16	16
BLAKE2s	32	32	8	8

注解: BLAKE2 specification defines constant lengths for salt and personalization parameters, however, for convenience, this implementation accepts byte strings of any size up to the specified length. If the length of the parameter is less than specified, it is padded with zeros, thus, for example, `b'salt'` and `b'salt\x00'` is the same value. (This is not the case for *key*.)

These sizes are available as module *constants* described below.

Constructor functions also accept the following tree hashing parameters:

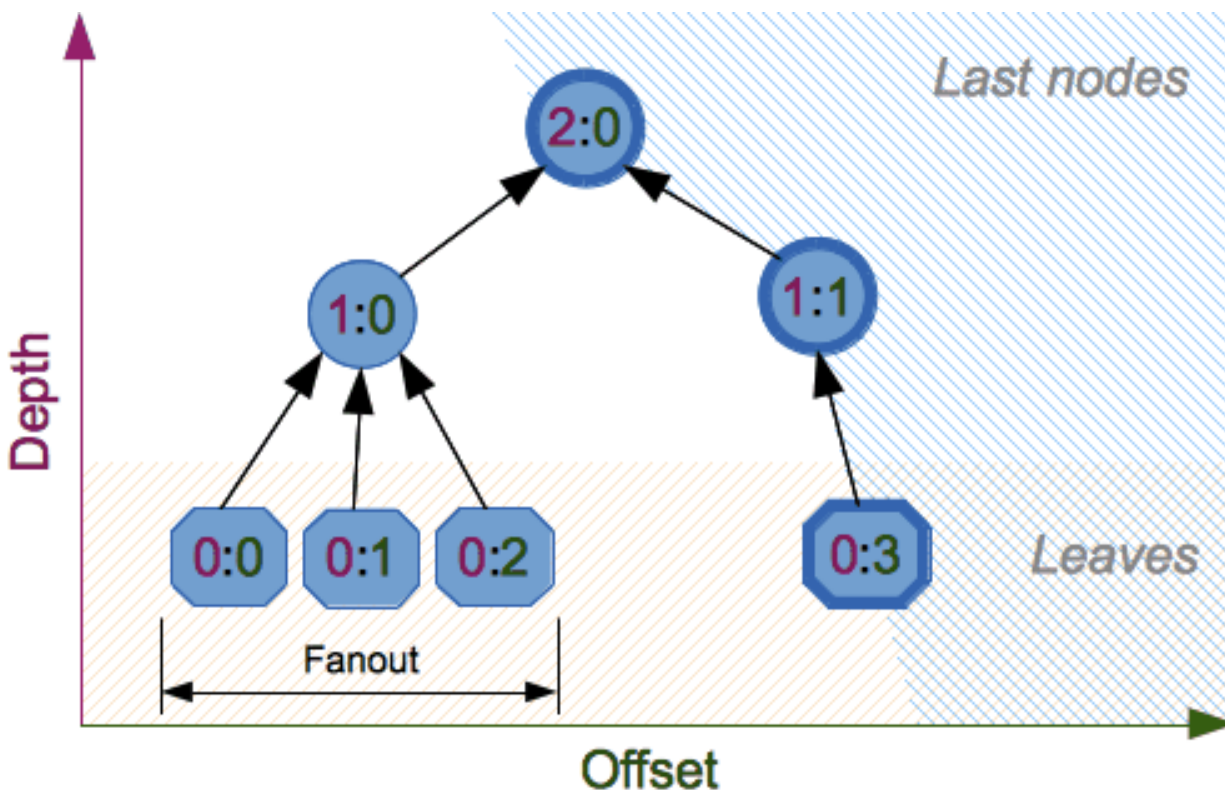
- *fanout*: fanout (0 to 255, 0 if unlimited, 1 in sequential mode).
- *depth*: maximal depth of tree (1 to 255, 255 if unlimited, 1 in sequential mode).
- *leaf_size*: maximal byte length of leaf (0 to $2^{**32}-1$, 0 if unlimited or in sequential mode).
- *node_offset*: node offset (0 to $2^{**64}-1$ for BLAKE2b, 0 to $2^{**48}-1$ for BLAKE2s, 0 for the first, leftmost, leaf, or in sequential mode).
- *node_depth*: node depth (0 to 255, 0 for leaves, or in sequential mode).
- *inner_size*: inner digest size (0 to 64 for BLAKE2b, 0 to 32 for BLAKE2s, 0 in sequential mode).
- *last_node*: boolean indicating whether the processed node is the last one (*False* for sequential mode).

See section 2.10 in [BLAKE2 specification](#) for comprehensive review of tree hashing.

常数

```
blake2b.SALT_SIZE
```

```
blake2s.SALT_SIZE
```



Salt length (maximum length accepted by constructors).

`blake2b.PERSON_SIZE`

`blake2s.PERSON_SIZE`

Personalization string length (maximum length accepted by constructors).

`blake2b.MAX_KEY_SIZE`

`blake2s.MAX_KEY_SIZE`

Maximum key size.

`blake2b.MAX_DIGEST_SIZE`

`blake2s.MAX_DIGEST_SIZE`

Maximum digest size that the hash function can output.

示例

Simple hashing

To calculate hash of some data, you should first construct a hash object by calling the appropriate constructor function (`blake2b()` or `blake2s()`), then update it with the data by calling `update()` on the object, and, finally, get the digest out of the object by calling `digest()` (or `hexdigest()` for hex-encoded string).

```
>>> from hashlib import blake2b
>>> h = blake2b()
```

(下页继续)

(续上页)

```
>>> h.update(b'Hello world')
>>> h.hexdigest()

↪ '6ff843ba685842aa82031d3f53c48b66326df7639a63d128974c5c14f31a0f33343a8c65551134ed1ae0f2b0dd2bb495d'
↪ '
```

As a shortcut, you can pass the first chunk of data to update directly to the constructor as the positional argument:

```
>>> from hashlib import blake2b
>>> blake2b(b'Hello world').hexdigest()

↪ '6ff843ba685842aa82031d3f53c48b66326df7639a63d128974c5c14f31a0f33343a8c65551134ed1ae0f2b0dd2bb495d'
↪ '
```

You can call `hash.update()` as many times as you need to iteratively update the hash:

```
>>> from hashlib import blake2b
>>> items = [b'Hello', b' ', b'world']
>>> h = blake2b()
>>> for item in items:
...     h.update(item)
>>> h.hexdigest()

↪ '6ff843ba685842aa82031d3f53c48b66326df7639a63d128974c5c14f31a0f33343a8c65551134ed1ae0f2b0dd2bb495d'
↪ '
```

Using different digest sizes

BLAKE2 has configurable size of digests up to 64 bytes for BLAKE2b and up to 32 bytes for BLAKE2s. For example, to replace SHA-1 with BLAKE2b without changing the size of output, we can tell BLAKE2b to produce 20-byte digests:

```
>>> from hashlib import blake2b
>>> h = blake2b(digest_size=20)
>>> h.update(b'Replacing SHA1 with the more secure function')
>>> h.hexdigest()
'd24f26cf8de66472d58d4e1b1774b4c9158b1f4c'
>>> h.digest_size
20
>>> len(h.digest())
20
```

Hash objects with different digest sizes have completely different outputs (shorter hashes are *not* prefixes of longer hashes); BLAKE2b and BLAKE2s produce different outputs even if the output length is the same:

```
>>> from hashlib import blake2b, blake2s
>>> blake2b(digest_size=10).hexdigest()
'6fa1d8fcfd719046d762'
>>> blake2b(digest_size=11).hexdigest()
'eb6ec15daf9546254f0809'
>>> blake2s(digest_size=10).hexdigest()
'1bf21a98c78a1c376ae9'
>>> blake2s(digest_size=11).hexdigest()
'567004bf96e4a25773ebf4'
```

Keyed hashing

Keyed hashing can be used for authentication as a faster and simpler replacement for [Hash-based message authentication code](#) (HMAC). BLAKE2 can be securely used in prefix-MAC mode thanks to the indistinguishability property inherited from BLAKE.

This example shows how to get a (hex-encoded) 128-bit authentication code for message `b'message data'` with key `b'pseudorandom key'`:

```
>>> from hashlib import blake2b
>>> h = blake2b(key=b'pseudorandom key', digest_size=16)
>>> h.update(b'message data')
>>> h.hexdigest()
'3d363ff7401e02026f4a4687d4863ced'
```

As a practical example, a web application can symmetrically sign cookies sent to users and later verify them to make sure they weren't tampered with:

```
>>> from hashlib import blake2b
>>> from hmac import compare_digest
>>>
>>> SECRET_KEY = b'pseudorandomly generated server secret key'
>>> AUTH_SIZE = 16
>>>
>>> def sign(cookie):
...     h = blake2b(digest_size=AUTH_SIZE, key=SECRET_KEY)
...     h.update(cookie)
...     return h.hexdigest().encode('utf-8')
>>>
>>> def verify(cookie, sig):
...     good_sig = sign(cookie)
...     return compare_digest(good_sig, sig)
>>>
>>> cookie = b'user-alice'
>>> sig = sign(cookie)
>>> print("{0},{1}".format(cookie.decode('utf-8'), sig))
user-alice,b'43b3c982cf697e0c5ab22172d1ca7421'
>>> verify(cookie, sig)
True
>>> verify(b'user-bob', sig)
False
>>> verify(cookie, b'0102030405060708090a0b0c0d0e0f00')
False
```

Even though there's a native keyed hashing mode, BLAKE2 can, of course, be used in HMAC construction with `hmac` module:

```
>>> import hmac, hashlib
>>> m = hmac.new(b'secret key', digestmod=hashlib.blake2s)
>>> m.update(b'message')
>>> m.hexdigest()
'e3c8102868d28b5ff85fc35dda07329970d1a01e273c37481326fe0c861c8142'
```

Randomized hashing

By setting *salt* parameter users can introduce randomization to the hash function. Randomized hashing is useful for protecting against collision attacks on the hash function used in digital signatures.

Randomized hashing is designed for situations where one party, the message preparer, generates all or part of a message to be signed by a second party, the message signer. If the message preparer is able to find cryptographic hash function collisions (i.e., two messages producing the same hash value), then they might prepare meaningful versions of the message that would produce the same hash value and digital signature, but with different results (e.g., transferring \$1,000,000 to an account, rather than \$10). Cryptographic hash functions have been designed with collision resistance as a major goal, but the current concentration on attacking cryptographic hash functions may result in a given cryptographic hash function providing less collision resistance than expected. Randomized hashing offers the signer additional protection by reducing the likelihood that a preparer can generate two or more messages that ultimately yield the same hash value during the digital signature generation process—even if it is practical to find collisions for the hash function. However, the use of randomized hashing may reduce the amount of security provided by a digital signature when all portions of the message are prepared by the signer.

(NIST SP-800-106 “Randomized Hashing for Digital Signatures”)

In BLAKE2 the salt is processed as a one-time input to the hash function during initialization, rather than as an input to each compression function.

警告: *Salted hashing* (or just hashing) with BLAKE2 or any other general-purpose cryptographic hash function, such as SHA-256, is not suitable for hashing passwords. See [BLAKE2 FAQ](#) for more information.

```
>>> import os
>>> from hashlib import blake2b
>>> msg = b'some message'
>>> # Calculate the first hash with a random salt.
>>> salt1 = os.urandom(blake2b.SALT_SIZE)
>>> h1 = blake2b(salt=salt1)
>>> h1.update(msg)
>>> # Calculate the second hash with a different random salt.
>>> salt2 = os.urandom(blake2b.SALT_SIZE)
>>> h2 = blake2b(salt=salt2)
>>> h2.update(msg)
>>> # The digests are different.
>>> h1.digest() != h2.digest()
True
```

Personalization

Sometimes it is useful to force hash function to produce different digests for the same input for different purposes. Quoting the authors of the Skein hash function:

We recommend that all application designers seriously consider doing this; we have seen many protocols where a hash that is computed in one part of the protocol can be used in an entirely different part because two hash computations were done on similar or related data, and the attacker can force the application to make the hash inputs the same. Personalizing each hash function used in the protocol summarily stops this type of attack.

(The Skein Hash Function Family, p. 21)

BLAKE2 can be personalized by passing bytes to the *person* argument:

```
>>> from hashlib import blake2b
>>> FILES_HASH_PERSON = b'MyApp Files Hash'
>>> BLOCK_HASH_PERSON = b'MyApp Block Hash'
>>> h = blake2b(digest_size=32, person=FILES_HASH_PERSON)
>>> h.update(b'the same content')
>>> h.hexdigest()
'20d9cd024d4fb086aae819a1432dd2466de12947831b75c5a30cf2676095d3b4'
>>> h = blake2b(digest_size=32, person=BLOCK_HASH_PERSON)
>>> h.update(b'the same content')
>>> h.hexdigest()
'cf68fb5761b9c44e7878bf2c4c9aea52264a80b75005e65619778de59f383a3'
```

Personalization together with the keyed mode can also be used to derive different keys from a single one.

```
>>> from hashlib import blake2s
>>> from base64 import b64decode, b64encode
>>> orig_key = b64decode(b'Rm5EPJai72qcK3RGBpW3vPNfZy5OZothY+kHY6h21KM=')
>>> enc_key = blake2s(key=orig_key, person=b'kEncrypt').digest()
>>> mac_key = blake2s(key=orig_key, person=b'kMAC').digest()
>>> print(b64encode(enc_key).decode('utf-8'))
rbPb15S/Z9t+agffno5wuhB77VbRi6F9Iv2qIxU7WHw=
>>> print(b64encode(mac_key).decode('utf-8'))
G9GtHFE1YluXY1zWPlYk1e/nWfu0WSEb0KRcjhDeP/o=
```

Tree mode

Here's an example of hashing a minimal tree with two leaf nodes:

```
  10
 /  \
00  01
```

This example uses 64-byte internal digests, and returns the 32-byte final digest:

```
>>> from hashlib import blake2b
>>>
>>> FANOUT = 2
>>> DEPTH = 2
>>> LEAF_SIZE = 4096
>>> INNER_SIZE = 64
>>>
>>> buf = bytearray(6000)
>>>
>>> # Left leaf
... h00 = blake2b(buf[0:LEAF_SIZE], fanout=FANOUT, depth=DEPTH,
...               leaf_size=LEAF_SIZE, inner_size=INNER_SIZE,
...               node_offset=0, node_depth=0, last_node=False)
>>> # Right leaf
... h01 = blake2b(buf[LEAF_SIZE:], fanout=FANOUT, depth=DEPTH,
...               leaf_size=LEAF_SIZE, inner_size=INNER_SIZE,
...               node_offset=1, node_depth=0, last_node=True)
>>> # Root node
... h10 = blake2b(digest_size=32, fanout=FANOUT, depth=DEPTH,
...               leaf_size=LEAF_SIZE, inner_size=INNER_SIZE,
...               node_offset=0, node_depth=1, last_node=True)
```

(下页继续)

(续上页)

```
>>> h10.update(h00.digest())
>>> h10.update(h01.digest())
>>> h10.hexdigest()
'3ad2a9b37c6070e374c7a8c508fe20ca86b6ed54e286e93a0318e95e881db5aa'
```

Credits

BLAKE2 was designed by *Jean-Philippe Aumasson, Samuel Neves, Zooko Wilcox-O' Hearn, and Christian Winnerlein* based on SHA-3 finalist BLAKE created by *Jean-Philippe Aumasson, Luca Henzen, Willi Meier, and Raphael C.-W. Phan*.

It uses core algorithm from ChaCha cipher designed by *Daniel J. Bernstein*.

The stdlib implementation is based on `pyblake2` module. It was written by *Dmitry Chestnykh* based on C implementation written by *Samuel Neves*. The documentation was copied from `pyblake2` and written by *Dmitry Chestnykh*.

The C code was partly rewritten for Python by *Christian Heimes*.

The following public domain dedication applies for both C hash function implementation, extension code, and this documentation:

To the extent possible under law, the author(s) have dedicated all copyright and related and neighboring rights to this software to the public domain worldwide. This software is distributed without any warranty.

You should have received a copy of the CC0 Public Domain Dedication along with this software. If not, see <http://creativecommons.org/publicdomain/zero/1.0/>.

The following people have helped with development or contributed their changes to the project and the public domain according to the Creative Commons Public Domain Dedication 1.0 Universal:

- *Alexandr Sokolovskiy*

参见:

Module `hmac` A module to generate message authentication codes using hashes.

模块 `base64` Another way to encode binary hashes for non-binary environments.

<https://blake2.net> Official BLAKE2 website.

<http://csrc.nist.gov/publications/fips/fips180-2/fips180-2.pdf> The FIPS 180-2 publication on Secure Hash Algorithms.

https://en.wikipedia.org/wiki/Cryptographic_hash_function#Cryptographic_hash_algorithms Wikipedia article with information on which algorithms have known issues and what that means regarding their use.

<https://www.ietf.org/rfc/rfc2898.txt> PKCS #5: Password-Based Cryptography Specification Version 2.0

15.2 hmac — 基于密钥的消息验证

Source code: `Lib/hmac.py`

This module implements the HMAC algorithm as described by **RFC 2104**.

`hmac.new(key, msg=None, digestmod=None)`

Return a new hmac object. `key` is a bytes or bytearray object giving the secret key. If `msg` is present, the method call `update(msg)` is made. `digestmod` is the digest name, digest constructor or module for the HMAC object to use. It supports any name suitable to `hashlib.new()` and defaults to the `hashlib.md5` constructor.

在 3.4 版更改: Parameter *key* can be a bytes or bytearray object. Parameter *msg* can be of any type supported by *hashlib*. Parameter *digestmod* can be the name of a hash algorithm.

3.4 版后已移除: MD5 as implicit default digest for *digestmod* is deprecated.

An HMAC object has the following methods:

HMAC.update(*msg*)

Update the hmac object with *msg*. Repeated calls are equivalent to a single call with the concatenation of all the arguments: `m.update(a); m.update(b)` is equivalent to `m.update(a + b)`.

在 3.4 版更改: Parameter *msg* can be of any type supported by *hashlib*.

HMAC.digest()

Return the digest of the bytes passed to the *update()* method so far. This bytes object will be the same length as the *digest_size* of the digest given to the constructor. It may contain non-ASCII bytes, including NUL bytes.

警告: When comparing the output of *digest()* to an externally-supplied digest during a verification routine, it is recommended to use the *compare_digest()* function instead of the `==` operator to reduce the vulnerability to timing attacks.

HMAC.hexdigest()

Like *digest()* except the digest is returned as a string twice the length containing only hexadecimal digits. This may be used to exchange the value safely in email or other non-binary environments.

警告: When comparing the output of *hexdigest()* to an externally-supplied digest during a verification routine, it is recommended to use the *compare_digest()* function instead of the `==` operator to reduce the vulnerability to timing attacks.

HMAC.copy()

Return a copy (“clone”) of the hmac object. This can be used to efficiently compute the digests of strings that share a common initial substring.

A hash object has the following attributes:

HMAC.digest_size

The size of the resulting HMAC digest in bytes.

HMAC.block_size

The internal block size of the hash algorithm in bytes.

3.4 新版功能.

HMAC.name

The canonical name of this HMAC, always lowercase, e.g. `hmac-md5`.

3.4 新版功能.

This module also provides the following helper function:

hmac.compare_digest(*a*, *b*)

Return `a == b`. This function uses an approach designed to prevent timing analysis by avoiding content-based short circuiting behaviour, making it appropriate for cryptography. *a* and *b* must both be of the same type: either *str* (ASCII only, as e.g. returned by *HMAC.hexdigest()*), or a *bytes-like object*.

注解: If *a* and *b* are of different lengths, or if an error occurs, a timing attack could theoretically reveal information about the types and lengths of *a* and *b*—but not their values.

3.3 新版功能.

参见:

Module `hashlib` The Python module providing secure hash functions.

15.3 `secrets` —Generate secure random numbers for managing secrets

3.6 新版功能.

Source code: [Lib/secrets.py](#)

The `secrets` module is used for generating cryptographically strong random numbers suitable for managing data such as passwords, account authentication, security tokens, and related secrets.

In particular, `secrets` should be used in preference to the default pseudo-random number generator in the `random` module, which is designed for modelling and simulation, not security or cryptography.

参见:

PEP 506

15.3.1 Random numbers

The `secrets` module provides access to the most secure source of randomness that your operating system provides.

class `secrets.SystemRandom`

A class for generating random numbers using the highest-quality sources provided by the operating system. See `random.SystemRandom` for additional details.

`secrets.choice(sequence)`

Return a randomly-chosen element from a non-empty sequence.

`secrets.randbelow(n)`

Return a random int in the range $[0, n)$.

`secrets.randbits(k)`

Return an int with *k* random bits.

15.3.2 Generating tokens

The `secrets` module provides functions for generating secure tokens, suitable for applications such as password resets, hard-to-guess URLs, and similar.

`secrets.token_bytes([nbytes=None])`

Return a random byte string containing *nbytes* number of bytes. If *nbytes* is `None` or not supplied, a reasonable default is used.

```
>>> token_bytes(16)
b'\xeb\r\x17D*t\xae\xd4\xe3S\xb6\xe2\xebP1\x8b'
```

`secrets.token_hex([nbytes=None])`

Return a random text string, in hexadecimal. The string has *nbytes* random bytes, each byte converted to two hex digits. If *nbytes* is `None` or not supplied, a reasonable default is used.

```
>>> token_hex(16)
'f9bf78b9a18ce6d46a0cd2b0b86df9da'
```

`secrets.token_urlsafe([nbytes=None])`

Return a random URL-safe text string, containing *nbytes* random bytes. The text is Base64 encoded, so on average each byte results in approximately 1.3 characters. If *nbytes* is `None` or not supplied, a reasonable default is used.

```
>>> token_urlsafe(16)
'Drmhze6EPcv0fN_81Bj-nA'
```

How many bytes should tokens use?

To be secure against [brute-force attacks](#), tokens need to have sufficient randomness. Unfortunately, what is considered sufficient will necessarily increase as computers get more powerful and able to make more guesses in a shorter period. As of 2015, it is believed that 32 bytes (256 bits) of randomness is sufficient for the typical use-case expected for the `secrets` module.

For those who want to manage their own token length, you can explicitly specify how much randomness is used for tokens by giving an `int` argument to the various `token_*` functions. That argument is taken as the number of bytes of randomness to use.

Otherwise, if no argument is provided, or if the argument is `None`, the `token_*` functions will use a reasonable default instead.

注解: That default is subject to change at any time, including during maintenance releases.

15.3.3 其他功能

`secrets.compare_digest(a, b)`

Return `True` if strings *a* and *b* are equal, otherwise `False`, in such a way as to reduce the risk of [timing attacks](#). See `hmac.compare_digest()` for additional details.

15.3.4 Recipes and best practices

This section shows recipes and best practices for using `secrets` to manage a basic level of security.

Generate an eight-character alphanumeric password:

```
import string
alphabet = string.ascii_letters + string.digits
password = ''.join(choice(alphabet) for i in range(8))
```

注解: Applications should not [store passwords in a recoverable format](#), whether plain text or encrypted. They should be salted and hashed using a cryptographically-strong one-way (irreversible) hash function.

Generate a ten-character alphanumeric password with at least one lowercase character, at least one uppercase character, and at least three digits:

```
import string
alphabet = string.ascii_letters + string.digits
while True:
    password = ''.join(choice(alphabet) for i in range(10))
    if (any(c.islower() for c in password)
        and any(c.isupper() for c in password)
        and sum(c.isdigit() for c in password) >= 3):
        break
```

Generate an XKCD-style passphrase:

```
# On standard Linux systems, use a convenient dictionary file.
# Other platforms may need to provide their own word-list.
with open('/usr/share/dict/words') as f:
    words = [word.strip() for word in f]
    password = ' '.join(choice(words) for i in range(4))
```

Generate a hard-to-guess temporary URL containing a security token suitable for password recovery applications:

```
url = 'https://mydomain.com/reset=' + token_urlsafe()
```

通用操作系统服务

本章中描述的各模块提供了在（几乎）所有的操作系统上可用的操作系统特性的接口，例如文件和时钟。这些接口通常以 Unix 或 C 接口为参照对象设计，不过在大多数其他系统上也可用。这里有一个概述：

16.1 `os` — 操作系统接口模块

源代码： [Lib/os.py](#)

该模块提供了一些方便使用操作系统相关功能的函数。如果你想读写一个文件，请参阅 `open()`，如果你想操作路径，请参阅 `os.path` 模块，如果你想在命令行上读取所有文件中的所有行请参阅 `fileinput` 模块。有关创建临时文件和目录的方法，请参阅 `tempfile` 模块，对于高级文件目录处理，请参阅 `shutil` 模块。

关于这些函数的可用性的说明：

- 所有 Python 内建的操作系统相关的模块的设计都是为了使得在同一功能可用的情况下，保持接口的一致性；例如，函数 `os.stat(path)` 以相同的格式返回关于 `path` 的统计信息（这个函数同时也是起源于 POSIX 接口）。
- 针对特定的操作的拓展同样在可用于 `os` 模块，但是使用它们必然会对可移植性产生威胁。
- 所有接受路径或文件名的函数都同时支持字节串和字符串对象，并在返回路径或文件名时使用相应类型的对象作为结果。
- An “Availability: Unix” note means that this function is commonly found on Unix systems. It does not make any claims about its existence on a specific operating system.
- If not separately noted, all functions that claim “Availability: Unix” are supported on Mac OS X, which builds on a Unix core.

注解：此模块的所有函数在遇到不可用或不可访问的文件名或路径，以及其他类型正确，但不被操作系统接受的参数时，会引发 `OSError` 异常。

exception `os.error`

内建的`OSError` 异常的一个别名。

`os.name`

导入的依赖特定操作系统的模块的名称。以下名称目前已注册: 'posix', 'nt', 'java'.

参见:

`sys.platform` 有更详细的描述. `os.uname()` 只给出系统提供的版本信息。

`platform` 模块对系统的标识有更详细的检查。

16.1.1 文件名，命令行参数，以及环境变量。

在 Python 中，使用字符串类型表示文件名、命令行参数和环境变量。在某些系统上，在将这些字符串传递给操作系统之前，必须将这些字符串解码为字节。Python 使用文件系统编码来执行此转换（请参阅`sys.getfilesystemencoding()`）。

在 3.1 版更改: 在某些系统上，使用文件系统编码进行转换可能会失败。在这种情况下，Python 会使用代理转义编码错误处理器，这意味着在解码时，不可解码的字节被 Unicode 字符 U+DCxx 替换，并且这些字节在编码时再次转换为原始字节。

文件系统编码必须保证成功解码小于 128 的所有字节。如果文件系统编码无法提供此保证，API 函数可能会引发 `UnicodeErrors`。

16.1.2 进程参数

这些函数和数据项提供了操作当前进程和用户的信息。

`os.ctermid()`

返回与进程控制终端对应的文件名。

Availability: Unix.

`os.environ`

一个表示字符串环境的 *mapping* 对象。例如，`environ['HOME']` 是你的主目录（在某些平台上）的路径名，相当于 C 中的 `getenv("HOME")`。

这个映射是在第一次导入 `os` 模块时捕获的，通常作为 Python 启动时处理 `site.py` 的一部分。除了通过直接修改 `os.environ` 之外，在此之后对环境所做的更改不会反映在 `os.environ` 中。

如果平台支持 `putenv()` 函数，这个映射除了可以用于查询环境外还能用于修改环境。当这个映射被修改时，`putenv()` 将被自动调用。

在 Unix 系统上，键和值会使用 `sys.getfilesystemencoding()` 和 'surrogateescape' 的错误处理。如果你想使用其他的编码，使用 `environb`。

注解: 直接调用 `putenv()` 并不会影响 `os.environ`，所以推荐直接修改“`os.environ`”。

注解: 在某些平台上，包括 FreeBSD 和 Mac OS X，设置 `environ` 可能导致内存泄露。参阅 `putenv()` 的系统文档。

如果平台没有提供 `putenv()`，为了使启动的子进程使用修改后的环境，一份修改后的映射会被传给合适的进程创建函数。

如果平台支持`unsetenv()` 函数, 你可以通过删除映射中元素的方式来删除对应的环境变量。当一个元素被从`os.environ` 删除时, 以及`pop()` 或`clear()` 被调用时, `unsetenv()` 会被自动调用。

`os.environb`

字节版本的`environ`: 一个以字节串表示环境的`mapping` 对象。`environ` 和`environb` 是同步的 (修改`environb` 会更新`environ`, 反之亦然)。

只有在`supports_bytes_environ` 为 `True` 的时候`environb` 才是可用的。

3.2 新版功能.

`os.chdir(path)`

`os.fchdir(fd)`

`os.getcwd()`

以上函数请参阅[Files and Directories](#) 。

`os.fsencode(filename)`

编码路径类 文件名为文件系统接受的形式, 使用 `'surrogateescape'` 代理转义编码错误处理器, 在 Windows 系统上会使用 `'strict'`; 返回`bytes` 字节类型不变。

`fsdecode()` 是此函数的逆向函数。

3.2 新版功能.

在 3.6 版更改: 增加对实现了`os.PathLike` 接口的对象的支持。

`os.fsdecode(filename)`

从文件系统编码方式解码为路径类 文件名, 使用 `'surrogateescape'` 代理转义编码错误处理器, 在 Windows 系统上会使用 `'strict'`; 返回`str` 字符串不变。

`fsencode()` 是此函数的逆向函数。

3.2 新版功能.

在 3.6 版更改: 增加对实现了`os.PathLike` 接口的对象的支持。

`os.fspath(path)`

返回路径的文件系统表示。

如果传入的是`str` 或`bytes` 类型的字符串, 将原样返回。否则`__fspath__()` 将被调用, 如果得到的是一个`str` 或`bytes` 类型的对象, 那就返回这个值。其他所有情况则会抛出`TypeError` 异常。

3.6 新版功能.

`class os.PathLike`

描述表示一个文件系统路径的`abstract base class` , 如`pathlib.PurePath` 。

3.6 新版功能.

`abstractmethod __fspath__()`

返回当前对象的文件系统表示。

这个方法只应该返回一个`str` 字符串或`bytes` 字节串, 请优先选择`str` 字符串。

`os.getenv(key, default=None)`

如果存在, 返回环境变量 `key` 的值, 否则返回 `default`。 `key` , `default` 和返回值均为 `str` 字符串类型。

在 Unix 系统上, 键和值会使用`sys.getfilesystemencoding()` 和 `'surrogateescape'` “ 错误处理进行解码。如果你想使用其他的编码, 使用`os.getenvb()` 。

Availability: most flavors of Unix, Windows.

`os.getenvb(key, default=None)`

如果存在环境变量 `key` 那么返回其值, 否则返回 `default`。 `key` , `default` 和返回值均为 `bytes` 字节串类型。

`getenvb()` 仅在`supports_bytes_environ` 为 `True` 时可用

Availability: most flavors of Unix.

3.2 新版功能.

`os.get_exec_path(env=None)`

返回将用于搜索可执行文件的目录列表，与在外壳程序中启动一个进程时相似。指定的 *env* 应为用于搜索 PATH 的环境变量字典。默认情况下，当 *env* 为 None 时，将会使用 *environ*。

3.2 新版功能.

`os.getegid()`

返回当前进程的有效组 ID。对应当前进程执行文件的 “set id” 位。

Availability: Unix.

`os.geteuid()`

返回当前进程的有效用户 ID。

Availability: Unix.

`os.getgid()`

返回当前进程的实际组 ID。

Availability: Unix.

`os.getgrouplist(user, group)`

返回该用户所在的组 ID 列表。可能 *group* 参数没有在返回的列表中，实际上用户应该也是属于该 *group*。*group* 参数一般可以从储存账户信息的密码记录文件中找到。

Availability: Unix.

3.3 新版功能.

`os.getgroups()`

返回当前进程对应的组 ID 列表

Availability: Unix.

注解： 在 Mac OS X 系统中，*getgroups()* 会和其他 Unix 平台有些不同。如果 Python 解释器是在 10.5 或更早版本中部署，*getgroups()* 返回当前用户进程相关的有效组 ID 列表。该列表长度由于系统预设的接口限制，最长为 16。而且在适当的权限下，返回结果还会因 *getgroups()* 而发生变化；如果 Python 解释器是在 10.5 以上版本中部署，*getgroups()* 返回进程所属有效用户 ID 所对应的用户的组 ID 列表，组用户列表可能因为进程的生存周期而发生变动，而且也不会因为 *setgroups()* 的调用而发生，返回的组用户列表长度也没有长度 16 的限制。在部署中，Python 解释器用到的变量 `MACOSX_DEPLOYMENT_TARGET` 可以用 *sysconfig.get_config_var()*。

`os.getlogin()`

返回通过控制终端进程进行登录的用户名。在多数情况下，使用 *getpass.getuser()* 会更有效，因为后者会通过检查环境变量 LOGNAME 或 USERNAME 来查找用户，再由 *pwd.getpwuid(os.getuid())[0]* 来获取当前用户 ID 的登录名。

Availability: Unix, Windows.

`os.getpgid(pid)`

根据进程 id *pid* 返回进程的组 ID 列表。如果 *pid* 为 0，则返回当前进程的进程组 ID 列表

Availability: Unix.

`os.getpgrp()`

返回当时进程组的 ID

Availability: Unix.

`os.getpid()`

返回当前进程 ID

`os.getppid()`

返回父进程 ID。当父进程已经结束，在 Unix 中返回的 ID 是初始进程 (1) 中的一个，在 Windows 中仍然是同一个进程 ID，该进程 ID 有可能已经被进行进程所占用。

Availability: Unix, Windows.

在 3.2 版更改: 添加 Windows 的支持。

`os.getpriority(which, who)`

获取程序调度优先级。*which* 参数值可以是 `PRIO_PROCESS`、`PRIO_PGRP`，或 `PRIO_USER` 中的一个，*who* 是相对于 *which* (`PRIO_PROCESS` 的进程标识符，`PRIO_PGRP` 的进程组标识符和 `PRIO_USER` 的用户 ID)。当 *who* 为 0 时（分别）表示调用的进程，调用进程的进程组或调用进程所属的真实用户 ID。

Availability: Unix.

3.3 新版功能.

`os.PRIO_PROCESS`

`os.PRIO_PGRP`

`os.PRIO_USER`

函数 `getpriority()` 和 `setpriority()` 的参数。

Availability: Unix.

3.3 新版功能.

`os.getresuid()`

返回一个由 (ruid, euid, suid) 所组成的元组，分别表示当前进程的真实用户 ID，有效用户 ID 和甲暂存用户 ID。

Availability: Unix.

3.2 新版功能.

`os.getresgid()`

返回一个由 (rgid, egid, sgid) 所组成的元组，分别表示当前进程的真实组 ID，有效组 ID 和暂存组 ID。

Availability: Unix.

3.2 新版功能.

`os.getuid()`

返回当前进程的真实用户 ID。

Availability: Unix.

`os.initgroups(username, gid)`

调用系统 `initgroups()`，使用指定用户所在的所有值来初始化组访问列表，包括指定的组 ID。

Availability: Unix.

3.2 新版功能.

`os.putenv(key, value)`

将名为 *key* 的环境变量值设置为 *value*。该变量名修改会影响由 `os.system()`、`popen()`、`fork()` 和 `execv()` 发起的子进程。

Availability: most flavors of Unix, Windows.

注解： 在一些平台，包括 FreeBSD 和 Mac OS X，设置 `environ` 可能导致内存泄露。详情参考 `putenv` 相关系统文档。

当系统支持 `putenv()` 时，`os.environ` 中的参数赋值会自动转换为对 `putenv()` 的调用。不过 `putenv()` 的调用不会更新 `os.environ`，因此最好使用 `os.environ` 对变量赋值。

`os.setegid(egid)`

设置当前进程的有效组 ID。

Availability: Unix.

`os.seteuid(euid)`

设置当前进程的有效用户 ID。

Availability: Unix.

`os.setgid(gid)`

设置当前进程的组 ID。

Availability: Unix.

`os.setgroups(groups)`

将 `group` 参数值设置为与当进程相关联的附加组 ID 列表。`group` 参数必须为一个序列，每个元素应为每个组的数字 ID。该操作通常只适用于超级用户。

Availability: Unix.

注解： 在 Mac OS X 中，`groups` 的长度不能超过系统定义的最大有效组 ID 个数，一般为 16。如果它没有返回与调用 `setgroups()` 所设置的相同的组列表，请参阅 `getgroups()` 的文档。

`os.setpgrp()`

根据已实现的版本（如果有）来调用系统 `setpgrp()` 或 `setpgrp(0, 0)`。相关说明，请参考 Unix 手册。

Availability: Unix.

`os.setpgid(pid, pgrp)`

使用系统调用 `setpgid()`，将 `pid` 对应进程的组 ID 设置为 `pgrp`。相关说明，请参考 Unix 手册。

Availability: Unix.

`os.setpriority(which, who, priority)`

设置程序调度优先级。`which` 的值为 `PRIO_PROCESS`、`PRIO_PGRP` 或 `PRIO_USER` 之一，而 `who` 会相对于 `which` (`PRIO_PROCESS` 的进程标识符，`PRIO_PGRP` 的进程组标识符和 `PRIO_USER` 的用户 ID) 被解析。`who` 值为零 (分别) 表示调用进程，调用进程的进程组或调用进程的真实用户 ID。`priority` 是范围在 -20 至 19 的值。默认优先级为 0；较小的优先级数值会更优先被调度。

Availability: Unix

3.3 新版功能.

`os.setregid(rgid, egid)`

设置当前进程的真实和有效组 ID。

Availability: Unix.

`os.setresgid(rgid, egid, sgid)`

设置当前进程的真实，有效和暂存组 ID。

Availability: Unix.

3.2 新版功能.

`os.setresuid(ruid, euid, suid)`

设置当前进程的真实，有效和暂存用户 ID。

Availability: Unix.

3.2 新版功能.

`os.setreuid(ruid, euid)`

设置当前进程的真实和有效用户 ID。

Availability: Unix.

`os.getsid(pid)`

调用系统调用 `getsid()`。相关语义请参阅 Unix 手册。

Availability: Unix.

`os.setsid()`

使用系统调用 `getsid()`。相关说明，请参考 Unix 手册。

Availability: Unix.

`os.setuid(uid)`

设置当前进程的用户 ID。

Availability: Unix.

`os.strerror(code)`

根据 `code` 中的错误码返回错误消息。在某些平台上当给出未知错误码时 `strerror()` 将返回 `NULL` 并会引发 `ValueError`。

`os.supports_bytes_environ`

如果操作系统上原生环境类型是字节型则为 `True` (例如在 Windows 上为 `False`)。

3.2 新版功能.

`os.umask(mask)`

设定当前数值掩码并返回之前的掩码。

`os.uname()`

返回当前操作系统的识别信息。返回值是一个有 5 个属性的对象：

- `sysname` - 操作系统名
- `nodename` - 机器在网络上的名称（需要先设定）
- `release` - 操作系统发行信息
- `version` - 操作系统版本信息
- `machine` - 硬件标识符

为了向后兼容，该对象也是可迭代的，像是一个按照 `sysname`, `nodename`, `release`, `version`, 和 `machine` 顺序组成的元组。

有些系统会将 `nodename` 截短为 8 个字符或截短至前缀部分；获取主机名的一个更好方式是 `socket.gethostname()` 或甚至可以用 `socket.gethostbyaddr(socket.gethostname())`。

Availability: recent flavors of Unix.

在 3.3 版更改: 返回结果的类型由元组变成一个类似元组的对象，同时具有命名的属性。

`os.unsetenv(key)`

取消设置（删除）名为 `key` 的环境变量。变量名的改变会影响由 `os.system()`, `popen()`, `fork()` 和 `execv()` 触发的子进程。

当系统支持`unsetenv()`，删除在`os.environ`中的变量会自动转换为对`unsetenv()`的调用。但是`unsetenv()`不能更新`os.environ`，因此最好直接删除`os.environ`中的变量。

Availability: most flavors of Unix, Windows.

16.1.3 创建文件对象

该函数创建新的文件对象。（另见`open()`关于打开文件的说明。）

`os.fdupen(fd, *args, **kwargs)`

返回打开文件描述符`fd`对应文件的对象。类似内建`open()`函数，二者接受同样的参数。不同之处在于`fdopen()`第一个参数应该为整数。

16.1.4 文件描述符操作

这些函数对文件描述符所引用的 I/O 流进行操作。

文件描述符是一些小的整数，对应于当前进程所打开的文件。例如，标准输入的文件描述符通常是 0，标准输出是 1，标准错误是 2。之后被进程打开的文件的文件描述符会被依次指定为 3, 4, 5 等。“文件描述符”这个词有点误导性，在 Unix 平台中套接字和管道也被文件描述符所引用。

当需要时，可以用`fileno()`可以获得`file object`所对应的文件描述符。需要注意的是，直接使用文件描述符会绕过文件对象的方法，会忽略如数据内部缓冲等情况。

`os.close(fd)`

关闭文件描述符`fd`。

注解：该功能适用于低级 I/O 操作，必须用于`os.open()`或`pipe()`返回的文件描述符。关闭由内建函数`open()`，`popen()`或`fdopen()`返回的“文件对象”，则使用其相应的`close()`方法。

`os.closerange(fd_low, fd_high)`

关闭从`fd_low`（包括）到`fd_high`（排除）间的文件描述符，并忽略错误。类似（但快于）：

```
for fd in range(fd_low, fd_high):
    try:
        os.close(fd)
    except OSError:
        pass
```

`os.device_encoding(fd)`

如果连接到终端，则返回一个与`fd`关联的设备描述字符，否则返回`None`。

`os.dup(fd)`

返回一个文件描述符`fd`的副本。该文件描述符的副本是**不可继承的**。

在 Windows 中，当复制一个标准流（0: stdin, 1: stdout, 2: stderr）时，新的文件描述符是**可继承的**。

在 3.4 版更改：新的文件描述符现在是**不可继承的**。

`os.dup2(fd, fd2, inheritable=True)`

Duplicate file descriptor `fd` to `fd2`, closing the latter first if necessary. The file descriptor `fd2` is *inheritable* by default, or non-inheritable if *inheritable* is `False`.

在 3.4 版更改：添加可选参数 *inheritable*。

os.fchmod(*fd*, *mode*)

将 *fd* 指定文件的权限状态修改为 *mode*。可以参考 `chmod()` 中列出 *mode* 的可用值。从 Python 3.3 开始, 这相当于 `os.chmod(fd, mode)`。

Availability: Unix.

os.fchown(*fd*, *uid*, *gid*)

分别将 *fd* 指定文件的所有者和组 ID 修改为 *uid* 和 *gid* 的值。若不想变更其中的某个 ID, 可将相应值设为 -1。参考 `chown()`。从 Python 3.3 开始, 这相当于 `os.chown(fd, uid, gid)`。

Availability: Unix.

os.fdatasync(*fd*)

强制将文件描述符 *fd* 指定文件写入磁盘。不强制更新元数据。

Availability: Unix.

注解: 该功能在 MacOS 中不可用。

os.fpathconf(*fd*, *name*)

Return system configuration information relevant to an open file. *name* specifies the configuration value to retrieve; it may be a string which is the name of a defined system value; these names are specified in a number of standards (POSIX.1, Unix 95, Unix 98, and others). Some platforms define additional names as well. The names known to the host operating system are given in the `pathconf_names` dictionary. For configuration variables not included in that mapping, passing an integer for *name* is also accepted.

If *name* is a string and is not known, `ValueError` is raised. If a specific value for *name* is not supported by the host system, even if it is included in `pathconf_names`, an `OSError` is raised with `errno.EINVAL` for the error number.

从 Python 3.3 起, 此功能等价于 `os.pathconf(fd, name)`。

Availability: Unix.

os.fstat(*fd*)

获取文件描述符 *fd* 的状态。返回一个 `stat_result` 对象。

从 Python 3.3 起, 此功能等价于 `os.stat(fd)`。

参见:

`stat()` 函数。

os.fstatvfs(*fd*)

Return information about the filesystem containing the file associated with file descriptor *fd*, like `statvfs()`. As of Python 3.3, this is equivalent to `os.statvfs(fd)`.

Availability: Unix.

os.fsync(*fd*)

Force write of file with filedescriptor *fd* to disk. On Unix, this calls the native `fsync()` function; on Windows, the `MS_commit()` function.

If you're starting with a buffered Python *file object* *f*, first do `f.flush()`, and then do `os.fsync(f.fileno())`, to ensure that all internal buffers associated with *f* are written to disk.

Availability: Unix, Windows.

os.ftruncate(*fd*, *length*)

Truncate the file corresponding to file descriptor *fd*, so that it is at most *length* bytes in size. As of Python 3.3, this is equivalent to `os.truncate(fd, length)`.

Availability: Unix, Windows.

在 3.5 版更改: 添加了 Windows 支持

`os.get_blocking(fd)`

Get the blocking mode of the file descriptor: `False` if the `O_NONBLOCK` flag is set, `True` if the flag is cleared.

See also `set_blocking()` and `socket.socket.setblocking()`.

Availability: Unix.

3.5 新版功能.

`os.isatty(fd)`

Return `True` if the file descriptor `fd` is open and connected to a tty(-like) device, else `False`.

`os.lockf(fd, cmd, len)`

Apply, test or remove a POSIX lock on an open file descriptor. `fd` is an open file descriptor. `cmd` specifies the command to use - one of `F_LOCK`, `F_TLOCK`, `F_ULOCK` or `F_TEST`. `len` specifies the section of the file to lock.

Availability: Unix.

3.3 新版功能.

`os.F_LOCK`

`os.F_TLOCK`

`os.F_ULOCK`

`os.F_TEST`

Flags that specify what action `lockf()` will take.

Availability: Unix.

3.3 新版功能.

`os.lseek(fd, pos, how)`

Set the current position of file descriptor `fd` to position `pos`, modified by `how`: `SEEK_SET` or 0 to set the position relative to the beginning of the file; `SEEK_CUR` or 1 to set it relative to the current position; `SEEK_END` or 2 to set it relative to the end of the file. Return the new cursor position in bytes, starting from the beginning.

`os.SEEK_SET`

`os.SEEK_CUR`

`os.SEEK_END`

Parameters to the `lseek()` function. Their values are 0, 1, and 2, respectively.

3.3 新版功能: Some operating systems could support additional values, like `os.SEEK_HOLE` or `os.SEEK_DATA`.

`os.open(path, flags, mode=0o777, *, dir_fd=None)`

Open the file `path` and set various flags according to `flags` and possibly its mode according to `mode`. When computing `mode`, the current umask value is first masked out. Return the file descriptor for the newly opened file. The new file descriptor is *non-inheritable*.

For a description of the flag and mode values, see the C run-time documentation; flag constants (like `O_RDONLY` and `O_WRONLY`) are defined in the `os` module. In particular, on Windows adding `O_BINARY` is needed to open files in binary mode.

This function can support *paths relative to directory descriptors* with the `dir_fd` parameter.

在 3.4 版更改: 新的文件描述符现在是不可继承的。

注解: This function is intended for low-level I/O. For normal usage, use the built-in function `open()`, which returns a *file object* with `read()` and `write()` methods (and many more). To wrap a file descriptor in a file

object, use `fdopen()`.

3.3 新版功能: `dir_fd` 参数。

在 3.5 版更改: 如果系统调用被中断, 但信号处理程序没有触发异常, 此函数现在会重试系统调用, 而不是触发 `InterruptedError` 异常 (原因详见 [PEP 475](#))。

在 3.6 版更改: 接受一个类路径对象。

The following constants are options for the `flags` parameter to the `open()` function. They can be combined using the bitwise OR operator `|`. Some of them are not available on all platforms. For descriptions of their availability and use, consult the `open(2)` manual page on Unix or [the MSDN](#) on Windows.

```
os.O_RDONLY
os.O_WRONLY
os.O_RDWR
os.O_APPEND
os.O_CREAT
os.O_EXCL
os.O_TRUNC
```

The above constants are available on Unix and Windows.

```
os.O_DSYNC
os.O_RSYNC
os.O_SYNC
os.O_NDELAY
os.O_NONBLOCK
os.O_NOCTTY
os.O_CLOEXEC
```

这个常数仅在 Unix 系统中可用。

在 3.3 版更改: Add `O_CLOEXEC` constant.

```
os.O_BINARY
os.O_NOINHERIT
os.O_SHORT_LIVED
os.O_TEMPORARY
os.O_RANDOM
os.O_SEQUENTIAL
os.O_TEXT
```

这个常数仅在 Windows 系统中可用。

```
os.O_ASYNC
os.O_DIRECT
os.O_DIRECTORY
os.O_NOFOLLOW
os.O_NOATIME
os.O_PATH
os.O_TMPFILE
os.O_SHLOCK
os.O_EXLOCK
```

The above constants are extensions and not present if they are not defined by the C library.

在 3.4 版更改: Add `O_PATH` on systems that support it. Add `O_TMPFILE`, only available on Linux Kernel 3.11 or newer.

```
os.openpty()
```

Open a new pseudo-terminal pair. Return a pair of file descriptors (`master`, `slave`) for the pty and the tty,

respectively. The new file descriptors are *non-inheritable*. For a (slightly) more portable approach, use the *pty* module.

Availability: some flavors of Unix.

在 3.4 版更改: 新的文件描述符不再可继承。

os.**pipe** ()

Create a pipe. Return a pair of file descriptors (*r*, *w*) usable for reading and writing, respectively. The new file descriptor is *non-inheritable*.

Availability: Unix, Windows.

在 3.4 版更改: 新的文件描述符不再可继承。

os.**pipe2** (*flags*)

Create a pipe with *flags* set atomically. *flags* can be constructed by ORing together one or more of these values: *O_NONBLOCK*, *O_CLOEXEC*. Return a pair of file descriptors (*r*, *w*) usable for reading and writing, respectively.

Availability: some flavors of Unix.

3.3 新版功能.

os.**posix_fallocate** (*fd*, *offset*, *len*)

Ensures that enough disk space is allocated for the file specified by *fd* starting from *offset* and continuing for *len* bytes.

Availability: Unix.

3.3 新版功能.

os.**posix_fadvise** (*fd*, *offset*, *len*, *advice*)

Announces an intention to access data in a specific pattern thus allowing the kernel to make optimizations. The advice applies to the region of the file specified by *fd* starting at *offset* and continuing for *len* bytes. *advice* is one of *POSIX_FADV_NORMAL*, *POSIX_FADV_SEQUENTIAL*, *POSIX_FADV_RANDOM*, *POSIX_FADV_NOREUSE*, *POSIX_FADV_WILLNEED* or *POSIX_FADV_DONTNEED*.

Availability: Unix.

3.3 新版功能.

os.**POSIX_FADV_NORMAL**

os.**POSIX_FADV_SEQUENTIAL**

os.**POSIX_FADV_RANDOM**

os.**POSIX_FADV_NOREUSE**

os.**POSIX_FADV_WILLNEED**

os.**POSIX_FADV_DONTNEED**

Flags that can be used in *advice* in *posix_fadvise()* that specify the access pattern that is likely to be used.

Availability: Unix.

3.3 新版功能.

os.**pread** (*fd*, *bufferize*, *offset*)

Read from a file descriptor, *fd*, at a position of *offset*. It will read up to *bufferize* number of bytes. The file offset remains unchanged.

Availability: Unix.

3.3 新版功能.

os.**pwrite** (*fd*, *str*, *offset*)

Write *bytestring* to a file descriptor, *fd*, from *offset*, leaving the file offset unchanged.

Availability: Unix.

3.3 新版功能.

`os.read(fd, n)`

Read at most *n* bytes from file descriptor *fd*. Return a bytearray containing the bytes read. If the end of the file referred to by *fd* has been reached, an empty bytes object is returned.

注解: This function is intended for low-level I/O and must be applied to a file descriptor as returned by `os.open()` or `pipe()`. To read a “file object” returned by the built-in function `open()` or by `popen()` or `fdopen()`, or `sys.stdin`, use its `read()` or `readline()` methods.

在 3.5 版更改: 如果系统调用被中断, 但信号处理程序没有触发异常, 此函数现在会重试系统调用, 而不是触发 `InterruptedError` 异常 (原因详见 [PEP 475](#)).

`os.sendfile(out, in, offset, count)`

`os.sendfile(out, in, offset, count[, headers][, trailers], flags=0)`

Copy *count* bytes from file descriptor *in* to file descriptor *out* starting at *offset*. Return the number of bytes sent. When EOF is reached return 0.

The first function notation is supported by all platforms that define `sendfile()`.

On Linux, if *offset* is given as `None`, the bytes are read from the current position of *in* and the position of *in* is updated.

The second case may be used on Mac OS X and FreeBSD where *headers* and *trailers* are arbitrary sequences of buffers that are written before and after the data from *in* is written. It returns the same as the first case.

On Mac OS X and FreeBSD, a value of 0 for *count* specifies to send until the end of *in* is reached.

All platforms support sockets as *out* file descriptor, and some platforms allow other types (e.g. regular file, pipe) as well.

Cross-platform applications should not use *headers*, *trailers* and *flags* arguments.

Availability: Unix.

注解: For a higher-level wrapper of `sendfile()`, see `socket.socket.sendfile()`.

3.3 新版功能.

`os.set_blocking(fd, blocking)`

Set the blocking mode of the specified file descriptor. Set the `O_NONBLOCK` flag if *blocking* is `False`, clear the flag otherwise.

See also `get_blocking()` and `socket.socket.setblocking()`.

Availability: Unix.

3.5 新版功能.

`os.SF_NODISKIO`

`os.SF_MNOWAIT`

`os.SF_SYNC`

Parameters to the `sendfile()` function, if the implementation supports them.

Availability: Unix.

3.3 新版功能.

os.readv(*fd*, *buffers*)

Read from a file descriptor *fd* into a number of mutable *bytes-like objects buffers*. *readv()* will transfer data into each buffer until it is full and then move on to the next buffer in the sequence to hold the rest of the data. *readv()* returns the total number of bytes read (which may be less than the total capacity of all the objects).

Availability: Unix.

3.3 新版功能.

os.tcgetpgrp(*fd*)

Return the process group associated with the terminal given by *fd* (an open file descriptor as returned by *os.open()*).

Availability: Unix.

os.tcsetpgrp(*fd*, *pg*)

Set the process group associated with the terminal given by *fd* (an open file descriptor as returned by *os.open()*) to *pg*.

Availability: Unix.

os.ttyname(*fd*)

Return a string which specifies the terminal device associated with file descriptor *fd*. If *fd* is not associated with a terminal device, an exception is raised.

Availability: Unix.

os.write(*fd*, *str*)

Write the bytestring in *str* to file descriptor *fd*. Return the number of bytes actually written.

注解: This function is intended for low-level I/O and must be applied to a file descriptor as returned by *os.open()* or *pipe()*. To write a “file object” returned by the built-in function *open()* or by *popen()* or *fdopen()*, or *sys.stdout* or *sys.stderr*, use its *write()* method.

在 3.5 版更改: 如果系统调用被中断, 但信号处理程序没有触发异常, 此函数现在会重试系统调用, 而不是触发 *InterruptedError* 异常 (原因详见 [PEP 475](#)).

os.writev(*fd*, *buffers*)

Write the contents of *buffers* to file descriptor *fd*. *buffers* must be a sequence of *bytes-like objects*. Buffers are processed in array order. Entire contents of first buffer is written before proceeding to second, and so on. The operating system may set a limit (*sysconf()* value *SC_IOV_MAX*) on the number of buffers that can be used.

writev() writes the contents of each object to the file descriptor and returns the total number of bytes written.

Availability: Unix.

3.3 新版功能.

Querying the size of a terminal

3.3 新版功能.

os.get_terminal_size(*fd=STDOUT_FILENO*)

Return the size of the terminal window as (*columns*, *lines*), tuple of type *terminal_size*.

The optional argument *fd* (default *STDOUT_FILENO*, or standard output) specifies which file descriptor should be queried.

If the file descriptor is not connected to a terminal, an *OSError* is raised.

`shutil.get_terminal_size()` is the high-level function which should normally be used, `os.get_terminal_size` is the low-level implementation.

Availability: Unix, Windows.

class `os.terminal_size`

A subclass of tuple, holding (`columns`, `lines`) of the terminal window size.

`columns`

Width of the terminal window in characters.

`lines`

Height of the terminal window in characters.

Inheritance of File Descriptors

3.4 新版功能.

A file descriptor has an “inheritable” flag which indicates if the file descriptor can be inherited by child processes. Since Python 3.4, file descriptors created by Python are non-inheritable by default.

On UNIX, non-inheritable file descriptors are closed in child processes at the execution of a new program, other file descriptors are inherited.

On Windows, non-inheritable handles and file descriptors are closed in child processes, except for standard streams (file descriptors 0, 1 and 2: `stdin`, `stdout` and `stderr`), which are always inherited. Using `spawn*` functions, all inheritable handles and all inheritable file descriptors are inherited. Using the `subprocess` module, all file descriptors except standard streams are closed, and inheritable handles are only inherited if the `close_fds` parameter is `False`.

`os.get_inheritable(fd)`

Get the “inheritable” flag of the specified file descriptor (a boolean).

`os.set_inheritable(fd, inheritable)`

Set the “inheritable” flag of the specified file descriptor.

`os.get_handle_inheritable(handle)`

Get the “inheritable” flag of the specified handle (a boolean).

Availability: Windows.

`os.set_handle_inheritable(handle, inheritable)`

Set the “inheritable” flag of the specified handle.

Availability: Windows.

16.1.5 Files and Directories

On some Unix platforms, many of these functions support one or more of these features:

- **specifying a file descriptor:** For some functions, the `path` argument can be not only a string giving a path name, but also a file descriptor. The function will then operate on the file referred to by the descriptor. (For POSIX systems, Python will call the `f...f` version of the function.)

You can check whether or not `path` can be specified as a file descriptor on your platform using `os.supports_fd`. If it is unavailable, using it will raise a `NotImplementedError`.

If the function also supports `dir_fd` or `follow_symlinks` arguments, it is an error to specify one of those when supplying `path` as a file descriptor.

- **paths relative to directory descriptors:** If `dir_fd` is not `None`, it should be a file descriptor referring to a directory, and the path to operate on should be relative; path will then be relative to that directory. If the path is absolute, `dir_fd` is ignored. (For POSIX systems, Python will call the `...at` or `f...at` version of the function.)

You can check whether or not `dir_fd` is supported on your platform using `os.supports_dir_fd`. If it is unavailable, using it will raise a `NotImplementedError`.

- **not following symlinks:** If `follow_symlinks` is `False`, and the last element of the path to operate on is a symbolic link, the function will operate on the symbolic link itself instead of the file the link points to. (For POSIX systems, Python will call the `l...at` version of the function.)

You can check whether or not `follow_symlinks` is supported on your platform using `os.supports_follow_symlinks`. If it is unavailable, using it will raise a `NotImplementedError`.

`os.access(path, mode, *, dir_fd=None, effective_ids=False, follow_symlinks=True)`

Use the real uid/gid to test for access to `path`. Note that most operations will use the effective uid/gid, therefore this routine can be used in a `suid/sgid` environment to test if the invoking user has the specified access to `path`. `mode` should be `F_OK` to test the existence of `path`, or it can be the inclusive OR of one or more of `R_OK`, `W_OK`, and `X_OK` to test permissions. Return `True` if access is allowed, `False` if not. See the Unix man page `access(2)` for more information.

This function can support specifying *paths relative to directory descriptors* and *not following symlinks*.

If `effective_ids` is `True`, `access()` will perform its access checks using the effective uid/gid instead of the real uid/gid. `effective_ids` may not be supported on your platform; you can check whether or not it is available using `os.supports_effective_ids`. If it is unavailable, using it will raise a `NotImplementedError`.

注解: Using `access()` to check if a user is authorized to e.g. open a file before actually doing so using `open()` creates a security hole, because the user might exploit the short time interval between checking and opening the file to manipulate it. It's preferable to use *EAFP* techniques. For example:

```
if os.access("myfile", os.R_OK):
    with open("myfile") as fp:
        return fp.read()
return "some default data"
```

is better written as:

```
try:
    fp = open("myfile")
except PermissionError:
    return "some default data"
else:
    with fp:
        return fp.read()
```

注解: I/O operations may fail even when `access()` indicates that they would succeed, particularly for operations on network filesystems which may have permissions semantics beyond the usual POSIX permission-bit model.

在 3.3 版更改: Added the `dir_fd`, `effective_ids`, and `follow_symlinks` parameters.

在 3.6 版更改: 接受一个类路径对象。

`os.F_OK`

`os.R_OK`

`os.W_OK`

os.X_OK

Values to pass as the *mode* parameter of `access()` to test the existence, readability, writability and executability of *path*, respectively.

os.chdir(*path*)

Change the current working directory to *path*.

This function can support *specifying a file descriptor*. The descriptor must refer to an opened directory, not an open file.

3.3 新版功能: Added support for specifying *path* as a file descriptor on some platforms.

在 3.6 版更改: 接受一个类路径对象。

os.chflags(*path*, *flags*, *, *follow_symlinks*=True)

Set the flags of *path* to the numeric *flags*. *flags* may take a combination (bitwise OR) of the following values (as defined in the `stat` module):

- `stat.UF_NODUMP`
- `stat.UF_IMMUTABLE`
- `stat.UF_APPEND`
- `stat.UF_OPAQUE`
- `stat.UF_NOUNLINK`
- `stat.UF_COMPRESSED`
- `stat.UF_HIDDEN`
- `stat.SF_ARCHIVED`
- `stat.SF_IMMUTABLE`
- `stat.SF_APPEND`
- `stat.SF_NOUNLINK`
- `stat.SF_SNAPSHOT`

This function can support *not following symlinks*.

Availability: Unix.

3.3 新版功能: The *follow_symlinks* argument.

在 3.6 版更改: 接受一个类路径对象。

os.chmod(*path*, *mode*, *, *dir_fd*=None, *follow_symlinks*=True)

Change the mode of *path* to the numeric *mode*. *mode* may take one of the following values (as defined in the `stat` module) or bitwise ORed combinations of them:

- `stat.S_ISUID`
- `stat.S_ISGID`
- `stat.S_ENFMT`
- `stat.S_ISVTX`
- `stat.S_IREAD`
- `stat.S_IWRITE`
- `stat.S_IEXEC`
- `stat.S_IRWXU`

- `stat.S_IRUSR`
- `stat.S_IWUSR`
- `stat.S_IXUSR`
- `stat.S_IRWXG`
- `stat.S_IRGRP`
- `stat.S_IWGRP`
- `stat.S_IXGRP`
- `stat.S_IRWXO`
- `stat.S_IROTH`
- `stat.S_IWOTH`
- `stat.S_IXOTH`

This function can support *specifying a file descriptor, paths relative to directory descriptors* and *not following symlinks*.

注解: Although Windows supports `chmod()`, you can only set the file's read-only flag with it (via the `stat.S_IWRITE` and `stat.S_IREAD` constants or a corresponding integer value). All other bits are ignored.

3.3 新版功能: Added support for specifying *path* as an open file descriptor, and the *dir_fd* and *follow_symlinks* arguments.

在 3.6 版更改: 接受一个类路径对象。

○ `os.chown(path, uid, gid, *, dir_fd=None, follow_symlinks=True)`

Change the owner and group id of *path* to the numeric *uid* and *gid*. To leave one of the ids unchanged, set it to -1.

This function can support *specifying a file descriptor, paths relative to directory descriptors* and *not following symlinks*.

See `shutil.chown()` for a higher-level function that accepts names in addition to numeric ids.

Availability: Unix.

3.3 新版功能: Added support for specifying an open file descriptor for *path*, and the *dir_fd* and *follow_symlinks* arguments.

在 3.6 版更改: Supports a *path-like object*.

○ `os.chroot(path)`

Change the root directory of the current process to *path*.

Availability: Unix.

在 3.6 版更改: 接受一个类路径对象。

○ `os.fchdir(fd)`

Change the current working directory to the directory represented by the file descriptor *fd*. The descriptor must refer to an opened directory, not an open file. As of Python 3.3, this is equivalent to `os.chdir(fd)`.

Availability: Unix.

○ `os.getcwd()`

Return a string representing the current working directory.

`os.getcwd()`

Return a bytestring representing the current working directory.

`os.lchflags(path, flags)`

Set the flags of *path* to the numeric *flags*, like `chflags()`, but do not follow symbolic links. As of Python 3.3, this is equivalent to `os.chflags(path, flags, follow_symlinks=False)`.

Availability: Unix.

在 3.6 版更改: 接受一个类路径对象。

`os.lchmod(path, mode)`

Change the mode of *path* to the numeric *mode*. If *path* is a symlink, this affects the symlink rather than the target. See the docs for `chmod()` for possible values of *mode*. As of Python 3.3, this is equivalent to `os.chmod(path, mode, follow_symlinks=False)`.

Availability: Unix.

在 3.6 版更改: 接受一个类路径对象。

`os.lchown(path, uid, gid)`

Change the owner and group id of *path* to the numeric *uid* and *gid*. This function will not follow symbolic links. As of Python 3.3, this is equivalent to `os.chown(path, uid, gid, follow_symlinks=False)`.

Availability: Unix.

在 3.6 版更改: 接受一个类路径对象。

`os.link(src, dst, *, src_dir_fd=None, dst_dir_fd=None, follow_symlinks=True)`

Create a hard link pointing to *src* named *dst*.

This function can support specifying *src_dir_fd* and/or *dst_dir_fd* to supply *paths relative to directory descriptors*, and *not following symlinks*.

Availability: Unix, Windows.

在 3.2 版更改: Added Windows support.

3.3 新版功能: Added the *src_dir_fd*, *dst_dir_fd*, and *follow_symlinks* arguments.

在 3.6 版更改: Accepts a *path-like object* for *src* and *dst*.

`os.listdir(path='.')`

Return a list containing the names of the entries in the directory given by *path*. The list is in arbitrary order, and does not include the special entries `'.'` and `'..'` even if they are present in the directory.

path may be a *path-like object*. If *path* is of type `bytes` (directly or indirectly through the *PathLike* interface), the filenames returned will also be of type `bytes`; in all other circumstances, they will be of type `str`.

This function can also support *specifying a file descriptor*; the file descriptor must refer to a directory.

注解: To encode `str` filenames to `bytes`, use `fsencode()`.

参见:

The `scandir()` function returns directory entries along with file attribute information, giving better performance for many common use cases.

在 3.2 版更改: The *path* parameter became optional.

3.3 新版功能: Added support for specifying an open file descriptor for *path*.

在 3.6 版更改: 接受一个类路径对象。

`os.lstat (path, *, dir_fd=None)`

Perform the equivalent of an `lstat()` system call on the given path. Similar to `stat()`, but does not follow symbolic links. Return a `stat_result` object.

On platforms that do not support symbolic links, this is an alias for `stat()`.

As of Python 3.3, this is equivalent to `os.stat(path, dir_fd=dir_fd, follow_symlinks=False)`.

This function can also support *paths relative to directory descriptors*.

参见:

`stat()` 函数。

在 3.2 版更改: Added support for Windows 6.0 (Vista) symbolic links.

在 3.3 版更改: Added the `dir_fd` parameter.

在 3.6 版更改: Accepts a *path-like object* for `src` and `dst`.

`os.mkdir (path, mode=0o777, *, dir_fd=None)`

Create a directory named `path` with numeric mode `mode`.

If the directory already exists, `FileExistsError` is raised.

On some systems, `mode` is ignored. Where it is used, the current umask value is first masked out. If bits other than the last 9 (i.e. the last 3 digits of the octal representation of the `mode`) are set, their meaning is platform-dependent. On some platforms, they are ignored and you should call `chmod()` explicitly to set them.

This function can also support *paths relative to directory descriptors*.

It is also possible to create temporary directories; see the `tempfile` module's `tempfile.mkdtemp()` function.

3.3 新版功能: `dir_fd` 参数。

在 3.6 版更改: 接受一个类路径对象。

`os.makedirs (name, mode=0o777, exist_ok=False)`

Recursive directory creation function. Like `mkdir()`, but makes all intermediate-level directories needed to contain the leaf directory.

The `mode` parameter is passed to `mkdir()`; see *the mkdir() description* for how it is interpreted.

If `exist_ok` is `False` (the default), an `OSError` is raised if the target directory already exists.

注解: `makedirs()` will become confused if the path elements to create include *pardir* (eg. “..” on UNIX systems).

This function handles UNC paths correctly.

3.2 新版功能: The `exist_ok` parameter.

在 3.4.1 版更改: Before Python 3.4.1, if `exist_ok` was `True` and the directory existed, `makedirs()` would still raise an error if `mode` did not match the mode of the existing directory. Since this behavior was impossible to implement safely, it was removed in Python 3.4.1. See [bpo-21082](#).

在 3.6 版更改: 接受一个类路径对象。

`os.mkfifo (path, mode=0o666, *, dir_fd=None)`

Create a FIFO (a named pipe) named `path` with numeric mode `mode`. The current umask value is first masked out from the mode.

This function can also support *paths relative to directory descriptors*.

FIFOs are pipes that can be accessed like regular files. FIFOs exist until they are deleted (for example with `os.unlink()`). Generally, FIFOs are used as rendezvous between “client” and “server” type processes: the server opens the FIFO for reading, and the client opens it for writing. Note that `mkfifo()` doesn’t open the FIFO — it just creates the rendezvous point.

Availability: Unix.

3.3 新版功能: `dir_fd` 参数。

在 3.6 版更改: 接受一个类路径对象。

`os.mknod(path, mode=0o600, device=0, *, dir_fd=None)`

Create a filesystem node (file, device special file or named pipe) named *path*. *mode* specifies both the permissions to use and the type of node to be created, being combined (bitwise OR) with one of `stat.S_IFREG`, `stat.S_IFCHR`, `stat.S_IFBLK`, and `stat.S_IFIFO` (those constants are available in `stat`). For `stat.S_IFCHR` and `stat.S_IFBLK`, *device* defines the newly created device special file (probably using `os.makedev()`), otherwise it is ignored.

This function can also support *paths relative to directory descriptors*.

Availability: Unix.

3.3 新版功能: `dir_fd` 参数。

在 3.6 版更改: 接受一个类路径对象。

`os.major(device)`

Extract the device major number from a raw device number (usually the `st_dev` or `st_rdev` field from `stat`).

`os.minor(device)`

Extract the device minor number from a raw device number (usually the `st_dev` or `st_rdev` field from `stat`).

`os.makedev(major, minor)`

Compose a raw device number from the major and minor device numbers.

`os.pathconf(path, name)`

Return system configuration information relevant to a named file. *name* specifies the configuration value to retrieve; it may be a string which is the name of a defined system value; these names are specified in a number of standards (POSIX.1, Unix 95, Unix 98, and others). Some platforms define additional names as well. The names known to the host operating system are given in the `pathconf_names` dictionary. For configuration variables not included in that mapping, passing an integer for *name* is also accepted.

If *name* is a string and is not known, `ValueError` is raised. If a specific value for *name* is not supported by the host system, even if it is included in `pathconf_names`, an `OSError` is raised with `errno.EINVAL` for the error number.

This function can support *specifying a file descriptor*.

Availability: Unix.

在 3.6 版更改: 接受一个类路径对象。

`os.pathconf_names`

Dictionary mapping names accepted by `pathconf()` and `fpathconf()` to the integer values defined for those names by the host operating system. This can be used to determine the set of names known to the system.

Availability: Unix.

`os.readlink(path, *, dir_fd=None)`

Return a string representing the path to which the symbolic link points. The result may be either an absolute or relative pathname; if it is relative, it may be converted to an absolute pathname using `os.path.join(os.path.dirname(path), result)`.

If the *path* is a string object (directly or indirectly through a *PathLike* interface), the result will also be a string object, and the call may raise a `UnicodeDecodeError`. If the *path* is a bytes object (direct or indirectly), the result will be a bytes object.

This function can also support *paths relative to directory descriptors*.

Availability: Unix, Windows

在 3.2 版更改: Added support for Windows 6.0 (Vista) symbolic links.

3.3 新版功能: *dir_fd* 参数。

在 3.6 版更改: 接受一个类路径对象。

`os.remove(path, *, dir_fd=None)`

Remove (delete) the file *path*. If *path* is a directory, `OSError` is raised. Use `rmdir()` to remove directories.

This function can support *paths relative to directory descriptors*.

On Windows, attempting to remove a file that is in use causes an exception to be raised; on Unix, the directory entry is removed but the storage allocated to the file is not made available until the original file is no longer in use.

This function is semantically identical to `unlink()`.

3.3 新版功能: *dir_fd* 参数。

在 3.6 版更改: 接受一个类路径对象。

`os.removedirs(name)`

Remove directories recursively. Works like `rmdir()` except that, if the leaf directory is successfully removed, `removedirs()` tries to successively remove every parent directory mentioned in *path* until an error is raised (which is ignored, because it generally means that a parent directory is not empty). For example, `os.removedirs('foo/bar/baz')` will first remove the directory 'foo/bar/baz', and then remove 'foo/bar' and 'foo' if they are empty. Raises `OSError` if the leaf directory could not be successfully removed.

在 3.6 版更改: 接受一个类路径对象。

`os.rename(src, dst, *, src_dir_fd=None, dst_dir_fd=None)`

Rename the file or directory *src* to *dst*. If *dst* is a directory, `OSError` will be raised. On Unix, if *dst* exists and is a file, it will be replaced silently if the user has permission. The operation may fail on some Unix flavors if *src* and *dst* are on different filesystems. If successful, the renaming will be an atomic operation (this is a POSIX requirement). On Windows, if *dst* already exists, `OSError` will be raised even if it is a file.

This function can support specifying *src_dir_fd* and/or *dst_dir_fd* to supply *paths relative to directory descriptors*.

If you want cross-platform overwriting of the destination, use `replace()`.

3.3 新版功能: The *src_dir_fd* and *dst_dir_fd* arguments.

在 3.6 版更改: Accepts a *path-like object* for *src* and *dst*.

`os.rename(old, new)`

Recursive directory or file renaming function. Works like `rename()`, except creation of any intermediate directories needed to make the new pathname good is attempted first. After the rename, directories corresponding to rightmost path segments of the old name will be pruned away using `removedirs()`.

注解: This function can fail with the new directory structure made if you lack permissions needed to remove the leaf directory or file.

在 3.6 版更改: Accepts a *path-like object* for *old* and *new*.

`os.replace(src, dst, *, src_dir_fd=None, dst_dir_fd=None)`

Rename the file or directory *src* to *dst*. If *dst* is a directory, *OSError* will be raised. If *dst* exists and is a file, it will be replaced silently if the user has permission. The operation may fail if *src* and *dst* are on different filesystems. If successful, the renaming will be an atomic operation (this is a POSIX requirement).

This function can support specifying *src_dir_fd* and/or *dst_dir_fd* to supply *paths relative to directory descriptors*.

3.3 新版功能.

在 3.6 版更改: Accepts a *path-like object* for *src* and *dst*.

`os.rmdir(path, *, dir_fd=None)`

Remove (delete) the directory *path*. Only works when the directory is empty, otherwise, *OSError* is raised. In order to remove whole directory trees, *shutil.rmtree()* can be used.

This function can support *paths relative to directory descriptors*.

3.3 新版功能: The *dir_fd* parameter.

在 3.6 版更改: 接受一个类路径对象。

`os.scandir(path='.')`

Return an iterator of *os.DirEntry* objects corresponding to the entries in the directory given by *path*. The entries are yielded in arbitrary order, and the special entries *'.'* and *'..'* are not included.

Using *scandir()* instead of *listdir()* can significantly increase the performance of code that also needs file type or file attribute information, because *os.DirEntry* objects expose this information if the operating system provides it when scanning a directory. All *os.DirEntry* methods may perform a system call, but *is_dir()* and *is_file()* usually only require a system call for symbolic links; *os.DirEntry.stat()* always requires a system call on Unix but only requires one for symbolic links on Windows.

path may be a *path-like object*. If *path* is of type *bytes* (directly or indirectly through the *PathLike* interface), the type of the *name* and *path* attributes of each *os.DirEntry* will be *bytes*; in all other circumstances, they will be of type *str*.

The *scandir()* iterator supports the *context manager* protocol and has the following method:

`scandir.close()`

Close the iterator and free acquired resources.

This is called automatically when the iterator is exhausted or garbage collected, or when an error happens during iterating. However it is advisable to call it explicitly or use the *with* statement.

3.6 新版功能.

The following example shows a simple use of *scandir()* to display all the files (excluding directories) in the given *path* that don't start with *'.'*. The *entry.is_file()* call will generally not make an additional system call:

```
with os.scandir(path) as it:
    for entry in it:
        if not entry.name.startswith('.') and entry.is_file():
            print(entry.name)
```

注解: On Unix-based systems, *scandir()* uses the system's *opendir()* and *readdir()* functions. On Windows, it uses the Win32 *FindFirstFileW* and *FindNextFileW* functions.

3.5 新版功能.

3.6 新版功能: Added support for the *context manager* protocol and the *close()* method. If a *scandir()* iterator is neither exhausted nor explicitly closed a *ResourceWarning* will be emitted in its destructor.

The function accepts a *path-like object*.

class `os.DirEntry`

Object yielded by `scandir()` to expose the file path and other file attributes of a directory entry.

`scandir()` will provide as much of this information as possible without making additional system calls. When a `stat()` or `lstat()` system call is made, the `os.DirEntry` object will cache the result.

`os.DirEntry` instances are not intended to be stored in long-lived data structures; if you know the file metadata has changed or if a long time has elapsed since calling `scandir()`, call `os.stat(entry.path)` to fetch up-to-date information.

Because the `os.DirEntry` methods can make operating system calls, they may also raise `OSError`. If you need very fine-grained control over errors, you can catch `OSError` when calling one of the `os.DirEntry` methods and handle as appropriate.

To be directly usable as a *path-like object*, `os.DirEntry` implements the `PathLike` interface.

Attributes and methods on a `os.DirEntry` instance are as follows:

name

The entry's base filename, relative to the `scandir()` *path* argument.

The `name` attribute will be `bytes` if the `scandir()` *path* argument is of type `bytes` and `str` otherwise. Use `fsdecode()` to decode byte filenames.

path

The entry's full path name: equivalent to `os.path.join(scandir_path, entry.name)` where `scandir_path` is the `scandir()` *path* argument. The path is only absolute if the `scandir()` *path* argument was absolute.

The `path` attribute will be `bytes` if the `scandir()` *path* argument is of type `bytes` and `str` otherwise. Use `fsdecode()` to decode byte filenames.

inode()

Return the inode number of the entry.

The result is cached on the `os.DirEntry` object. Use `os.stat(entry.path, follow_symlinks=False).st_ino` to fetch up-to-date information.

On the first, uncached call, a system call is required on Windows but not on Unix.

is_dir (*, *follow_symlinks=True*)

Return `True` if this entry is a directory or a symbolic link pointing to a directory; return `False` if the entry is or points to any other kind of file, or if it doesn't exist anymore.

If *follow_symlinks* is `False`, return `True` only if this entry is a directory (without following symlinks); return `False` if the entry is any other kind of file or if it doesn't exist anymore.

The result is cached on the `os.DirEntry` object, with a separate cache for *follow_symlinks* `True` and `False`. Call `os.stat()` along with `stat.S_ISDIR()` to fetch up-to-date information.

On the first, uncached call, no system call is required in most cases. Specifically, for non-symlinks, neither Windows or Unix require a system call, except on certain Unix file systems, such as network file systems, that return `dirent.d_type == DT_UNKNOWN`. If the entry is a symlink, a system call will be required to follow the symlink unless *follow_symlinks* is `False`.

This method can raise `OSError`, such as `PermissionError`, but `FileNotFoundError` is caught and not raised.

is_file (*, *follow_symlinks=True*)

Return `True` if this entry is a file or a symbolic link pointing to a file; return `False` if the entry is or points to a directory or other non-file entry, or if it doesn't exist anymore.

If `follow_symlinks` is `False`, return `True` only if this entry is a file (without following symlinks); return `False` if the entry is a directory or other non-file entry, or if it doesn't exist anymore.

The result is cached on the `os.DirEntry` object. Caching, system calls made, and exceptions raised are as per `is_dir()`.

`is_symlink()`

Return `True` if this entry is a symbolic link (even if broken); return `False` if the entry points to a directory or any kind of file, or if it doesn't exist anymore.

The result is cached on the `os.DirEntry` object. Call `os.path.islink()` to fetch up-to-date information.

On the first, uncached call, no system call is required in most cases. Specifically, neither Windows or Unix require a system call, except on certain Unix file systems, such as network file systems, that return `dirent.d_type == DT_UNKNOWN`.

This method can raise `OSError`, such as `PermissionError`, but `FileNotFoundError` is caught and not raised.

`stat(*, follow_symlinks=True)`

Return a `stat_result` object for this entry. This method follows symbolic links by default; to stat a symbolic link add the `follow_symlinks=False` argument.

On Unix, this method always requires a system call. On Windows, it only requires a system call if `follow_symlinks` is `True` and the entry is a symbolic link.

On Windows, the `st_ino`, `st_dev` and `st_nlink` attributes of the `stat_result` are always set to zero. Call `os.stat()` to get these attributes.

The result is cached on the `os.DirEntry` object, with a separate cache for `follow_symlinks` `True` and `False`. Call `os.stat()` to fetch up-to-date information.

Note that there is a nice correspondence between several attributes and methods of `os.DirEntry` and of `pathlib.Path`. In particular, the `name` attribute has the same meaning, as do the `is_dir()`, `is_file()`, `is_symlink()` and `stat()` methods.

3.5 新版功能.

在 3.6 版更改: Added support for the `PathLike` interface. Added support for `bytes` paths on Windows.

`os.stat(path, *, dir_fd=None, follow_symlinks=True)`

Get the status of a file or a file descriptor. Perform the equivalent of a `stat()` system call on the given path. `path` may be specified as either a string or bytes – directly or indirectly through the `PathLike` interface – or as an open file descriptor. Return a `stat_result` object.

This function normally follows symlinks; to stat a symlink add the argument `follow_symlinks=False`, or use `lstat()`.

This function can support *specifying a file descriptor* and *not following symlinks*.

示例:

```
>>> import os
>>> statinfo = os.stat('somefile.txt')
>>> statinfo
os.stat_result(st_mode=33188, st_ino=7876932, st_dev=234881026,
st_nlink=1, st_uid=501, st_gid=501, st_size=264, st_atime=1297230295,
st_mtime=1297230027, st_ctime=1297230027)
>>> statinfo.st_size
264
```

参见:

`fstat()` and `lstat()` functions.

3.3 新版功能: Added the `dir_fd` and `follow_symlinks` arguments, specifying a file descriptor instead of a path.

在 3.6 版更改: 接受一个类路径对象。

class `os.stat_result`

Object whose attributes correspond roughly to the members of the `stat` structure. It is used for the result of `os.stat()`, `os.fstat()` and `os.lstat()`.

Attributes:

st_mode

File mode: file type and file mode bits (permissions).

st_ino

Platform dependent, but if non-zero, uniquely identifies the file for a given value of `st_dev`. Typically:

- the inode number on Unix,
- the [file index](#) on Windows

st_dev

Identifier of the device on which this file resides.

st_nlink

Number of hard links.

st_uid

User identifier of the file owner.

st_gid

Group identifier of the file owner.

st_size

Size of the file in bytes, if it is a regular file or a symbolic link. The size of a symbolic link is the length of the pathname it contains, without a terminating null byte.

Timestamps:

st_atime

Time of most recent access expressed in seconds.

st_mtime

Time of most recent content modification expressed in seconds.

st_ctime

Platform dependent:

- the time of most recent metadata change on Unix,
- the time of creation on Windows, expressed in seconds.

st_atime_ns

Time of most recent access expressed in nanoseconds as an integer.

st_mtime_ns

Time of most recent content modification expressed in nanoseconds as an integer.

st_ctime_ns

Platform dependent:

- the time of most recent metadata change on Unix,

- the time of creation on Windows, expressed in nanoseconds as an integer.

See also the `stat_float_times()` function.

注解: The exact meaning and resolution of the `st_atime`, `st_mtime`, and `st_ctime` attributes depend on the operating system and the file system. For example, on Windows systems using the FAT or FAT32 file systems, `st_mtime` has 2-second resolution, and `st_atime` has only 1-day resolution. See your operating system documentation for details.

Similarly, although `st_atime_ns`, `st_mtime_ns`, and `st_ctime_ns` are always expressed in nanoseconds, many systems do not provide nanosecond precision. On systems that do provide nanosecond precision, the floating-point object used to store `st_atime`, `st_mtime`, and `st_ctime` cannot preserve all of it, and as such will be slightly inexact. If you need the exact timestamps you should always use `st_atime_ns`, `st_mtime_ns`, and `st_ctime_ns`.

On some Unix systems (such as Linux), the following attributes may also be available:

st_blocks

Number of 512-byte blocks allocated for file. This may be smaller than `st_size/512` when the file has holes.

st_blksize

“Preferred” blocksize for efficient file system I/O. Writing to a file in smaller chunks may cause an inefficient read-modify-rewrite.

st_rdev

Type of device if an inode device.

st_flags

User defined flags for file.

On other Unix systems (such as FreeBSD), the following attributes may be available (but may be only filled out if root tries to use them):

st_gen

File generation number.

st_birthtime

Time of file creation.

On Mac OS systems, the following attributes may also be available:

st_rsize

Real size of the file.

st_creator

Creator of the file.

st_type

File type.

On Windows systems, the following attribute is also available:

st_file_attributes

Windows file attributes: `dwFileAttributes` member of the `BY_HANDLE_FILE_INFORMATION` structure returned by `GetFileInformationByHandle()`. See the `FILE_ATTRIBUTE_*` constants in the `stat` module.

The standard module `stat` defines functions and constants that are useful for extracting information from a `stat` structure. (On Windows, some items are filled with dummy values.)

For backward compatibility, a `stat_result` instance is also accessible as a tuple of at least 10 integers giving the most important (and portable) members of the `stat` structure, in the order `st_mode`, `st_ino`, `st_dev`, `st_nlink`, `st_uid`, `st_gid`, `st_size`, `st_atime`, `st_mtime`, `st_ctime`. More items may be added at the end by some implementations. For compatibility with older Python versions, accessing `stat_result` as a tuple always returns integers.

3.3 新版功能: Added the `st_atime_ns`, `st_mtime_ns`, and `st_ctime_ns` members.

3.5 新版功能: Added the `st_file_attributes` member on Windows.

在 3.5 版更改: Windows now returns the file index as `st_ino` when available.

`os.stat_float_times([newvalue])`

Determine whether `stat_result` represents time stamps as float objects. If `newvalue` is `True`, future calls to `stat()` return floats, if it is `False`, future calls return ints. If `newvalue` is omitted, return the current setting.

For compatibility with older Python versions, accessing `stat_result` as a tuple always returns integers.

Python now returns float values by default. Applications which do not work correctly with floating point time stamps can use this function to restore the old behaviour.

The resolution of the timestamps (that is the smallest possible fraction) depends on the system. Some systems only support second resolution; on these systems, the fraction will always be zero.

It is recommended that this setting is only changed at program startup time in the `__main__` module; libraries should never change this setting. If an application uses a library that works incorrectly if floating point time stamps are processed, this application should turn the feature off until the library has been corrected.

3.3 版后已移除.

`os.statvfs(path)`

Perform a `statvfs()` system call on the given path. The return value is an object whose attributes describe the filesystem on the given path, and correspond to the members of the `statvfs` structure, namely: `f_bsize`, `f_frsize`, `f_blocks`, `f_bfree`, `f_bavail`, `f_files`, `f_ffree`, `f_favail`, `f_flag`, `f_namemax`.

Two module-level constants are defined for the `f_flag` attribute's bit-flags: if `ST_RDONLY` is set, the filesystem is mounted read-only, and if `ST_NOSUID` is set, the semantics of `setuid`/`setgid` bits are disabled or not supported.

Additional module-level constants are defined for GNU/glibc based systems. These are `ST_NODEV` (disallow access to device special files), `ST_NOEXEC` (disallow program execution), `ST_SYNCHRONOUS` (writes are synced at once), `ST_MANDLOCK` (allow mandatory locks on an FS), `ST_WRITE` (write on file/directory/symlink), `ST_APPEND` (append-only file), `ST_IMMUTABLE` (immutable file), `ST_NOATIME` (do not update access times), `ST_NODIRATIME` (do not update directory access times), `ST_RELATIME` (update atime relative to mtime/ctime).

This function can support *specifying a file descriptor*.

Availability: Unix.

在 3.2 版更改: The `ST_RDONLY` and `ST_NOSUID` constants were added.

3.3 新版功能: Added support for specifying an open file descriptor for `path`.

在 3.4 版更改: The `ST_NODEV`, `ST_NOEXEC`, `ST_SYNCHRONOUS`, `ST_MANDLOCK`, `ST_WRITE`, `ST_APPEND`, `ST_IMMUTABLE`, `ST_NOATIME`, `ST_NODIRATIME`, and `ST_RELATIME` constants were added.

在 3.6 版更改: 接受一个类路径对象。

`os.supports_dir_fd`

A `Set` object indicating which functions in the `os` module permit use of their `dir_fd` parameter. Different platforms provide different functionality, and an option that might work on one might be unsupported on another.

For consistency's sakes, functions that support *dir_fd* always allow specifying the parameter, but will raise an exception if the functionality is not actually available.

To check whether a particular function permits use of its *dir_fd* parameter, use the `in` operator on `supports_dir_fd`. As an example, this expression determines whether the *dir_fd* parameter of `os.stat()` is locally available:

```
os.stat in os.supports_dir_fd
```

Currently *dir_fd* parameters only work on Unix platforms; none of them work on Windows.

3.3 新版功能.

`os.supports_effective_ids`

A *Set* object indicating which functions in the `os` module permit use of the *effective_ids* parameter for `os.access()`. If the local platform supports it, the collection will contain `os.access()`, otherwise it will be empty.

To check whether you can use the *effective_ids* parameter for `os.access()`, use the `in` operator on `supports_effective_ids`, like so:

```
os.access in os.supports_effective_ids
```

Currently *effective_ids* only works on Unix platforms; it does not work on Windows.

3.3 新版功能.

`os.supports_fd`

A *Set* object indicating which functions in the `os` module permit specifying their *path* parameter as an open file descriptor. Different platforms provide different functionality, and an option that might work on one might be unsupported on another. For consistency's sakes, functions that support *fd* always allow specifying the parameter, but will raise an exception if the functionality is not actually available.

To check whether a particular function permits specifying an open file descriptor for its *path* parameter, use the `in` operator on `supports_fd`. As an example, this expression determines whether `os.chdir()` accepts open file descriptors when called on your local platform:

```
os.chdir in os.supports_fd
```

3.3 新版功能.

`os.supports_follow_symlinks`

A *Set* object indicating which functions in the `os` module permit use of their *follow_symlinks* parameter. Different platforms provide different functionality, and an option that might work on one might be unsupported on another. For consistency's sakes, functions that support *follow_symlinks* always allow specifying the parameter, but will raise an exception if the functionality is not actually available.

To check whether a particular function permits use of its *follow_symlinks* parameter, use the `in` operator on `supports_follow_symlinks`. As an example, this expression determines whether the *follow_symlinks* parameter of `os.stat()` is locally available:

```
os.stat in os.supports_follow_symlinks
```

3.3 新版功能.

`os.symlink(src, dst, target_is_directory=False, *, dir_fd=None)`

Create a symbolic link pointing to *src* named *dst*.

On Windows, a symlink represents either a file or a directory, and does not morph to the target dynamically. If the target is present, the type of the symlink will be created to match. Otherwise, the symlink will be created as

a directory if *target_is_directory* is `True` or a file symlink (the default) otherwise. On non-Windows platforms, *target_is_directory* is ignored.

Symbolic link support was introduced in Windows 6.0 (Vista). `symlink()` will raise a `NotImplementedError` on Windows versions earlier than 6.0.

This function can support *paths relative to directory descriptors*.

注解: On Windows, the `SeCreateSymbolicLinkPrivilege` is required in order to successfully create symlinks. This privilege is not typically granted to regular users but is available to accounts which can escalate privileges to the administrator level. Either obtaining the privilege or running your application as an administrator are ways to successfully create symlinks.

`OSError` is raised when the function is called by an unprivileged user.

Availability: Unix, Windows.

在 3.2 版更改: Added support for Windows 6.0 (Vista) symbolic links.

3.3 新版功能: Added the *dir_fd* argument, and now allow *target_is_directory* on non-Windows platforms.

在 3.6 版更改: Accepts a *path-like object* for *src* and *dst*.

os.sync()

Force write of everything to disk.

Availability: Unix.

3.3 新版功能.

os.truncate(path, length)

Truncate the file corresponding to *path*, so that it is at most *length* bytes in size.

This function can support *specifying a file descriptor*.

Availability: Unix, Windows.

3.3 新版功能.

在 3.5 版更改: 添加了 Windows 支持

在 3.6 版更改: 接受一个类路径对象。

os.unlink(path, *, dir_fd=None)

Remove (delete) the file *path*. This function is semantically identical to `remove()`; the `unlink` name is its traditional Unix name. Please see the documentation for `remove()` for further information.

3.3 新版功能: The *dir_fd* parameter.

在 3.6 版更改: 接受一个类路径对象。

os.utime(path, times=None, *[, ns], dir_fd=None, follow_symlinks=True)

Set the access and modified times of the file specified by *path*.

`utime()` takes two optional parameters, *times* and *ns*. These specify the times set on *path* and are used as follows:

- If *ns* is specified, it must be a 2-tuple of the form `(atime_ns, mtime_ns)` where each member is an int expressing nanoseconds.
- If *times* is not `None`, it must be a 2-tuple of the form `(atime, mtime)` where each member is an int or float expressing seconds.
- If *times* is `None` and *ns* is unspecified, this is equivalent to specifying `ns=(atime_ns, mtime_ns)` where both times are the current time.

It is an error to specify tuples for both *times* and *ns*.

Whether a directory can be given for *path* depends on whether the operating system implements directories as files (for example, Windows does not). Note that the exact times you set here may not be returned by a subsequent `stat()` call, depending on the resolution with which your operating system records access and modification times; see `stat()`. The best way to preserve exact times is to use the `st_atime_ns` and `st_mtime_ns` fields from the `os.stat()` result object with the *ns* parameter to `utime`.

This function can support *specifying a file descriptor*, *paths relative to directory descriptors* and *not following symlinks*.

3.3 新版功能: Added support for specifying an open file descriptor for *path*, and the *dir_fd*, *follow_symlinks*, and *ns* parameters.

在 3.6 版更改: 接受一个类路径对象。

`os.walk(top, topdown=True, onerror=None, followlinks=False)`

Generate the file names in a directory tree by walking the tree either top-down or bottom-up. For each directory in the tree rooted at directory *top* (including *top* itself), it yields a 3-tuple (*dirpath*, *dirnames*, *filenames*).

dirpath is a string, the path to the directory. *dirnames* is a list of the names of the subdirectories in *dirpath* (excluding `'.'` and `'..'`). *filenames* is a list of the names of the non-directory files in *dirpath*. Note that the names in the lists contain no path components. To get a full path (which begins with *top*) to a file or directory in *dirpath*, do `os.path.join(dirpath, name)`.

If optional argument *topdown* is `True` or not specified, the triple for a directory is generated before the triples for any of its subdirectories (directories are generated top-down). If *topdown* is `False`, the triple for a directory is generated after the triples for all of its subdirectories (directories are generated bottom-up). No matter the value of *topdown*, the list of subdirectories is retrieved before the tuples for the directory and its subdirectories are generated.

When *topdown* is `True`, the caller can modify the *dirnames* list in-place (perhaps using `del` or slice assignment), and `walk()` will only recurse into the subdirectories whose names remain in *dirnames*; this can be used to prune the search, impose a specific order of visiting, or even to inform `walk()` about directories the caller creates or renames before it resumes `walk()` again. Modifying *dirnames* when *topdown* is `False` has no effect on the behavior of the walk, because in bottom-up mode the directories in *dirnames* are generated before *dirpath* itself is generated.

By default, errors from the `scandir()` call are ignored. If optional argument *onerror* is specified, it should be a function; it will be called with one argument, an `OSError` instance. It can report the error to continue with the walk, or raise the exception to abort the walk. Note that the filename is available as the *filename* attribute of the exception object.

By default, `walk()` will not walk down into symbolic links that resolve to directories. Set *followlinks* to `True` to visit directories pointed to by symlinks, on systems that support them.

注解: Be aware that setting *followlinks* to `True` can lead to infinite recursion if a link points to a parent directory of itself. `walk()` does not keep track of the directories it visited already.

注解: If you pass a relative pathname, don't change the current working directory between resumptions of `walk()`. `walk()` never changes the current directory, and assumes that its caller doesn't either.

This example displays the number of bytes taken by non-directory files in each directory under the starting directory, except that it doesn't look under any CVS subdirectory:

```
import os
from os.path import join, getsize
for root, dirs, files in os.walk('python/Lib/email'):
    print(root, "consumes", end=" ")
    print(sum(getsize(join(root, name)) for name in files), end=" ")
    print("bytes in", len(files), "non-directory files")
    if 'CVS' in dirs:
        dirs.remove('CVS') # don't visit CVS directories
```

In the next example (simple implementation of `shutil.rmtree()`), walking the tree bottom-up is essential, `rmdir()` doesn't allow deleting a directory before the directory is empty:

```
# Delete everything reachable from the directory named in "top",
# assuming there are no symbolic links.
# CAUTION: This is dangerous! For example, if top == '/', it
# could delete all your disk files.
import os
for root, dirs, files in os.walk(top, topdown=False):
    for name in files:
        os.remove(os.path.join(root, name))
    for name in dirs:
        os.rmdir(os.path.join(root, name))
```

在 3.5 版更改: This function now calls `os.scandir()` instead of `os.listdir()`, making it faster by reducing the number of calls to `os.stat()`.

在 3.6 版更改: 接受一个类路径对象。

`os.fwalk` (`top='.'`, `topdown=True`, `onerror=None`, *, `follow_symlinks=False`, `dir_fd=None`)

This behaves exactly like `walk()`, except that it yields a 4-tuple (`dirpath`, `dirnames`, `filenames`, `dirfd`), and it supports `dir_fd`.

`dirpath`, `dirnames` and `filenames` are identical to `walk()` output, and `dirfd` is a file descriptor referring to the directory `dirpath`.

This function always supports *paths relative to directory descriptors* and *not following symlinks*. Note however that, unlike other functions, the `fwalk()` default value for `follow_symlinks` is `False`.

注解: Since `fwalk()` yields file descriptors, those are only valid until the next iteration step, so you should duplicate them (e.g. with `dup()`) if you want to keep them longer.

This example displays the number of bytes taken by non-directory files in each directory under the starting directory, except that it doesn't look under any CVS subdirectory:

```
import os
for root, dirs, files, rootfd in os.fwalk('python/Lib/email'):
    print(root, "consumes", end=" ")
    print(sum([os.stat(name, dir_fd=rootfd).st_size for name in files]),
          end=" ")
    print("bytes in", len(files), "non-directory files")
    if 'CVS' in dirs:
        dirs.remove('CVS') # don't visit CVS directories
```

In the next example, walking the tree bottom-up is essential: `rmdir()` doesn't allow deleting a directory before the directory is empty:

```
# Delete everything reachable from the directory named in "top",
# assuming there are no symbolic links.
# CAUTION: This is dangerous! For example, if top == '/', it
# could delete all your disk files.
import os
for root, dirs, files, rootfd in os.fwalk(top, topdown=False):
    for name in files:
        os.unlink(name, dir_fd=rootfd)
    for name in dirs:
        os.rmdir(name, dir_fd=rootfd)
```

Availability: Unix.

3.3 新版功能.

在 3.6 版更改: 接受一个类路径对象。

Linux extended attributes

3.3 新版功能.

These functions are all available on Linux only.

`os.getxattr(path, attribute, *, follow_symlinks=True)`

Return the value of the extended filesystem attribute *attribute* for *path*. *attribute* can be bytes or str (directly or indirectly through the *PathLike* interface). If it is str, it is encoded with the filesystem encoding.

This function can support *specifying a file descriptor* and *not following symlinks*.

在 3.6 版更改: Accepts a *path-like object* for *path* and *attribute*.

`os.listdirxattr(path=None, *, follow_symlinks=True)`

Return a list of the extended filesystem attributes on *path*. The attributes in the list are represented as strings decoded with the filesystem encoding. If *path* is None, `listxattr()` will examine the current directory.

This function can support *specifying a file descriptor* and *not following symlinks*.

在 3.6 版更改: 接受一个类路径对象。

`os.removexattr(path, attribute, *, follow_symlinks=True)`

Removes the extended filesystem attribute *attribute* from *path*. *attribute* should be bytes or str (directly or indirectly through the *PathLike* interface). If it is a string, it is encoded with the filesystem encoding.

This function can support *specifying a file descriptor* and *not following symlinks*.

在 3.6 版更改: Accepts a *path-like object* for *path* and *attribute*.

`os.setxattr(path, attribute, value, flags=0, *, follow_symlinks=True)`

Set the extended filesystem attribute *attribute* on *path* to *value*. *attribute* must be a bytes or str with no embedded NULs (directly or indirectly through the *PathLike* interface). If it is a str, it is encoded with the filesystem encoding. *flags* may be `XATTR_REPLACE` or `XATTR_CREATE`. If `XATTR_REPLACE` is given and the attribute does not exist, `EEXIST` will be raised. If `XATTR_CREATE` is given and the attribute already exists, the attribute will not be created and `ENODATA` will be raised.

This function can support *specifying a file descriptor* and *not following symlinks*.

注解: A bug in Linux kernel versions less than 2.6.39 caused the flags argument to be ignored on some filesystems.

在 3.6 版更改: Accepts a *path-like object* for *path* and *attribute*.

os.XATTR_SIZE_MAX

The maximum size the value of an extended attribute can be. Currently, this is 64 KiB on Linux.

os.XATTR_CREATE

This is a possible value for the flags argument in `setxattr()`. It indicates the operation must create an attribute.

os.XATTR_REPLACE

This is a possible value for the flags argument in `setxattr()`. It indicates the operation must replace an existing attribute.

16.1.6 Process Management

These functions may be used to create and manage processes.

The various `exec*` functions take a list of arguments for the new program loaded into the process. In each case, the first of these arguments is passed to the new program as its own name rather than as an argument a user may have typed on a command line. For the C programmer, this is the `argv[0]` passed to a program's `main()`. For example, `os.execv('/bin/echo', ['foo', 'bar'])` will only print `bar` on standard output; `foo` will seem to be ignored.

os.abort()

Generate a SIGABRT signal to the current process. On Unix, the default behavior is to produce a core dump; on Windows, the process immediately returns an exit code of 3. Be aware that calling this function will not call the Python signal handler registered for SIGABRT with `signal.signal()`.

os.execl(path, arg0, arg1, ...)**os.execle(path, arg0, arg1, ..., env)****os.execlp(file, arg0, arg1, ...)****os.execlpe(file, arg0, arg1, ..., env)****os.execv(path, args)****os.execve(path, args, env)****os.execvp(file, args)****os.execvpe(file, args, env)**

These functions all execute a new program, replacing the current process; they do not return. On Unix, the new executable is loaded into the current process, and will have the same process id as the caller. Errors will be reported as `OSError` exceptions.

The current process is replaced immediately. Open file objects and descriptors are not flushed, so if there may be data buffered on these open files, you should flush them using `sys.stdout.flush()` or `os.fsync()` before calling an `exec*` function.

The “l” and “v” variants of the `exec*` functions differ in how command-line arguments are passed. The “l” variants are perhaps the easiest to work with if the number of parameters is fixed when the code is written; the individual parameters simply become additional parameters to the `execl*()` functions. The “v” variants are good when the number of parameters is variable, with the arguments being passed in a list or tuple as the `args` parameter. In either case, the arguments to the child process should start with the name of the command being run, but this is not enforced.

The variants which include a “p” near the end (`execlp()`, `execlpe()`, `execvp()`, and `execvpe()`) will use the `PATH` environment variable to locate the program *file*. When the environment is being replaced (using one of the `exec*e` variants, discussed in the next paragraph), the new environment is used as the source of the `PATH` variable. The other variants, `execl()`, `execle()`, `execv()`, and `execve()`, will not use the `PATH` variable to locate the executable; *path* must contain an appropriate absolute or relative path.

For `execle()`, `execlpe()`, `execve()`, and `execvpe()` (note that these all end in “e”), the *env* parameter must be a mapping which is used to define the environment variables for the new process (these are used instead of

the current process' environment); the functions `execl()`, `execvp()`, and `execvp()` all cause the new process to inherit the environment of the current process.

For `execve()` on some platforms, *path* may also be specified as an open file descriptor. This functionality may not be supported on your platform; you can check whether or not it is available using `os.supports_fd`. If it is unavailable, using it will raise a `NotImplementedError`.

Availability: Unix, Windows.

3.3 新版功能: Added support for specifying an open file descriptor for *path* for `execve()`.

在 3.6 版更改: 接受一个类路径对象。

`os._exit(n)`

Exit the process with status *n*, without calling cleanup handlers, flushing stdio buffers, etc.

注解: The standard way to exit is `sys.exit(n)`. `_exit()` should normally only be used in the child process after a `fork()`.

The following exit codes are defined and can be used with `_exit()`, although they are not required. These are typically used for system programs written in Python, such as a mail server's external command delivery program.

注解: Some of these may not be available on all Unix platforms, since there is some variation. These constants are defined where they are defined by the underlying platform.

`os.EX_OK`

Exit code that means no error occurred.

Availability: Unix.

`os.EX_USAGE`

Exit code that means the command was used incorrectly, such as when the wrong number of arguments are given.

Availability: Unix.

`os.EX_DATAERR`

Exit code that means the input data was incorrect.

Availability: Unix.

`os.EX_NOINPUT`

Exit code that means an input file did not exist or was not readable.

Availability: Unix.

`os.EX_NOUSER`

Exit code that means a specified user did not exist.

Availability: Unix.

`os.EX_NOHOST`

Exit code that means a specified host did not exist.

Availability: Unix.

`os.EX_UNAVAILABLE`

Exit code that means that a required service is unavailable.

Availability: Unix.

os.**EX_SOFTWARE**

Exit code that means an internal software error was detected.

Availability: Unix.

os.**EX_OSERR**

Exit code that means an operating system error was detected, such as the inability to fork or create a pipe.

Availability: Unix.

os.**EX_OSFILE**

Exit code that means some system file did not exist, could not be opened, or had some other kind of error.

Availability: Unix.

os.**EX_CANTCREAT**

Exit code that means a user specified output file could not be created.

Availability: Unix.

os.**EX_IOERR**

Exit code that means that an error occurred while doing I/O on some file.

Availability: Unix.

os.**EX_TEMPFAIL**

Exit code that means a temporary failure occurred. This indicates something that may not really be an error, such as a network connection that couldn't be made during a retryable operation.

Availability: Unix.

os.**EX_PROTOCOL**

Exit code that means that a protocol exchange was illegal, invalid, or not understood.

Availability: Unix.

os.**EX_NOPERM**

Exit code that means that there were insufficient permissions to perform the operation (but not intended for file system problems).

Availability: Unix.

os.**EX_CONFIG**

Exit code that means that some kind of configuration error occurred.

Availability: Unix.

os.**EX_NOTFOUND**

Exit code that means something like “an entry was not found”.

Availability: Unix.

os.**fork()**

Fork a child process. Return 0 in the child and the child's process id in the parent. If an error occurs *OSError* is raised.

Note that some platforms including FreeBSD <= 6.3 and Cygwin have known issues when using `fork()` from a thread.

警告: See *ssl* for applications that use the SSL module with `fork()`.

Availability: Unix.

os.forkpty()

Fork a child process, using a new pseudo-terminal as the child's controlling terminal. Return a pair of (*pid*, *fd*), where *pid* is 0 in the child, the new child's process id in the parent, and *fd* is the file descriptor of the master end of the pseudo-terminal. For a more portable approach, use the *pty* module. If an error occurs *OSError* is raised.

Availability: some flavors of Unix.

os.kill(pid, sig)

Send signal *sig* to the process *pid*. Constants for the specific signals available on the host platform are defined in the *signal* module.

Windows: The *signal.CTRL_C_EVENT* and *signal.CTRL_BREAK_EVENT* signals are special signals which can only be sent to console processes which share a common console window, e.g., some subprocesses. Any other value for *sig* will cause the process to be unconditionally killed by the *TerminateProcess* API, and the exit code will be set to *sig*. The Windows version of *kill()* additionally takes process handles to be killed.

See also *signal.thread_kill()*.

3.2 新版功能: Windows support.

os.killpg(pgid, sig)

Send the signal *sig* to the process group *pgid*.

Availability: Unix.

os.nice(increment)

Add *increment* to the process's "niceness". Return the new niceness.

Availability: Unix.

os.plock(op)

Lock program segments into memory. The value of *op* (defined in *<sys/lock.h>*) determines which segments are locked.

Availability: Unix.

os.popen(cmd, mode='r', buffering=-1)

Open a pipe to or from command *cmd*. The return value is an open file object connected to the pipe, which can be read or written depending on whether *mode* is 'r' (default) or 'w'. The *buffering* argument has the same meaning as the corresponding argument to the built-in *open()* function. The returned file object reads or writes text strings rather than bytes.

The *close* method returns *None* if the subprocess exited successfully, or the subprocess's return code if there was an error. On POSIX systems, if the return code is positive it represents the return value of the process left-shifted by one byte. If the return code is negative, the process was terminated by the signal given by the negated value of the return code. (For example, the return value might be $-\text{signal.SIGKILL}$ if the subprocess was killed.) On Windows systems, the return value contains the signed integer return code from the child process.

This is implemented using *subprocess.Popen*; see that class's documentation for more powerful ways to manage and communicate with subprocesses.

os.spawnl(mode, path, ...)**os.spawnle(mode, path, ..., env)****os.spawnlp(mode, file, ...)****os.spawnlpe(mode, file, ..., env)****os.spawnnv(mode, path, args)****os.spawnve(mode, path, args, env)****os.spawnvp(mode, file, args)****os.spawnvpe(mode, file, args, env)**

Execute the program *path* in a new process.

(Note that the `subprocess` module provides more powerful facilities for spawning new processes and retrieving their results; using that module is preferable to using these functions. Check especially the [Replacing Older Functions with the subprocess Module](#) section.)

If `mode` is `P_NOWAIT`, this function returns the process id of the new process; if `mode` is `P_WAIT`, returns the process' s exit code if it exits normally, or `-signal`, where `signal` is the signal that killed the process. On Windows, the process id will actually be the process handle, so can be used with the `waitpid()` function.

The “l” and “v” variants of the `spawn*` functions differ in how command-line arguments are passed. The “l” variants are perhaps the easiest to work with if the number of parameters is fixed when the code is written; the individual parameters simply become additional parameters to the `spawnl*()` functions. The “v” variants are good when the number of parameters is variable, with the arguments being passed in a list or tuple as the `args` parameter. In either case, the arguments to the child process must start with the name of the command being run.

The variants which include a second “p” near the end (`spawnlp()`, `spawnlpe()`, `spawnvp()`, and `spawnvpe()`) will use the `PATH` environment variable to locate the program *file*. When the environment is being replaced (using one of the `spawn*e` variants, discussed in the next paragraph), the new environment is used as the source of the `PATH` variable. The other variants, `spawnl()`, `spawnle()`, `spawnv()`, and `spawnve()`, will not use the `PATH` variable to locate the executable; *path* must contain an appropriate absolute or relative path.

For `spawnle()`, `spawnlpe()`, `spawnve()`, and `spawnvpe()` (note that these all end in “e”), the `env` parameter must be a mapping which is used to define the environment variables for the new process (they are used instead of the current process' environment); the functions `spawnl()`, `spawnlp()`, `spawnv()`, and `spawnvp()` all cause the new process to inherit the environment of the current process. Note that keys and values in the `env` dictionary must be strings; invalid keys or values will cause the function to fail, with a return value of 127.

As an example, the following calls to `spawnlp()` and `spawnvpe()` are equivalent:

```
import os
os.spawnlp(os.P_WAIT, 'cp', 'cp', 'index.html', '/dev/null')

L = ['cp', 'index.html', '/dev/null']
os.spawnvpe(os.P_WAIT, 'cp', L, os.environ)
```

Availability: Unix, Windows. `spawnlp()`, `spawnlpe()`, `spawnvp()` and `spawnvpe()` are not available on Windows. `spawnle()` and `spawnve()` are not thread-safe on Windows; we advise you to use the `subprocess` module instead.

在 3.6 版更改: 接受一个类路径对象。

os.P_NOWAIT

os.P_NOWAITO

Possible values for the `mode` parameter to the `spawn*` family of functions. If either of these values is given, the `spawn*()` functions will return as soon as the new process has been created, with the process id as the return value.

Availability: Unix, Windows.

os.P_WAIT

Possible value for the `mode` parameter to the `spawn*` family of functions. If this is given as `mode`, the `spawn*()` functions will not return until the new process has run to completion and will return the exit code of the process the run is successful, or `-signal` if a signal kills the process.

Availability: Unix, Windows.

os.P_DETACH

os.P_OVERLAY

Possible values for the `mode` parameter to the `spawn*` family of functions. These are less portable than those

listed above. `P_DETACH` is similar to `P_NOWAIT`, but the new process is detached from the console of the calling process. If `P_OVERLAY` is used, the current process will be replaced; the `spawn*` function will not return.

Availability: Windows.

`os.startfile(path[, operation])`

Start a file with its associated application.

When `operation` is not specified or `'open'`, this acts like double-clicking the file in Windows Explorer, or giving the file name as an argument to the `start` command from the interactive command shell: the file is opened with whatever application (if any) its extension is associated.

When another `operation` is given, it must be a “command verb” that specifies what should be done with the file. Common verbs documented by Microsoft are `'print'` and `'edit'` (to be used on files) as well as `'explore'` and `'find'` (to be used on directories).

`startfile()` returns as soon as the associated application is launched. There is no option to wait for the application to close, and no way to retrieve the application’s exit status. The `path` parameter is relative to the current directory. If you want to use an absolute path, make sure the first character is not a slash (`'/'`); the underlying Win32 `ShellExecute()` function doesn’t work if it is. Use the `os.path.normpath()` function to ensure that the path is properly encoded for Win32.

To reduce interpreter startup overhead, the Win32 `ShellExecute()` function is not resolved until this function is first called. If the function cannot be resolved, `NotImplementedError` will be raised.

Availability: Windows.

`os.system(command)`

Execute the command (a string) in a subshell. This is implemented by calling the Standard C function `system()`, and has the same limitations. Changes to `sys.stdin`, etc. are not reflected in the environment of the executed command. If `command` generates any output, it will be sent to the interpreter standard output stream.

On Unix, the return value is the exit status of the process encoded in the format specified for `wait()`. Note that POSIX does not specify the meaning of the return value of the C `system()` function, so the return value of the Python function is system-dependent.

On Windows, the return value is that returned by the system shell after running `command`. The shell is given by the Windows environment variable `COMSPEC`; it is usually `cmd.exe`, which returns the exit status of the command run; on systems using a non-native shell, consult your shell documentation.

The `subprocess` module provides more powerful facilities for spawning new processes and retrieving their results; using that module is preferable to using this function. See the [Replacing Older Functions with the subprocess Module](#) section in the `subprocess` documentation for some helpful recipes.

Availability: Unix, Windows.

`os.times()`

Returns the current global process times. The return value is an object with five attributes:

- `user` - user time
- `system` - system time
- `children_user` - user time of all child processes
- `children_system` - system time of all child processes
- `elapsed` - elapsed real time since a fixed point in the past

For backwards compatibility, this object also behaves like a five-tuple containing `user`, `system`, `children_user`, `children_system`, and `elapsed` in that order.

See the Unix manual page `times(2)` or the corresponding Windows Platform API documentation. On Windows, only `user` and `system` are known; the other attributes are zero.

Availability: Unix, Windows.

在 3.3 版更改: 返回结果的类型由元组变成一个类似元组的对象, 同时具有命名的属性。

`os.wait()`

Wait for completion of a child process, and return a tuple containing its pid and exit status indication: a 16-bit number, whose low byte is the signal number that killed the process, and whose high byte is the exit status (if the signal number is zero); the high bit of the low byte is set if a core file was produced.

Availability: Unix.

`os.waitid(idtype, id, options)`

Wait for the completion of one or more child processes. *idtype* can be `P_PID`, `P_PGID` or `P_ALL`. *id* specifies the pid to wait on. *options* is constructed from the ORing of one or more of `WEXITED`, `WSTOPPED` or `WCONTINUED` and additionally may be ORed with `WNOHANG` or `WNOWAIT`. The return value is an object representing the data contained in the `siginfo_t` structure, namely: `si_pid`, `si_uid`, `si_signo`, `si_status`, `si_code` or `None` if `WNOHANG` is specified and there are no children in a waitable state.

Availability: Unix.

3.3 新版功能.

`os.P_PID`

`os.P_PGID`

`os.P_ALL`

These are the possible values for *idtype* in `waitid()`. They affect how *id* is interpreted.

Availability: Unix.

3.3 新版功能.

`os.WEXITED`

`os.WSTOPPED`

`os.WNOWAIT`

Flags that can be used in *options* in `waitid()` that specify what child signal to wait for.

Availability: Unix.

3.3 新版功能.

`os.CLD_EXITED`

`os.CLD_DUMPED`

`os.CLD_TRAPPED`

`os.CLD_CONTINUED`

These are the possible values for `si_code` in the result returned by `waitid()`.

Availability: Unix.

3.3 新版功能.

`os.waitpid(pid, options)`

The details of this function differ on Unix and Windows.

On Unix: Wait for completion of a child process given by process id *pid*, and return a tuple containing its process id and exit status indication (encoded as for `wait()`). The semantics of the call are affected by the value of the integer *options*, which should be 0 for normal operation.

If *pid* is greater than 0, `waitpid()` requests status information for that specific process. If *pid* is 0, the request is for the status of any child in the process group of the current process. If *pid* is -1, the request pertains to any child of the current process. If *pid* is less than -1, status is requested for any process in the process group -*pid* (the absolute value of *pid*).

An `OSError` is raised with the value of `errno` when the syscall returns -1.

On Windows: Wait for completion of a process given by process handle *pid*, and return a tuple containing *pid*, and its exit status shifted left by 8 bits (shifting makes cross-platform use of the function easier). A *pid* less than or equal to 0 has no special meaning on Windows, and raises an exception. The value of integer *options* has no effect. *pid* can refer to any process whose id is known, not necessarily a child process. The *spawn** functions called with *P_NOWAIT* return suitable process handles.

在 3.5 版更改: 如果系统调用被中断, 但信号处理程序没有触发异常, 此函数现在会重试系统调用, 而不是触发 *InterruptedError* 异常 (原因详见 [PEP 475](#))。

os.wait3(*options*)

Similar to *waitpid()*, except no process id argument is given and a 3-element tuple containing the child's process id, exit status indication, and resource usage information is returned. Refer to *resource.getrusage()* for details on resource usage information. The option argument is the same as that provided to *waitpid()* and *wait4()*.

Availability: Unix.

os.wait4(*pid, options*)

Similar to *waitpid()*, except a 3-element tuple, containing the child's process id, exit status indication, and resource usage information is returned. Refer to *resource.getrusage()* for details on resource usage information. The arguments to *wait4()* are the same as those provided to *waitpid()*.

Availability: Unix.

os.WNOHANG

The option for *waitpid()* to return immediately if no child process status is available immediately. The function returns (0, 0) in this case.

Availability: Unix.

os.WCONTINUED

This option causes child processes to be reported if they have been continued from a job control stop since their status was last reported.

Availability: some Unix systems.

os.WUNTRACED

This option causes child processes to be reported if they have been stopped but their current state has not been reported since they were stopped.

Availability: Unix.

The following functions take a process status code as returned by *system()*, *wait()*, or *waitpid()* as a parameter. They may be used to determine the disposition of a process.

os.WCOREDUMP(*status*)

Return True if a core dump was generated for the process, otherwise return False.

Availability: Unix.

os.WIFCONTINUED(*status*)

Return True if the process has been continued from a job control stop, otherwise return False.

Availability: Unix.

os.WIFSTOPPED(*status*)

Return True if the process has been stopped, otherwise return False.

Availability: Unix.

os.WIFSIGNALED(*status*)

Return True if the process exited due to a signal, otherwise return False.

Availability: Unix.

`os.WIFEXITED(status)`

Return `True` if the process exited using the `exit(2)` system call, otherwise return `False`.

Availability: Unix.

`os.WEXITSTATUS(status)`

If `WIFEXITED(status)` is true, return the integer parameter to the `exit(2)` system call. Otherwise, the return value is meaningless.

Availability: Unix.

`os.WSTOPSIG(status)`

Return the signal which caused the process to stop.

Availability: Unix.

`os.WTERMSIG(status)`

Return the signal which caused the process to exit.

Availability: Unix.

16.1.7 Interface to the scheduler

These functions control how a process is allocated CPU time by the operating system. They are only available on some Unix platforms. For more detailed information, consult your Unix manpages.

3.3 新版功能.

The following scheduling policies are exposed if they are supported by the operating system.

`os.SCHED_OTHER`

The default scheduling policy.

`os.SCHED_BATCH`

Scheduling policy for CPU-intensive processes that tries to preserve interactivity on the rest of the computer.

`os.SCHED_IDLE`

Scheduling policy for extremely low priority background tasks.

`os.SCHED_SPORADIC`

Scheduling policy for sporadic server programs.

`os.SCHED_FIFO`

A First In First Out scheduling policy.

`os.SCHED_RR`

A round-robin scheduling policy.

`os.SCHED_RESET_ON_FORK`

This flag can be OR'ed with any other scheduling policy. When a process with this flag set forks, its child's scheduling policy and priority are reset to the default.

class `os.sched_param(sched_priority)`

This class represents tunable scheduling parameters used in `sched_setparam()`, `sched_setscheduler()`, and `sched_getparam()`. It is immutable.

At the moment, there is only one possible parameter:

sched_priority

The scheduling priority for a scheduling policy.

`os.sched_get_priority_min(policy)`

Get the minimum priority value for *policy*. *policy* is one of the scheduling policy constants above.

- `os.sched_get_priority_max(policy)`
Get the maximum priority value for *policy*. *policy* is one of the scheduling policy constants above.
- `os.sched_setscheduler(pid, policy, param)`
Set the scheduling policy for the process with PID *pid*. A *pid* of 0 means the calling process. *policy* is one of the scheduling policy constants above. *param* is a *sched_param* instance.
- `os.sched_getscheduler(pid)`
Return the scheduling policy for the process with PID *pid*. A *pid* of 0 means the calling process. The result is one of the scheduling policy constants above.
- `os.sched_setparam(pid, param)`
Set a scheduling parameters for the process with PID *pid*. A *pid* of 0 means the calling process. *param* is a *sched_param* instance.
- `os.sched_getparam(pid)`
Return the scheduling parameters as a *sched_param* instance for the process with PID *pid*. A *pid* of 0 means the calling process.
- `os.sched_rr_get_interval(pid)`
Return the round-robin quantum in seconds for the process with PID *pid*. A *pid* of 0 means the calling process.
- `os.sched_yield()`
Voluntarily relinquish the CPU.
- `os.sched_setaffinity(pid, mask)`
Restrict the process with PID *pid* (or the current process if zero) to a set of CPUs. *mask* is an iterable of integers representing the set of CPUs to which the process should be restricted.
- `os.sched_getaffinity(pid)`
Return the set of CPUs the process with PID *pid* (or the current process if zero) is restricted to.

16.1.8 Miscellaneous System Information

- `os.confstr(name)`
Return string-valued system configuration values. *name* specifies the configuration value to retrieve; it may be a string which is the name of a defined system value; these names are specified in a number of standards (POSIX, Unix 95, Unix 98, and others). Some platforms define additional names as well. The names known to the host operating system are given as the keys of the `confstr_names` dictionary. For configuration variables not included in that mapping, passing an integer for *name* is also accepted.

If the configuration value specified by *name* isn't defined, `None` is returned.

If *name* is a string and is not known, `ValueError` is raised. If a specific value for *name* is not supported by the host system, even if it is included in `confstr_names`, an `OSError` is raised with `errno.EINVAL` for the error number.

Availability: Unix.
- `os.confstr_names`
Dictionary mapping names accepted by `confstr()` to the integer values defined for those names by the host operating system. This can be used to determine the set of names known to the system.

Availability: Unix.
- `os.cpu_count()`
Return the number of CPUs in the system. Returns `None` if undetermined.

该数量不同于当前进程可以使用的 CPU 数量。可用的 CPU 数量可以由 `len(os.sched_getaffinity(0))` 方法获得。

3.4 新版功能.

os.getloadavg()

Return the number of processes in the system run queue averaged over the last 1, 5, and 15 minutes or raises *OSError* if the load average was unobtainable.

Availability: Unix.

os.sysconf(name)

Return integer-valued system configuration values. If the configuration value specified by *name* isn't defined, -1 is returned. The comments regarding the *name* parameter for *confstr()* apply here as well; the dictionary that provides information on the known names is given by *sysconf_names*.

Availability: Unix.

os.sysconf_names

Dictionary mapping names accepted by *sysconf()* to the integer values defined for those names by the host operating system. This can be used to determine the set of names known to the system.

Availability: Unix.

The following data values are used to support path manipulation operations. These are defined for all platforms.

Higher-level operations on pathnames are defined in the *os.path* module.

os.curdir

The constant string used by the operating system to refer to the current directory. This is *'.'* for Windows and POSIX. Also available via *os.path*.

os.pardir

The constant string used by the operating system to refer to the parent directory. This is *'..'* for Windows and POSIX. Also available via *os.path*.

os.sep

The character used by the operating system to separate pathname components. This is *'/'* for POSIX and *'\\'* for Windows. Note that knowing this is not sufficient to be able to parse or concatenate pathnames —use *os.path.split()* and *os.path.join()* —but it is occasionally useful. Also available via *os.path*.

os.altsep

An alternative character used by the operating system to separate pathname components, or *None* if only one separator character exists. This is set to *'/'* on Windows systems where *sep* is a backslash. Also available via *os.path*.

os.extsep

The character which separates the base filename from the extension; for example, the *'.'* in *os.py*. Also available via *os.path*.

os.pathsep

The character conventionally used by the operating system to separate search path components (as in *PATH*), such as *':'* for POSIX or *';'* for Windows. Also available via *os.path*.

os.defpath

The default search path used by *exec*p** and *spawn*p** if the environment doesn't have a *'PATH'* key. Also available via *os.path*.

os.linesep

The string used to separate (or, rather, terminate) lines on the current platform. This may be a single character, such as *'\n'* for POSIX, or multiple characters, for example, *'\r\n'* for Windows. Do not use *os.linesep* as a line terminator when writing files opened in text mode (the default); use a single *'\n'* instead, on all platforms.

os.devnull

The file path of the null device. For example: `'/dev/null'` for POSIX, `'nul'` for Windows. Also available via `os.path`.

os.RTLD_LAZY**os.RTLD_NOW****os.RTLD_GLOBAL****os.RTLD_LOCAL****os.RTLD_NODELETE****os.RTLD_NOLOAD****os.RTLD_DEEPBIND**

Flags for use with the `setdlopenflags()` and `getdlopenflags()` functions. See the Unix manual page `dlopen(3)` for what the different flags mean.

3.3 新版功能.

16.1.9 Random numbers

os.getrandom(size, flags=0)

Get up to *size* random bytes. The function can return less bytes than requested.

These bytes can be used to seed user-space random number generators or for cryptographic purposes.

`getrandom()` relies on entropy gathered from device drivers and other sources of environmental noise. Unnecessarily reading large quantities of data will have a negative impact on other users of the `/dev/random` and `/dev/urandom` devices.

The flags argument is a bit mask that can contain zero or more of the following values ORed together: `os.GRND_RANDOM` and `GRND_NONBLOCK`.

See also the [Linux getrandom\(\) manual page](#).

Availability: Linux 3.17 and newer.

3.6 新版功能.

os.urandom(size)

Return a string of *size* random bytes suitable for cryptographic use.

This function returns random bytes from an OS-specific randomness source. The returned data should be unpredictable enough for cryptographic applications, though its exact quality depends on the OS implementation.

On Linux, if the `getrandom()` syscall is available, it is used in blocking mode: block until the system urandom entropy pool is initialized (128 bits of entropy are collected by the kernel). See the [PEP 524](#) for the rationale. On Linux, the `getrandom()` function can be used to get random bytes in non-blocking mode (using the `GRND_NONBLOCK` flag) or to poll until the system urandom entropy pool is initialized.

On a Unix-like system, random bytes are read from the `/dev/urandom` device. If the `/dev/urandom` device is not available or not readable, the `NotImplementedError` exception is raised.

On Windows, it will use `CryptGenRandom()`.

参见:

The `secrets` module provides higher level functions. For an easy-to-use interface to the random number generator provided by your platform, please see `random.SystemRandom`.

在 3.6.0 版更改: On Linux, `getrandom()` is now used in blocking mode to increase the security.

在 3.5.2 版更改: On Linux, if the `getrandom()` syscall blocks (the urandom entropy pool is not initialized yet), fall back on reading `/dev/urandom`.

在 3.5 版更改: On Linux 3.17 and newer, the `getrandom()` syscall is now used when available. On OpenBSD 5.6 and newer, the C `getentropy()` function is now used. These functions avoid the usage of an internal file descriptor.

`os.GRND_NONBLOCK`

By default, when reading from `/dev/random`, `getrandom()` blocks if no random bytes are available, and when reading from `/dev/urandom`, it blocks if the entropy pool has not yet been initialized.

If the `GRND_NONBLOCK` flag is set, then `getrandom()` does not block in these cases, but instead immediately raises `BlockingIOError`.

3.6 新版功能.

`os.GRND_RANDOM`

If this bit is set, then random bytes are drawn from the `/dev/random` pool instead of the `/dev/urandom` pool.

3.6 新版功能.

16.2 `io` —处理流的核心工具

源代码: [Lib/io.py](#)

16.2.1 概述

`io` 模块提供了 Python 用于处理各种 I/O 类型的主要工具。三种主要的 I/O 类型分别为: 文本 I/O, 二进制 I/O 和 原始 I/O。这些是泛型类型, 有很多种后端存储可以用在他们上面。一个隶属于任何这些类型的具体对象被称作 *file object*。其他同类的术语还有 流和 类文件对象。

独立于其类别, 每个具体流对象也将具有各种功能: 它可以是只读, 只写或读写。它还可以允许任意随机访问 (向前或向后寻找任何位置), 或仅允许顺序访问 (例如在套接字或管道的情况下)。

All streams are careful about the type of data you give to them. For example giving a `str` object to the `write()` method of a binary stream will raise a `TypeError`. So will giving a `bytes` object to the `write()` method of a text stream.

在 3.3 版更改: Operations that used to raise `IOError` now raise `OSError`, since `IOError` is now an alias of `OSError`.

Text I/O

Text I/O expects and produces `str` objects. This means that whenever the backing store is natively made of bytes (such as in the case of a file), encoding and decoding of data is made transparently as well as optional translation of platform-specific newline characters.

The easiest way to create a text stream is with `open()`, optionally specifying an encoding:

```
f = open("myfile.txt", "r", encoding="utf-8")
```

In-memory text streams are also available as `StringIO` objects:

```
f = io.StringIO("some initial text data")
```

The text stream API is described in detail in the documentation of `TextIOBase`.

Binary I/O

Binary I/O (also called *buffered I/O*) expects *bytes-like objects* and produces *bytes* objects. No encoding, decoding, or newline translation is performed. This category of streams can be used for all kinds of non-text data, and also when manual control over the handling of text data is desired.

The easiest way to create a binary stream is with `open()` with 'b' in the mode string:

```
f = open("myfile.jpg", "rb")
```

In-memory binary streams are also available as *BytesIO* objects:

```
f = io.BytesIO(b"some initial binary data: \x00\x01")
```

The binary stream API is described in detail in the docs of *BufferedIOBase*.

Other library modules may provide additional ways to create text or binary streams. See `socket.socket.makefile()` for example.

Raw I/O

Raw I/O (also called *unbuffered I/O*) is generally used as a low-level building-block for binary and text streams; it is rarely useful to directly manipulate a raw stream from user code. Nevertheless, you can create a raw stream by opening a file in binary mode with buffering disabled:

```
f = open("myfile.jpg", "rb", buffering=0)
```

The raw stream API is described in detail in the docs of *RawIOBase*.

16.2.2 High-level Module Interface

`io.DEFAULT_BUFFER_SIZE`

An int containing the default buffer size used by the module's buffered I/O classes. `open()` uses the file's blksize (as obtained by `os.stat()`) if possible.

`io.open(file, mode='r', buffering=-1, encoding=None, errors=None, newline=None, closefd=True, opener=None)`

This is an alias for the builtin `open()` function.

`exception io.BlockingIOError`

This is a compatibility alias for the builtin *BlockingIOError* exception.

`exception io.UnsupportedOperation`

An exception inheriting *OSError* and *ValueError* that is raised when an unsupported operation is called on a stream.

In-memory streams

It is also possible to use a *str* or *bytes-like object* as a file for both reading and writing. For strings *StringIO* can be used like a file opened in text mode. *BytesIO* can be used like a file opened in binary mode. Both provide full read-write capabilities with random access.

参见:

sys contains the standard IO streams: `sys.stdin`, `sys.stdout`, and `sys.stderr`.

16.2.3 Class hierarchy

The implementation of I/O streams is organized as a hierarchy of classes. First *abstract base classes* (ABCs), which are used to specify the various categories of streams, then concrete classes providing the standard stream implementations.

注解: The abstract base classes also provide default implementations of some methods in order to help implementation of concrete stream classes. For example, *BufferedIOBase* provides unoptimized implementations of `readinto()` and `readline()`.

At the top of the I/O hierarchy is the abstract base class *IOBase*. It defines the basic interface to a stream. Note, however, that there is no separation between reading and writing to streams; implementations are allowed to raise *UnsupportedOperation* if they do not support a given operation.

The *RawIOBase* ABC extends *IOBase*. It deals with the reading and writing of bytes to a stream. *FileIO* subclasses *RawIOBase* to provide an interface to files in the machine's file system.

The *BufferedIOBase* ABC deals with buffering on a raw byte stream (*RawIOBase*). Its subclasses, *BufferedWriter*, *BufferedReader*, and *BufferedRWPair* buffer streams that are readable, writable, and both readable and writable. *BufferedRandom* provides a buffered interface to random access streams. Another *BufferedIOBase* subclass, *BytesIO*, is a stream of in-memory bytes.

The *TextIOBase* ABC, another subclass of *IOBase*, deals with streams whose bytes represent text, and handles encoding and decoding to and from strings. *TextIOWrapper*, which extends it, is a buffered text interface to a buffered raw stream (*BufferedIOBase*). Finally, *StringIO* is an in-memory stream for text.

Argument names are not part of the specification, and only the arguments of `open()` are intended to be used as keyword arguments.

The following table summarizes the ABCs provided by the *io* module:

ABC	Inherits	Stub Methods	Mixin Methods and Properties
<i>IOBase</i>		<code>fileno</code> , <code>seek</code> , 和 <code>truncate</code>	<code>close</code> , <code>closed</code> , <code>__enter__</code> , <code>__exit__</code> , <code>flush</code> , <code>isatty</code> , <code>__iter__</code> , <code>__next__</code> , <code>readable</code> , <code>readline</code> , <code>readlines</code> , <code>seekable</code> , <code>tell</code> , <code>writable</code> , 和 <code>writelines</code>
<i>RawIOBase</i>	<i>IOBase</i>	<code>readinto</code> 和 <code>write</code>	Inherited <i>IOBase</i> methods, <code>read</code> , and <code>readall</code>
<i>BufferedIOBase</i>	<i>IOBase</i>	<code>detach</code> , <code>read</code> , <code>read1</code> , 和 <code>write</code>	Inherited <i>IOBase</i> methods, <code>readinto</code> , and <code>readinto1</code>
<i>TextIOBase</i>	<i>IOBase</i>	<code>detach</code> , <code>read</code> , <code>readline</code> , 和 <code>write</code>	Inherited <i>IOBase</i> methods, <code>encoding</code> , <code>errors</code> , and <code>newlines</code>

I/O Base Classes

class `io.IOBase`

The abstract base class for all I/O classes, acting on streams of bytes. There is no public constructor.

This class provides empty abstract implementations for many methods that derived classes can override selectively; the default implementations represent a file that cannot be read, written or seeked.

Even though `IOBase` does not declare `read()`, `readinto()`, or `write()` because their signatures will vary, implementations and clients should consider those methods part of the interface. Also, implementations may raise a `ValueError` (or `UnsupportedOperation`) when operations they do not support are called.

The basic type used for binary data read from or written to a file is `bytes`. Other *bytes-like objects* are accepted as method arguments too. In some cases, such as `readinto()`, a writable object such as `bytearray` is required. Text I/O classes work with `str` data.

Note that calling any method (even inquiries) on a closed stream is undefined. Implementations may raise `ValueError` in this case.

`IOBase` (and its subclasses) supports the iterator protocol, meaning that an `IOBase` object can be iterated over yielding the lines in a stream. Lines are defined slightly differently depending on whether the stream is a binary stream (yielding bytes), or a text stream (yielding character strings). See `readline()` below.

`IOBase` is also a context manager and therefore supports the `with` statement. In this example, `file` is closed after the `with` statement's suite is finished—even if an exception occurs:

```
with open('spam.txt', 'w') as file:
    file.write('Spam and eggs!')
```

`IOBase` provides these data attributes and methods:

close()

Flush and close this stream. This method has no effect if the file is already closed. Once the file is closed, any operation on the file (e.g. reading or writing) will raise a `ValueError`.

As a convenience, it is allowed to call this method more than once; only the first call, however, will have an effect.

closed

True if the stream is closed.

fileno()

Return the underlying file descriptor (an integer) of the stream if it exists. An `OSError` is raised if the IO object does not use a file descriptor.

flush()

Flush the write buffers of the stream if applicable. This does nothing for read-only and non-blocking streams.

isatty()

Return True if the stream is interactive (i.e., connected to a terminal/tty device).

readable()

Return True if the stream can be read from. If False, `read()` will raise `OSError`.

readline (size=-1)

Read and return one line from the stream. If `size` is specified, at most `size` bytes will be read.

The line terminator is always `b'\n'` for binary files; for text files, the `newline` argument to `open()` can be used to select the line terminator(s) recognized.

readlines (hint=-1)

Read and return a list of lines from the stream. `hint` can be specified to control the number of lines read: no more lines will be read if the total size (in bytes/characters) of all lines so far exceeds `hint`.

Note that it's already possible to iterate on file objects using `for line in file: ...` without calling `file.readlines()`.

seek (offset[, whence])

Change the stream position to the given byte `offset`. `offset` is interpreted relative to the position indicated by `whence`. The default value for `whence` is `SEEK_SET`. Values for `whence` are:

- `SEEK_SET` or 0 – start of the stream (the default); *offset* should be zero or positive
- `SEEK_CUR` or 1 – current stream position; *offset* may be negative
- `SEEK_END` or 2 – end of the stream; *offset* is usually negative

Return the new absolute position.

3.1 新版功能: The `SEEK_*` constants.

3.3 新版功能: Some operating systems could support additional values, like `os.SEEK_HOLE` or `os.SEEK_DATA`. The valid values for a file could depend on it being open in text or binary mode.

seekable()

Return True if the stream supports random access. If False, `seek()`, `tell()` and `truncate()` will raise `OSError`.

tell()

Return the current stream position.

truncate(size=None)

Resize the stream to the given *size* in bytes (or the current position if *size* is not specified). The current stream position isn't changed. This resizing can extend or reduce the current file size. In case of extension, the contents of the new file area depend on the platform (on most systems, additional bytes are zero-filled). The new file size is returned.

在 3.5 版更改: Windows will now zero-fill files when extending.

writable()

Return True if the stream supports writing. If False, `write()` and `truncate()` will raise `OSError`.

writelines(lines)

Write a list of lines to the stream. Line separators are not added, so it is usual for each of the lines provided to have a line separator at the end.

__del__()

Prepare for object destruction. `IOBase` provides a default implementation of this method that calls the instance's `close()` method.

class io.RawIOBase

Base class for raw binary I/O. It inherits `IOBase`. There is no public constructor.

Raw binary I/O typically provides low-level access to an underlying OS device or API, and does not try to encapsulate it in high-level primitives (this is left to Buffered I/O and Text I/O, described later in this page).

In addition to the attributes and methods from `IOBase`, `RawIOBase` provides the following methods:

read(size=-1)

Read up to *size* bytes from the object and return them. As a convenience, if *size* is unspecified or -1, all bytes until EOF are returned. Otherwise, only one system call is ever made. Fewer than *size* bytes may be returned if the operating system call returns fewer than *size* bytes.

If 0 bytes are returned, and *size* was not 0, this indicates end of file. If the object is in non-blocking mode and no bytes are available, `None` is returned.

The default implementation defers to `readall()` and `readinto()`.

readall()

Read and return all the bytes from the stream until EOF, using multiple calls to the stream if necessary.

readinto(b)

Read bytes into a pre-allocated, writable *bytes-like object* *b*, and return the number of bytes read. If the object is in non-blocking mode and no bytes are available, `None` is returned.

write(*b*)

Write the given *bytes-like object*, *b*, to the underlying raw stream, and return the number of bytes written. This can be less than the length of *b* in bytes, depending on specifics of the underlying raw stream, and especially if it is in non-blocking mode. *None* is returned if the raw stream is set not to block and no single byte could be readily written to it. The caller may release or mutate *b* after this method returns, so the implementation should only access *b* during the method call.

class io.BufferedIOBase

Base class for binary streams that support some kind of buffering. It inherits *IOBase*. There is no public constructor.

The main difference with *RawIOBase* is that methods *read()*, *readinto()* and *write()* will try (respectively) to read as much input as requested or to consume all given output, at the expense of making perhaps more than one system call.

In addition, those methods can raise *BlockingIOError* if the underlying raw stream is in non-blocking mode and cannot take or give enough data; unlike their *RawIOBase* counterparts, they will never return *None*.

Besides, the *read()* method does not have a default implementation that defers to *readinto()*.

A typical *BufferedIOBase* implementation should not inherit from a *RawIOBase* implementation, but wrap one, like *BufferedWriter* and *BufferedReader* do.

BufferedIOBase provides or overrides these methods and attribute in addition to those from *IOBase*:

raw

The underlying raw stream (a *RawIOBase* instance) that *BufferedIOBase* deals with. This is not part of the *BufferedIOBase* API and may not exist on some implementations.

detach()

Separate the underlying raw stream from the buffer and return it.

After the raw stream has been detached, the buffer is in an unusable state.

Some buffers, like *BytesIO*, do not have the concept of a single raw stream to return from this method. They raise *UnsupportedOperation*.

3.1 新版功能.

read(*size=-1*)

Read and return up to *size* bytes. If the argument is omitted, *None*, or negative, data is read and returned until EOF is reached. An empty *bytes* object is returned if the stream is already at EOF.

If the argument is positive, and the underlying raw stream is not interactive, multiple raw reads may be issued to satisfy the byte count (unless EOF is reached first). But for interactive raw streams, at most one raw read will be issued, and a short result does not imply that EOF is imminent.

A *BlockingIOError* is raised if the underlying raw stream is in non blocking-mode, and has no data available at the moment.

read1(*size=-1*)

Read and return up to *size* bytes, with at most one call to the underlying raw stream's *read()* (or *readinto()*) method. This can be useful if you are implementing your own buffering on top of a *BufferedIOBase* object.

readinto(*b*)

Read bytes into a pre-allocated, writable *bytes-like object* *b* and return the number of bytes read.

Like *read()*, multiple reads may be issued to the underlying raw stream, unless the latter is interactive.

A *BlockingIOError* is raised if the underlying raw stream is in non blocking-mode, and has no data available at the moment.

readinto1(b)

Read bytes into a pre-allocated, writable *bytes-like object* *b*, using at most one call to the underlying raw stream's *read()* (or *readinto()*) method. Return the number of bytes read.

A *BlockingIOError* is raised if the underlying raw stream is in non blocking-mode, and has no data available at the moment.

3.5 新版功能.

write(b)

Write the given *bytes-like object*, *b*, and return the number of bytes written (always equal to the length of *b* in bytes, since if the write fails an *OSError* will be raised). Depending on the actual implementation, these bytes may be readily written to the underlying stream, or held in a buffer for performance and latency reasons.

When in non-blocking mode, a *BlockingIOError* is raised if the data needed to be written to the raw stream but it couldn't accept all the data without blocking.

The caller may release or mutate *b* after this method returns, so the implementation should only access *b* during the method call.

Raw File I/O

class `io.FileIO(name, mode='r', closefd=True, opener=None)`

FileIO represents an OS-level file containing bytes data. It implements the *RawIOBase* interface (and therefore the *IOBase* interface, too).

The *name* can be one of two things:

- a character string or *bytes* object representing the path to the file which will be opened. In this case *closefd* must be *True* (the default) otherwise an error will be raised.
- an integer representing the number of an existing OS-level file descriptor to which the resulting *FileIO* object will give access. When the *FileIO* object is closed this *fd* will be closed as well, unless *closefd* is set to *False*.

The *mode* can be 'r', 'w', 'x' or 'a' for reading (default), writing, exclusive creation or appending. The file will be created if it doesn't exist when opened for writing or appending; it will be truncated when opened for writing. *FileExistsError* will be raised if it already exists when opened for creating. Opening a file for creating implies writing, so this mode behaves in a similar way to 'w'. Add a '+' to the mode to allow simultaneous reading and writing.

The *read()* (when called with a positive argument), *readinto()* and *write()* methods on this class will only make one system call.

A custom opener can be used by passing a callable as *opener*. The underlying file descriptor for the file object is then obtained by calling *opener* with (*name*, *flags*). *opener* must return an open file descriptor (passing *os.open* as *opener* results in functionality similar to passing *None*).

新创建的文件是不可继承的。

See the *open()* built-in function for examples on using the *opener* parameter.

在 3.3 版更改: The *opener* parameter was added. The 'x' mode was added.

在 3.4 版更改: 文件现在禁止继承。

In addition to the attributes and methods from *IOBase* and *RawIOBase*, *FileIO* provides the following data attributes:

mode

The mode as given in the constructor.

name

The file name. This is the file descriptor of the file when no name is given in the constructor.

Buffered Streams

Buffered I/O streams provide a higher-level interface to an I/O device than raw I/O does.

class `io.BytesIO([initial_bytes])`

A stream implementation using an in-memory bytes buffer. It inherits `BufferedIOBase`. The buffer is discarded when the `close()` method is called.

The optional argument `initial_bytes` is a *bytes-like object* that contains initial data.

`BytesIO` provides or overrides these methods in addition to those from `BufferedIOBase` and `IOBase`:

getbuffer()

Return a readable and writable view over the contents of the buffer without copying them. Also, mutating the view will transparently update the contents of the buffer:

```
>>> b = io.BytesIO(b"abcdef")
>>> view = b.getbuffer()
>>> view[2:4] = b"56"
>>> b.getvalue()
b'ab56ef'
```

注解: As long as the view exists, the `BytesIO` object cannot be resized or closed.

3.2 新版功能.

getvalue()

Return *bytes* containing the entire contents of the buffer.

read1()

In `BytesIO`, this is the same as `read()`.

readinto1()

In `BytesIO`, this is the same as `readinto()`.

3.5 新版功能.

class `io.BufferedReader(raw, buffer_size=DEFAULT_BUFFER_SIZE)`

A buffer providing higher-level access to a readable, sequential `RawIOBase` object. It inherits `BufferedIOBase`. When reading data from this object, a larger amount of data may be requested from the underlying raw stream, and kept in an internal buffer. The buffered data can then be returned directly on subsequent reads.

The constructor creates a `BufferedReader` for the given readable `raw` stream and `buffer_size`. If `buffer_size` is omitted, `DEFAULT_BUFFER_SIZE` is used.

`BufferedReader` provides or overrides these methods in addition to those from `BufferedIOBase` and `IOBase`:

peek([size])

Return bytes from the stream without advancing the position. At most one single read on the raw stream is done to satisfy the call. The number of bytes returned may be less or more than requested.

read([size])

Read and return `size` bytes, or if `size` is not given or negative, until EOF or if the read call would block in non-blocking mode.

read1 (*size*)

Read and return up to *size* bytes with only one call on the raw stream. If at least one byte is buffered, only buffered bytes are returned. Otherwise, one raw stream read call is made.

class `io.BufferedWriter` (*raw*, *buffer_size*=`DEFAULT_BUFFER_SIZE`)

A buffer providing higher-level access to a writeable, sequential `RawIOBase` object. It inherits `BufferedIOBase`. When writing to this object, data is normally placed into an internal buffer. The buffer will be written out to the underlying `RawIOBase` object under various conditions, including:

- when the buffer gets too small for all pending data;
- when `flush()` is called;
- when a `seek()` is requested (for `BufferedRandom` objects);
- when the `BufferedWriter` object is closed or destroyed.

The constructor creates a `BufferedWriter` for the given writeable *raw* stream. If the *buffer_size* is not given, it defaults to `DEFAULT_BUFFER_SIZE`.

`BufferedWriter` provides or overrides these methods in addition to those from `BufferedIOBase` and `IOBase`:

flush ()

Force bytes held in the buffer into the raw stream. A `BlockingIOError` should be raised if the raw stream blocks.

write (*b*)

Write the *bytes-like object*, *b*, and return the number of bytes written. When in non-blocking mode, a `BlockingIOError` is raised if the buffer needs to be written out but the raw stream blocks.

class `io.BufferedRandom` (*raw*, *buffer_size*=`DEFAULT_BUFFER_SIZE`)

A buffered interface to random access streams. It inherits `BufferedReader` and `BufferedWriter`, and further supports `seek()` and `tell()` functionality.

The constructor creates a reader and writer for a seekable raw stream, given in the first argument. If the *buffer_size* is omitted it defaults to `DEFAULT_BUFFER_SIZE`.

`BufferedRandom` is capable of anything `BufferedReader` or `BufferedWriter` can do.

class `io.BufferedRWPair` (*reader*, *writer*, *buffer_size*=`DEFAULT_BUFFER_SIZE`)

A buffered I/O object combining two unidirectional `RawIOBase` objects – one readable, the other writeable – into a single bidirectional endpoint. It inherits `BufferedIOBase`.

reader and *writer* are `RawIOBase` objects that are readable and writeable respectively. If the *buffer_size* is omitted it defaults to `DEFAULT_BUFFER_SIZE`.

`BufferedRWPair` implements all of `BufferedIOBase`'s methods except for `detach()`, which raises `UnsupportedOperation`.

警告: `BufferedRWPair` does not attempt to synchronize accesses to its underlying raw streams. You should not pass it the same object as reader and writer; use `BufferedRandom` instead.

Text I/O

class `io.TextIOBase`

Base class for text streams. This class provides a character and line based interface to stream I/O. There is no `readinto()` method because Python's character strings are immutable. It inherits `IOBase`. There is no public constructor.

TextIOBase provides or overrides these data attributes and methods in addition to those from *IOBase*:

encoding

The name of the encoding used to decode the stream's bytes into strings, and to encode strings into bytes.

errors

The error setting of the decoder or encoder.

newlines

A string, a tuple of strings, or *None*, indicating the newlines translated so far. Depending on the implementation and the initial constructor flags, this may not be available.

buffer

The underlying binary buffer (a *BufferedIOBase* instance) that *TextIOBase* deals with. This is not part of the *TextIOBase* API and may not exist in some implementations.

detach()

Separate the underlying binary buffer from the *TextIOBase* and return it.

After the underlying buffer has been detached, the *TextIOBase* is in an unusable state.

Some *TextIOBase* implementations, like *StringIO*, may not have the concept of an underlying buffer and calling this method will raise *UnsupportedOperation*.

3.1 新版功能.

read(size=-1)

Read and return at most *size* characters from the stream as a single *str*. If *size* is negative or *None*, reads until EOF.

readline(size=-1)

Read until newline or EOF and return a single *str*. If the stream is already at EOF, an empty string is returned.

If *size* is specified, at most *size* characters will be read.

seek(offset[, whence])

Change the stream position to the given *offset*. Behaviour depends on the *whence* parameter. The default value for *whence* is *SEEK_SET*.

- *SEEK_SET* or 0: seek from the start of the stream (the default); *offset* must either be a number returned by *TextIOBase.tell()*, or zero. Any other *offset* value produces undefined behaviour.
- *SEEK_CUR* or 1: “seek” to the current position; *offset* must be zero, which is a no-operation (all other values are unsupported).
- *SEEK_END* or 2: seek to the end of the stream; *offset* must be zero (all other values are unsupported).

Return the new absolute position as an opaque number.

3.1 新版功能: The *SEEK_** constants.

tell()

Return the current stream position as an opaque number. The number does not usually represent a number of bytes in the underlying binary storage.

write(s)

Write the string *s* to the stream and return the number of characters written.

class io.TextIOWrapper(buffer, encoding=None, errors=None, newline=None, line_buffering=False, write_through=False)

A buffered text stream over a *BufferedIOBase* binary stream. It inherits *TextIOBase*.

encoding gives the name of the encoding that the stream will be decoded or encoded with. It defaults to *locale.getpreferredencoding(False)*.

errors is an optional string that specifies how encoding and decoding errors are to be handled. Pass `'strict'` to raise a `ValueError` exception if there is an encoding error (the default of `None` has the same effect), or pass `'ignore'` to ignore errors. (Note that ignoring encoding errors can lead to data loss.) `'replace'` causes a replacement marker (such as `'?'`) to be inserted where there is malformed data. `'backslashreplace'` causes malformed data to be replaced by a backslashed escape sequence. When writing, `'xmlcharrefreplace'` (replace with the appropriate XML character reference) or `'namereplace'` (replace with `\N{...}` escape sequences) can be used. Any other error handling name that has been registered with `codecs.register_error()` is also valid.

newline controls how line endings are handled. It can be `None`, `''`, `'\n'`, `'\r'`, and `'\r\n'`. It works as follows:

- When reading input from the stream, if *newline* is `None`, *universal newlines* mode is enabled. Lines in the input can end in `'\n'`, `'\r'`, or `'\r\n'`, and these are translated into `'\n'` before being returned to the caller. If it is `''`, universal newlines mode is enabled, but line endings are returned to the caller untranslated. If it has any of the other legal values, input lines are only terminated by the given string, and the line ending is returned to the caller untranslated.
- 将输出写入流时，如果 *newline* 为 `None`，则写入的任何 `'\n'` 字符都将转换为系统默认行分隔符 `os.linesep`。如果 *newline* 是 `''` 或 `'\n'`，则不进行翻译。如果 *newline* 是任何其他合法值，则写入的任何 `'\n'` 字符将被转换为给定的字符串。

If *line_buffering* is `True`, `flush()` is implied when a call to write contains a newline character or a carriage return.

If *write_through* is `True`, calls to `write()` are guaranteed not to be buffered: any data written on the `TextIOWrapper` object is immediately handled to its underlying binary *buffer*.

在 3.3 版更改: The *write_through* argument has been added.

在 3.3 版更改: The default *encoding* is now `locale.getpreferredencoding(False)` instead of `locale.getpreferredencoding()`. Don't change temporary the locale encoding using `locale.setlocale()`, use the current locale encoding instead of the user preferred encoding.

`TextIOWrapper` provides one attribute in addition to those of `TextIOBase` and its parents:

line_buffering

Whether line buffering is enabled.

class `io.StringIO` (*initial_value*="", *newline*='\n')

An in-memory stream for text I/O. The text buffer is discarded when the `close()` method is called.

The initial value of the buffer can be set by providing *initial_value*. If newline translation is enabled, newlines will be encoded as if by `write()`. The stream is positioned at the start of the buffer.

The *newline* argument works like that of `TextIOWrapper`. The default is to consider only `\n` characters as ends of lines and to do no newline translation. If *newline* is set to `None`, newlines are written as `\n` on all platforms, but universal newline decoding is still performed when reading.

`StringIO` provides this method in addition to those from `TextIOBase` and its parents:

getvalue()

Return a `str` containing the entire contents of the buffer. Newlines are decoded as if by `read()`, although the stream position is not changed.

Example usage:

```
import io

output = io.StringIO()
output.write('First line.\n')
```

(下页继续)

(续上页)

```
print('Second line.', file=output)

# Retrieve file contents -- this will be
# 'First line.\nSecond line.\n'
contents = output.getvalue()

# Close object and discard memory buffer --
# .getvalue() will now raise an exception.
output.close()
```

class io.IncrementalNewlineDecoder

A helper codec that decodes newlines for *universal newlines* mode. It inherits *codecs.IncrementalDecoder*.

16.2.4 性能

This section discusses the performance of the provided concrete I/O implementations.

Binary I/O

By reading and writing only large chunks of data even when the user asks for a single byte, buffered I/O hides any inefficiency in calling and executing the operating system's unbuffered I/O routines. The gain depends on the OS and the kind of I/O which is performed. For example, on some modern OSes such as Linux, unbuffered disk I/O can be as fast as buffered I/O. The bottom line, however, is that buffered I/O offers predictable performance regardless of the platform and the backing device. Therefore, it is almost always preferable to use buffered I/O rather than unbuffered I/O for binary data.

Text I/O

Text I/O over a binary storage (such as a file) is significantly slower than binary I/O over the same storage, because it requires conversions between unicode and binary data using a character codec. This can become noticeable handling huge amounts of text data like large log files. Also, `TextIOWrapper.tell()` and `TextIOWrapper.seek()` are both quite slow due to the reconstruction algorithm used.

StringIO, however, is a native in-memory unicode container and will exhibit similar speed to *BytesIO*.

多线程

FileIO objects are thread-safe to the extent that the operating system calls (such as `read(2)` under Unix) they wrap are thread-safe too.

Binary buffered objects (instances of *BufferedReader*, *BufferedWriter*, *BufferedRandom* and *BufferedRWPair*) protect their internal structures using a lock; it is therefore safe to call them from multiple threads at once.

TextIOWrapper objects are not thread-safe.

Reentrancy

Binary buffered objects (instances of *BufferedReader*, *BufferedWriter*, *BufferedRandom* and *BufferedRWPair*) are not reentrant. While reentrant calls will not happen in normal situations, they can arise

from doing I/O in a `signal` handler. If a thread tries to re-enter a buffered object which it is already accessing, a `RuntimeError` is raised. Note this doesn't prohibit a different thread from entering the buffered object.

The above implicitly extends to text files, since the `open()` function will wrap a buffered object inside a `TextIOWrapper`. This includes standard streams and therefore affects the built-in function `print()` as well.

16.3 `time` — 时间的访问和转换

该模块提供了各种时间相关的函数。相关功能还可以参阅 `datetime` 和 `calendar` 模块。

尽管此模块始终可用，但并非所有平台上都提供所有功能。此模块中定义的大多数函数调用都具有相同名称的平台 C 库函数。因为这些函数的语义因平台而异，所以使用时最好查阅平台相关文档。

下面是一些术语和惯例的解释。

- **epoch** 是时间开始的点，并且取决于平台。对于 Unix，epoch 是 1970 年 1 月 1 日 00:00:00 (UTC)。要找出给定平台上的 epoch，请查看 `time.gmtime(0)`。
- 术语 **Unix 纪元秒数** 是指自国际标准时间 1970 年 1 月 1 日零时以来经过的总秒数，通常不包括 闰秒。在所有符合 POSIX 标准的平台上，闰秒都会从总秒数中被扣除。
- 此模块中的功能可能无法处理纪元之前或将来的远期日期和时间。未来的截止点由 C 库决定；对于 32 位系统，它通常在 2038 年。
- **2000 年 (Y2K) 问题**：Python 依赖于平台的 C 库，它通常没有 2000 年问题，因为所有日期和时间都在内部表示为自纪元以来的秒数。函数 `strptime()` 在给出 `%Y` 格式代码时可以解析 2 位数年份。当解析 2 位数年份时，它们将根据 POSIX 和 ISO C 标准进行转换：值 69–99 映射到 1969–1999，值 0–68 映射到 2000–2068。
- **UTC** 是协调世界时（以前称为格林威治标准时间，或 GMT）。缩写 UTC 不是错误，而是英语和法语之间的妥协。
- **DST** 是夏令时，在一年中的一部分时间（通常）调整时区一小时。DST 规则很神奇（由当地法律确定），并且每年都会发生变化。C 库有一个包含本地规则的表（通常是从系统文件中读取以获得灵活性），并且在这方面是 True Wisdom 的唯一来源。
- 各种实时函数的精度可能低于表示其值或参数的单位所建议的精度。例如，在大多数 Unix 系统上，时钟 “ticks” 仅为每秒 50 或 100 次。
- 另一方面，`time()` 和 `sleep()` 的精度优于它们的 Unix 等价物：时间表示为浮点数，`time()` 返回最准确的时间（使用 Unix `gettimeofday()` 如果可用），并且 `sleep()` 将接受非零分数的时间（Unix `select()` 用于实现此功能，如果可用）。
- 时间值由 `gmtime()`、`localtime()` 和 `strptime()` 返回，并被 `asctime()`、`mktime()` 和 `strftime()` 接受，是一个 9 个整数的序列。`gmtime()`、`localtime()` 和 `strptime()` 的返回值还提供各个字段的属性名称。

请参阅 `struct_time` 以获取这些对象的描述。

在 3.3 版更改：在平台支持相应的 `struct tm` 成员时，`struct_time` 类型被扩展提供 `tm_gmtoff` 和 `tm_zone` 属性。

在 3.6 版更改：`struct_time` 的属性 `tm_gmtoff` 和 `tm_zone` 现在可在所有平台上使用。

- 使用以下函数在时间表示之间进行转换：

从	到	使用
seconds since the epoch	UTC 的 <code>struct_time</code>	<code>gmtime()</code>
seconds since the epoch	本地时间的 <code>struct_time</code>	<code>localtime()</code>
UTC 的 <code>struct_time</code>	seconds since the epoch	<code>calendar.timegm()</code>
本地时间的 <code>struct_time</code>	seconds since the epoch	<code>mktime()</code>

16.3.1 函数

`time.asctime([t])`
转换一个元组或 `struct_time` 表示的时间，由 `gmtime()` 或 `localtime()` 返回为以下形式的字符串：'Sun Jun 20 23:21:05 1993'。如果未提供 `t`，则使用由 `localtime()` 返回的当前时间。区域信息不被函数 `asctime()` 使用。

注解：与同名的 C 函数不同，`asctime()` 不添加尾随换行符。

`time.clock()`
在 Unix 上，将当前处理器时间返回为以秒为单位的浮点数。精确度，实际上是“处理器时间”含义的定义，取决于同名 C 函数的精度。
在 Windows 上，此函数返回自第一次调用此函数以来经过的 `wallclock` 秒数，作为浮点数，基于 Win32 函数 `QueryPerformanceCounter()`。分辨率通常优于 1 微秒。
3.3 版后已移除：此函数的行为取决于平台：根据你的需求，使用 `perf_counter()` 或 `process_time()` 获得具有明确定义的行为。

`time.clock_getres(clk_id)`
返回指定时钟 `clk_id` 的分辨率（精度）。有关 `clk_id` 的可接受值列表，请参阅 `Clock ID 常量`。
Availability: Unix.
3.3 新版功能.

`time.clock_gettime(clk_id)`
返回指定 `clk_id` 时钟的时间。有关 `clk_id` 的可接受值列表，请参阅 `Clock ID 常量`。
Availability: Unix.
3.3 新版功能.

`time.clock_settime(clk_id, time)`
设置指定 `clk_id` 时钟的时间。目前，`CLOCK_REALTIME` 是 `clk_id` 唯一可接受的值。
Availability: Unix.
3.3 新版功能.

`time.ctime([secs])`
将自 seconds since the epoch 表示的时间转换为表示本地时间的字符串。如果未提供 `secs` 或为 `None`，则使用由 `time()` 返回的当前时间。`ctime(secs)` 相当于 `asctime(localtime(secs))`。区域信息不被 `ctime()` 使用。

`time.get_clock_info(name)`
获取有关指定时钟的信息作为命名空间对象。支持的时钟名称和读取其值的相应函数是：

- 'clock': `time.clock()`
- 'monotonic': `time.monotonic()`
- 'perf_counter': `time.perf_counter()`

- `'process_time': time.process_time()`
- `'time': time.time()`

结果具有以下属性:

- *adjustable*: 如果时钟可以自动更改 (例如通过 NTP 守护程序) 或由系统管理员手动更改, 则为 `True`, 否则为 `False`。
- *implementation*: 用于获取时钟值的基础 C 函数的名称。有关可能的值, 请参阅 *Clock ID 常量*。
- *monotonic*: 如果时钟不能倒退, 则为 `True`, 否则为 `False`。
- *resolution*: 以秒为单位的时钟分辨率 (*float*)

3.3 新版功能.

`time.gmtime([secs])`

将 seconds since the epoch 为单位的时间转换为 UTC 的 *struct_time*, 其中 *dst* 标志始终为零。如果未提供 *secs* 或为 `None`, 则使用由 *time()* 返回的当前时间。忽略一秒的分数。有关 *struct_time* 对象的说明, 请参见上文。有关此函数的反函数, 请参阅 *calendar.timegm()*。

`time.localtime([secs])`

与 *gmtime()* 相似但转换为当地时间。如果未提供 *secs* 或为 `None`, 则使用由 *time()* 返回的当前时间。当 DST 适用于给定时间时, *dst* 标志设置为 1。

`time.mktime(t)`

这是 *localtime()* 的反函数。它的参数是 *struct_time* 或者完整的 9 元组 (因为需要 *dst* 标志; 如果它是未知的则使用 -1 作为 *dst* 标志), 它表示 *local* 的时间, 而不是 UTC。它返回一个浮点数, 以便与 *time()* 兼容。如果输入值不能表示为有效时间, 则 *OverflowError* 或 *ValueError* 将被引发 (这取决于 Python 或底层 C 库是否捕获到无效值)。它可以生成时间的最早日期取决于平台。

`time.monotonic()`

返回单调时钟的值 (以小数秒为单位), 即不能倒退的时钟。时钟不受系统时钟更新的影响。返回值的参考点未定义, 因此只有连续调用结果之间的差异才有效。

On Windows versions older than Vista, *monotonic()* detects `GetTickCount()` integer overflow (32 bits, roll-over after 49.7 days). It increases an internal epoch (reference time) by 2^{32} each time that an overflow is detected. The epoch is stored in the process-local state and so the value of *monotonic()* may be different in two Python processes running for more than 49 days. On more recent versions of Windows and on other operating systems, *monotonic()* is system-wide.

3.3 新版功能.

在 3.5 版更改: The function is now always available.

`time.perf_counter()`

返回性能计数器的值 (以小数秒为单位), 即具有最高可用分辨率的时钟, 以测量短持续时间。它确实包括睡眠期间经过的时间, 并且是系统范围的。返回值的参考点未定义, 因此只有连续调用结果之间的差异才有效。

3.3 新版功能.

`time.process_time()`

返回当前进程的系统和用户 CPU 时间总和的值 (以小数秒为单位)。它不包括睡眠期间经过的时间。根据定义, 它在整个进程范围中。返回值的参考点未定义, 因此只有连续调用结果之间的差异才有效。

3.3 新版功能.

`time.sleep(secs)`

暂停执行调用线程达到给定的秒数。参数可以是浮点数, 以指示更精确的睡眠时间。实际的暂停时间可能小于请求的时间, 因为任何捕获的信号将在执行该信号的捕获例程后终止 *sleep()*。此外, 由于系统中其他活动的安排, 暂停时间可能比请求的时间长任意量。

在 3.5 版更改: 即使睡眠被信号中断, 该函数现在至少睡眠 *secs*, 除非信号处理程序引发异常 (参见 [PEP 475](#) 作为基本原理)。

`time.strptime(format[, t])`
转换一个元组或 `struct_time` 表示的由 `gmtime()` 或 `localtime()` 返回的时间到由 `format` 参数指定的字符串。如果未提供 `t`, 则使用由 `localtime()` 返回的当前时间。`format` 必须是一个字符串。如果 `t` 中的任何字段超出允许范围, 则引发 `ValueError`。

0 是时间元组中任何位置的合法参数; 如果它通常是非法的, 则该值被强制改为正确的值。

以下指令可以嵌入 `format` 字符串中。它们显示时没有可选的字段宽度和精度规范, 并被 `strptime()` 结果中的指示字符替换:

指令	意义	注释
%a	本地化的缩写星期中每日的名称。	
%A	本地化的星期中每日的完整名称。	
%b	本地化的月缩写名称。	
%B	本地化的月完整名称。	
%c	本地化的适当日期和时间表示。	
%d	十进制数 [01,31] 表示的月中日。	
%H	十进制数 [00,23] 表示的小时 (24 小时制)。	
%I	十进制数 [01,12] 表示的小时 (12 小时制)。	
%j	十进制数 [001,366] 表示的年中日。	
%m	十进制数 [01,12] 表示的月。	
%M	十进制数 [00,59] 表示的分钟。	
%p	本地化的 AM 或 PM。	(1)
%S	十进制数 [00,61] 表示的秒。	(2)
%U	十进制数 [00,53] 表示的一年中的周数 (星期日作为一周的第一天) 作为。在第一个星期日之前的新年中的所有日子都被认为是在第 0 周。	(3)
%w	十进制数 [0(星期日),6] 表示的周中日。	
%W	十进制数 [00,53] 表示的一年中的周数 (星期一作为一周的第一天) 作为。在第一个星期一之前的新年中的所有日子被认为是在第 0 周。	(3)
%x	本地化的适当日期表示。	
%X	本地化的适当时间表示。	
%y	十进制数 [00,99] 表示的没有世纪的年份。	
%Y	十进制数表示的带世纪的年份。	
%z	时区偏移以格式 +HHMM 或 -HHMM 形式的 UTC/GMT 的正或负时差指示, 其中 H 表示十进制小时数字, M 表示小数分钟数字 [-23:59, +23:59]。	
%Z	时区名称 (如果不存在时区, 则不包含字符)。	
%%	字面的 '%' 字符。	

注释:

- (1) 当与 `strptime()` 函数一起使用时, 如果使用 `%I` 指令来解析小时, `%p` 指令只影响输出小时字段。
- (2) 范围真的是 0 到 61; 值 60 在表示 `leap seconds` 的时间戳中有效, 并且由于历史原因支持值 61。
- (3) 当与 `strptime()` 函数一起使用时, `%U` 和 `%W` 仅用于指定星期几和年份的计算。

下面是一个示例, 一个与 [RFC 2822](#) Internet 电子邮件标准以兼容的日期格式。¹

¹ 现在不推荐使用 `%z`, 但是所有 ANSI C 库都不支持扩展为首选小时/分钟偏移量的 “%z” 转义符。此外, 严格的 1982 年原始 [RFC 822](#) 标准要求两位数的年份 (`%y` 而不是 `%Y`), 但是实际在 2000 年之前很久就转移到了 4 位数年。之后, [RFC 822](#) 已经废弃了, 4 位数的年份首先被推荐 [RFC 1123](#), 然后被 [RFC 2822](#) 强制执行。

```
>>> from time import gmtime, strftime
>>> strftime("%a, %d %b %Y %H:%M:%S +0000", gmtime())
'Thu, 28 Jun 2001 14:17:15 +0000'
```

某些平台可能支持其他指令，但只有此处列出的指令具有 ANSI C 标准化的含义。要查看平台支持的完整格式代码集，请参阅 `strftime(3)` 文档。

在某些平台上，可选的字段宽度和精度规范可以按照以下顺序紧跟在指令的初始 '%' 之后；这也不可移植。字段宽度通常为 2，除了 %j，它是 3。

`time.strptime(string[, format])`

根据格式解析表示时间的字符串。返回值为一个被 `gmtime()` 或 `localtime()` 返回的 `struct_time`。

`format` 参数使用与 `strftime()`；使用的指令相同的指令。它默认为匹配 `ctime()` 返回格式的 "%a %b %d %H:%M:%S %Y"。如果 `string` 不能根据 `*format` 解析，或者解析后它有多余的数据，则引发 `ValueError`。当无法推断出更准确的值时，用于填充任何缺失数据的默认值是 (1900, 1, 1, 0, 0, 0, 0, 1, -1)。`string` 和 `format` 都必须是字符串。

例如：

```
>>> import time
>>> time.strptime("30 Nov 00", "%d %b %y") # doctest: +NORMALIZE_WHITESPACE
time.struct_time(tm_year=2000, tm_mon=11, tm_mday=30, tm_hour=0, tm_min=0,
                  tm_sec=0, tm_wday=3, tm_yday=335, tm_isdst=-1)
```

支持 %Z 指令是基于 `tzname` 中包含的值以及 `daylight` 是否为真。因此，它是特定于平台的，除了识别始终已知的 UTC 和 GMT（并且被认为是非夏令时时区）。

仅支持文档中指定的指令。因为每个平台都实现了 `strftime()`，它有时会提供比列出的指令更多的指令。但是 `strptime()` 独立于任何平台，因此不一定支持所有未记录为支持的可用指令。

class `time.struct_time`

返回的时间值序列的类型为 `gmtime()`、`localtime()` 和 `strptime()`。它是一个带有 *named tuple* 接口的对象：可以通过索引和属性名访问值。存在以下值：

索引	属性	值
0	<code>tm_year</code>	(例如, 1993)
1	<code>tm_mon</code>	range [1, 12]
2	<code>tm_mday</code>	range [1, 31]
3	<code>tm_hour</code>	range [0, 23]
4	<code>tm_min</code>	range [0, 59]
5	<code>tm_sec</code>	range [0, 61]; 见 <code>strftime()</code> 介绍中的 (2)
6	<code>tm_wday</code>	range [0, 6], 周一为 0
7	<code>tm_yday</code>	range [1, 366]
8	<code>tm_isdst</code>	0, 1 或 -1; 如下所示
N/A	<code>tm_zone</code>	时区名称的缩写
N/A	<code>tm_gmtoff</code>	以秒为单位的 UTC 以东偏离

请注意，与 C 结构不同，月份值是 [1,12] 的范围，而不是 [0,11]。

在调用 `mktime()` 时，`tm_isdst` 可以在夏令时生效时设置为 1，而在夏令时不生效时设置为 0。值 -1 表示这是未知的，并且通常会导致填写正确的状态。

当一个长度不正确的元组被传递给期望 `struct_time` 的函数，或者具有错误类型的元素时，会引发 `TypeError`。

`time.time()`

返回浮点数的 seconds since the *epoch*。*epoch* 的具体日期和 *leap seconds* 的处理取决于平台。在 Windows 和大多数 Unix 系统上，*epoch* 是 1970 年 1 月 1 日 00:00:00 (UTC)，并且闰秒不计入 seconds since the *epoch*。这通常被称为 *Unix time*。要找出给定平台上的 *epoch*，请查看 `gmtime(0)`。

请注意，即使时间总是作为浮点数返回，但并非所有系统都提供高于 1 秒的精度。虽然此函数通常返回非递减值，但如果在两次调用之间设置了系统时钟，则它可以返回比先前调用更低的值。

返回的数字 `time()` 可以通过将其传递给 `gmtime()` 函数或转换为 UTC 中更常见的时间格式（即年、月、日、小时等）或通过将它传递给 `localtime()` 函数获得本地时间。在这两种情况下都返回一个 `struct_time` 对象，日历日期组件可以从中作为属性访问。

`time.tzset()`

Resets the time conversion rules used by the library routines. The environment variable TZ specifies how this is done.

Availability: Unix.

注解： 虽然在很多情况下，更改 TZ 环境变量而不调用 `tzset()` 可能会影响函数的输出，例如 `localtime()`，不应该依赖此行为。

TZ 不应该包含空格。

TZ 环境变量的标准格式是（为了清晰起见，添加了空格）：

```
std offset [dst [offset [,start[/time], end[/time]]]]
```

组件的位置是：

std 和 **dst** 三个或更多字母数字，给出时区缩写。这些将传到 `time.tzname`

offset 偏移量的形式为： $\pm hh[:mm[:ss]]$ 。这表示添加到达 UTC 的本地时间的值。如果前面有 ‘-’，则时区位于本初子午线的东边；否则，在它是西边。如果 **dst** 之后没有偏移，则假设夏令时比标准时间提前一小时。

start[/time], end[/time] 指示何时更改为 DST 和从 DST 返回。开始日期和结束日期的格式为以下之一：

Jn Julian 日 n ($1 \leq n \leq 365$)。闰日不计算在内，因此在所有年份中，2 月 28 日是第 59 天，3 月 1 日是第 60 天。

n 从零开始的 Julian 日 ($0 \leq n \leq 365$)。闰日计入，可以引用 2 月 29 日。

Mm.n.d 一年中 m 月的第 n 周 ($1 \leq n \leq 5$ ， $1 \leq m \leq 12$ ，第 5 周表示“可能在 m 月第 4 周或第 5 周出现的最后第 d 日”) 的第 d 天 ($0 \leq d \leq 6$)。第 1 周是第 d 天发生的第一周。第 0 天是星期天。

`time` 的格式与 `offset` 的格式相同，但不允许使用前导符号（‘-’或‘+’）。如果没有给出时间，则默认值为 02:00:00。

```
>>> os.environ['TZ'] = 'EST+05EDT,M4.1.0,M10.5.0'
>>> time.tzset()
>>> time.strftime('%X %x %Z')
'02:07:36 05/08/03 EDT'
>>> os.environ['TZ'] = 'AEST-10AEDT-11,M10.5.0,M3.5.0'
>>> time.tzset()
>>> time.strftime('%X %x %Z')
'16:08:12 05/08/03 AEST'
```

在许多 Unix 系统（包括 *BSD，Linux，Solaris 和 Darwin 上），使用系统的区域信息（*tzfile(5)*）数据库来指定时区规则会更方便。为此，将 TZ 环境变量设置为所需时区数据文件的路径，相对于系统 ‘zoneinfo’ 时区数据库的根目录，通常位于 /usr/share/zoneinfo。例如，‘US/Eastern’、‘Australia/Melbourne’、‘Egypt’ 或 ‘Europe/Amsterdam’。

```
>>> os.environ['TZ'] = 'US/Eastern'
>>> time.tzset()
>>> time.tzname
('EST', 'EDT')
>>> os.environ['TZ'] = 'Egypt'
>>> time.tzset()
>>> time.tzname
('EET', 'EEST')
```

16.3.2 Clock ID 常量

这些常量用作 *clock_getres()* 和 *clock_gettime()* 的参数。

time.CLOCK_HIGHRES

Solaris OS 有一个 CLOCK_HIGHRES 计时器，试图使用最佳硬件源，并可能提供接近纳秒的分辨率。CLOCK_HIGHRES 是不可调节的高分辨率时钟。

Availability: Solaris.

3.3 新版功能.

time.CLOCK_MONOTONIC

无法设置的时钟，表示自某些未指定的起点以来的单调时间。

Availability: Unix.

3.3 新版功能.

time.CLOCK_MONOTONIC_RAW

类似于 *CLOCK_MONOTONIC*，但可以访问不受 NTP 调整影响的原始硬件时间。

Availability: Linux 2.6.28 or later.

3.3 新版功能.

time.CLOCK_PROCESS_CPUTIME_ID

来自 CPU 的高分辨率每进程计时器。

Availability: Unix.

3.3 新版功能.

time.CLOCK_THREAD_CPUTIME_ID

特定于线程的 CPU 时钟。

Availability: Unix.

3.3 新版功能.

以下常量是唯一可以发送到 *clock_settime()* 的参数。

time.CLOCK_REALTIME

系统范围的实时时钟。设置此时钟需要适当的权限。

Availability: Unix.

3.3 新版功能.

16.3.3 时区常量

`time.altzone`

本地 DST 时区的偏移量，以 UTC 为单位的秒数，如果已定义。如果当地 DST 时区在 UTC 以东（如在西欧，包括英国），则是负数。只有当 `daylight` 非零时才使用它。见下面的注释。

`time.daylight`

如果定义了 DST 时区，则为非零。见下面的注释。

`time.timezone`

本地（非 DST）时区的偏移量，UTC 以西的秒数（西欧大部分地区为负，美国为正，英国为零）。见下面的注释。

`time.tzname`

两个字符串的元组：第一个是本地非 DST 时区的名称，第二个是本地 DST 时区的名称。如果未定义 DST 时区，则不应使用第二个字符串。见下面的注释。

注解：对于上述时区常量（`altzone`、`daylight`、`timezone` 和 `tzname`），该值由模块加载时有效的时区规则确定，或者最后一次 `tzset()` 被调用时，并且在过去的时间可能不正确。建议使用来自 `localtime()` 结果的 `tm_gmtoff` 和 `tm_zone` 来获取时区信息。

参见：

模块 `datetime` 更多面向对象的日期和时间接口。

模块 `locale` 国际化服务。区域设置会影响 `strftime()` 和 `strptime()` 中许多格式说明符的解析。

模块 `calendar` 一般日历相关功能。这个模块的 `timegm()` 是函数 `gmtime()` 的反函数。

16.4 argparse — 命令行选项、参数和子命令解析器

3.2 新版功能.

源代码： [Lib/argparse.py](#)

教程

此页面包含该 API 的参考信息。有关 Python 命令行解析更细致的介绍，请参阅 `argparse` 教程。

`argparse` 模块可以让人轻松编写用户友好的命令行接口。程序定义它需要的参数，然后 `argparse` 将弄清如何从 `sys.argv` 解析出那些参数。`argparse` 模块还会自动生成帮助和使用手册，并在用户给程序传入无效参数时报出错误信息。

16.4.1 示例

以下代码是一个 Python 程序，它获取一个整数列表并计算总和或者最大值：

```
import argparse

parser = argparse.ArgumentParser(description='Process some integers.')
parser.add_argument('integers', metavar='N', type=int, nargs='+',
```

(下页继续)

(续上页)

```

        help='an integer for the accumulator')
parser.add_argument('--sum', dest='accumulate', action='store_const',
                    const=sum, default=max,
                    help='sum the integers (default: find the max)')

args = parser.parse_args()
print(args.accumulate(args.integers))

```

假设上面的 Python 代码保存在名为 prog.py 的文件中，它可以在命令行运行并提供有用的帮助消息：

```

$ python prog.py -h
usage: prog.py [-h] [--sum] N [N ...]

Process some integers.

positional arguments:
  N              an integer for the accumulator

optional arguments:
  -h, --help    show this help message and exit
  --sum         sum the integers (default: find the max)

```

当使用适当的参数运行时，它会输出命令行传入整数的总和或者最大值：

```

$ python prog.py 1 2 3 4
4

$ python prog.py 1 2 3 4 --sum
10

```

如果传入无效参数，则会报出错误：

```

$ python prog.py a b c
usage: prog.py [-h] [--sum] N [N ...]
prog.py: error: argument N: invalid int value: 'a'

```

以下部分将引导你完成这个示例。

创建一个解析器

使用 `argparse` 的第一步是创建一个 `ArgumentParser` 对象：

```
>>> parser = argparse.ArgumentParser(description='Process some integers.')
```

`ArgumentParser` 对象包含将命令行解析成 Python 数据类型所需的全部信息。

添加参数

给一个 `ArgumentParser` 添加程序参数信息是通过调用 `add_argument()` 方法完成的。通常，这些调用指定 `ArgumentParser` 如何获取命令行字符串并将其转换为对象。这些信息在 `parse_args()` 调用时被存储和使用。例如：

```
>>> parser.add_argument('integers', metavar='N', type=int, nargs='+',
...                       help='an integer for the accumulator')
>>> parser.add_argument('--sum', dest='accumulate', action='store_const',
...                       const=sum, default=max,
...                       help='sum the integers (default: find the max)')
```

稍后，调用 `parse_args()` 将返回一个具有 `integers` 和 `accumulate` 两个属性的对象。`integers` 属性将是一个包含一个或多个整数的列表，而 `accumulate` 属性当命令行中指定了 `--sum` 参数时将是 `sum()` 函数，否则则是 `max()` 函数。

解析参数

`ArgumentParser` 通过 `parse_args()` 方法解析参数。它将检查命令行，把每个参数转换为适当的类型然后调用相应的操作。在大多数情况下，这意味着一个简单的 `Namespace` 对象将从命令行参数中解析出的属性构建：

```
>>> parser.parse_args(['--sum', '7', '-1', '42'])
Namespace(accumulate=<built-in function sum>, integers=[7, -1, 42])
```

在脚本中，通常 `parse_args()` 会被不带参数调用，而 `ArgumentParser` 将自动从 `sys.argv` 中确定命令行参数。

16.4.2 ArgumentParser 对象

```
class argparse.ArgumentParser(prog=None, usage=None, description=None, epilog=None,
                               parents=[], formatter_class=argparse.HelpFormatter, prefix_chars='-',
                               fromfile_prefix_chars=None, argument_default=None, conflict_handler='error', add_help=True, allow_abbrev=True)
```

创建一个新的 `ArgumentParser` 对象。所有的参数都应当作为关键字参数传入。每个参数在下面都有它更详细的描述，但简而言之，它们是：

- `prog` - 程序的名称（默认： `sys.argv[0]`）
- `usage` - 描述程序用途的字符串（默认值：从添加到解析器的参数生成）
- `description` - 在参数帮助文档之前显示的文本（默认值：无）
- `epilog` - 在参数帮助文档之后显示的文本（默认值：无）
- `parents` - 一个 `ArgumentParser` 对象的列表，它们的参数也应包含在内
- `formatter_class` - 用于自定义帮助文档输出格式的类
- `prefix_chars` - 可选参数的前缀字符集合（默认值：' - '）
- `fromfile_prefix_chars` - 当需要从文件中读取其他参数时，用于标识文件名的前缀字符集合（默认值：None）
- `argument_default` - 参数的全局默认值（默认值：None）
- `conflict_handler` - 解决冲突选项的策略（通常是不必要的）
- `add_help` - 为解析器添加一个 `-h/--help` 选项（默认值：True）
- `allow_abbrev` - 如果缩写是无歧义的，则允许缩写长选项（默认值：True）

在 3.5 版更改：添加 `allow_abbrev` 参数。

以下部分描述这些参数如何使用。

prog

默认情况下, *ArgumentParser* 对象使用 `sys.argv[0]` 来确定如何在帮助消息中显示程序名称。这一默认值几乎总是可取的, 因为它将使帮助消息与从命令行调用此程序的方式相匹配。例如, 对于有如下代码的名为 `myprogram.py` 的文件:

```
import argparse
parser = argparse.ArgumentParser()
parser.add_argument('--foo', help='foo help')
args = parser.parse_args()
```

该程序的帮助信息将显示 `myprogram.py` 作为程序名称 (无论程序从何处被调用):

```
$ python myprogram.py --help
usage: myprogram.py [-h] [--foo FOO]

optional arguments:
  -h, --help  show this help message and exit
  --foo FOO   foo help
$ cd ..
$ python subdir/myprogram.py --help
usage: myprogram.py [-h] [--foo FOO]

optional arguments:
  -h, --help  show this help message and exit
  --foo FOO   foo help
```

要更改这样的默认行为, 可以使用 `prog=` 参数为 *ArgumentParser* 提供另一个值:

```
>>> parser = argparse.ArgumentParser(prog='myprogram')
>>> parser.print_help()
usage: myprogram [-h]

optional arguments:
  -h, --help  show this help message and exit
```

需要注意的是, 无论是从 `sys.argv[0]` 或是从 `prog=` 参数确定的程序名称, 都可以在帮助消息里通过 `%(prog)s` 格式串来引用。

```
>>> parser = argparse.ArgumentParser(prog='myprogram')
>>> parser.add_argument('--foo', help='foo of the %(prog)s program')
>>> parser.print_help()
usage: myprogram [-h] [--foo FOO]

optional arguments:
  -h, --help  show this help message and exit
  --foo FOO   foo of the myprogram program
```

usage

默认情况下, *ArgumentParser* 根据它包含的参数来构建用法消息:

```
>>> parser = argparse.ArgumentParser(prog='PROG')
>>> parser.add_argument('--foo', nargs='?', help='foo help')
>>> parser.add_argument('bar', nargs='+', help='bar help')
>>> parser.print_help()
```

(下页继续)

(续上页)

```
usage: PROG [-h] [--foo [FOO]] bar [bar ...]

positional arguments:
  bar                bar help

optional arguments:
  -h, --help        show this help message and exit
  --foo [FOO]       foo help
```

可以通过 `usage=` 关键字参数覆盖这一默认消息：

```
>>> parser = argparse.ArgumentParser(prog='PROG', usage='% (prog)s [options]')
>>> parser.add_argument('--foo', nargs='?', help='foo help')
>>> parser.add_argument('bar', nargs='+', help='bar help')
>>> parser.print_help()
usage: PROG [options]

positional arguments:
  bar                bar help

optional arguments:
  -h, --help        show this help message and exit
  --foo [FOO]       foo help
```

在用法消息中可以使用 `%(prog)s` 格式说明符来填入程序名称。

description

大多数对 `ArgumentParser` 构造方法的调用都会使用 `description=` 关键字参数。这个参数简要描述这个程度做什么以及怎么做。在帮助消息中，这个描述会显示在命令行用法字符串和各种参数的帮助消息之间：

```
>>> parser = argparse.ArgumentParser(description='A foo that bars')
>>> parser.print_help()
usage: argparse.py [-h]

A foo that bars

optional arguments:
  -h, --help  show this help message and exit
```

在默认情况下，`description` 将被换行以便适应给定的空间。如果想改变这种行为，见 `formatter_class` 参数。

epilog

一些程序喜欢在 `description` 参数后显示额外的对程序的描述。这种文字能够通过给 `ArgumentParser::` 提供 `epilog=` 参数而被指定。

```
>>> parser = argparse.ArgumentParser(
...     description='A foo that bars',
...     epilog="And that's how you'd foo a bar")
>>> parser.print_help()
usage: argparse.py [-h]

A foo that bars

And that's how you'd foo a bar
```

(下页继续)

(续上页)

```
optional arguments:
  -h, --help  show this help message and exit
```

```
And that's how you'd foo a bar
```

和 *description* 参数一样, *epilog= text* 在默认情况下会换行, 但是这种行为能够被调整通过提供 *formatter_class* 参数给 *ArgumentParser*.

parents

有些时候, 少数解析器会使用同一系列参数。单个解析器能够通过提供 *parents=* 参数给 *ArgumentParser* 而使用相同的参数而不是重复这些参数的定义。*parents=* 参数使用 *ArgumentParser* 对象的列表, 从它们那里收集所有的位置和可选的行为, 然后将这写行为加到正在构建的 *ArgumentParser* 对象。

```
>>> parent_parser = argparse.ArgumentParser(add_help=False)
>>> parent_parser.add_argument('--parent', type=int)

>>> foo_parser = argparse.ArgumentParser(parents=[parent_parser])
>>> foo_parser.add_argument('foo')
>>> foo_parser.parse_args(['--parent', '2', 'XXX'])
Namespace(foo='XXX', parent=2)

>>> bar_parser = argparse.ArgumentParser(parents=[parent_parser])
>>> bar_parser.add_argument('--bar')
>>> bar_parser.parse_args(['--bar', 'YYY'])
Namespace(bar='YYY', parent=None)
```

请注意大多数父解析器会指定 *add_help=False*. 否则, *ArgumentParser* 将会看到两个 *-h/--help* 选项 (一个在父参数中一个在子参数中) 并且产生一个错误。

注解: 你在传 “*parents=*” 给那些解析器时必须完全初始化它们。如果你在子解析器之后改变父解析器是, 这些改变不会反映在子解析器上。

formatter_class

ArgumentParser 对象允许通过指定备用格式化类来自定义帮助格式。目前, 有四种这样的类。

```
class argparse.RawDescriptionHelpFormatter
class argparse.RawTextHelpFormatter
class argparse.ArgumentDefaultsHelpFormatter
class argparse.MetavarTypeHelpFormatter
```

RawDescriptionHelpFormatter 和 *RawTextHelpFormatter* 在正文的描述和展示上给与了更多的控制。*ArgumentParser* 对象会将 *description* 和 *epilog* 的文字在命令行中自动换行。

```
>>> parser = argparse.ArgumentParser(
...     prog='PROG',
...     description='''this description
...         was indented weird
...         but that is okay''',
...     epilog='')
```

(下页继续)

(续上页)

```

...         likewise for this epilog whose whitespace will
...         be cleaned up and whose words will be wrapped
...         across a couple lines'''
>>> parser.print_help()
usage: PROG [-h]

this description was indented weird but that is okay

optional arguments:
  -h, --help  show this help message and exit

likewise for this epilog whose whitespace will be cleaned up and whose words
will be wrapped across a couple lines

```

传 `RawDescriptionHelpFormatter` 给 `formatter_class=` 表示 `description` 和 `epilog` 已经被正确的格式化了, 不能在命令行中被自动换行:

```

>>> parser = argparse.ArgumentParser(
...     prog='PROG',
...     formatter_class=argparse.RawDescriptionHelpFormatter,
...     description=textwrap.dedent('''\
...         Please do not mess up this text!
...         -----
...         I have indented it
...         exactly the way
...         I want it
...     '''))
>>> parser.print_help()
usage: PROG [-h]

Please do not mess up this text!
-----
    I have indented it
    exactly the way
    I want it

optional arguments:
  -h, --help  show this help message and exit

```

`RawTextHelpFormatter` 保留所有种类文字的空格, 包括参数的描述。然而, 多重的新行会被替换成一行。如果你想保留多重的空白行, 可以在新行之间加空格。

`ArgumentDefaultsHelpFormatter` 自动添加默认的值的信息到每一个帮助信息的参数中:

```

>>> parser = argparse.ArgumentParser(
...     prog='PROG',
...     formatter_class=argparse.ArgumentDefaultsHelpFormatter)
>>> parser.add_argument('--foo', type=int, default=42, help='FOO!')
>>> parser.add_argument('bar', nargs='*', default=[1, 2, 3], help='BAR!')
>>> parser.print_help()
usage: PROG [-h] [--foo FOO] [bar [bar ...]]

positional arguments:
  bar          BAR! (default: [1, 2, 3])

optional arguments:
  -h, --help  show this help message and exit

```

(下页继续)

(续上页)

```
--foo FOO    FOO! (default: 42)
```

`MetavarTypeHelpFormatter` 为它的值在每一个参数中使用 *type* 的参数名当作它的显示名（而不是使用通常的格式 *dest*）:

```
>>> parser = argparse.ArgumentParser(
...     prog='PROG',
...     formatter_class=argparse.MetavarTypeHelpFormatter)
>>> parser.add_argument('--foo', type=int)
>>> parser.add_argument('bar', type=float)
>>> parser.print_help()
usage: PROG [-h] [--foo int] float

positional arguments:
  float

optional arguments:
  -h, --help  show this help message and exit
  --foo int
```

prefix_chars

许多命令行会使用 `-` 当作前缀，比如 `-f/--foo`。如果解析器需要支持不同的或者额外的字符，比如像 `+f` 或者 `/foo` 的选项，可以在参数解析构建器中使用 `prefix_chars=` 参数。

```
>>> parser = argparse.ArgumentParser(prog='PROG', prefix_chars='-+')
>>> parser.add_argument('+f')
>>> parser.add_argument('++bar')
>>> parser.parse_args('+f X ++bar Y'.split())
Namespace(bar='Y', f='X')
```

The `prefix_chars=` 参数默认使用 `'-'`。支持一系列字符，但是不包括 `-`，这样会产生不被允许的 `-f/--foo` 选项。

fromfile_prefix_chars

有些时候，先举个例子，当处理一个特别长的参数列表的时候，把它存入一个文件中而不是在命令行打出来会很有意义。如果 `fromfile_prefix_chars=` 参数提供给 `ArgumentParser` 构造函数，之后所有类型的字符的参数都会被当成文件处理，并且会被文件包含的参数替代。举个例子:

```
>>> with open('args.txt', 'w') as fp:
...     fp.write('-f\nbar')
>>> parser = argparse.ArgumentParser(fromfile_prefix_chars='@')
>>> parser.add_argument('-f')
>>> parser.parse_args(['-f', 'foo', '@args.txt'])
Namespace(f='bar')
```

从文件读取的参数在默认情况下必须一个一行（但是可参见 `convert_arg_line_to_args()`）并且它们被视为与命令行上的原始文件引用参数位于同一位置。所以在以上例子中，`['-f', 'foo', '@args.txt']` 的表示和 `['-f', 'foo', '-f', 'bar']` 的表示相同。

`fromfile_prefix_chars=` 参数默认为 `None`，意味着参数不会被当作文件对待。

argument_default

一般情况下，参数默认会通过设置一个默认到 `add_argument()` 或者调用带一组指定键值对的 `ArgumentParser.set_defaults()` 方法。但是有些时候，为参数指定一个普遍适用的解析器会更有用。这能够通过传输 `argument_default=` 关键词参数给 `ArgumentParser` 来完成。举个例子，要全局禁止在 `parse_args()` 中创建属性，我们提供 `argument_default=SUPPRESS`：

```
>>> parser = argparse.ArgumentParser(argument_default=argparse.SUPPRESS)
>>> parser.add_argument('--foo')
>>> parser.add_argument('bar', nargs='?')
>>> parser.parse_args(['--foo', '1', 'BAR'])
Namespace(bar='BAR', foo='1')
>>> parser.parse_args([])
Namespace()
```

allow_abbrev

正常情况下，当你向 `ArgumentParser` 的 `parse_args()` 方法传入一个参数列表时，它会 *recognizes abbreviations*。

这个特性可以设置 `allow_abbrev` 为 `False` 来关闭：

```
>>> parser = argparse.ArgumentParser(prog='PROG', allow_abbrev=False)
>>> parser.add_argument('--foobar', action='store_true')
>>> parser.add_argument('--foonley', action='store_false')
>>> parser.parse_args(['--foon'])
usage: PROG [-h] [--foobar] [--foonley]
PROG: error: unrecognized arguments: --foon
```

3.5 新版功能.

conflict_handler

`ArgumentParser` 对象不允许在相同选项字符串下有两种行为。默认情况下，`ArgumentParser` 对象会产生一个异常如果去创建一个正在使用的选项字符串参数。

```
>>> parser = argparse.ArgumentParser(prog='PROG')
>>> parser.add_argument('-f', '--foo', help='old foo help')
>>> parser.add_argument('--foo', help='new foo help')
Traceback (most recent call last):
..
ArgumentError: argument --foo: conflicting option string(s): --foo
```

有些时候（例如：使用 *parents*），重写旧的有相同选项字符串的参数会更有用。为了产生这种行为，`'resolve'` 值可以提供给 `ArgumentParser` 的 `conflict_handler=` 参数：

```
>>> parser = argparse.ArgumentParser(prog='PROG', conflict_handler='resolve')
>>> parser.add_argument('-f', '--foo', help='old foo help')
>>> parser.add_argument('--foo', help='new foo help')
>>> parser.print_help()
usage: PROG [-h] [-f FOO] [--foo FOO]

optional arguments:
  -h, --help  show this help message and exit
```

(下页继续)

(续上页)

```
-f FOO      old foo help
--foo FOO   new foo help
```

注意 `ArgumentParser` 对象只能移除一个行为如果它所有的选项字符串都被重写。所以，在上面的例子中，旧的 `-f/--foo` 行为回合 `-f` 行为保持一致，因为只有 `--foo` 选项字符串被重写。

add_help

默认情况下，`ArgumentParser` 对象添加一个简单的显示解析器帮助信息的选项。举个栗子，考虑一个名为 `myprogram.py` 的文件包含如下代码：

```
import argparse
parser = argparse.ArgumentParser()
parser.add_argument('--foo', help='foo help')
args = parser.parse_args()
```

如果 `-h` or `--help` 在命令行中被提供，参数解析器帮助信息会打印：

```
$ python myprogram.py --help
usage: myprogram.py [-h] [--foo FOO]

optional arguments:
  -h, --help  show this help message and exit
  --foo FOO   foo help
```

有时候可能会需要关闭额外的帮助信息。这可以通过在 `ArgumentParser` 中设置 `add_help=` 参数为 `False` 来实现。

```
>>> parser = argparse.ArgumentParser(prog='PROG', add_help=False)
>>> parser.add_argument('--foo', help='foo help')
>>> parser.print_help()
usage: PROG [--foo FOO]

optional arguments:
  --foo FOO   foo help
```

帮助选项一般为 `-h/--help`。如果 `prefix_chars=` 被指定并且没有包含 `-` 字符，在这种情况下，`-h` `--help` 不是有效的选项。此时，`prefix_chars` 的第一个字符将用作帮助选项的前缀。

```
>>> parser = argparse.ArgumentParser(prog='PROG', prefix_chars='+/')
>>> parser.print_help()
usage: PROG [+h]

optional arguments:
  +h, ++help  show this help message and exit
```

16.4.3 add_argument() 方法

`ArgumentParser.add_argument` (*name or flags...* [, *action*] [, *nargs*] [, *const*] [, *default*] [, *type*] [, *choices*] [, *required*] [, *help*] [, *metavar*] [, *dest*])

定义单个的命令行参数应当如何解析。每个形参都在下面有它自己更多的描述，长话短说有：

- *name or flags* - 一个命名或者一个选项字符串的列表，例如 `foo` 或 `-f`，`--foo`。

- *action* - 当参数在命令行中出现时使用的动作基本类型。
- *nargs* - 命令行参数应当消耗的数目。
- *const* - 被一些*action* 和*nargs* 选择所需求的常数。
- *default* - 当参数未在命令行中出现时使用的值。
- *type* - 命令行参数应当被转换成的类型。
- *choices* - 可用的参数的容器。
- *required* - 此命令行选项是否可省略（仅选项可用）。
- *help* - 一个此选项作用的简单描述。
- *metavar* - 在使用方法消息中使用的参数值示例。
- *dest* - 被添加到`parse_args()` 所返回对象上的属性名。

以下部分描述这些参数如何使用。

name or flags

`add_argument()` 方法必须知道它是否是一个选项，例如 `-f` 或 `--foo`，或是一个位置参数，例如一组文件名。第一个传递给`add_argument()` 的参数必须是一系列 **flags** 或者是一个简单的参数名。例如，可以选项可以被这样创建：

```
>>> parser.add_argument('-f', '--foo')
```

而位置参数可以这么创建：

```
>>> parser.add_argument('bar')
```

当`parse_args()` 被调用，选项会以 `-` 前缀识别，剩下的参数则会被假定为位置参数：

```
>>> parser = argparse.ArgumentParser(prog='PROG')
>>> parser.add_argument('-f', '--foo')
>>> parser.add_argument('bar')
>>> parser.parse_args(['BAR'])
Namespace(bar='BAR', foo=None)
>>> parser.parse_args(['BAR', '--foo', 'FOO'])
Namespace(bar='BAR', foo='FOO')
>>> parser.parse_args(['--foo', 'FOO'])
usage: PROG [-h] [-f FOO] bar
PROG: error: the following arguments are required: bar
```

action

`ArgumentParser` 对象将命令行参数与动作相关联。这些动作可以做与它们相关联的命令行参数的任何事，尽管大多数动作只是简单的向`parse_args()` 返回的对象上添加属性。`action` 命名参数指定了这个命令行参数应当如何处理。供应的动作有：

- `'store'` - 存储参数的值。这是默认的动作。例如：

```
>>> parser = argparse.ArgumentParser()
>>> parser.add_argument('--foo')
>>> parser.parse_args('--foo 1'.split())
Namespace(foo='1')
```

- 'store_const' - 存储被`const` 命名参数指定的值。'store_const' 动作通常用在选项中来指定一些标志。例如:

```
>>> parser = argparse.ArgumentParser()
>>> parser.add_argument('--foo', action='store_const', const=42)
>>> parser.parse_args(['--foo'])
Namespace(foo=42)
```

- 'store_true' and 'store_false' - 这些是 'store_const' 分别用作存储 True 和 False 值的特殊用例。另外，它们的默认值分别为 False 和 True。例如:

```
>>> parser = argparse.ArgumentParser()
>>> parser.add_argument('--foo', action='store_true')
>>> parser.add_argument('--bar', action='store_false')
>>> parser.add_argument('--baz', action='store_false')
>>> parser.parse_args('--foo --bar'.split())
Namespace(foo=True, bar=False, baz=True)
```

- 'append' - 存储一个列表，并且将每个参数值追加到列表中。在允许多次使用选项时很有用。例如:

```
>>> parser = argparse.ArgumentParser()
>>> parser.add_argument('--foo', action='append')
>>> parser.parse_args('--foo 1 --foo 2'.split())
Namespace(foo=['1', '2'])
```

- 'append_const' - 这存储一个列表，并将`const` 命名参数指定的值追加到列表中。(注意`const` 命名参数默认为 None。)“append_const” 动作一般在多个参数需要在同一列表中存储常数时会有用。例如:

```
>>> parser = argparse.ArgumentParser()
>>> parser.add_argument('--str', dest='types', action='append_const', const=str)
>>> parser.add_argument('--int', dest='types', action='append_const', const=int)
>>> parser.parse_args('--str --int'.split())
Namespace(types=[<class 'str'>, <class 'int'>])
```

- 'count' - 计算一个关键字参数出现的数目或次数。例如，对于一个增长的详情等级来说有用:

```
>>> parser = argparse.ArgumentParser()
>>> parser.add_argument('--verbose', '-v', action='count')
>>> parser.parse_args(['-vvv'])
Namespace(verbose=3)
```

- 'help' - 打印所有当前解析器中的选项和参数的完整帮助信息，然后退出。默认情况下，一个 help 动作会被自动加入解析器。关于输出是如何创建的，参与`ArgumentParser`。
- 'version' - 期望有一个 version= 命名参数在`add_argument()` 调用中，并打印版本信息并在调用后退出:

```
>>> import argparse
>>> parser = argparse.ArgumentParser(prog='PROG')
>>> parser.add_argument('--version', action='version', version='%s 2.0' % (prog))
>>> parser.parse_args(['--version'])
PROG 2.0
```

您还可以通过传递 Action 子类或实现相同接口的其他对象来指定任意操作。建议的方法是扩展`Action`，覆盖 `__call__` 方法和可选的 `__init__` 方法。

一个自定义动作的例子:

```
>>> class FooAction(argparse.Action):
...     def __init__(self, option_strings, dest, nargs=None, **kwargs):
...         if nargs is not None:
...             raise ValueError("nargs not allowed")
...         super(FooAction, self).__init__(option_strings, dest, **kwargs)
...     def __call__(self, parser, namespace, values, option_string=None):
...         print('%r %r %r' % (namespace, values, option_string))
...         setattr(namespace, self.dest, values)
...
>>> parser = argparse.ArgumentParser()
>>> parser.add_argument('--foo', action=FooAction)
>>> parser.add_argument('bar', action=FooAction)
>>> args = parser.parse_args('1 --foo 2'.split())
Namespace(bar=None, foo=None) '1' None
Namespace(bar='1', foo=None) '2' '--foo'
>>> args
Namespace(bar='1', foo='2')
```

更多描述，见 [Action](#)。

nargs

`ArgumentParser` 对象通常关联一个单独的命令行参数到一个单独的被执行的动作。`nargs` 命名参数关联不同数目的命令行参数到单一动作。支持的值有：

- `N`（一个整数）。命令行中的 `N` 个参数会被聚集到一个列表中。例如：

```
>>> parser = argparse.ArgumentParser()
>>> parser.add_argument('--foo', nargs=2)
>>> parser.add_argument('bar', nargs=1)
>>> parser.parse_args('c --foo a b'.split())
Namespace(bar=['c'], foo=['a', 'b'])
```

注意 `nargs=1` 会产生一个单元素列表。这和默认的元素本身是不同的。

- `'?'`。如果可能的话，会从命令行中消耗一个参数，并产生一个单一项。如果当前没有命令行参数，则会产生 *default* 值。注意，对于选项，有另外的用例 - 选项字符串出现但没有跟随命令行参数，则会产生 *const* 值。一些说用用例：

```
>>> parser = argparse.ArgumentParser()
>>> parser.add_argument('--foo', nargs='?', const='c', default='d')
>>> parser.add_argument('bar', nargs='?', default='d')
>>> parser.parse_args(['XX', '--foo', 'YY'])
Namespace(bar='XX', foo='YY')
>>> parser.parse_args(['XX', '--foo'])
Namespace(bar='XX', foo='c')
>>> parser.parse_args([])
Namespace(bar='d', foo='d')
```

`nargs='?'` 的一个更普遍用法是允许可选的输入或输出文件：

```
>>> parser = argparse.ArgumentParser()
>>> parser.add_argument('infile', nargs='?', type=argparse.FileType('r'),
...                     default=sys.stdin)
>>> parser.add_argument('outfile', nargs='?', type=argparse.FileType('w'),
...                     default=sys.stdout)
```

(下页继续)

(续上页)

```
>>> parser.parse_args(['input.txt', 'output.txt'])
Namespace(infile=<_io.TextIOWrapper name='input.txt' encoding='UTF-8'>,
          outfile=<_io.TextIOWrapper name='output.txt' encoding='UTF-8'>)
>>> parser.parse_args([])
Namespace(infile=<_io.TextIOWrapper name='<stdin>' encoding='UTF-8'>,
          outfile=<_io.TextIOWrapper name='<stdout>' encoding='UTF-8'>)
```

- `'*'`。所有当前命令行参数被聚集到一个列表中。注意通过 `nargs='*'` 来实现多个位置参数通常没有意义，但是多个选项是可能的。例如：

```
>>> parser = argparse.ArgumentParser()
>>> parser.add_argument('--foo', nargs='*')
>>> parser.add_argument('--bar', nargs='*')
>>> parser.add_argument('baz', nargs='*')
>>> parser.parse_args('a b --foo x y --bar 1 2'.split())
Namespace(bar=['1', '2'], baz=['a', 'b'], foo=['x', 'y'])
```

- `'+'`。和 `'*'` 类似，所有当前命令行参数被聚集到一个列表中。另外，当前没有至少一个命令行参数时会产生一个错误信息。例如：

```
>>> parser = argparse.ArgumentParser(prog='PROG')
>>> parser.add_argument('foo', nargs='+')
>>> parser.parse_args(['a', 'b'])
Namespace(foo=['a', 'b'])
>>> parser.parse_args([])
usage: PROG [-h] foo [foo ...]
PROG: error: the following arguments are required: foo
```

- `argparse.REMAINDER`。所有剩余的命令行参数被聚集到一个列表中。这通常在从一个命令行功能传递参数到另一个命令行功能中时有用：

```
>>> parser = argparse.ArgumentParser(prog='PROG')
>>> parser.add_argument('--foo')
>>> parser.add_argument('command')
>>> parser.add_argument('args', nargs=argparse.REMAINDER)
>>> print(parser.parse_args('--foo B cmd --arg1 XX ZZ'.split()))
Namespace(args=['--arg1', 'XX', 'ZZ'], command='cmd', foo='B')
```

如果不提供 `nargs` 命名参数，则消耗参数的数目将被 *action* 决定。通常这意味着单一项目（非列表）消耗单一命令行参数。

const

`add_argument()` 的 `"const"` 参数用于保存不从命令行中读取但被各种 *ArgumentParser* 动作需求的常数值。最常用的两例为：

- 当 `add_argument()` 通过 `action='store_const'` 或 `action='append_const'` 调用时。这些动作将 `const` 值添加到 `parse_args()` 返回的对象的属性中。在 *action* 的描述中查看案例。
- 当 `add_argument()` 通过选项（例如 `-f` 或 `--foo`）调用并且 `nargs='?'` 时。这会创建一个可以跟随零个或一个命令行参数的选项。当解析命令行时，如果选项后没有参数，则用 `const` 代替。在 *nargs* 描述中查看案例。

对 `'store_const'` 和 `'append_const'` 动作，`const` 命名参数必须给出。对其他动作，默认为 `None`。

default

所有选项和一些位置参数可能在命令行中被忽略。`add_argument()` 的命名参数 `default`，默认值为 `None`，指定了在命令行参数未出现时应当使用的值。对于选项，`default` 值在选项未在命令行中出现时使用：

```
>>> parser = argparse.ArgumentParser()
>>> parser.add_argument('--foo', default=42)
>>> parser.parse_args(['--foo', '2'])
Namespace(foo='2')
>>> parser.parse_args([])
Namespace(foo=42)
```

如果 `default` 值是一个字符串，解析器解析此值就像一个命令行参数。特别是，在将属性设置在 `Namespace` 的返回值之前，解析器应用任何提供的 `type` 转换参数。否则解析器使用原值：

```
>>> parser = argparse.ArgumentParser()
>>> parser.add_argument('--length', default='10', type=int)
>>> parser.add_argument('--width', default=10.5, type=int)
>>> parser.parse_args()
Namespace(length=10, width=10.5)
```

对于 `nargs` 等于 `?` 或 `*` 的位置参数，`default` 值在没有命令行参数出现时使用。

```
>>> parser = argparse.ArgumentParser()
>>> parser.add_argument('foo', nargs='?', default=42)
>>> parser.parse_args(['a'])
Namespace(foo='a')
>>> parser.parse_args([])
Namespace(foo=42)
```

提供 `default=argparse.SUPPRESS` 导致命令行参数未出现时没有属性被添加：

```
>>> parser = argparse.ArgumentParser()
>>> parser.add_argument('--foo', default=argparse.SUPPRESS)
>>> parser.parse_args([])
Namespace()
>>> parser.parse_args(['--foo', '1'])
Namespace(foo='1')
```

type

By default, `ArgumentParser` objects read command-line arguments in as simple strings. However, quite often the command-line string should instead be interpreted as another type, like a `float` or `int`. The `type` keyword argument of `add_argument()` allows any necessary type-checking and type conversions to be performed. Common built-in types and functions can be used directly as the value of the `type` argument:

```
>>> parser = argparse.ArgumentParser()
>>> parser.add_argument('foo', type=int)
>>> parser.add_argument('bar', type=open)
>>> parser.parse_args('2 temp.txt'.split())
Namespace(bar=<_io.TextIOWrapper name='temp.txt' encoding='UTF-8'>, foo=2)
```

See the section on the `default` keyword argument for information on when the `type` argument is applied to default arguments.

To ease the use of various types of files, the `argparse` module provides the factory `FileType` which takes the `mode=`, `bufsize=`, `encoding=` and `errors=` arguments of the `open()` function. For example, `FileType('w')` can be used to create a writable file:

```
>>> parser = argparse.ArgumentParser()
>>> parser.add_argument('bar', type=argparse.FileType('w'))
>>> parser.parse_args(['out.txt'])
Namespace(bar=<_io.TextIOWrapper name='out.txt' encoding='UTF-8'>)
```

`type=` can take any callable that takes a single string argument and returns the converted value:

```
>>> def perfect_square(string):
...     value = int(string)
...     sqrt = math.sqrt(value)
...     if sqrt != int(sqrt):
...         msg = "%r is not a perfect square" % string
...         raise argparse.ArgumentTypeError(msg)
...     return value
...
>>> parser = argparse.ArgumentParser(prog='PROG')
>>> parser.add_argument('foo', type=perfect_square)
>>> parser.parse_args(['9'])
Namespace(foo=9)
>>> parser.parse_args(['7'])
usage: PROG [-h] foo
PROG: error: argument foo: '7' is not a perfect square
```

The `choices` keyword argument may be more convenient for type checkers that simply check against a range of values:

```
>>> parser = argparse.ArgumentParser(prog='PROG')
>>> parser.add_argument('foo', type=int, choices=range(5, 10))
>>> parser.parse_args(['7'])
Namespace(foo=7)
>>> parser.parse_args(['11'])
usage: PROG [-h] {5,6,7,8,9}
PROG: error: argument foo: invalid choice: 11 (choose from 5, 6, 7, 8, 9)
```

See the `choices` section for more details.

选择

Some command-line arguments should be selected from a restricted set of values. These can be handled by passing a container object as the `choices` keyword argument to `add_argument()`. When the command line is parsed, argument values will be checked, and an error message will be displayed if the argument was not one of the acceptable values:

```
>>> parser = argparse.ArgumentParser(prog='game.py')
>>> parser.add_argument('move', choices=['rock', 'paper', 'scissors'])
>>> parser.parse_args(['rock'])
Namespace(move='rock')
>>> parser.parse_args(['fire'])
usage: game.py [-h] {rock,paper,scissors}
game.py: error: argument move: invalid choice: 'fire' (choose from 'rock',
'paper', 'scissors')
```

Note that inclusion in the `choices` container is checked after any `type` conversions have been performed, so the type of the objects in the `choices` container should match the `type` specified:

```
>>> parser = argparse.ArgumentParser(prog='doors.py')
>>> parser.add_argument('door', type=int, choices=range(1, 4))
>>> print(parser.parse_args(['3']))
Namespace(door=3)
>>> parser.parse_args(['4'])
usage: doors.py [-h] {1,2,3}
doors.py: error: argument door: invalid choice: 4 (choose from 1, 2, 3)
```

Any object that supports the `in` operator can be passed as the *choices* value, so *dict* objects, *set* objects, custom containers, etc. are all supported.

required

In general, the *argparse* module assumes that flags like `-f` and `--bar` indicate *optional* arguments, which can always be omitted at the command line. To make an option *required*, `True` can be specified for the `required=` keyword argument to *add_argument()*:

```
>>> parser = argparse.ArgumentParser()
>>> parser.add_argument('--foo', required=True)
>>> parser.parse_args(['--foo', 'BAR'])
Namespace(foo='BAR')
>>> parser.parse_args([])
usage: argparse.py [-h] [--foo FOO]
argparse.py: error: option --foo is required
```

As the example shows, if an option is marked as required, *parse_args()* will report an error if that option is not present at the command line.

注解： Required options are generally considered bad form because users expect *options* to be *optional*, and thus they should be avoided when possible.

help

The *help* value is a string containing a brief description of the argument. When a user requests help (usually by using `-h` or `--help` at the command line), these help descriptions will be displayed with each argument:

```
>>> parser = argparse.ArgumentParser(prog='frobble')
>>> parser.add_argument('--foo', action='store_true',
...                     help='foo the bars before frobbling')
>>> parser.add_argument('bar', nargs='+',
...                     help='one of the bars to be frobbled')
>>> parser.parse_args(['-h'])
usage: frobble [-h] [--foo] bar [bar ...]

positional arguments:
  bar                one of the bars to be frobbled

optional arguments:
  -h, --help        show this help message and exit
  --foo            foo the bars before frobbling
```

The *help* strings can include various format specifiers to avoid repetition of things like the program name or the argument *default*. The available specifiers include the program name, `%(prog)s` and most keyword arguments to

`add_argument()`, e.g. `%(default)s`, `%(type)s`, etc.:

```
>>> parser = argparse.ArgumentParser(prog='frobble')
>>> parser.add_argument('bar', nargs='?', type=int, default=42,
...                     help='the bar to %(prog)s (default: %(default)s)')
>>> parser.print_help()
usage: frobble [-h] [bar]

positional arguments:
  bar      the bar to frobble (default: 42)

optional arguments:
  -h, --help  show this help message and exit
```

As the help string supports %-formatting, if you want a literal % to appear in the help string, you must escape it as %%.

`argparse` supports silencing the help entry for certain options, by setting the help value to `argparse.SUPPRESS`:

```
>>> parser = argparse.ArgumentParser(prog='frobble')
>>> parser.add_argument('--foo', help=argparse.SUPPRESS)
>>> parser.print_help()
usage: frobble [-h]

optional arguments:
  -h, --help  show this help message and exit
```

metavar

When `ArgumentParser` generates help messages, it needs some way to refer to each expected argument. By default, `ArgumentParser` objects use the `dest` value as the “name” of each object. By default, for positional argument actions, the `dest` value is used directly, and for optional argument actions, the `dest` value is uppercased. So, a single positional argument with `dest='bar'` will be referred to as `bar`. A single optional argument `--foo` that should be followed by a single command-line argument will be referred to as `FOO`. An example:

```
>>> parser = argparse.ArgumentParser()
>>> parser.add_argument('--foo')
>>> parser.add_argument('bar')
>>> parser.parse_args('X --foo Y'.split())
Namespace(bar='X', foo='Y')
>>> parser.print_help()
usage: [-h] [--foo FOO] bar

positional arguments:
  bar

optional arguments:
  -h, --help  show this help message and exit
  --foo FOO
```

An alternative name can be specified with `metavar`:

```
>>> parser = argparse.ArgumentParser()
>>> parser.add_argument('--foo', metavar='YYY')
>>> parser.add_argument('bar', metavar='XXX')
>>> parser.parse_args('X --foo Y'.split())
Namespace(bar='X', foo='Y')
```

(下页继续)

(续上页)

```
>>> parser.print_help()
usage:  [-h] [--foo YYY] XXX

positional arguments:
  XXX

optional arguments:
  -h, --help  show this help message and exit
  --foo YYY
```

Note that `metavar` only changes the *displayed* name - the name of the attribute on the `parse_args()` object is still determined by the *dest* value.

Different values of `nargs` may cause the metavar to be used multiple times. Providing a tuple to `metavar` specifies a different display for each of the arguments:

```
>>> parser = argparse.ArgumentParser(prog='PROG')
>>> parser.add_argument('-x', nargs=2)
>>> parser.add_argument('--foo', nargs=2, metavar=('bar', 'baz'))
>>> parser.print_help()
usage: PROG [-h] [-x X X] [--foo bar baz]

optional arguments:
  -h, --help            show this help message and exit
  -x X X
  --foo bar baz
```

dest

Most `ArgumentParser` actions add some value as an attribute of the object returned by `parse_args()`. The name of this attribute is determined by the `dest` keyword argument of `add_argument()`. For positional argument actions, `dest` is normally supplied as the first argument to `add_argument()`:

```
>>> parser = argparse.ArgumentParser()
>>> parser.add_argument('bar')
>>> parser.parse_args(['XXX'])
Namespace(bar='XXX')
```

For optional argument actions, the value of `dest` is normally inferred from the option strings. `ArgumentParser` generates the value of `dest` by taking the first long option string and stripping away the initial `--` string. If no long option strings were supplied, `dest` will be derived from the first short option string by stripping the initial `-` character. Any internal `-` characters will be converted to `_` characters to make sure the string is a valid attribute name. The examples below illustrate this behavior:

```
>>> parser = argparse.ArgumentParser()
>>> parser.add_argument('-f', '--foo-bar', '--foo')
>>> parser.add_argument('-x', '-y')
>>> parser.parse_args('-f 1 -x 2'.split())
Namespace(foo_bar='1', x='2')
>>> parser.parse_args('--foo 1 -y 2'.split())
Namespace(foo_bar='1', x='2')
```

`dest` allows a custom attribute name to be provided:

```
>>> parser = argparse.ArgumentParser()
>>> parser.add_argument('--foo', dest='bar')
>>> parser.parse_args('--foo XXX'.split())
Namespace(bar='XXX')
```

Action classes

Action classes implement the Action API, a callable which returns a callable which processes arguments from the command-line. Any object which follows this API may be passed as the `action` parameter to `add_argument()`.

class `argparse.Action`(*option_strings*, *dest*, *nargs=None*, *const=None*, *default=None*, *type=None*, *choices=None*, *required=False*, *help=None*, *metavar=None*)

Action objects are used by an `ArgumentParser` to represent the information needed to parse a single argument from one or more strings from the command line. The Action class must accept the two positional arguments plus any keyword arguments passed to `ArgumentParser.add_argument()` except for the `action` itself.

Instances of Action (or return value of any callable to the `action` parameter) should have attributes “`dest`”, “`option_strings`”, “`default`”, “`type`”, “`required`”, “`help`”, etc. defined. The easiest way to ensure these attributes are defined is to call `Action.__init__`.

Action instances should be callable, so subclasses must override the `__call__` method, which should accept four parameters:

- `parser` - The `ArgumentParser` object which contains this action.
- `namespace` - The `Namespace` object that will be returned by `parse_args()`. Most actions add an attribute to this object using `setattr()`.
- `values` - The associated command-line arguments, with any type conversions applied. Type conversions are specified with the `type` keyword argument to `add_argument()`.
- `option_string` - The option string that was used to invoke this action. The `option_string` argument is optional, and will be absent if the action is associated with a positional argument.

The `__call__` method may perform arbitrary actions, but will typically set attributes on the `namespace` based on `dest` and `values`.

16.4.4 The `parse_args()` method

`ArgumentParser.parse_args`(*args=None*, *namespace=None*)

Convert argument strings to objects and assign them as attributes of the namespace. Return the populated namespace.

Previous calls to `add_argument()` determine exactly what objects are created and how they are assigned. See the documentation for `add_argument()` for details.

- *args* - List of strings to parse. The default is taken from `sys.argv`.
- *namespace* - An object to take the attributes. The default is a new empty `Namespace` object.

Option value syntax

The `parse_args()` method supports several ways of specifying the value of an option (if it takes one). In the simplest case, the option and its value are passed as two separate arguments:

```
>>> parser = argparse.ArgumentParser(prog='PROG')
>>> parser.add_argument('-x')
>>> parser.add_argument('--foo')
>>> parser.parse_args(['-x', 'X'])
Namespace(foo=None, x='X')
>>> parser.parse_args(['--foo', 'FOO'])
Namespace(foo='FOO', x=None)
```

For long options (options with names longer than a single character), the option and value can also be passed as a single command-line argument, using `=` to separate them:

```
>>> parser.parse_args(['--foo=FOO'])
Namespace(foo='FOO', x=None)
```

For short options (options only one character long), the option and its value can be concatenated:

```
>>> parser.parse_args(['-xX'])
Namespace(foo=None, x='X')
```

Several short options can be joined together, using only a single `-` prefix, as long as only the last option (or none of them) requires a value:

```
>>> parser = argparse.ArgumentParser(prog='PROG')
>>> parser.add_argument('-x', action='store_true')
>>> parser.add_argument('-y', action='store_true')
>>> parser.add_argument('-z')
>>> parser.parse_args(['-xyzZ'])
Namespace(x=True, y=True, z='Z')
```

Invalid arguments

While parsing the command line, `parse_args()` checks for a variety of errors, including ambiguous options, invalid types, invalid options, wrong number of positional arguments, etc. When it encounters such an error, it exits and prints the error along with a usage message:

```
>>> parser = argparse.ArgumentParser(prog='PROG')
>>> parser.add_argument('--foo', type=int)
>>> parser.add_argument('bar', nargs='?')

>>> # invalid type
>>> parser.parse_args(['--foo', 'spam'])
usage: PROG [-h] [--foo FOO] [bar]
PROG: error: argument --foo: invalid int value: 'spam'

>>> # invalid option
>>> parser.parse_args(['--bar'])
usage: PROG [-h] [--foo FOO] [bar]
PROG: error: no such option: --bar

>>> # wrong number of arguments
>>> parser.parse_args(['spam', 'badger'])
usage: PROG [-h] [--foo FOO] [bar]
PROG: error: extra arguments found: badger
```

Arguments containing –

The `parse_args()` method attempts to give errors whenever the user has clearly made a mistake, but some situations are inherently ambiguous. For example, the command-line argument `-1` could either be an attempt to specify an option or an attempt to provide a positional argument. The `parse_args()` method is cautious here: positional arguments may only begin with `-` if they look like negative numbers and there are no options in the parser that look like negative numbers:

```
>>> parser = argparse.ArgumentParser(prog='PROG')
>>> parser.add_argument('-x')
>>> parser.add_argument('foo', nargs='?')

>>> # no negative number options, so -1 is a positional argument
>>> parser.parse_args(['-x', '-1'])
Namespace(foo=None, x='-1')

>>> # no negative number options, so -1 and -5 are positional arguments
>>> parser.parse_args(['-x', '-1', '-5'])
Namespace(foo='-5', x='-1')

>>> parser = argparse.ArgumentParser(prog='PROG')
>>> parser.add_argument('-1', dest='one')
>>> parser.add_argument('foo', nargs='?')

>>> # negative number options present, so -1 is an option
>>> parser.parse_args(['-1', 'X'])
Namespace(foo=None, one='X')

>>> # negative number options present, so -2 is an option
>>> parser.parse_args(['-2'])
usage: PROG [-h] [-1 ONE] [foo]
PROG: error: no such option: -2

>>> # negative number options present, so both -1s are options
>>> parser.parse_args(['-1', '-1'])
usage: PROG [-h] [-1 ONE] [foo]
PROG: error: argument -1: expected one argument
```

If you have positional arguments that must begin with `-` and don't look like negative numbers, you can insert the pseudo-argument `--` which tells `parse_args()` that everything after that is a positional argument:

```
>>> parser.parse_args(['--', '-f'])
Namespace(foo='-f', one=None)
```

Argument abbreviations (prefix matching)

The `parse_args()` method *by default* allows long options to be abbreviated to a prefix, if the abbreviation is unambiguous (the prefix matches a unique option):

```
>>> parser = argparse.ArgumentParser(prog='PROG')
>>> parser.add_argument('-bacon')
>>> parser.add_argument('-badger')
>>> parser.parse_args(['-bac MMM'].split())
Namespace(bacon='MMM', badger=None)
>>> parser.parse_args(['-bad WOOD'].split())
Namespace(bacon=None, badger='WOOD')
```

(下页继续)

(续上页)

```
>>> parser.parse_args('-ba BA'.split())
usage: PROG [-h] [-bacon BACON] [-badger BADGER]
PROG: error: ambiguous option: -ba could match -badger, -bacon
```

An error is produced for arguments that could produce more than one options. This feature can be disabled by setting `allow_abbrev` to `False`.

Beyond `sys.argv`

Sometimes it may be useful to have an `ArgumentParser` parse arguments other than those of `sys.argv`. This can be accomplished by passing a list of strings to `parse_args()`. This is useful for testing at the interactive prompt:

```
>>> parser = argparse.ArgumentParser()
>>> parser.add_argument(
...     'integers', metavar='int', type=int, choices=range(10),
...     nargs='+', help='an integer in the range 0..9')
>>> parser.add_argument(
...     '--sum', dest='accumulate', action='store_const', const=sum,
...     default=max, help='sum the integers (default: find the max)')
>>> parser.parse_args(['1', '2', '3', '4'])
Namespace(accumulate=<built-in function max>, integers=[1, 2, 3, 4])
>>> parser.parse_args(['1', '2', '3', '4', '--sum'])
Namespace(accumulate=<built-in function sum>, integers=[1, 2, 3, 4])
```

The Namespace object

`class` `argparse.Namespace`

Simple class used by default by `parse_args()` to create an object holding attributes and return it.

This class is deliberately simple, just an `object` subclass with a readable string representation. If you prefer to have dict-like view of the attributes, you can use the standard Python idiom, `vars()`:

```
>>> parser = argparse.ArgumentParser()
>>> parser.add_argument('--foo')
>>> args = parser.parse_args(['--foo', 'BAR'])
>>> vars(args)
{'foo': 'BAR'}
```

It may also be useful to have an `ArgumentParser` assign attributes to an already existing object, rather than a new `Namespace` object. This can be achieved by specifying the `namespace=` keyword argument:

```
>>> class C:
...     pass
...
>>> c = C()
>>> parser = argparse.ArgumentParser()
>>> parser.add_argument('--foo')
>>> parser.parse_args(args=['--foo', 'BAR'], namespace=c)
>>> c.foo
'BAR'
```

16.4.5 Other utilities

Sub-commands

`ArgumentParser.add_subparsers` (`[title][, description][, prog][, parser_class][, action][, option_string][, dest][, help][, metavar]`)

Many programs split up their functionality into a number of sub-commands, for example, the `svn` program can invoke sub-commands like `svn checkout`, `svn update`, and `svn commit`. Splitting up functionality this way can be a particularly good idea when a program performs several different functions which require different kinds of command-line arguments. `ArgumentParser` supports the creation of such sub-commands with the `add_subparsers()` method. The `add_subparsers()` method is normally called with no arguments and returns a special action object. This object has a single method, `add_parser()`, which takes a command name and any `ArgumentParser` constructor arguments, and returns an `ArgumentParser` object that can be modified as usual.

Description of parameters:

- `title` - title for the sub-parser group in help output; by default “subcommands” if description is provided, otherwise uses title for positional arguments
- `description` - description for the sub-parser group in help output, by default `None`
- `prog` - usage information that will be displayed with sub-command help, by default the name of the program and any positional arguments before the subparser argument
- `parser_class` - class which will be used to create sub-parser instances, by default the class of the current parser (e.g. `ArgumentParser`)
- `action` - the basic type of action to be taken when this argument is encountered at the command line
- `dest` - name of the attribute under which sub-command name will be stored; by default `None` and no value is stored
- `help` - help for sub-parser group in help output, by default `None`
- `metavar` - string presenting available sub-commands in help; by default it is `None` and presents sub-commands in form `{cmd1, cmd2, ..}`

Some example usage:

```
>>> # create the top-level parser
>>> parser = argparse.ArgumentParser(prog='PROG')
>>> parser.add_argument('--foo', action='store_true', help='foo help')
>>> subparsers = parser.add_subparsers(help='sub-command help')
>>>
>>> # create the parser for the "a" command
>>> parser_a = subparsers.add_parser('a', help='a help')
>>> parser_a.add_argument('bar', type=int, help='bar help')
>>>
>>> # create the parser for the "b" command
>>> parser_b = subparsers.add_parser('b', help='b help')
>>> parser_b.add_argument('--baz', choices='XYZ', help='baz help')
>>>
>>> # parse some argument lists
>>> parser.parse_args(['a', '12'])
Namespace(bar=12, foo=False)
>>> parser.parse_args(['--foo', 'b', '--baz', 'Z'])
Namespace(baz='Z', foo=True)
```

Note that the object returned by `parse_args()` will only contain attributes for the main parser and the subparser that was selected by the command line (and not any other subparsers). So in the example above, when the `a` command is specified, only the `foo` and `bar` attributes are present, and when the `b` command is specified, only the `foo` and `baz` attributes are present.

Similarly, when a help message is requested from a subparser, only the help for that particular parser will be printed. The help message will not include parent parser or sibling parser messages. (A help message for each subparser command, however, can be given by supplying the `help=` argument to `add_parser()` as above.)

```
>>> parser.parse_args(['--help'])
usage: PROG [-h] [--foo] {a,b} ...

positional arguments:
  {a,b}    sub-command help
  a        a help
  b        b help

optional arguments:
  -h, --help  show this help message and exit
  --foo      foo help

>>> parser.parse_args(['a', '--help'])
usage: PROG a [-h] bar

positional arguments:
  bar      bar help

optional arguments:
  -h, --help  show this help message and exit

>>> parser.parse_args(['b', '--help'])
usage: PROG b [-h] [--baz {X,Y,Z}]

optional arguments:
  -h, --help      show this help message and exit
  --baz {X,Y,Z}  baz help
```

The `add_subparsers()` method also supports `title` and `description` keyword arguments. When either is present, the subparser's commands will appear in their own group in the help output. For example:

```
>>> parser = argparse.ArgumentParser()
>>> subparsers = parser.add_subparsers(title='subcommands',
...                                   description='valid subcommands',
...                                   help='additional help')
>>> subparsers.add_parser('foo')
>>> subparsers.add_parser('bar')
>>> parser.parse_args(['-h'])
usage: [-h] {foo,bar} ...

optional arguments:
  -h, --help  show this help message and exit

subcommands:
  valid subcommands

  {foo,bar}  additional help
```

Furthermore, `add_parser` supports an additional `aliases` argument, which allows multiple strings to refer to the same subparser. This example, like `svn`, aliases `co` as a shorthand for `checkout`:

```
>>> parser = argparse.ArgumentParser()
>>> subparsers = parser.add_subparsers()
>>> checkout = subparsers.add_parser('checkout', aliases=['co'])
```

(下页继续)

(续上页)

```
>>> checkout.add_argument('foo')
>>> parser.parse_args(['co', 'bar'])
Namespace(foo='bar')
```

One particularly effective way of handling sub-commands is to combine the use of the `add_subparsers()` method with calls to `set_defaults()` so that each subparser knows which Python function it should execute. For example:

```
>>> # sub-command functions
>>> def foo(args):
...     print(args.x * args.y)
...
>>> def bar(args):
...     print('((%s))' % args.z)
...
>>> # create the top-level parser
>>> parser = argparse.ArgumentParser()
>>> subparsers = parser.add_subparsers()
>>>
>>> # create the parser for the "foo" command
>>> parser_foo = subparsers.add_parser('foo')
>>> parser_foo.add_argument('-x', type=int, default=1)
>>> parser_foo.add_argument('y', type=float)
>>> parser_foo.set_defaults(func=foo)
>>>
>>> # create the parser for the "bar" command
>>> parser_bar = subparsers.add_parser('bar')
>>> parser_bar.add_argument('z')
>>> parser_bar.set_defaults(func=bar)
>>>
>>> # parse the args and call whatever function was selected
>>> args = parser.parse_args('foo 1 -x 2'.split())
>>> args.func(args)
2.0
>>>
>>> # parse the args and call whatever function was selected
>>> args = parser.parse_args('bar XYZYX'.split())
>>> args.func(args)
((XYZYX))
```

This way, you can let `parse_args()` do the job of calling the appropriate function after argument parsing is complete. Associating functions with actions like this is typically the easiest way to handle the different actions for each of your subparsers. However, if it is necessary to check the name of the subparser that was invoked, the `dest` keyword argument to the `add_subparsers()` call will work:

```
>>> parser = argparse.ArgumentParser()
>>> subparsers = parser.add_subparsers(dest='subparser_name')
>>> subparser1 = subparsers.add_parser('1')
>>> subparser1.add_argument('-x')
>>> subparser2 = subparsers.add_parser('2')
>>> subparser2.add_argument('y')
>>> parser.parse_args(['2', 'frobble'])
Namespace(subparser_name='2', y='frobble')
```

FileType objects

class `argparse.FileType` (*mode='r', bufsize=-1, encoding=None, errors=None*)

The `FileType` factory creates objects that can be passed to the `type` argument of `ArgumentParser.add_argument()`. Arguments that have `FileType` objects as their type will open command-line arguments as files with the requested modes, buffer sizes, encodings and error handling (see the `open()` function for more details):

```
>>> parser = argparse.ArgumentParser()
>>> parser.add_argument('--raw', type=argparse.FileType('wb', 0))
>>> parser.add_argument('out', type=argparse.FileType('w', encoding='UTF-8'))
>>> parser.parse_args(['--raw', 'raw.dat', 'file.txt'])
Namespace(out=<_io.TextIOWrapper name='file.txt' mode='w' encoding='UTF-8'>, raw=
↳ <_io.FileIO name='raw.dat' mode='wb'>)
```

`FileType` objects understand the pseudo-argument `'-'` and automatically convert this into `sys.stdin` for readable `FileType` objects and `sys.stdout` for writable `FileType` objects:

```
>>> parser = argparse.ArgumentParser()
>>> parser.add_argument('infile', type=argparse.FileType('r'))
>>> parser.parse_args(['-'])
Namespace(infile=<_io.TextIOWrapper name='<stdin>' encoding='UTF-8'>)
```

3.4 新版功能: The *encodings* and *errors* keyword arguments.

Argument groups

`ArgumentParser.add_argument_group` (*title=None, description=None*)

By default, `ArgumentParser` groups command-line arguments into “positional arguments” and “optional arguments” when displaying help messages. When there is a better conceptual grouping of arguments than this default one, appropriate groups can be created using the `add_argument_group()` method:

```
>>> parser = argparse.ArgumentParser(prog='PROG', add_help=False)
>>> group = parser.add_argument_group('group')
>>> group.add_argument('--foo', help='foo help')
>>> group.add_argument('bar', help='bar help')
>>> parser.print_help()
usage: PROG [--foo FOO] bar

group:
  bar      bar help
  --foo FOO  foo help
```

The `add_argument_group()` method returns an argument group object which has an `add_argument()` method just like a regular `ArgumentParser`. When an argument is added to the group, the parser treats it just like a normal argument, but displays the argument in a separate group for help messages. The `add_argument_group()` method accepts *title* and *description* arguments which can be used to customize this display:

```
>>> parser = argparse.ArgumentParser(prog='PROG', add_help=False)
>>> group1 = parser.add_argument_group('group1', 'group1 description')
>>> group1.add_argument('foo', help='foo help')
>>> group2 = parser.add_argument_group('group2', 'group2 description')
>>> group2.add_argument('--bar', help='bar help')
>>> parser.print_help()
```

(下页继续)

(续上页)

```
usage: PROG [--bar BAR] foo

group1:
  group1 description

  foo      foo help

group2:
  group2 description

  --bar BAR  bar help
```

Note that any arguments not in your user-defined groups will end up back in the usual “positional arguments” and “optional arguments” sections.

Mutual exclusion

`ArgumentParser.add_mutually_exclusive_group(required=False)`

创建一个互斥组。`argparse` 将会确保互斥组中只有一个参数在命令行中可用:

```
>>> parser = argparse.ArgumentParser(prog='PROG')
>>> group = parser.add_mutually_exclusive_group()
>>> group.add_argument('--foo', action='store_true')
>>> group.add_argument('--bar', action='store_false')
>>> parser.parse_args(['--foo'])
Namespace(bar=True, foo=True)
>>> parser.parse_args(['--bar'])
Namespace(bar=False, foo=False)
>>> parser.parse_args(['--foo', '--bar'])
usage: PROG [-h] [--foo | --bar]
PROG: error: argument --bar: not allowed with argument --foo
```

`add_mutually_exclusive_group()` 方法也接受一个 `required` 参数, 表示在互斥组中至少有一个参数是需要的:

```
>>> parser = argparse.ArgumentParser(prog='PROG')
>>> group = parser.add_mutually_exclusive_group(required=True)
>>> group.add_argument('--foo', action='store_true')
>>> group.add_argument('--bar', action='store_false')
>>> parser.parse_args([])
usage: PROG [-h] (--foo | --bar)
PROG: error: one of the arguments --foo --bar is required
```

注意, 目前互斥参数组不支持 `add_argument_group()` 的 `title` 和 `description` 参数。

Parser defaults

`ArgumentParser.set_defaults(**kwargs)`

Most of the time, the attributes of the object returned by `parse_args()` will be fully determined by inspecting the command-line arguments and the argument actions. `set_defaults()` allows some additional attributes that are determined without any inspection of the command line to be added:

```
>>> parser = argparse.ArgumentParser()
>>> parser.add_argument('foo', type=int)
>>> parser.set_defaults(bar=42, baz='badger')
>>> parser.parse_args(['736'])
Namespace(bar=42, baz='badger', foo=736)
```

Note that parser-level defaults always override argument-level defaults:

```
>>> parser = argparse.ArgumentParser()
>>> parser.add_argument('--foo', default='bar')
>>> parser.set_defaults(foo='spam')
>>> parser.parse_args([])
Namespace(foo='spam')
```

Parser-level defaults can be particularly useful when working with multiple parsers. See the `add_subparsers()` method for an example of this type.

`ArgumentParser.get_default(dest)`

Get the default value for a namespace attribute, as set by either `add_argument()` or by `set_defaults()`:

```
>>> parser = argparse.ArgumentParser()
>>> parser.add_argument('--foo', default='badger')
>>> parser.get_default('foo')
'badger'
```

Printing help

In most typical applications, `parse_args()` will take care of formatting and printing any usage or error messages. However, several formatting methods are available:

`ArgumentParser.print_usage(file=None)`

Print a brief description of how the `ArgumentParser` should be invoked on the command line. If `file` is `None`, `sys.stdout` is assumed.

`ArgumentParser.print_help(file=None)`

Print a help message, including the program usage and information about the arguments registered with the `ArgumentParser`. If `file` is `None`, `sys.stdout` is assumed.

There are also variants of these methods that simply return a string instead of printing it:

`ArgumentParser.format_usage()`

Return a string containing a brief description of how the `ArgumentParser` should be invoked on the command line.

`ArgumentParser.format_help()`

Return a string containing a help message, including the program usage and information about the arguments registered with the `ArgumentParser`.

Partial parsing

`ArgumentParser.parse_known_args(args=None, namespace=None)`

Sometimes a script may only parse a few of the command-line arguments, passing the remaining arguments on to another script or program. In these cases, the `parse_known_args()` method can be useful. It works much like `parse_args()` except that it does not produce an error when extra arguments are present. Instead, it returns a two item tuple containing the populated namespace and the list of remaining argument strings.

```
>>> parser = argparse.ArgumentParser()
>>> parser.add_argument('--foo', action='store_true')
>>> parser.add_argument('bar')
>>> parser.parse_known_args(['--foo', '--badger', 'BAR', 'spam'])
(Namespace(bar='BAR', foo=True), ['--badger', 'spam'])
```

警告: *Prefix matching* rules apply to `parse_known_args()`. The parser may consume an option even if it's just a prefix of one of its known options, instead of leaving it in the remaining arguments list.

Customizing file parsing

`ArgumentParser.convert_arg_line_to_args(arg_line)`

Arguments that are read from a file (see the *fromfile_prefix_chars* keyword argument to the *ArgumentParser* constructor) are read one argument per line. *convert_arg_line_to_args()* can be overridden for fancier reading.

This method takes a single argument *arg_line* which is a string read from the argument file. It returns a list of arguments parsed from this string. The method is called once per line read from the argument file, in order.

A useful override of this method is one that treats each space-separated word as an argument. The following example demonstrates how to do this:

```
class MyArgumentParser(argparse.ArgumentParser):
    def convert_arg_line_to_args(self, arg_line):
        return arg_line.split()
```

Exiting methods

`ArgumentParser.exit(status=0, message=None)`

This method terminates the program, exiting with the specified *status* and, if given, it prints a *message* before that.

`ArgumentParser.error(message)`

This method prints a usage message including the *message* to the standard error and terminates the program with a status code of 2.

16.4.6 Upgrading optparse code

Originally, the *argparse* module had attempted to maintain compatibility with *optparse*. However, *optparse* was difficult to extend transparently, particularly with the changes required to support the new *nargs=* specifiers and better usage messages. When most everything in *optparse* had either been copy-pasted over or monkey-patched, it no longer seemed practical to try to maintain the backwards compatibility.

The *argparse* module improves on the standard library *optparse* module in a number of ways including:

- Handling positional arguments.
- Supporting sub-commands.
- Allowing alternative option prefixes like `+` and `/`.
- Handling zero-or-more and one-or-more style arguments.
- Producing more informative usage messages.

- Providing a much simpler interface for custom type and action.

A partial upgrade path from *optparse* to *argparse*:

- Replace all `optparse.OptionParser.add_option()` calls with `ArgumentParser.add_argument()` calls.
- Replace `(options, args) = parser.parse_args()` with `args = parser.parse_args()` and add additional `ArgumentParser.add_argument()` calls for the positional arguments. Keep in mind that what was previously called `options`, now in the *argparse* context is called `args`.
- Replace `optparse.OptionParser.disable_interspersed_args()` by setting `nargs` of a positional argument to `argparse.REMAINDER`, or use `parse_known_args()` to collect unparsed argument strings in a separate list.
- Replace callback actions and the `callback_*` keyword arguments with `type` or `action` arguments.
- Replace string names for `type` keyword arguments with the corresponding type objects (e.g. `int`, `float`, `complex`, etc).
- Replace `optparse.Values` with `Namespace` and `optparse.OptionError` and `optparse.OptionValueError` with `ArgumentError`.
- Replace strings with implicit arguments such as `%default` or `%prog` with the standard Python syntax to use dictionaries to format strings, that is, `%(default)s` and `%(prog)s`.
- Replace the `OptionParser` constructor `version` argument with a call to `parser.add_argument('--version', action='version', version='<the version>')`.

16.5 getopt —C-style parser for command line options

Source code: [Lib/getopt.py](#)

注解: The *getopt* module is a parser for command line options whose API is designed to be familiar to users of the C `getopt()` function. Users who are unfamiliar with the C `getopt()` function or who would like to write less code and get better help and error messages should consider using the *argparse* module instead.

This module helps scripts to parse the command line arguments in `sys.argv`. It supports the same conventions as the Unix `getopt()` function (including the special meanings of arguments of the form `'-'` and `'--'`). Long options similar to those supported by GNU software may be used as well via an optional third argument.

This module provides two functions and an exception:

`getopt.getopt(args, shortopts, longopts=[])`

Parses command line options and parameter list. *args* is the argument list to be parsed, without the leading reference to the running program. Typically, this means `sys.argv[1:]`. *shortopts* is the string of option letters that the script wants to recognize, with options that require an argument followed by a colon (':'); i.e., the same format that Unix `getopt()` uses).

注解: Unlike GNU `getopt()`, after a non-option argument, all further arguments are considered also non-options. This is similar to the way non-GNU Unix systems work.

longopts, if specified, must be a list of strings with the names of the long options which should be supported. The leading `'--'` characters should not be included in the option name. Long options which require an argument

should be followed by an equal sign ('='). Optional arguments are not supported. To accept only long options, *shortopts* should be an empty string. Long options on the command line can be recognized so long as they provide a prefix of the option name that matches exactly one of the accepted options. For example, if *longopts* is ['foo', 'frob'], the option --fo will match as --foo, but --f will not match uniquely, so *GetoptError* will be raised.

The return value consists of two elements: the first is a list of (option, value) pairs; the second is the list of program arguments left after the option list was stripped (this is a trailing slice of *args*). Each option-and-value pair returned has the option as its first element, prefixed with a hyphen for short options (e.g., '-x') or two hyphens for long options (e.g., '--long-option'), and the option argument as its second element, or an empty string if the option has no argument. The options occur in the list in the same order in which they were found, thus allowing multiple occurrences. Long and short options may be mixed.

`getopt.gnu_getopt (args, shortopts, longopts=[])`

This function works like *getopt()*, except that GNU style scanning mode is used by default. This means that option and non-option arguments may be intermixed. The *getopt()* function stops processing options as soon as a non-option argument is encountered.

If the first character of the option string is '+', or if the environment variable `POSIIXLY_CORRECT` is set, then option processing stops as soon as a non-option argument is encountered.

exception `getopt.GetoptError`

This is raised when an unrecognized option is found in the argument list or when an option requiring an argument is given none. The argument to the exception is a string indicating the cause of the error. For long options, an argument given to an option which does not require one will also cause this exception to be raised. The attributes *msg* and *opt* give the error message and related option; if there is no specific option to which the exception relates, *opt* is an empty string.

exception `getopt.error`

Alias for *GetoptError*; for backward compatibility.

An example using only Unix style options:

```
>>> import getopt
>>> args = '-a -b -cfoo -d bar a1 a2'.split()
>>> args
['-a', '-b', '-cfoo', '-d', 'bar', 'a1', 'a2']
>>> optlist, args = getopt.getopt(args, 'abc:d:')
>>> optlist
[('-a', ''), ('-b', ''), ('-c', 'foo'), ('-d', 'bar')]
>>> args
['a1', 'a2']
```

Using long option names is equally easy:

```
>>> s = '--condition=foo --testing --output-file abc.def -x a1 a2'
>>> args = s.split()
>>> args
['--condition=foo', '--testing', '--output-file', 'abc.def', '-x', 'a1', 'a2']
>>> optlist, args = getopt.getopt(args, 'x', [
...     'condition=', 'output-file=', 'testing'])
>>> optlist
[('--condition', 'foo'), ('--testing', ''), ('--output-file', 'abc.def'), ('-x', '')]
>>> args
['a1', 'a2']
```

In a script, typical usage is something like this:

```

import getopt, sys

def main():
    try:
        opts, args = getopt.getopt(sys.argv[1:], "ho:v", ["help", "output="])
    except getopt.GetoptError as err:
        # print help information and exit:
        print(err) # will print something like "option -a not recognized"
        usage()
        sys.exit(2)
    output = None
    verbose = False
    for o, a in opts:
        if o == "-v":
            verbose = True
        elif o in ("-h", "--help"):
            usage()
            sys.exit()
        elif o in ("-o", "--output"):
            output = a
        else:
            assert False, "unhandled option"
    # ...

if __name__ == "__main__":
    main()

```

Note that an equivalent command line interface could be produced with less code and more informative help and error messages by using the [argparse](#) module:

```

import argparse

if __name__ == '__main__':
    parser = argparse.ArgumentParser()
    parser.add_argument('-o', '--output')
    parser.add_argument('-v', dest='verbose', action='store_true')
    args = parser.parse_args()
    # ... do something with args.output ...
    # ... do something with args.verbose ..

```

参见:

Module [argparse](#) Alternative command line option and argument parsing library.

16.6 模块 logging —Python 的日志记录工具

源代码: [Lib/logging/__init__.py](#)

Important

此页面仅包含 API 参考信息。有关更多高级主题的教程信息和讨论，请参阅

- 基础教程
- 进阶教程

这个模块为应用与库定义了实现灵活的事件日志系统的函数与类。

使用标准库提供的 logging API 最主要的好处是，所有的 Python 模块都可能参与日志输出，包括你的日志消息和第三方模块的日志消息。

这个模块提供许多强大而灵活的功能。如果你对 logging 不太熟悉的话，掌握它最好的方式就是查看它对应的教程（详见右侧的链接）。

该模块定义的基础类和函数都列在下面。

- 记录器暴露了应用程序代码直接使用的接口。
- 处理程序将日志记录（由记录器创建）发送到适当的目标。
- 过滤器提供了更精细的设施，用于确定要输出的日志记录。
- 格式化程序指定最终输出中日志记录的样式。

16.6.1 Logger 对象

Loggers 有下面的属性和方法。要注意的是，不要直接实例化 Loggers，但是可以通过模块级的函数“logging.getLogger(name)”。多次使用相同的名字调用 `getLogger()` 函数只会返回指向同一个 Logger 对象。

The name is potentially a period-separated hierarchical value, like `foo.bar.baz` (though it could also be just plain `foo`, for example). Loggers that are further down in the hierarchical list are children of loggers higher up in the list. For example, given a logger with a name of `foo`, loggers with names of `foo.bar`, `foo.bar.baz`, and `foo.bam` are all descendants of `foo`. The logger name hierarchy is analogous to the Python package hierarchy, and identical to it if you organise your loggers on a per-module basis using the recommended construction `logging.getLogger(__name__)`. That's because in a module, `__name__` is the module's name in the Python package namespace.

```
class logging.Logger
```

propagate

If this attribute evaluates to true, events logged to this logger will be passed to the handlers of higher level (ancestor) loggers, in addition to any handlers attached to this logger. Messages are passed directly to the ancestor loggers' handlers - neither the level nor filters of the ancestor loggers in question are considered.

If this evaluates to false, logging messages are not passed to the handlers of ancestor loggers.

构造器将这个属性初始化为 `True`。

注解： If you attach a handler to a logger *and* one or more of its ancestors, it may emit the same record multiple times. In general, you should not need to attach a handler to more than one logger - if you just attach it to the appropriate logger which is highest in the logger hierarchy, then it will see all events logged by all descendant loggers, provided that their propagate setting is left set to `True`. A common scenario is to attach handlers only to the root logger, and to let propagation take care of the rest.

setLevel (level)

Sets the threshold for this logger to *level*. Logging messages which are less severe than *level* will be ignored; logging messages which have severity *level* or higher will be emitted by whichever handler or handlers service this logger, unless a handler's level has been set to a higher severity level than *level*.

When a logger is created, the level is set to `NOTSET` (which causes all messages to be processed when the logger is the root logger, or delegation to the parent when the logger is a non-root logger). Note that the root logger is created with level `WARNING`.

The term ‘delegation to the parent’ means that if a logger has a level of `NOTSET`, its chain of ancestor loggers is traversed until either an ancestor with a level other than `NOTSET` is found, or the root is reached.

If an ancestor is found with a level other than `NOTSET`, then that ancestor’s level is treated as the effective level of the logger where the ancestor search began, and is used to determine how a logging event is handled.

If the root is reached, and it has a level of `NOTSET`, then all messages will be processed. Otherwise, the root’s level will be used as the effective level.

See [日志级别](#) for a list of levels.

在 3.2 版更改: The *level* parameter now accepts a string representation of the level such as ‘`INFO`’ as an alternative to the integer constants such as `INFO`. Note, however, that levels are internally stored as integers, and methods such as e.g. `getEffectiveLevel()` and `isEnabledFor()` will return/expect to be passed integers.

`isEnabledFor(lvl)`

Indicates if a message of severity *lvl* would be processed by this logger. This method checks first the module-level level set by `logging.disable(lvl)` and then the logger’s effective level as determined by `getEffectiveLevel()`.

`getEffectiveLevel()`

Indicates the effective level for this logger. If a value other than `NOTSET` has been set using `setLevel()`, it is returned. Otherwise, the hierarchy is traversed towards the root until a value other than `NOTSET` is found, and that value is returned. The value returned is an integer, typically one of `logging.DEBUG`, `logging.INFO` etc.

`getChild(suffix)`

Returns a logger which is a descendant to this logger, as determined by the suffix. Thus, `logging.getLogger('abc').getChild('def.ghi')` would return the same logger as would be returned by `logging.getLogger('abc.def.ghi')`. This is a convenience method, useful when the parent logger is named using e.g. `__name__` rather than a literal string.

3.2 新版功能.

`debug(msg, *args, **kwargs)`

Logs a message with level `DEBUG` on this logger. The *msg* is the message format string, and the *args* are the arguments which are merged into *msg* using the string formatting operator. (Note that this means that you can use keywords in the format string, together with a single dictionary argument.)

There are three keyword arguments in *kwargs* which are inspected: *exc_info*, *stack_info*, and *extra*.

If *exc_info* does not evaluate as false, it causes exception information to be added to the logging message. If an exception tuple (in the format returned by `sys.exc_info()`) or an exception instance is provided, it is used; otherwise, `sys.exc_info()` is called to get the exception information.

The second optional keyword argument is *stack_info*, which defaults to `False`. If true, stack information is added to the logging message, including the actual logging call. Note that this is not the same stack information as that displayed through specifying *exc_info*: The former is stack frames from the bottom of the stack up to the logging call in the current thread, whereas the latter is information about stack frames which have been unwound, following an exception, while searching for exception handlers.

You can specify *stack_info* independently of *exc_info*, e.g. to just show how you got to a certain point in your code, even when no exceptions were raised. The stack frames are printed following a header line which says:

```
Stack (most recent call last):
```

This mimics the `Traceback (most recent call last):` which is used when displaying exception frames.

The third keyword argument is *extra* which can be used to pass a dictionary which is used to populate the `__dict__` of the `LogRecord` created for the logging event with user-defined attributes. These custom attributes can then be used as you like. For example, they could be incorporated into logged messages. For example:

```
FORMAT = '%(asctime)-15s %(clientip)s %(user)-8s %(message)s'
logging.basicConfig(format=FORMAT)
d = {'clientip': '192.168.0.1', 'user': 'fbloggs'}
logger = logging.getLogger('tcpserver')
logger.warning('Protocol problem: %s', 'connection reset', extra=d)
```

would print something like

```
2006-02-08 22:20:02,165 192.168.0.1 fbloggs Protocol problem: connection_
↪reset
```

The keys in the dictionary passed in *extra* should not clash with the keys used by the logging system. (See the *Formatter* documentation for more information on which keys are used by the logging system.)

If you choose to use these attributes in logged messages, you need to exercise some care. In the above example, for instance, the *Formatter* has been set up with a format string which expects ‘clientip’ and ‘user’ in the attribute dictionary of the `LogRecord`. If these are missing, the message will not be logged because a string formatting exception will occur. So in this case, you always need to pass the *extra* dictionary with these keys.

While this might be annoying, this feature is intended for use in specialized circumstances, such as multi-threaded servers where the same code executes in many contexts, and interesting conditions which arise are dependent on this context (such as remote client IP address and authenticated user name, in the above example). In such circumstances, it is likely that specialized *Formatters* would be used with particular *Handlers*.

3.2 新版功能: The *stack_info* parameter was added.

在 3.5 版更改: The *exc_info* parameter can now accept exception instances.

info (*msg*, **args*, ***kwargs*)

Logs a message with level INFO on this logger. The arguments are interpreted as for *debug()*.

warning (*msg*, **args*, ***kwargs*)

Logs a message with level WARNING on this logger. The arguments are interpreted as for *debug()*.

注解: There is an obsolete method `warn` which is functionally identical to `warning`. As `warn` is deprecated, please do not use it - use `warning` instead.

error (*msg*, **args*, ***kwargs*)

Logs a message with level ERROR on this logger. The arguments are interpreted as for *debug()*.

critical (*msg*, **args*, ***kwargs*)

Logs a message with level CRITICAL on this logger. The arguments are interpreted as for *debug()*.

log (*lvl*, *msg*, **args*, ***kwargs*)

Logs a message with integer level *lvl* on this logger. The other arguments are interpreted as for *debug()*.

exception (*msg*, **args*, ***kwargs*)

Logs a message with level ERROR on this logger. The arguments are interpreted as for *debug()*. Exception info is added to the logging message. This method should only be called from an exception handler.

addFilter (*filter*)

Adds the specified filter *filter* to this logger.

removeFilter (*filter*)

Removes the specified filter *filter* from this logger.

filter (*record*)

Applies this logger's filters to the record and returns a true value if the record is to be processed. The filters are consulted in turn, until one of them returns a false value. If none of them return a false value, the record will be processed (passed to handlers). If one returns a false value, no further processing of the record occurs.

addHandler (*hdlr*)

Adds the specified handler *hdlr* to this logger.

removeHandler (*hdlr*)

Removes the specified handler *hdlr* from this logger.

findCaller (*stack_info=False*)

Finds the caller's source filename and line number. Returns the filename, line number, function name and stack information as a 4-element tuple. The stack information is returned as `None` unless *stack_info* is `True`.

handle (*record*)

Handles a record by passing it to all handlers associated with this logger and its ancestors (until a false value of *propagate* is found). This method is used for unpickled records received from a socket, as well as those created locally. Logger-level filtering is applied using *filter()*.

makeRecord (*name, lvl, fn, lno, msg, args, exc_info, func=None, extra=None, sinfo=None*)

This is a factory method which can be overridden in subclasses to create specialized *LogRecord* instances.

hasHandlers ()

Checks to see if this logger has any handlers configured. This is done by looking for handlers in this logger and its parents in the logger hierarchy. Returns `True` if a handler was found, else `False`. The method stops searching up the hierarchy whenever a logger with the 'propagate' attribute set to false is found - that will be the last logger which is checked for the existence of handlers.

3.2 新版功能.

在 3.7 版更改: Loggers can now be pickled and unpickled.

16.6.2 日志级别

日志记录级别的数值在下表中给出。如果你想要定义自己的级别，并且需要它们具有相对于预定义级别的特定值，那么这些内容可能是你感兴趣的。如果你定义具有相同数值的级别，它将覆盖预定义的值；预定义的名称丢失。

级别	数值
CRITICAL	50
ERROR	40
WARNING	30
INFO	20
DEBUG	10
NOTSET	0

16.6.3 Handler Objects

Handlers have the following attributes and methods. Note that *Handler* is never instantiated directly; this class acts as a base for more useful subclasses. However, the `__init__()` method in subclasses needs to call *Handler*.
`__init__()`.

class `logging.Handler`

`__init__` (*level*=NOTSET)

Initializes the *Handler* instance by setting its level, setting the list of filters to the empty list and creating a lock (using `createLock()`) for serializing access to an I/O mechanism.

`createLock` ()

Initializes a thread lock which can be used to serialize access to underlying I/O functionality which may not be threadsafe.

`acquire` ()

Acquires the thread lock created with `createLock()`.

`release` ()

Releases the thread lock acquired with `acquire()`.

`setLevel` (*level*)

Sets the threshold for this handler to *level*. Logging messages which are less severe than *level* will be ignored. When a handler is created, the level is set to NOTSET (which causes all messages to be processed).

See [日志级别](#) for a list of levels.

在 3.2 版更改: The *level* parameter now accepts a string representation of the level such as ‘INFO’ as an alternative to the integer constants such as INFO.

`setFormatter` (*fmt*)

Sets the *Formatter* for this handler to *fmt*.

`addFilter` (*filter*)

Adds the specified filter *filter* to this handler.

`removeFilter` (*filter*)

Removes the specified filter *filter* from this handler.

`filter` (*record*)

Applies this handler’s filters to the record and returns a true value if the record is to be processed. The filters are consulted in turn, until one of them returns a false value. If none of them return a false value, the record will be emitted. If one returns a false value, the handler will not emit the record.

`flush` ()

Ensure all logging output has been flushed. This version does nothing and is intended to be implemented by subclasses.

`close` ()

Tidy up any resources used by the handler. This version does no output but removes the handler from an internal list of handlers which is closed when `shutdown()` is called. Subclasses should ensure that this gets called from overridden `close()` methods.

`handle` (*record*)

Conditionally emits the specified logging record, depending on filters which may have been added to the handler. Wraps the actual emission of the record with acquisition/release of the I/O thread lock.

`handleError` (*record*)

This method should be called from handlers when an exception is encountered during an `emit()` call. If the module-level attribute `raiseExceptions` is False, exceptions get silently ignored. This is what is

mostly wanted for a logging system - most users will not care about errors in the logging system, they are more interested in application errors. You could, however, replace this with a custom handler if you wish. The specified record is the one which was being processed when the exception occurred. (The default value of `raiseExceptions` is `True`, as that is more useful during development).

format (*record*)

Do formatting for a record - if a formatter is set, use it. Otherwise, use the default formatter for the module.

emit (*record*)

Do whatever it takes to actually log the specified logging record. This version is intended to be implemented by subclasses and so raises a `NotImplementedError`.

For a list of handlers included as standard, see `logging.handlers`.

16.6.4 Formatter Objects

`Formatter` objects have the following attributes and methods. They are responsible for converting a `LogRecord` to (usually) a string which can be interpreted by either a human or an external system. The base `Formatter` allows a formatting string to be specified. If none is supplied, the default value of `'%(message)s'` is used, which just includes the message in the logging call. To have additional items of information in the formatted output (such as a timestamp), keep reading.

A Formatter can be initialized with a format string which makes use of knowledge of the `LogRecord` attributes - such as the default value mentioned above making use of the fact that the user's message and arguments are pre-formatted into a `LogRecord`'s `message` attribute. This format string contains standard Python %-style mapping keys. See section `printf 风格的字符串格式化` for more information on string formatting.

The useful mapping keys in a `LogRecord` are given in the section on `LogRecord attributes`.

class `logging.Formatter` (*fmt=None, datefmt=None, style='%*)

Returns a new instance of the `Formatter` class. The instance is initialized with a format string for the message as a whole, as well as a format string for the date/time portion of a message. If no *fmt* is specified, `'%(message)s'` is used. If no *datefmt* is specified, a format is used which is described in the `formatTime()` documentation.

The *style* parameter can be one of `'%'`, `'{'` or `'$'` and determines how the format string will be merged with its data: using one of %-formatting, `str.format()` or `string.Template`. See formatting-styles for more information on using `{- and $-formatting` for log messages.

在 3.2 版更改: The *style* parameter was added.

format (*record*)

The record's attribute dictionary is used as the operand to a string formatting operation. Returns the resulting string. Before formatting the dictionary, a couple of preparatory steps are carried out. The `message` attribute of the record is computed using `msg % args`. If the formatting string contains `'(asctime)'`, `formatTime()` is called to format the event time. If there is exception information, it is formatted using `formatException()` and appended to the message. Note that the formatted exception information is cached in attribute `exc_text`. This is useful because the exception information can be pickled and sent across the wire, but you should be careful if you have more than one `Formatter` subclass which customizes the formatting of exception information. In this case, you will have to clear the cached value after a formatter has done its formatting, so that the next formatter to handle the event doesn't use the cached value but recalculates it afresh.

If stack information is available, it's appended after the exception information, using `formatStack()` to transform it if necessary.

formatTime (*record, datefmt=None*)

This method should be called from `format()` by a formatter which wants to make use of a formatted time. This method can be overridden in formatters to provide for any specific requirement, but the basic behavior

is as follows: if *datefmt* (a string) is specified, it is used with *time.strftime()* to format the creation time of the record. Otherwise, the format `'%Y-%m-%d %H:%M:%S,uuu'` is used, where the *uuu* part is a millisecond value and the other letters are as per the *time.strftime()* documentation. An example time in this format is `2003-01-23 00:29:50,411`. The resulting string is returned.

This function uses a user-configurable function to convert the creation time to a tuple. By default, *time.localtime()* is used; to change this for a particular formatter instance, set the *converter* attribute to a function with the same signature as *time.localtime()* or *time.gmtime()*. To change it for all formatters, for example if you want all logging times to be shown in GMT, set the *converter* attribute in the *Formatter* class.

在 3.3 版更改: Previously, the default format was hard-coded as in this example: `2010-09-06 22:38:15,292` where the part before the comma is handled by a *strftime* format string (`'%Y-%m-%d %H:%M:%S'`), and the part after the comma is a millisecond value. Because *strftime* does not have a format placeholder for milliseconds, the millisecond value is appended using another format string, `'%s,%03d'` —and both of these format strings have been hardcoded into this method. With the change, these strings are defined as class-level attributes which can be overridden at the instance level when desired. The names of the attributes are *default_time_format* (for the *strftime* format string) and *default_msec_format* (for appending the millisecond value).

formatException (*exc_info*)

Formats the specified exception information (a standard exception tuple as returned by *sys.exc_info()*) as a string. This default implementation just uses *traceback.print_exception()*. The resulting string is returned.

formatStack (*stack_info*)

Formats the specified stack information (a string as returned by *traceback.print_stack()*, but with the last newline removed) as a string. This default implementation just returns the input value.

16.6.5 Filter Objects

Filters can be used by *Handlers* and *Loggers* for more sophisticated filtering than is provided by levels. The base filter class only allows events which are below a certain point in the logger hierarchy. For example, a filter initialized with `'A.B'` will allow events logged by loggers `'A.B'`, `'A.B.C'`, `'A.B.C.D'`, `'A.B.D'` etc. but not `'A.BB'`, `'B.A.B'` etc. If initialized with the empty string, all events are passed.

class *logging.Filter* (*name=""*)

Returns an instance of the *Filter* class. If *name* is specified, it names a logger which, together with its children, will have its events allowed through the filter. If *name* is the empty string, allows every event.

filter (*record*)

Is the specified record to be logged? Returns zero for no, nonzero for yes. If deemed appropriate, the record may be modified in-place by this method.

Note that filters attached to handlers are consulted before an event is emitted by the handler, whereas filters attached to loggers are consulted whenever an event is logged (using *debug()*, *info()*, etc.), before sending an event to handlers. This means that events which have been generated by descendant loggers will not be filtered by a logger's filter setting, unless the filter has also been applied to those descendant loggers.

You don't actually need to subclass *Filter*: you can pass any instance which has a *filter* method with the same semantics.

在 3.2 版更改: You don't need to create specialized *Filter* classes, or use other classes with a *filter* method: you can use a function (or other callable) as a filter. The filtering logic will check to see if the filter object has a *filter* attribute: if it does, it's assumed to be a *Filter* and its *filter()* method is called. Otherwise, it's assumed to be a callable and called with the record as the single parameter. The returned value should conform to that returned by *filter()*.

Although filters are used primarily to filter records based on more sophisticated criteria than levels, they get to see every record which is processed by the handler or logger they're attached to: this can be useful if you want to do things like counting how many records were processed by a particular logger or handler, or adding, changing or removing attributes in the `LogRecord` being processed. Obviously changing the `LogRecord` needs to be done with some care, but it does allow the injection of contextual information into logs (see `filters-contextual`).

16.6.6 LogRecord Objects

`LogRecord` instances are created automatically by the `Logger` every time something is logged, and can be created manually via `makeLogRecord()` (for example, from a pickled event received over the wire).

class `logging.LogRecord` (*name, level, pathname, lineno, msg, args, exc_info, func=None, sinfo=None*)
Contains all the information pertinent to the event being logged.

The primary information is passed in `msg` and `args`, which are combined using `msg % args` to create the message field of the record.

参数

- **name** – The name of the logger used to log the event represented by this `LogRecord`. Note that this name will always have this value, even though it may be emitted by a handler attached to a different (ancestor) logger.
- **level** – The numeric level of the logging event (one of `DEBUG`, `INFO` etc.) Note that this is converted to *two* attributes of the `LogRecord`: `lineno` for the numeric value and `levelname` for the corresponding level name.
- **pathname** – The full pathname of the source file where the logging call was made.
- **lineno** – The line number in the source file where the logging call was made.
- **msg** – The event description message, possibly a format string with placeholders for variable data.
- **args** – Variable data to merge into the `msg` argument to obtain the event description.
- **exc_info** – An exception tuple with the current exception information, or `None` if no exception information is available.
- **func** – The name of the function or method from which the logging call was invoked.
- **sinfo** – A text string representing stack information from the base of the stack in the current thread, up to the logging call.

`getMessage()`

Returns the message for this `LogRecord` instance after merging any user-supplied arguments with the message. If the user-supplied message argument to the logging call is not a string, `str()` is called on it to convert it to a string. This allows use of user-defined classes as messages, whose `__str__` method can return the actual format string to be used.

在 3.2 版更改: The creation of a `LogRecord` has been made more configurable by providing a factory which is used to create the record. The factory can be set using `getLogRecordFactory()` and `setLogRecordFactory()` (see this for the factory's signature).

This functionality can be used to inject your own values into a `LogRecord` at creation time. You can use the following pattern:

```
old_factory = logging.getLogRecordFactory()

def record_factory(*args, **kwargs):
```

(下页继续)

(续上页)

```
record = old_factory(*args, **kwargs)
record.custom_attribute = 0xdecafbad
return record

logging.setLogRecordFactory(record_factory)
```

With this pattern, multiple factories could be chained, and as long as they don't overwrite each other's attributes or unintentionally overwrite the standard attributes listed above, there should be no surprises.

16.6.7 LogRecord attributes

The LogRecord has a number of attributes, most of which are derived from the parameters to the constructor. (Note that the names do not always correspond exactly between the LogRecord constructor parameters and the LogRecord attributes.) These attributes can be used to merge data from the record into the format string. The following table lists (in alphabetical order) the attribute names, their meanings and the corresponding placeholder in a %-style format string.

If you are using {}-formatting (*str.format()*), you can use {attrname} as the placeholder in the format string. If you are using \$-formatting (*string.Template*), use the form \${attrname}. In both cases, of course, replace attrname with the actual attribute name you want to use.

In the case of {}-formatting, you can specify formatting flags by placing them after the attribute name, separated from it with a colon. For example: a placeholder of {msecs:03d} would format a millisecond value of 4 as 004. Refer to the *str.format()* documentation for full details on the options available to you.

属性名称	格式	描述
args	You shouldn't need to format this yourself.	The tuple of arguments merged into <code>msg</code> to produce <code>message</code> , or a dict whose values are used for the merge (when there is only one argument, and it is a dictionary).
asctime	<code>%(asctime)s</code>	Human-readable time when the <code>LogRecord</code> was created. By default this is of the form <code>'2003-07-08 16:49:45,896'</code> (the numbers after the comma are millisecond portion of the time).
created	<code>%(created)f</code>	Time when the <code>LogRecord</code> was created (as returned by <code>time.time()</code>).
exc_info	You shouldn't need to format this yourself.	Exception tuple (à la <code>sys.exc_info</code>) or, if no exception has occurred, <code>None</code> .
filename	<code>%(filename)s</code>	Filename portion of <code>pathname</code> .
funcName	<code>%(funcName)s</code>	Name of function containing the logging call.
levelname	<code>%(levelname)s</code>	Text logging level for the message (<code>'DEBUG'</code> , <code>'INFO'</code> , <code>'WARNING'</code> , <code>'ERROR'</code> , <code>'CRITICAL'</code>).
levelno	<code>%(levelno)s</code>	Numeric logging level for the message (<code>DEBUG</code> , <code>INFO</code> , <code>WARNING</code> , <code>ERROR</code> , <code>CRITICAL</code>).
lineno	<code>%(lineno)d</code>	Source line number where the logging call was issued (if available).
message	<code>%(message)s</code>	The logged message, computed as <code>msg % args</code> . This is set when <code>Formatter.format()</code> is invoked.
module	<code>%(module)s</code>	Module (name portion of <code>filename</code>).
msecs	<code>%(msecs)d</code>	Millisecond portion of the time when the <code>LogRecord</code> was created.
msg	You shouldn't need to format this yourself.	The format string passed in the original logging call. Merged with <code>args</code> to produce <code>message</code> , or an arbitrary object (see arbitrary-object-messages).
name	<code>%(name)s</code>	Name of the logger used to log the call.
pathname	<code>%(pathname)s</code>	Full pathname of the source file where the logging call was issued (if available).
process	<code>%(process)d</code>	Process ID (if available).
processName	<code>%(processName)s</code>	Process name (if available).
relativeCreated	<code>%(relativeCreated)d</code>	Time in milliseconds when the <code>LogRecord</code> was created, relative to the time the logging module was loaded.
stack_info	You shouldn't need to format this yourself.	Stack frame information (where available) from the bottom of the stack in the current thread, up to and including the stack frame of the logging call which resulted in the creation of this record.
thread	<code>%(thread)d</code>	Thread ID (if available).
threadName	<code>%(threadName)s</code>	Thread name (if available).

在 3.1 版更改: `processName` was added.

16.6.8 LoggerAdapter Objects

LoggerAdapter instances are used to conveniently pass contextual information into logging calls. For a usage example, see the section on adding contextual information to your logging output.

class `logging.LoggerAdapter` (*logger, extra*)

Returns an instance of *LoggerAdapter* initialized with an underlying *Logger* instance and a dict-like object.

process (*msg, kwargs*)

Modifies the message and/or keyword arguments passed to a logging call in order to insert contextual information. This implementation takes the object passed as *extra* to the constructor and adds it to *kwargs* using key 'extra'. The return value is a (*msg, kwargs*) tuple which has the (possibly modified) versions of the arguments passed in.

In addition to the above, *LoggerAdapter* supports the following methods of *Logger*: *debug()*, *info()*, *warning()*, *error()*, *exception()*, *critical()*, *log()*, *isEnabledFor()*, *getEffectiveLevel()*, *setLevel()* and *hasHandlers()*. These methods have the same signatures as their counterparts in *Logger*, so you can use the two types of instances interchangeably.

在 3.2 版更改: The *isEnabledFor()*, *getEffectiveLevel()*, *setLevel()* and *hasHandlers()* methods were added to *LoggerAdapter*. These methods delegate to the underlying logger.

16.6.9 线程安全

The logging module is intended to be thread-safe without any special work needing to be done by its clients. It achieves this though using threading locks; there is one lock to serialize access to the module's shared data, and each handler also creates a lock to serialize access to its underlying I/O.

If you are implementing asynchronous signal handlers using the *signal* module, you may not be able to use logging from within such handlers. This is because lock implementations in the *threading* module are not always re-entrant, and so cannot be invoked from such signal handlers.

16.6.10 模块级别函数

In addition to the classes described above, there are a number of module-level functions.

`logging.getLogger` (*name=None*)

Return a logger with the specified name or, if *name* is *None*, return a logger which is the root logger of the hierarchy. If specified, the name is typically a dot-separated hierarchical name like 'a', 'a.b' or 'a.b.c.d'. Choice of these names is entirely up to the developer who is using logging.

All calls to this function with a given name return the same logger instance. This means that logger instances never need to be passed between different parts of an application.

`logging.getLoggerClass` ()

Return either the standard *Logger* class, or the last class passed to *setLoggerClass()*. This function may be called from within a new class definition, to ensure that installing a customized *Logger* class will not undo customizations already applied by other code. For example:

```
class MyLogger(logging.getLoggerClass()):
    # ... override behaviour here
```

`logging.getLogRecordFactory` ()

Return a callable which is used to create a *LogRecord*.

3.2 新版功能: This function has been provided, along with *setLogRecordFactory()*, to allow developers more control over how the *LogRecord* representing a logging event is constructed.

See `setLogRecordFactory()` for more information about the how the factory is called.

`logging.debug(msg, *args, **kwargs)`

Logs a message with level `DEBUG` on the root logger. The `msg` is the message format string, and the `args` are the arguments which are merged into `msg` using the string formatting operator. (Note that this means that you can use keywords in the format string, together with a single dictionary argument.)

There are three keyword arguments in `kwargs` which are inspected: `exc_info` which, if it does not evaluate as false, causes exception information to be added to the logging message. If an exception tuple (in the format returned by `sys.exc_info()`) or an exception instance is provided, it is used; otherwise, `sys.exc_info()` is called to get the exception information.

The second optional keyword argument is `stack_info`, which defaults to `False`. If true, stack information is added to the logging message, including the actual logging call. Note that this is not the same stack information as that displayed through specifying `exc_info`: The former is stack frames from the bottom of the stack up to the logging call in the current thread, whereas the latter is information about stack frames which have been unwound, following an exception, while searching for exception handlers.

You can specify `stack_info` independently of `exc_info`, e.g. to just show how you got to a certain point in your code, even when no exceptions were raised. The stack frames are printed following a header line which says:

```
Stack (most recent call last):
```

This mimics the `Traceback (most recent call last):` which is used when displaying exception frames.

The third optional keyword argument is `extra` which can be used to pass a dictionary which is used to populate the `__dict__` of the `LogRecord` created for the logging event with user-defined attributes. These custom attributes can then be used as you like. For example, they could be incorporated into logged messages. For example:

```
FORMAT = '%(asctime)-15s %(clientip)s %(user)-8s %(message)s'
logging.basicConfig(format=FORMAT)
d = {'clientip': '192.168.0.1', 'user': 'fbloggs'}
logging.warning('Protocol problem: %s', 'connection reset', extra=d)
```

would print something like:

```
2006-02-08 22:20:02,165 192.168.0.1 fbloggs Protocol problem: connection reset
```

The keys in the dictionary passed in `extra` should not clash with the keys used by the logging system. (See the `Formatter` documentation for more information on which keys are used by the logging system.)

If you choose to use these attributes in logged messages, you need to exercise some care. In the above example, for instance, the `Formatter` has been set up with a format string which expects 'clientip' and 'user' in the attribute dictionary of the `LogRecord`. If these are missing, the message will not be logged because a string formatting exception will occur. So in this case, you always need to pass the `extra` dictionary with these keys.

While this might be annoying, this feature is intended for use in specialized circumstances, such as multi-threaded servers where the same code executes in many contexts, and interesting conditions which arise are dependent on this context (such as remote client IP address and authenticated user name, in the above example). In such circumstances, it is likely that specialized `Formatters` would be used with particular `Handlers`.

3.2 新版功能: The `stack_info` parameter was added.

`logging.info(msg, *args, **kwargs)`

Logs a message with level `INFO` on the root logger. The arguments are interpreted as for `debug()`.

`logging.warning(msg, *args, **kwargs)`

Logs a message with level `WARNING` on the root logger. The arguments are interpreted as for `debug()`.

注解: There is an obsolete function `warn` which is functionally identical to `warning`. As `warn` is deprecated, please do not use it - use `warning` instead.

`logging.error(msg, *args, **kwargs)`

Logs a message with level `ERROR` on the root logger. The arguments are interpreted as for `debug()`.

`logging.critical(msg, *args, **kwargs)`

Logs a message with level `CRITICAL` on the root logger. The arguments are interpreted as for `debug()`.

`logging.exception(msg, *args, **kwargs)`

Logs a message with level `ERROR` on the root logger. The arguments are interpreted as for `debug()`. Exception info is added to the logging message. This function should only be called from an exception handler.

`logging.log(level, msg, *args, **kwargs)`

Logs a message with level `level` on the root logger. The other arguments are interpreted as for `debug()`.

注解: The above module-level convenience functions, which delegate to the root logger, call `basicConfig()` to ensure that at least one handler is available. Because of this, they should *not* be used in threads, in versions of Python earlier than 2.7.1 and 3.2, unless at least one handler has been added to the root logger *before* the threads are started. In earlier versions of Python, due to a thread safety shortcoming in `basicConfig()`, this can (under rare circumstances) lead to handlers being added multiple times to the root logger, which can in turn lead to multiple messages for the same event.

`logging.disable(lvl=CRITICAL)`

Provides an overriding level `lvl` for all loggers which takes precedence over the logger's own level. When the need arises to temporarily throttle logging output down across the whole application, this function can be useful. Its effect is to disable all logging calls of severity `lvl` and below, so that if you call it with a value of `INFO`, then all `INFO` and `DEBUG` events would be discarded, whereas those of severity `WARNING` and above would be processed according to the logger's effective level. If `logging.disable(logging.NOTSET)` is called, it effectively removes this overriding level, so that logging output again depends on the effective levels of individual loggers.

Note that if you have defined any custom logging level higher than `CRITICAL` (this is not recommended), you won't be able to rely on the default value for the `lvl` parameter, but will have to explicitly supply a suitable value.

在 3.7 版更改: The `lvl` parameter was defaulted to level `CRITICAL`. See Issue #28524 for more information about this change.

`logging.addLevelName(lvl, levelName)`

Associates level `lvl` with text `levelName` in an internal dictionary, which is used to map numeric levels to a textual representation, for example when a `Formatter` formats a message. This function can also be used to define your own levels. The only constraints are that all levels used must be registered using this function, levels should be positive integers and they should increase in increasing order of severity.

注解: If you are thinking of defining your own levels, please see the section on custom-levels.

`logging.getLevelName(lvl)`

Returns the textual representation of logging level `lvl`. If the level is one of the predefined levels `CRITICAL`, `ERROR`, `WARNING`, `INFO` or `DEBUG` then you get the corresponding string. If you have associated levels with names using `addLevelName()` then the name you have associated with `lvl` is returned. If a numeric value corresponding to one of the defined levels is passed in, the corresponding string representation is returned. Otherwise, the string 'Level %s' % `lvl` is returned.

注解: Levels are internally integers (as they need to be compared in the logging logic). This function is used

to convert between an integer level and the level name displayed in the formatted log output by means of the `% (levelname)s` format specifier (see *LogRecord attributes*).

在 3.4 版更改: In Python versions earlier than 3.4, this function could also be passed a text level, and would return the corresponding numeric value of the level. This undocumented behaviour was considered a mistake, and was removed in Python 3.4, but reinstated in 3.4.2 due to retain backward compatibility.

```
logging.makeLogRecord(attrdict)
```

Creates and returns a new *LogRecord* instance whose attributes are defined by *attrdict*. This function is useful for taking a pickled *LogRecord* attribute dictionary, sent over a socket, and reconstituting it as a *LogRecord* instance at the receiving end.

```
logging.basicConfig(**kwargs)
```

Does basic configuration for the logging system by creating a *StreamHandler* with a default *Formatter* and adding it to the root logger. The functions *debug()*, *info()*, *warning()*, *error()* and *critical()* will call *basicConfig()* automatically if no handlers are defined for the root logger.

This function does nothing if the root logger already has handlers configured for it.

注解: This function should be called from the main thread before other threads are started. In versions of Python prior to 2.7.1 and 3.2, if this function is called from multiple threads, it is possible (in rare circumstances) that a handler will be added to the root logger more than once, leading to unexpected results such as messages being duplicated in the log.

The following keyword arguments are supported.

格式	描述
<i>filename</i>	Specifies that a <code>FileHandler</code> be created, using the specified filename, rather than a <code>StreamHandler</code> .
<i>filemode</i>	If <i>filename</i> is specified, open the file in this <i>mode</i> . Defaults to 'a'.
<i>format</i>	Use the specified format string for the handler.
<i>datefmt</i>	Use the specified date/time format, as accepted by <code>time.strftime()</code> .
<i>style</i>	If <i>format</i> is specified, use this style for the format string. One of '%', '{' or '\$' for <code>printf-style</code> , <code>str.format()</code> or <code>string.Template</code> respectively. Defaults to '% '.
<i>level</i>	Set the root logger level to the specified <i>level</i> .
<i>stream</i>	Use the specified stream to initialize the <code>StreamHandler</code> . Note that this argument is incompatible with <i>filename</i> - if both are present, a <code>ValueError</code> is raised.
<i>handlers</i>	If specified, this should be an iterable of already created handlers to add to the root logger. Any handlers which don't already have a formatter set will be assigned the default formatter created in this function. Note that this argument is incompatible with <i>filename</i> or <i>stream</i> - if both are present, a <code>ValueError</code> is raised.

在 3.2 版更改: The *style* argument was added.

In 3.3 版更改: The *handlers* argument was added. Additional checks were added to catch situations where incompatible arguments are specified (e.g. *handlers* together with *stream* or *filename*, or *stream* together with *filename*).

```
logging.shutdown()
```

Informs the logging system to perform an orderly shutdown by flushing and closing all handlers. This should be called at application exit and no further use of the logging system should be made after this call.

```
logging.setLoggerClass(klass)
```

Tells the logging system to use the class *klass* when instantiating a logger. The class should define `__init__()` such that only a name argument is required, and the `__init__()` should call `Logger.__init__()`. This function is typically called before any loggers are instantiated by applications which need to use custom logger behavior.

`logging.setLogRecordFactory(factory)`

Set a callable which is used to create a *LogRecord*.

参数 factory – The factory callable to be used to instantiate a log record.

3.2 新版功能: This function has been provided, along with *getLogRecordFactory()*, to allow developers more control over how the *LogRecord* representing a logging event is constructed.

The factory has the following signature:

```
factory(name, level, fn, lno, msg, args, exc_info, func=None, sinfo=None,
**kwargs)
```

name The logger name.

level The logging level (numeric).

fn The full pathname of the file where the logging call was made.

lno The line number in the file where the logging call was made.

msg The logging message.

args The arguments for the logging message.

exc_info An exception tuple, or None.

func The name of the function or method which invoked the logging call.

sinfo A stack traceback such as is provided by *traceback.print_stack()*, showing the call hierarchy.

kwargs Additional keyword arguments.

16.6.11 Module-Level Attributes

`logging.lastResort`

A “handler of last resort” is available through this attribute. This is a *StreamHandler* writing to `sys.stderr` with a level of `WARNING`, and is used to handle logging events in the absence of any logging configuration. The end result is to just print the message to `sys.stderr`. This replaces the earlier error message saying that “no handlers could be found for logger XYZ”. If you need the earlier behaviour for some reason, `lastResort` can be set to `None`.

3.2 新版功能.

16.6.12 Integration with the warnings module

The *captureWarnings()* function can be used to integrate *logging* with the *warnings* module.

`logging.captureWarnings(capture)`

This function is used to turn the capture of warnings by logging on and off.

If *capture* is `True`, warnings issued by the *warnings* module will be redirected to the logging system. Specifically, a warning will be formatted using *warnings.formatwarning()* and the resulting string logged to a logger named `'py.warnings'` with a severity of `WARNING`.

If *capture* is `False`, the redirection of warnings to the logging system will stop, and warnings will be redirected to their original destinations (i.e. those in effect before *captureWarnings(True)* was called).

参见:

模块 *logging.config* 日志记录模块的配置 API。

模块 `logging.handlers` 日志记录模块附带的有用处理程序。

PEP 282 - A Logging System 该提案描述了 Python 标准库中包含的这个特性。

Original Python logging package This is the original source for the `logging` package. The version of the package available from this site is suitable for use with Python 1.5.2, 2.1.x and 2.2.x, which do not include the `logging` package in the standard library.

16.7 logging.config — 日志记录配置

源代码: [Lib/logging/config.py](#)

Important

此页面仅包含参考信息。有关教程，请参阅

- 基础教程
- Advanced Tutorial
- Logging Cookbook

This section describes the API for configuring the logging module.

16.7.1 Configuration functions

The following functions configure the logging module. They are located in the `logging.config` module. Their use is optional—you can configure the logging module using these functions or by making calls to the main API (defined in `logging` itself) and defining handlers which are declared either in `logging` or `logging.handlers`.

`logging.config.dictConfig(config)`

Takes the logging configuration from a dictionary. The contents of this dictionary are described in *Configuration dictionary schema* below.

If an error is encountered during configuration, this function will raise a `ValueError`, `TypeError`, `AttributeError` or `ImportError` with a suitably descriptive message. The following is a (possibly incomplete) list of conditions which will raise an error:

- A `level` which is not a string or which is a string not corresponding to an actual logging level.
- A `propagate` value which is not a boolean.
- An `id` which does not have a corresponding destination.
- A non-existent handler `id` found during an incremental call.
- An invalid logger name.
- Inability to resolve to an internal or external object.

Parsing is performed by the `DictConfigurator` class, whose constructor is passed the dictionary used for configuration, and has a `configure()` method. The `logging.config` module has a callable attribute `dictConfigClass` which is initially set to `DictConfigurator`. You can replace the value of `dictConfigClass` with a suitable implementation of your own.

`dictConfig()` calls `dictConfigClass` passing the specified dictionary, and then calls the `configure()` method on the returned object to put the configuration into effect:

```
def dictConfig(config):
    dictConfigClass(config).configure()
```

For example, a subclass of `DictConfigurator` could call `DictConfigurator.__init__()` in its own `__init__()`, then set up custom prefixes which would be usable in the subsequent `configure()` call. `dictConfigClass` would be bound to this new subclass, and then `dictConfig()` could be called exactly as in the default, uncustomized state.

3.2 新版功能.

`logging.config.fileConfig(fname, defaults=None, disable_existing_loggers=True)`

Reads the logging configuration from a `configparser`-format file. The format of the file should be as described in [Configuration file format](#). This function can be called several times from an application, allowing an end user to select from various pre-canned configurations (if the developer provides a mechanism to present the choices and load the chosen configuration).

参数

- **fname** – A filename, or a file-like object, or an instance derived from `RawConfigParser`. If a `RawConfigParser`-derived instance is passed, it is used as is. Otherwise, a `Configparser` is instantiated, and the configuration read by it from the object passed in `fname`. If that has a `readline()` method, it is assumed to be a file-like object and read using `read_file()`; otherwise, it is assumed to be a filename and passed to `read()`.
- **defaults** – Defaults to be passed to the `ConfigParser` can be specified in this argument.
- **disable_existing_loggers** – If specified as `False`, loggers which exist when this call is made are left enabled. The default is `True` because this enables old behaviour in a backward-compatible way. This behaviour is to disable any existing loggers unless they or their ancestors are explicitly named in the logging configuration.

在 3.4 版更改: An instance of a subclass of `RawConfigParser` is now accepted as a value for `fname`. This facilitates:

- Use of a configuration file where logging configuration is just part of the overall application configuration.
- Use of a configuration read from a file, and then modified by the using application (e.g. based on command-line parameters or other aspects of the runtime environment) before being passed to `fileConfig`.

`logging.config.listen(port=DEFAULT_LOGGING_CONFIG_PORT, verify=None)`

Starts up a socket server on the specified port, and listens for new configurations. If no port is specified, the module's default `DEFAULT_LOGGING_CONFIG_PORT` is used. Logging configurations will be sent as a file suitable for processing by `dictConfig()` or `fileConfig()`. Returns a `Thread` instance on which you can call `start()` to start the server, and which you can `join()` when appropriate. To stop the server, call `stopListening()`.

The `verify` argument, if specified, should be a callable which should verify whether bytes received across the socket are valid and should be processed. This could be done by encrypting and/or signing what is sent across the socket, such that the `verify` callable can perform signature verification and/or decryption. The `verify` callable is called with a single argument - the bytes received across the socket - and should return the bytes to be processed, or `None` to indicate that the bytes should be discarded. The returned bytes could be the same as the passed in bytes (e.g. when only verification is done), or they could be completely different (perhaps if decryption were performed).

To send a configuration to the socket, read in the configuration file and send it to the socket as a sequence of bytes preceded by a four-byte length string packed in binary using `struct.pack('>L', n)`.

注解: Because portions of the configuration are passed through `eval()`, use of this function may open its users

to a security risk. While the function only binds to a socket on `localhost`, and so does not accept connections from remote machines, there are scenarios where untrusted code could be run under the account of the process which calls `listen()`. Specifically, if the process calling `listen()` runs on a multi-user machine where users cannot trust each other, then a malicious user could arrange to run essentially arbitrary code in a victim user's process, simply by connecting to the victim's `listen()` socket and sending a configuration which runs whatever code the attacker wants to have executed in the victim's process. This is especially easy to do if the default port is used, but not hard even if a different port is used). To avoid the risk of this happening, use the `verify` argument to `listen()` to prevent unrecognised configurations from being applied.

在 3.4 版更改: The `verify` argument was added.

注解: If you want to send configurations to the listener which don't disable existing loggers, you will need to use a JSON format for the configuration, which will use `dictConfig()` for configuration. This method allows you to specify `disable_existing_loggers` as `False` in the configuration you send.

`logging.config.stopListening()`

Stops the listening server which was created with a call to `listen()`. This is typically called before calling `join()` on the return value from `listen()`.

16.7.2 Configuration dictionary schema

Describing a logging configuration requires listing the various objects to create and the connections between them; for example, you may create a handler named 'console' and then say that the logger named 'startup' will send its messages to the 'console' handler. These objects aren't limited to those provided by the `logging` module because you might write your own formatter or handler class. The parameters to these classes may also need to include external objects such as `sys.stderr`. The syntax for describing these objects and connections is defined in *Object connections* below.

Dictionary Schema Details

The dictionary passed to `dictConfig()` must contain the following keys:

- *version* - to be set to an integer value representing the schema version. The only valid value at present is 1, but having this key allows the schema to evolve while still preserving backwards compatibility.

All other keys are optional, but if present they will be interpreted as described below. In all cases below where a 'configuring dict' is mentioned, it will be checked for the special '()' key to see if a custom instantiation is required. If so, the mechanism described in *User-defined objects* below is used to create an instance; otherwise, the context is used to determine what to instantiate.

- *formatters* - the corresponding value will be a dict in which each key is a formatter id and each value is a dict describing how to configure the corresponding *Formatter* instance.

The configuring dict is searched for keys `format` and `datefmt` (with defaults of `None`) and these are used to construct a *Formatter* instance.

- *filters* - the corresponding value will be a dict in which each key is a filter id and each value is a dict describing how to configure the corresponding *Filter* instance.

The configuring dict is searched for the key `name` (defaulting to the empty string) and this is used to construct a *logging.Filter* instance.

- *handlers* - the corresponding value will be a dict in which each key is a handler id and each value is a dict describing how to configure the corresponding *Handler* instance.

The configuring dict is searched for the following keys:

- `class` (mandatory). This is the fully qualified name of the handler class.
- `level` (optional). The level of the handler.
- `formatter` (optional). The id of the formatter for this handler.
- `filters` (optional). A list of ids of the filters for this handler.

All *other* keys are passed through as keyword arguments to the handler's constructor. For example, given the snippet:

```
handlers:
  console:
    class : logging.StreamHandler
    formatter: brief
    level  : INFO
    filters: [allow_foo]
    stream : ext://sys.stdout
  file:
    class : logging.handlers.RotatingFileHandler
    formatter: precise
    filename: logconfig.log
    maxBytes: 1024
    backupCount: 3
```

the handler with id `console` is instantiated as a `logging.StreamHandler`, using `sys.stdout` as the underlying stream. The handler with id `file` is instantiated as a `logging.handlers.RotatingFileHandler` with the keyword arguments `filename='logconfig.log'`, `maxBytes=1024`, `backupCount=3`.

- *loggers* - the corresponding value will be a dict in which each key is a logger name and each value is a dict describing how to configure the corresponding Logger instance.

The configuring dict is searched for the following keys:

- `level` (optional). The level of the logger.
- `propagate` (optional). The propagation setting of the logger.
- `filters` (optional). A list of ids of the filters for this logger.
- `handlers` (optional). A list of ids of the handlers for this logger.

The specified loggers will be configured according to the level, propagation, filters and handlers specified.

- *root* - this will be the configuration for the root logger. Processing of the configuration will be as for any logger, except that the `propagate` setting will not be applicable.
- *incremental* - whether the configuration is to be interpreted as incremental to the existing configuration. This value defaults to `False`, which means that the specified configuration replaces the existing configuration with the same semantics as used by the existing `fileConfig()` API.

If the specified value is `True`, the configuration is processed as described in the section on *Incremental Configuration*.

- *disable_existing_loggers* - whether any existing loggers are to be disabled. This setting mirrors the parameter of the same name in `fileConfig()`. If absent, this parameter defaults to `True`. This value is ignored if *incremental* is `True`.

Incremental Configuration

It is difficult to provide complete flexibility for incremental configuration. For example, because objects such as filters and formatters are anonymous, once a configuration is set up, it is not possible to refer to such anonymous objects when augmenting a configuration.

Furthermore, there is not a compelling case for arbitrarily altering the object graph of loggers, handlers, filters, formatters at run-time, once a configuration is set up; the verbosity of loggers and handlers can be controlled just by setting levels (and, in the case of loggers, propagation flags). Changing the object graph arbitrarily in a safe way is problematic in a multi-threaded environment; while not impossible, the benefits are not worth the complexity it adds to the implementation.

Thus, when the `incremental` key of a configuration dict is present and is `True`, the system will completely ignore any `formatters` and `filters` entries, and process only the `level` settings in the `handlers` entries, and the `level` and `propagate` settings in the `loggers` and `root` entries.

Using a value in the configuration dict lets configurations to be sent over the wire as pickled dicts to a socket listener. Thus, the logging verbosity of a long-running application can be altered over time with no need to stop and restart the application.

Object connections

The schema describes a set of logging objects - loggers, handlers, formatters, filters - which are connected to each other in an object graph. Thus, the schema needs to represent connections between the objects. For example, say that, once configured, a particular logger has attached to it a particular handler. For the purposes of this discussion, we can say that the logger represents the source, and the handler the destination, of a connection between the two. Of course in the configured objects this is represented by the logger holding a reference to the handler. In the configuration dict, this is done by giving each destination object an id which identifies it unambiguously, and then using the id in the source object's configuration to indicate that a connection exists between the source and the destination object with that id.

So, for example, consider the following YAML snippet:

```
formatters:
  brief:
    # configuration for formatter with id 'brief' goes here
  precise:
    # configuration for formatter with id 'precise' goes here
handlers:
  h1: #This is an id
    # configuration of handler with id 'h1' goes here
    formatter: brief
  h2: #This is another id
    # configuration of handler with id 'h2' goes here
    formatter: precise
loggers:
  foo.bar.baz:
    # other configuration for logger 'foo.bar.baz'
    handlers: [h1, h2]
```

(Note: YAML used here because it's a little more readable than the equivalent Python source form for the dictionary.)

The ids for loggers are the logger names which would be used programmatically to obtain a reference to those loggers, e.g. `foo.bar.baz`. The ids for Formatters and Filters can be any string value (such as `brief`, `precise` above) and they are transient, in that they are only meaningful for processing the configuration dictionary and used to determine connections between objects, and are not persisted anywhere when the configuration call is complete.

The above snippet indicates that logger named `foo.bar.baz` should have two handlers attached to it, which are described by the handler ids `h1` and `h2`. The formatter for `h1` is that described by id `brief`, and the formatter for `h2` is that described by id `precise`.

User-defined objects

The schema supports user-defined objects for handlers, filters and formatters. (Loggers do not need to have different types for different instances, so there is no support in this configuration schema for user-defined logger classes.)

Objects to be configured are described by dictionaries which detail their configuration. In some places, the logging system will be able to infer from the context how an object is to be instantiated, but when a user-defined object is to be instantiated, the system will not know how to do this. In order to provide complete flexibility for user-defined object instantiation, the user needs to provide a ‘factory’ - a callable which is called with a configuration dictionary and which returns the instantiated object. This is signalled by an absolute import path to the factory being made available under the special key ‘()’. Here’s a concrete example:

```
formatters:
  brief:
    format: '%(message)s'
  default:
    format: '%(asctime)s %(levelname)-8s %(name)-15s %(message)s'
    datefmt: '%Y-%m-%d %H:%M:%S'
  custom:
    (): my.package.customFormatterFactory
    bar: baz
    spam: 99.9
    answer: 42
```

The above YAML snippet defines three formatters. The first, with id `brief`, is a standard `logging.Formatter` instance with the specified format string. The second, with id `default`, has a longer format and also defines the time format explicitly, and will result in a `logging.Formatter` initialized with those two format strings. Shown in Python source form, the `brief` and `default` formatters have configuration sub-dictionaries:

```
{
  'format' : '%(message)s'
}
```

和:

```
{
  'format' : '%(asctime)s %(levelname)-8s %(name)-15s %(message)s',
  'datefmt' : '%Y-%m-%d %H:%M:%S'
}
```

respectively, and as these dictionaries do not contain the special key ‘()’, the instantiation is inferred from the context: as a result, standard `logging.Formatter` instances are created. The configuration sub-dictionary for the third formatter, with id `custom`, is:

```
{
  '()' : 'my.package.customFormatterFactory',
  'bar' : 'baz',
  'spam' : 99.9,
  'answer' : 42
}
```

and this contains the special key ‘()’, which means that user-defined instantiation is wanted. In this case, the specified factory callable will be used. If it is an actual callable it will be used directly - otherwise, if you specify a string (as in the example) the actual callable will be located using normal import mechanisms. The callable will be called with the **remaining** items in the configuration sub-dictionary as keyword arguments. In the above example, the formatter with id `custom` will be assumed to be returned by the call:

```
my.package.customFormatterFactory(bar='baz', spam=99.9, answer=42)
```

The key '()' has been used as the special key because it is not a valid keyword parameter name, and so will not clash with the names of the keyword arguments used in the call. The '()' also serves as a mnemonic that the corresponding value is a callable.

Access to external objects

There are times where a configuration needs to refer to objects external to the configuration, for example `sys.stderr`. If the configuration dict is constructed using Python code, this is straightforward, but a problem arises when the configuration is provided via a text file (e.g. JSON, YAML). In a text file, there is no standard way to distinguish `sys.stderr` from the literal string `'sys.stderr'`. To facilitate this distinction, the configuration system looks for certain special prefixes in string values and treat them specially. For example, if the literal string `'ext://sys.stderr'` is provided as a value in the configuration, then the `ext://` will be stripped off and the remainder of the value processed using normal import mechanisms.

The handling of such prefixes is done in a way analogous to protocol handling: there is a generic mechanism to look for prefixes which match the regular expression `^(?P<prefix>[a-z]+):/(?P<suffix>.*)$` whereby, if the `prefix` is recognised, the `suffix` is processed in a prefix-dependent manner and the result of the processing replaces the string value. If the prefix is not recognised, then the string value will be left as-is.

Access to internal objects

As well as external objects, there is sometimes also a need to refer to objects in the configuration. This will be done implicitly by the configuration system for things that it knows about. For example, the string value `'DEBUG'` for a level in a logger or handler will automatically be converted to the value `logging.DEBUG`, and the `handlers`, `filters` and `formatter` entries will take an object id and resolve to the appropriate destination object.

However, a more generic mechanism is needed for user-defined objects which are not known to the `logging` module. For example, consider `logging.handlers.MemoryHandler`, which takes a `target` argument which is another handler to delegate to. Since the system already knows about this class, then in the configuration, the given `target` just needs to be the object id of the relevant target handler, and the system will resolve to the handler from the id. If, however, a user defines a `my.package.MyHandler` which has an `alternate` handler, the configuration system would not know that the `alternate` referred to a handler. To cater for this, a generic resolution system allows the user to specify:

```
handlers:
  file:
    # configuration of file handler goes here

  custom:
    (): my.package.MyHandler
    alternate: cfg://handlers.file
```

The literal string `'cfg://handlers.file'` will be resolved in an analogous way to strings with the `ext://` prefix, but looking in the configuration itself rather than the import namespace. The mechanism allows access by dot or by index, in a similar way to that provided by `str.format`. Thus, given the following snippet:

```
handlers:
  email:
    class: logging.handlers.SMTPHandler
    mailhost: localhost
    fromaddr: my_app@domain.tld
    toaddrs:
      - support_team@domain.tld
```

(下页继续)

(续上页)

```
- dev_team@domain.tld
subject: Houston, we have a problem.
```

in the configuration, the string `'cfg://handlers'` would resolve to the dict with key `handlers`, the string `'cfg://handlers.email'` would resolve to the dict with key `email` in the `handlers` dict, and so on. The string `'cfg://handlers.email.toaddrs[1]'` would resolve to `'dev_team.domain.tld'` and the string `'cfg://handlers.email.toaddrs[0]'` would resolve to the value `'support_team@domain.tld'`. The subject value could be accessed using either `'cfg://handlers.email.subject'` or, equivalently, `'cfg://handlers.email[subject]'`. The latter form only needs to be used if the key contains spaces or non-alphanumeric characters. If an index value consists only of decimal digits, access will be attempted using the corresponding integer value, falling back to the string value if needed.

Given a string `cfg://handlers.myhandler.mykey.123`, this will resolve to `config_dict['handlers']['myhandler']['mykey']['123']`. If the string is specified as `cfg://handlers.myhandler.mykey[123]`, the system will attempt to retrieve the value from `config_dict['handlers']['myhandler']['mykey'][123]`, and fall back to `config_dict['handlers']['myhandler']['mykey']['123']` if that fails.

Import resolution and custom importers

Import resolution, by default, uses the builtin `__import__()` function to do its importing. You may want to replace this with your own importing mechanism: if so, you can replace the `importer` attribute of the `DictConfigurator` or its superclass, the `BaseConfigurator` class. However, you need to be careful because of the way functions are accessed from classes via descriptors. If you are using a Python callable to do your imports, and you want to define it at class level rather than instance level, you need to wrap it with `staticmethod()`. For example:

```
from importlib import import_module
from logging.config import BaseConfigurator

BaseConfigurator.importer = staticmethod(import_module)
```

You don't need to wrap with `staticmethod()` if you're setting the import callable on a configurator *instance*.

16.7.3 Configuration file format

The configuration file format understood by `fileConfig()` is based on `configparser` functionality. The file must contain sections called `[loggers]`, `[handlers]` and `[formatters]` which identify by name the entities of each type which are defined in the file. For each such entity, there is a separate section which identifies how that entity is configured. Thus, for a logger named `log01` in the `[loggers]` section, the relevant configuration details are held in a section `[logger_log01]`. Similarly, a handler called `hand01` in the `[handlers]` section will have its configuration held in a section called `[handler_hand01]`, while a formatter called `form01` in the `[formatters]` section will have its configuration specified in a section called `[formatter_form01]`. The root logger configuration must be specified in a section called `[logger_root]`.

注解: The `fileConfig()` API is older than the `dictConfig()` API and does not provide functionality to cover certain aspects of logging. For example, you cannot configure `Filter` objects, which provide for filtering of messages beyond simple integer levels, using `fileConfig()`. If you need to have instances of `Filter` in your logging configuration, you will need to use `dictConfig()`. Note that future enhancements to configuration functionality will be added to `dictConfig()`, so it's worth considering transitioning to this newer API when it's convenient to do so.

Examples of these sections in the file are given below.

```
[loggers]
keys=root,log02,log03,log04,log05,log06,log07

[handlers]
keys=hand01,hand02,hand03,hand04,hand05,hand06,hand07,hand08,hand09

[formatters]
keys=form01,form02,form03,form04,form05,form06,form07,form08,form09
```

The root logger must specify a level and a list of handlers. An example of a root logger section is given below.

```
[logger_root]
level=NOTSET
handlers=hand01
```

The `level` entry can be one of `DEBUG`, `INFO`, `WARNING`, `ERROR`, `CRITICAL` or `NOTSET`. For the root logger only, `NOTSET` means that all messages will be logged. Level values are *eval()*uated in the context of the logging package's namespace.

The `handlers` entry is a comma-separated list of handler names, which must appear in the `[handlers]` section. These names must appear in the `[handlers]` section and have corresponding sections in the configuration file.

For loggers other than the root logger, some additional information is required. This is illustrated by the following example.

```
[logger_parser]
level=DEBUG
handlers=hand01
propagate=1
qualname=compiler.parser
```

The `level` and `handlers` entries are interpreted as for the root logger, except that if a non-root logger's level is specified as `NOTSET`, the system consults loggers higher up the hierarchy to determine the effective level of the logger. The `propagate` entry is set to 1 to indicate that messages must propagate to handlers higher up the logger hierarchy from this logger, or 0 to indicate that messages are **not** propagated to handlers up the hierarchy. The `qualname` entry is the hierarchical channel name of the logger, that is to say the name used by the application to get the logger.

Sections which specify handler configuration are exemplified by the following.

```
[handler_hand01]
class=StreamHandler
level=NOTSET
formatter=form01
args=(sys.stdout,)
```

The `class` entry indicates the handler's class (as determined by *eval()* in the logging package's namespace). The `level` is interpreted as for loggers, and `NOTSET` is taken to mean 'log everything'.

The `formatter` entry indicates the key name of the formatter for this handler. If blank, a default formatter (`logging._defaultFormatter`) is used. If a name is specified, it must appear in the `[formatters]` section and have a corresponding section in the configuration file.

The `args` entry, when *eval()*uated in the context of the logging package's namespace, is the list of arguments to the constructor for the handler class. Refer to the constructors for the relevant handlers, or to the examples below, to see how typical entries are constructed.

```
[handler_hand02]
class=FileHandler
level=DEBUG
```

(下页继续)

(续上页)

```

formatter=form02
args=('python.log', 'w')

[handler_hand03]
class=handlers.SocketHandler
level=INFO
formatter=form03
args=('localhost', handlers.DEFAULT_TCP_LOGGING_PORT)

[handler_hand04]
class=handlers.DatagramHandler
level=WARN
formatter=form04
args=('localhost', handlers.DEFAULT_UDP_LOGGING_PORT)

[handler_hand05]
class=handlers.SysLogHandler
level=ERROR
formatter=form05
args= (('localhost', handlers.SYSLOG_UDP_PORT), handlers.SysLogHandler.LOG_USER)

[handler_hand06]
class=handlers.NTEventLogHandler
level=CRITICAL
formatter=form06
args=('Python Application', '', 'Application')

[handler_hand07]
class=handlers.SMTPHandler
level=WARN
formatter=form07
args=('localhost', 'from@abc', ['user1@abc', 'user2@xyz'], 'Logger Subject')

[handler_hand08]
class=handlers.MemoryHandler
level=NOTSET
formatter=form08
target=
args=(10, ERROR)

[handler_hand09]
class=handlers.HTTPHandler
level=NOTSET
formatter=form09
args=('localhost:9022', '/log', 'GET')

```

Sections which specify formatter configuration are typified by the following.

```

[formatter_form01]
format=F1 %(asctime)s %(levelname)s %(message)s
datefmt=
class=logging.Formatter

```

The `format` entry is the overall format string, and the `datefmt` entry is the `strftime()`-compatible date/time format string. If empty, the package substitutes something which is almost equivalent to specifying the date format string `'%Y-%m-%d %H:%M:%S'`. This format also specifies milliseconds, which are appended to the result of using the above format string, with a comma separator. An example time in this format is `2003-01-23 00:29:50,411`.

The `class` entry is optional. It indicates the name of the formatter's class (as a dotted module and class name.) This option is useful for instantiating a *Formatter* subclass. Subclasses of *Formatter* can present exception tracebacks in an expanded or condensed format.

注解： Due to the use of `eval()` as described above, there are potential security risks which result from using the `listen()` to send and receive configurations via sockets. The risks are limited to where multiple users with no mutual trust run code on the same machine; see the `listen()` documentation for more information.

参见：

模块 *logging* 日志记录模块的 API 参考。

模块 *logging.handlers* 日志记录模块附带的有用处理程序。

16.8 logging.handlers —Logging handlers

源代码: [Lib/logging/handlers.py](#)

Important

此页面仅包含参考信息。有关教程，请参阅

- Basic Tutorial
- Advanced Tutorial
- Logging Cookbook

The following useful handlers are provided in the package. Note that three of the handlers (*StreamHandler*, *FileHandler* and *NullHandler*) are actually defined in the *logging* module itself, but have been documented here along with the other handlers.

16.8.1 StreamHandler

The *StreamHandler* class, located in the core *logging* package, sends logging output to streams such as `sys.stdout`, `sys.stderr` or any file-like object (or, more precisely, any object which supports `write()` and `flush()` methods).

class `logging.StreamHandler` (*stream=None*)

Returns a new instance of the *StreamHandler* class. If *stream* is specified, the instance will use it for logging output; otherwise, `sys.stderr` will be used.

emit (*record*)

If a formatter is specified, it is used to format the record. The record is then written to the stream with a terminator. If exception information is present, it is formatted using `traceback.print_exception()` and appended to the stream.

flush ()

Flushes the stream by calling its `flush()` method. Note that the `close()` method is inherited from *Handler* and so does no output, so an explicit `flush()` call may be needed at times.

在 3.2 版更改: The `StreamHandler` class now has a `terminator` attribute, default value `'\n'`, which is used as the terminator when writing a formatted record to a stream. If you don't want this newline termination, you can set the handler instance's `terminator` attribute to the empty string. In earlier versions, the terminator was hardcoded as `'\n'`.

16.8.2 FileHandler

The `FileHandler` class, located in the core `logging` package, sends logging output to a disk file. It inherits the output functionality from `StreamHandler`.

class `logging.FileHandler` (*filename*, *mode*='a', *encoding*=None, *delay*=False)

Returns a new instance of the `FileHandler` class. The specified file is opened and used as the stream for logging. If *mode* is not specified, 'a' is used. If *encoding* is not None, it is used to open the file with that encoding. If *delay* is true, then file opening is deferred until the first call to `emit()`. By default, the file grows indefinitely.

在 3.6 版更改: As well as string values, `Path` objects are also accepted for the *filename* argument.

close()

Closes the file.

emit (*record*)

Outputs the record to the file.

16.8.3 NullHandler

3.1 新版功能.

The `NullHandler` class, located in the core `logging` package, does not do any formatting or output. It is essentially a 'no-op' handler for use by library developers.

class `logging.NullHandler`

Returns a new instance of the `NullHandler` class.

emit (*record*)

This method does nothing.

handle (*record*)

This method does nothing.

createLock ()

This method returns None for the lock, since there is no underlying I/O to which access needs to be serialized.

See library-config for more information on how to use `NullHandler`.

16.8.4 WatchedFileHandler

The `WatchedFileHandler` class, located in the `logging.handlers` module, is a `FileHandler` which watches the file it is logging to. If the file changes, it is closed and reopened using the file name.

A file change can happen because of usage of programs such as `newsyslog` and `logrotate` which perform log file rotation. This handler, intended for use under Unix/Linux, watches the file to see if it has changed since the last emit. (A file is deemed to have changed if its device or inode have changed.) If the file has changed, the old file stream is closed, and the file opened to get a new stream.

This handler is not appropriate for use under Windows, because under Windows open log files cannot be moved or renamed - logging opens the files with exclusive locks - and so there is no need for such a handler. Furthermore, `ST_INO` is not supported under Windows; `stat()` always returns zero for this value.

class `logging.handlers.WatchedFileHandler` (*filename*, *mode*='a', *encoding*=None, *delay*=False)

Returns a new instance of the `WatchedFileHandler` class. The specified file is opened and used as the stream for logging. If *mode* is not specified, 'a' is used. If *encoding* is not None, it is used to open the file with that encoding. If *delay* is true, then file opening is deferred until the first call to `emit()`. By default, the file grows indefinitely.

在 3.6 版更改: As well as string values, `Path` objects are also accepted for the *filename* argument.

reopenIfNeeded()

Checks to see if the file has changed. If it has, the existing stream is flushed and closed and the file opened again, typically as a precursor to outputting the record to the file.

3.6 新版功能.

emit (*record*)

Outputs the record to the file, but first calls `reopenIfNeeded()` to reopen the file if it has changed.

16.8.5 BaseRotatingHandler

The `BaseRotatingHandler` class, located in the `logging.handlers` module, is the base class for the rotating file handlers, `RotatingFileHandler` and `TimedRotatingFileHandler`. You should not need to instantiate this class, but it has attributes and methods you may need to override.

class `logging.handlers.BaseRotatingHandler` (*filename*, *mode*, *encoding*=None, *delay*=False)

The parameters are as for `FileHandler`. The attributes are:

namer

If this attribute is set to a callable, the `rotation_filename()` method delegates to this callable. The parameters passed to the callable are those passed to `rotation_filename()`.

注解: The namer function is called quite a few times during rollover, so it should be as simple and as fast as possible. It should also return the same output every time for a given input, otherwise the rollover behaviour may not work as expected.

3.3 新版功能.

rotator

If this attribute is set to a callable, the `rotate()` method delegates to this callable. The parameters passed to the callable are those passed to `rotate()`.

3.3 新版功能.

rotation_filename (*default_name*)

Modify the filename of a log file when rotating.

This is provided so that a custom filename can be provided.

The default implementation calls the 'namer' attribute of the handler, if it's callable, passing the default name to it. If the attribute isn't callable (the default is None), the name is returned unchanged.

参数 default_name – The default name for the log file.

3.3 新版功能.

rotate (*source*, *dest*)

When rotating, rotate the current log.

The default implementation calls the ‘rotator’ attribute of the handler, if it’s callable, passing the source and dest arguments to it. If the attribute isn’t callable (the default is `None`), the source is simply renamed to the destination.

参数

- **source** – The source filename. This is normally the base filename, e.g. ‘test.log’.
- **dest** – The destination filename. This is normally what the source is rotated to, e.g. ‘test.log.1’.

3.3 新版功能.

The reason the attributes exist is to save you having to subclass - you can use the same callables for instances of `RotatingFileHandler` and `TimedRotatingFileHandler`. If either the namer or rotator callable raises an exception, this will be handled in the same way as any other exception during an `emit()` call, i.e. via the `handleError()` method of the handler.

If you need to make more significant changes to rotation processing, you can override the methods.

For an example, see `cookbook-rotator-namer`.

16.8.6 RotatingFileHandler

The `RotatingFileHandler` class, located in the `logging.handlers` module, supports rotation of disk log files.

class `logging.handlers.RotatingFileHandler` (*filename*, *mode*='a', *maxBytes*=0, *backupCount*=0, *encoding*=None, *delay*=False)

Returns a new instance of the `RotatingFileHandler` class. The specified file is opened and used as the stream for logging. If *mode* is not specified, 'a' is used. If *encoding* is not None, it is used to open the file with that encoding. If *delay* is true, then file opening is deferred until the first call to `emit()`. By default, the file grows indefinitely.

You can use the *maxBytes* and *backupCount* values to allow the file to *rollover* at a predetermined size. When the size is about to be exceeded, the file is closed and a new file is silently opened for output. Rollover occurs whenever the current log file is nearly *maxBytes* in length; but if either of *maxBytes* or *backupCount* is zero, rollover never occurs, so you generally want to set *backupCount* to at least 1, and have a non-zero *maxBytes*. When *backupCount* is non-zero, the system will save old log files by appending the extensions ‘.1’, ‘.2’ etc., to the filename. For example, with a *backupCount* of 5 and a base file name of `app.log`, you would get `app.log`, `app.log.1`, `app.log.2`, up to `app.log.5`. The file being written to is always `app.log`. When this file is filled, it is closed and renamed to `app.log.1`, and if files `app.log.1`, `app.log.2`, etc. exist, then they are renamed to `app.log.2`, `app.log.3` etc. respectively.

在 3.6 版更改: As well as string values, `Path` objects are also accepted for the *filename* argument.

doRollover()

Does a rollover, as described above.

emit (*record*)

Outputs the record to the file, catering for rollover as described previously.

16.8.7 TimedRotatingFileHandler

The `TimedRotatingFileHandler` class, located in the `logging.handlers` module, supports rotation of disk log files at certain timed intervals.

class `logging.handlers.TimedRotatingFileHandler` (*filename*, *when*='h', *interval*=1, *backupCount*=0, *encoding*=None, *delay*=False, *utc*=False, *atTime*=None)

Returns a new instance of the `TimedRotatingFileHandler` class. The specified file is opened and used as the stream for logging. On rotating it also sets the filename suffix. Rotating happens based on the product of *when* and *interval*.

You can use the *when* to specify the type of *interval*. The list of possible values is below. Note that they are not case sensitive.

值	间隔类型	如果/如何使用 <i>atTime</i>
'S'	秒	忽略
'M'	分钟	忽略
'H'	小时	忽略
'D'	天	忽略
'W0'-'W6'	工作日 (0= 星期一)	Used to compute initial rollover time
'midnight'	Roll over at midnight, if <i>atTime</i> not specified, else at time <i>atTime</i>	Used to compute initial rollover time

When using weekday-based rotation, specify 'W0' for Monday, 'W1' for Tuesday, and so on up to 'W6' for Sunday. In this case, the value passed for *interval* isn't used.

The system will save old log files by appending extensions to the filename. The extensions are date-and-time based, using the strftime format `%Y-%m-%d_%H-%M-%S` or a leading portion thereof, depending on the rollover interval.

When computing the next rollover time for the first time (when the handler is created), the last modification time of an existing log file, or else the current time, is used to compute when the next rotation will occur.

If the *utc* argument is true, times in UTC will be used; otherwise local time is used.

If *backupCount* is nonzero, at most *backupCount* files will be kept, and if more would be created when rollover occurs, the oldest one is deleted. The deletion logic uses the interval to determine which files to delete, so changing the interval may leave old files lying around.

If *delay* is true, then file opening is deferred until the first call to `emit()`.

If *atTime* is not None, it must be a `datetime.time` instance which specifies the time of day when rollover occurs, for the cases where rollover is set to happen "at midnight" or "on a particular weekday". Note that in these cases, the *atTime* value is effectively used to compute the *initial* rollover, and subsequent rollovers would be calculated via the normal interval calculation.

注解: Calculation of the initial rollover time is done when the handler is initialised. Calculation of subsequent rollover times is done only when rollover occurs, and rollover occurs only when emitting output. If this is not kept in mind, it might lead to some confusion. For example, if an interval of "every minute" is set, that does not mean you will always see log files with times (in the filename) separated by a minute; if, during application execution, logging output is generated more frequently than once a minute, *then* you can expect to see log files with times separated by a minute. If, on the other hand, logging messages are only output once every five minutes (say), then there will be gaps in the file times corresponding to the minutes where no output (and hence no rollover) occurred.

在 3.4 版更改: *atTime* parameter was added.

在 3.6 版更改: As well as string values, `Path` objects are also accepted for the *filename* argument.

doRollover()

Does a rollover, as described above.

emit (*record*)

Outputs the record to the file, catering for rollover as described above.

16.8.8 SocketHandler

The *SocketHandler* class, located in the *logging.handlers* module, sends logging output to a network socket. The base class uses a TCP socket.

class *logging.handlers.SocketHandler* (*host*, *port*)

Returns a new instance of the *SocketHandler* class intended to communicate with a remote machine whose address is given by *host* and *port*.

在 3.4 版更改: If *port* is specified as *None*, a Unix domain socket is created using the value in *host* - otherwise, a TCP socket is created.

close ()

Closes the socket.

emit ()

Pickles the record's attribute dictionary and writes it to the socket in binary format. If there is an error with the socket, silently drops the packet. If the connection was previously lost, re-establishes the connection. To unpickle the record at the receiving end into a *LogRecord*, use the *makeLogRecord()* function.

handleError ()

Handles an error which has occurred during *emit()*. The most likely cause is a lost connection. Closes the socket so that we can retry on the next event.

makeSocket ()

This is a factory method which allows subclasses to define the precise type of socket they want. The default implementation creates a TCP socket (*socket.SOCK_STREAM*).

makePickle (*record*)

Pickles the record's attribute dictionary in binary format with a length prefix, and returns it ready for transmission across the socket.

Note that pickles aren't completely secure. If you are concerned about security, you may want to override this method to implement a more secure mechanism. For example, you can sign pickles using HMAC and then verify them on the receiving end, or alternatively you can disable unpickling of global objects on the receiving end.

send (*packet*)

Send a pickled string *packet* to the socket. This function allows for partial sends which can happen when the network is busy.

createSocket ()

Tries to create a socket; on failure, uses an exponential back-off algorithm. On initial failure, the handler will drop the message it was trying to send. When subsequent messages are handled by the same instance, it will not try connecting until some time has passed. The default parameters are such that the initial delay is one second, and if after that delay the connection still can't be made, the handler will double the delay each time up to a maximum of 30 seconds.

This behaviour is controlled by the following handler attributes:

- *retryStart* (initial delay, defaulting to 1.0 seconds).
- *retryFactor* (multiplier, defaulting to 2.0).
- *retryMax* (maximum delay, defaulting to 30.0 seconds).

This means that if the remote listener starts up *after* the handler has been used, you could lose messages (since the handler won't even attempt a connection until the delay has elapsed, but just silently drop messages during the delay period).

16.8.9 DatagramHandler

The `DatagramHandler` class, located in the `logging.handlers` module, inherits from `SocketHandler` to support sending logging messages over UDP sockets.

class `logging.handlers.DatagramHandler` (*host*, *port*)

Returns a new instance of the `DatagramHandler` class intended to communicate with a remote machine whose address is given by *host* and *port*.

在 3.4 版更改: If *port* is specified as `None`, a Unix domain socket is created using the value in *host* - otherwise, a UDP socket is created.

emit ()

Pickles the record's attribute dictionary and writes it to the socket in binary format. If there is an error with the socket, silently drops the packet. To unpickle the record at the receiving end into a `LogRecord`, use the `makeLogRecord()` function.

makeSocket ()

The factory method of `SocketHandler` is here overridden to create a UDP socket (`socket.SOCK_DGRAM`).

send (*s*)

Send a pickled string to a socket.

16.8.10 SysLogHandler

The `SysLogHandler` class, located in the `logging.handlers` module, supports sending logging messages to a remote or local Unix syslog.

class `logging.handlers.SysLogHandler` (*address*=(`'localhost'`, `SYSLOG_UDP_PORT`), *facility*=`LOG_USER`, *socktype*=`socket.SOCK_DGRAM`)

Returns a new instance of the `SysLogHandler` class intended to communicate with a remote Unix machine whose address is given by *address* in the form of a (*host*, *port*) tuple. If *address* is not specified, (`'localhost'`, `514`) is used. The address is used to open a socket. An alternative to providing a (*host*, *port*) tuple is providing an address as a string, for example `"/dev/log"`. In this case, a Unix domain socket is used to send the message to the syslog. If *facility* is not specified, `LOG_USER` is used. The type of socket opened depends on the *socktype* argument, which defaults to `socket.SOCK_DGRAM` and thus opens a UDP socket. To open a TCP socket (for use with the newer syslog daemons such as `rsyslog`), specify a value of `socket.SOCK_STREAM`.

Note that if your server is not listening on UDP port 514, `SysLogHandler` may appear not to work. In that case, check what address you should be using for a domain socket - it's system dependent. For example, on Linux it's usually `"/dev/log"` but on OS/X it's `"/var/run/syslog"`. You'll need to check your platform and use the appropriate address (you may need to do this check at runtime if your application needs to run on several platforms). On Windows, you pretty much have to use the UDP option.

在 3.2 版更改: *socktype* was added.

close ()

Closes the socket to the remote host.

emit (*record*)

The record is formatted, and then sent to the syslog server. If exception information is present, it is *not* sent to the server.

在 3.2.1 版更改: (See: [bpo-12168](#).) In earlier versions, the message sent to the syslog daemons was always terminated with a NUL byte, because early versions of these daemons expected a NUL terminated message - even though it's not in the relevant specification ([RFC 5424](#)). More recent versions of these daemons don't expect the NUL byte but strip it off if it's there, and even more recent daemons (which adhere more closely to RFC 5424) pass the NUL byte on as part of the message.

To enable easier handling of syslog messages in the face of all these differing daemon behaviours, the appending of the NUL byte has been made configurable, through the use of a class-level attribute, `append_nul`. This defaults to `True` (preserving the existing behaviour) but can be set to `False` on a `SysLogHandler` instance in order for that instance to *not* append the NUL terminator.

在 3.3 版更改: (See: [bpo-12419](#).) In earlier versions, there was no facility for an “ident” or “tag” prefix to identify the source of the message. This can now be specified using a class-level attribute, defaulting to `""` to preserve existing behaviour, but which can be overridden on a `SysLogHandler` instance in order for that instance to prepend the ident to every message handled. Note that the provided ident must be text, not bytes, and is prepended to the message exactly as is.

`encodePriority` (*facility, priority*)

Encodes the facility and priority into an integer. You can pass in strings or integers - if strings are passed, internal mapping dictionaries are used to convert them to integers.

The symbolic `LOG_` values are defined in `SysLogHandler` and mirror the values defined in the `sys/syslog.h` header file.

优先级

名称 (字符串)	符号值
alert	LOG_ALERT
crit 或 critical	LOG_CRIT
debug	LOG_DEBUG
emerg 或 panic	LOG_EMERG
err 或 error	LOG_ERR
info	LOG_INFO
notice	LOG_NOTICE
warn 或 warning	LOG_WARNING

设备

名称 (字符串)	符号值
auth	LOG_AUTH
authpriv	LOG_AUTHPRIV
cron	LOG_CRON
daemon	LOG_DAEMON
ftp	LOG_FTP
kern	LOG_KERN
lpr	LOG_LPR
mail	LOG_MAIL
news	LOG_NEWS
syslog	LOG_SYSLOG
user	LOG_USER
uucp	LOG_UUCP
local0	LOG_LOCAL0
local1	LOG_LOCAL1
local2	LOG_LOCAL2
local3	LOG_LOCAL3
local4	LOG_LOCAL4
local5	LOG_LOCAL5
local6	LOG_LOCAL6
local7	LOG_LOCAL7

mapPriority (*levelname*)

Maps a logging level name to a syslog priority name. You may need to override this if you are using custom levels, or if the default algorithm is not suitable for your needs. The default algorithm maps DEBUG, INFO, WARNING, ERROR and CRITICAL to the equivalent syslog names, and all other level names to ‘warning’.

16.8.11 NTEventLogHandler

The *NTEventLogHandler* class, located in the *logging.handlers* module, supports sending logging messages to a local Windows NT, Windows 2000 or Windows XP event log. Before you can use it, you need Mark Hammond’s Win32 extensions for Python installed.

class *logging.handlers.NTEventLogHandler* (*appname*, *dllname=None*, *logtype='Application'*)

Returns a new instance of the *NTEventLogHandler* class. The *appname* is used to define the application name as it appears in the event log. An appropriate registry entry is created using this name. The *dllname* should give the fully qualified pathname of a .dll or .exe which contains message definitions to hold in the log (if not specified, 'win32service.pyd' is used - this is installed with the Win32 extensions and contains some basic placeholder message definitions. Note that use of these placeholders will make your event logs big, as the entire message source is held in the log. If you want slimmer logs, you have to pass in the name of your own .dll or .exe which contains the message definitions you want to use in the event log). The *logtype* is one of 'Application', 'System' or 'Security', and defaults to 'Application'.

close ()

At this point, you can remove the application name from the registry as a source of event log entries. However, if you do this, you will not be able to see the events as you intended in the Event Log Viewer - it needs to be able to access the registry to get the .dll name. The current version does not do this.

emit (*record*)

Determines the message ID, event category and event type, and then logs the message in the NT event log.

getEventCategory (*record*)

Returns the event category for the record. Override this if you want to specify your own categories. This version returns 0.

getEventType (*record*)

Returns the event type for the record. Override this if you want to specify your own types. This version does a mapping using the handler's `typemap` attribute, which is set up in `__init__()` to a dictionary which contains mappings for `DEBUG`, `INFO`, `WARNING`, `ERROR` and `CRITICAL`. If you are using your own levels, you will either need to override this method or place a suitable dictionary in the handler's `typemap` attribute.

getMessageID (*record*)

Returns the message ID for the record. If you are using your own messages, you could do this by having the *msg* passed to the logger being an ID rather than a format string. Then, in here, you could use a dictionary lookup to get the message ID. This version returns 1, which is the base message ID in `win32service.pyd`.

16.8.12 SMTPHandler

The `SMTPHandler` class, located in the `logging.handlers` module, supports sending logging messages to an email address via SMTP.

class `logging.handlers.SMTPHandler` (*mailhost*, *fromaddr*, *toaddrs*, *subject*, *credentials=None*, *secure=None*, *timeout=1.0*)

Returns a new instance of the `SMTPHandler` class. The instance is initialized with the from and to addresses and subject line of the email. The *toaddrs* should be a list of strings. To specify a non-standard SMTP port, use the (host, port) tuple format for the *mailhost* argument. If you use a string, the standard SMTP port is used. If your SMTP server requires authentication, you can specify a (username, password) tuple for the *credentials* argument.

To specify the use of a secure protocol (TLS), pass in a tuple to the *secure* argument. This will only be used when authentication credentials are supplied. The tuple should be either an empty tuple, or a single-value tuple with the name of a keyfile, or a 2-value tuple with the names of the keyfile and certificate file. (This tuple is passed to the `smtplib.SMTP.starttls()` method.)

A timeout can be specified for communication with the SMTP server using the *timeout* argument.

3.3 新版功能: The *timeout* argument was added.

emit (*record*)

Formats the record and sends it to the specified addressees.

getSubject (*record*)

If you want to specify a subject line which is record-dependent, override this method.

16.8.13 MemoryHandler

The `MemoryHandler` class, located in the `logging.handlers` module, supports buffering of logging records in memory, periodically flushing them to a *target* handler. Flushing occurs whenever the buffer is full, or when an event of a certain severity or greater is seen.

`MemoryHandler` is a subclass of the more general `BufferingHandler`, which is an abstract class. This buffers logging records in memory. Whenever each record is added to the buffer, a check is made by calling `shouldFlush()` to see if the buffer should be flushed. If it should, then `flush()` is expected to do the flushing.

class `logging.handlers.BufferingHandler` (*capacity*)

Initializes the handler with a buffer of the specified capacity.

emit (*record*)

Appends the record to the buffer. If `shouldFlush()` returns true, calls `flush()` to process the buffer.

flush()

You can override this to implement custom flushing behavior. This version just zaps the buffer to empty.

shouldFlush(record)

Returns true if the buffer is up to capacity. This method can be overridden to implement custom flushing strategies.

class logging.handlers.**MemoryHandler**(*capacity*, *flushLevel*=ERROR, *target*=None, *flushOnClose*=True)

Returns a new instance of the *MemoryHandler* class. The instance is initialized with a buffer size of *capacity*. If *flushLevel* is not specified, ERROR is used. If no *target* is specified, the target will need to be set using *setTarget()* before this handler does anything useful. If *flushOnClose* is specified as False, then the buffer is *not* flushed when the handler is closed. If not specified or specified as True, the previous behaviour of flushing the buffer will occur when the handler is closed.

在 3.6 版更改: The *flushOnClose* parameter was added.

close()

Calls *flush()*, sets the target to None and clears the buffer.

flush()

For a *MemoryHandler*, flushing means just sending the buffered records to the target, if there is one. The buffer is also cleared when this happens. Override if you want different behavior.

setTarget(target)

Sets the target handler for this handler.

shouldFlush(record)

Checks for buffer full or a record at the *flushLevel* or higher.

16.8.14 HTTPHandler

The *HTTPHandler* class, located in the *logging.handlers* module, supports sending logging messages to a Web server, using either GET or POST semantics.

class logging.handlers.**HTTPHandler**(*host*, *url*, *method*='GET', *secure*=False, *credentials*=None, *context*=None)

Returns a new instance of the *HTTPHandler* class. The *host* can be of the form *host:port*, should you need to use a specific port number. If no *method* is specified, GET is used. If *secure* is true, a HTTPS connection will be used. The *context* parameter may be set to a *ssl.SSLContext* instance to configure the SSL settings used for the HTTPS connection. If *credentials* is specified, it should be a 2-tuple consisting of userid and password, which will be placed in a HTTP 'Authorization' header using Basic authentication. If you specify credentials, you should also specify *secure*=True so that your userid and password are not passed in cleartext across the wire.

在 3.5 版更改: The *context* parameter was added.

mapLogRecord(record)

Provides a dictionary, based on *record*, which is to be URL-encoded and sent to the web server. The default implementation just returns *record.__dict__*. This method can be overridden if e.g. only a subset of *LogRecord* is to be sent to the web server, or if more specific customization of what's sent to the server is required.

emit(record)

Sends the record to the Web server as a URL-encoded dictionary. The *mapLogRecord()* method is used to convert the record to the dictionary to be sent.

注解: Since preparing a record for sending it to a Web server is not the same as a generic formatting operation, using *setFormatter()* to specify a *Formatter* for a *HTTPHandler* has no effect. Instead of calling

`format()`, this handler calls `mapLogRecord()` and then `urllib.parse.urlencode()` to encode the dictionary in a form suitable for sending to a Web server.

16.8.15 QueueHandler

3.2 新版功能.

The `QueueHandler` class, located in the `logging.handlers` module, supports sending logging messages to a queue, such as those implemented in the `queue` or `multiprocessing` modules.

Along with the `QueueListener` class, `QueueHandler` can be used to let handlers do their work on a separate thread from the one which does the logging. This is important in Web applications and also other service applications where threads servicing clients need to respond as quickly as possible, while any potentially slow operations (such as sending an email via `SMTPHandler`) are done on a separate thread.

class `logging.handlers.QueueHandler` (*queue*)

Returns a new instance of the `QueueHandler` class. The instance is initialized with the queue to send messages to. The queue can be any queue-like object; it's used as-is by the `enqueue()` method, which needs to know how to send messages to it.

emit (*record*)

Enqueues the result of preparing the LogRecord.

prepare (*record*)

Prepares a record for queuing. The object returned by this method is enqueued.

The base implementation formats the record to merge the message and arguments, and removes unpickleable items from the record in-place.

You might want to override this method if you want to convert the record to a dict or JSON string, or send a modified copy of the record while leaving the original intact.

enqueue (*record*)

Enqueues the record on the queue using `put_nowait()`; you may want to override this if you want to use blocking behaviour, or a timeout, or a customized queue implementation.

16.8.16 QueueListener

3.2 新版功能.

The `QueueListener` class, located in the `logging.handlers` module, supports receiving logging messages from a queue, such as those implemented in the `queue` or `multiprocessing` modules. The messages are received from a queue in an internal thread and passed, on the same thread, to one or more handlers for processing. While `QueueListener` is not itself a handler, it is documented here because it works hand-in-hand with `QueueHandler`.

Along with the `QueueHandler` class, `QueueListener` can be used to let handlers do their work on a separate thread from the one which does the logging. This is important in Web applications and also other service applications where threads servicing clients need to respond as quickly as possible, while any potentially slow operations (such as sending an email via `SMTPHandler`) are done on a separate thread.

class `logging.handlers.QueueListener` (*queue*, **handlers*, *respect_handler_level=False*)

Returns a new instance of the `QueueListener` class. The instance is initialized with the queue to send messages to and a list of handlers which will handle entries placed on the queue. The queue can be any queue-like object; it's passed as-is to the `dequeue()` method, which needs to know how to get messages from it. If `respect_handler_level` is True, a handler's level is respected (compared with the level for the message) when deciding whether to pass messages to that handler; otherwise, the behaviour is as in previous Python versions - to always pass each message to each handler.

在 3.5 版更改: The `respect_handler_levels` argument was added.

dequeue (*block*)

Dequeues a record and return it, optionally blocking.

The base implementation uses `get()`. You may want to override this method if you want to use timeouts or work with custom queue implementations.

prepare (*record*)

Prepare a record for handling.

This implementation just returns the passed-in record. You may want to override this method if you need to do any custom marshalling or manipulation of the record before passing it to the handlers.

handle (*record*)

Handle a record.

This just loops through the handlers offering them the record to handle. The actual object passed to the handlers is that which is returned from `prepare()`.

start ()

Starts the listener.

This starts up a background thread to monitor the queue for LogRecords to process.

stop ()

Stops the listener.

This asks the thread to terminate, and then waits for it to do so. Note that if you don't call this before your application exits, there may be some records still left on the queue, which won't be processed.

enqueue_sentinel ()

Writes a sentinel to the queue to tell the listener to quit. This implementation uses `put_nowait()`. You may want to override this method if you want to use timeouts or work with custom queue implementations.

3.3 新版功能.

参见:

模块 `logging` 日志记录模块的 API 参考。

模块 `logging.config` 日志记录模块的配置 API。

16.9 getpass — 便携式密码输入工具

源代码: [Lib/getpass.py](#)

`getpass` 模块提供了两个函数:

`getpass.getpass` (*prompt='Password: ', stream=None*)

提示用户输入一个密码且不会回显。用户会看到字符串 *prompt* 作为提示, 其默认值为 'Password: '。在 Unix 上, 如有必要提示会使用替换错误句柄写入到文件类对象 *stream*。*stream* 默认指向控制终端 (/dev/tty), 如果不可用则指向 `sys.stderr` (此参数在 Windows 上会被忽略)。

如果回显自由输入不可用则 `getpass()` 将回退为打印一条警告消息到 *stream* 并且从 `sys.stdin` 读取同时发出 `GetPassWarning`。

注解：如果你从 IDLE 内部调用 `getpass`，输入可能是在你启动 IDLE 的终端中而非在 IDLE 窗口本身中完成。

exception `getpass.GetPassWarning`

一个当密码输入可能被回显时发出的 `UserWarning` 子类。

`getpass.getuser()`

返回用户的“登录名称”。

此函数会按顺序检查环境变量 `LOGNAME`, `USER`, `LNAME` 和 `USERNAME`，并返回其中第一个被设置为非空字符串的值。如果均未设置，则在支持 `pwd` 模块的系统上将返回来自密码数据库的登录名，否则将引发一个异常。

通常情况下，此函数应优先于 `os.getlogin()` 使用。

16.10 curses — 终端字符单元显示的处理

`curses` 模块提供了 `curses` 库的接口，这是可移植高级终端处理的事实标准。

虽然 `curses` 在 Unix 环境中使用最为广泛，但也有适用于 Windows, DOS 以及其他可能的系统的版本。此扩展模块旨在匹配 `ncurses` 的 API，这是一个部署在 Linux 和 Unix 的 BSD 变体上的开源 `curses` 库。

注解：每当文档提到 **字符**时，它可以被指定为一个整数，一个单字符 Unicode 字符串或者一个单字节的字节字符串。

每当此文档提到 **字符串**时，它可以被指定为一个 Unicode 字符串或者一个字节字符串。

注解：从 5.4 版本开始，`ncurses` 库使用 `nl_langinfo` 函数来决定如何解释非 ASCII 数据。这意味着你需要在程序中调用 `locale.setlocale()` 函数，并使用一种系统中可用的编码方法来编码 Unicode 字符串。这个例子使用了系统默认的编码：

```
import locale
locale.setlocale(locale.LC_ALL, '')
code = locale.getpreferredencoding()
```

然后使用 `code` 作为 `str.encode()` 调用的编码。

参见：

模块 `curses.ascii` 在 ASCII 字符上工作的工具，无论你的区域设置是什么。

模块 `curses.panel` A panel stack extension that adds depth to curses windows.

Module `curses.textpad` Editable text widget for curses supporting **Emacs**-like bindings.

curses-howto Tutorial material on using curses with Python, by Andrew Kuchling and Eric Raymond.

The `Tools/demo/` directory in the Python source distribution contains some example programs using the curses bindings provided by this module.

16.10.1 函数

`curses` 模块定义了以下异常:

exception `curses.error`

当 `curses` 库中函数返回一个错误时引发的异常。

注解: Whenever *x* or *y* arguments to a function or a method are optional, they default to the current cursor location. Whenever *attr* is optional, it defaults to `A_NORMAL`.

`curses` 模块定义了以下函数:

`curses.baudrate()`

Return the output speed of the terminal in bits per second. On software terminal emulators it will have a fixed high value. Included for historical reasons; in former times, it was used to write output loops for time delays and occasionally to change interfaces depending on the line speed.

`curses.beep()`

Emit a short attention sound.

`curses.can_change_color()`

Return True or False, depending on whether the programmer can change the colors displayed by the terminal.

`curses.cbreak()`

Enter cbreak mode. In cbreak mode (sometimes called “rare” mode) normal tty line buffering is turned off and characters are available to be read one by one. However, unlike raw mode, special characters (interrupt, quit, suspend, and flow control) retain their effects on the tty driver and calling program. Calling first `raw()` then `cbreak()` leaves the terminal in cbreak mode.

`curses.color_content(color_number)`

Return the intensity of the red, green, and blue (RGB) components in the color *color_number*, which must be between 0 and `COLORS`. Return a 3-tuple, containing the R,G,B values for the given color, which will be between 0 (no component) and 1000 (maximum amount of component).

`curses.color_pair(color_number)`

Return the attribute value for displaying text in the specified color. This attribute value can be combined with `A_STANDOUT`, `A_REVERSE`, and the other `A_*` attributes. `pair_number()` is the counterpart to this function.

`curses.curs_set(visibility)`

Set the cursor state. *visibility* can be set to 0, 1, or 2, for invisible, normal, or very visible. If the terminal supports the visibility requested, return the previous cursor state; otherwise raise an exception. On many terminals, the “visible” mode is an underline cursor and the “very visible” mode is a block cursor.

`curses.def_prog_mode()`

Save the current terminal mode as the “program” mode, the mode when the running program is using `curses`. (Its counterpart is the “shell” mode, for when the program is not in `curses`.) Subsequent calls to `reset_prog_mode()` will restore this mode.

`curses.def_shell_mode()`

Save the current terminal mode as the “shell” mode, the mode when the running program is not using `curses`. (Its counterpart is the “program” mode, when the program is using `curses` capabilities.) Subsequent calls to `reset_shell_mode()` will restore this mode.

`curses.delay_output(ms)`

Insert an *ms* millisecond pause in output.

`curses.doupdate()`

Update the physical screen. The `curses` library keeps two data structures, one representing the current physical

screen contents and a virtual screen representing the desired next state. The `doupdate()` ground updates the physical screen to match the virtual screen.

The virtual screen may be updated by a `noutrefresh()` call after write operations such as `addstr()` have been performed on a window. The normal `refresh()` call is simply `noutrefresh()` followed by `doupdate()`; if you have to update multiple windows, you can speed performance and perhaps reduce screen flicker by issuing `noutrefresh()` calls on all windows, followed by a single `doupdate()`.

`curses.echo()`

Enter echo mode. In echo mode, each character input is echoed to the screen as it is entered.

`curses.endwin()`

De-initialize the library, and return terminal to normal status.

`curses.erasechar()`

Return the user's current erase character as a one-byte bytes object. Under Unix operating systems this is a property of the controlling tty of the curses program, and is not set by the curses library itself.

`curses.filter()`

The `filter()` routine, if used, must be called before `initscr()` is called. The effect is that, during those calls, `LINES` is set to 1; the capabilities `clear`, `cup`, `cud`, `cud1`, `cuu1`, `cuu`, `vpa` are disabled; and the home string is set to the value of `cr`. The effect is that the cursor is confined to the current line, and so are screen updates. This may be used for enabling character-at-a-time line editing without touching the rest of the screen.

`curses.flash()`

Flash the screen. That is, change it to reverse-video and then change it back in a short interval. Some people prefer such as 'visible bell' to the audible attention signal produced by `beep()`.

`curses.flushinp()`

Flush all input buffers. This throws away any typeahead that has been typed by the user and has not yet been processed by the program.

`curses.getmouse()`

After `getch()` returns `KEY_MOUSE` to signal a mouse event, this method should be call to retrieve the queued mouse event, represented as a 5-tuple (`id`, `x`, `y`, `z`, `bstate`). `id` is an ID value used to distinguish multiple devices, and `x`, `y`, `z` are the event's coordinates. (`z` is currently unused.) `bstate` is an integer value whose bits will be set to indicate the type of event, and will be the bitwise OR of one or more of the following constants, where `n` is the button number from 1 to 4: `BUTTONn_PRESSED`, `BUTTONn_RELEASED`, `BUTTONn_CLICKED`, `BUTTONn_DOUBLE_CLICKED`, `BUTTONn_TRIPLE_CLICKED`, `BUTTON_SHIFT`, `BUTTON_CTRL`, `BUTTON_ALT`.

`curses.getsyx()`

Return the current coordinates of the virtual screen cursor as a tuple (`y`, `x`). If `leaveok` is currently `True`, then return `(-1, -1)`.

`curses.getwin(file)`

Read window related data stored in the file by an earlier `putwin()` call. The routine then creates and initializes a new window using that data, returning the new window object.

`curses.has_colors()`

Return `True` if the terminal can display colors; otherwise, return `False`.

`curses.has_ic()`

Return `True` if the terminal has insert- and delete-character capabilities. This function is included for historical reasons only, as all modern software terminal emulators have such capabilities.

`curses.has_il()`

Return `True` if the terminal has insert- and delete-line capabilities, or can simulate them using scrolling regions. This function is included for historical reasons only, as all modern software terminal emulators have such capabilities.

`curses.has_key(ch)`

Take a key value *ch*, and return True if the current terminal type recognizes a key with that value.

`curses.halfdelay(tenths)`

Used for half-delay mode, which is similar to cbreak mode in that characters typed by the user are immediately available to the program. However, after blocking for *tenths* tenths of seconds, raise an exception if nothing has been typed. The value of *tenths* must be a number between 1 and 255. Use `nocbreak()` to leave half-delay mode.

`curses.init_color(color_number, r, g, b)`

Change the definition of a color, taking the number of the color to be changed followed by three RGB values (for the amounts of red, green, and blue components). The value of *color_number* must be between 0 and COLORS. Each of *r*, *g*, *b*, must be a value between 0 and 1000. When `init_color()` is used, all occurrences of that color on the screen immediately change to the new definition. This function is a no-op on most terminals; it is active only if `can_change_color()` returns True.

`curses.init_pair(pair_number, fg, bg)`

Change the definition of a color-pair. It takes three arguments: the number of the color-pair to be changed, the foreground color number, and the background color number. The value of *pair_number* must be between 1 and COLOR_PAIRS - 1 (the 0 color pair is wired to white on black and cannot be changed). The value of *fg* and *bg* arguments must be between 0 and COLORS. If the color-pair was previously initialized, the screen is refreshed and all occurrences of that color-pair are changed to the new definition.

`curses.initscr()`

Initialize the library. Return a `Window` object which represents the whole screen.

注解: If there is an error opening the terminal, the underlying curses library may cause the interpreter to exit.

`curses.is_term_resized(nlines, ncols)`

Return True if `resize_term()` would modify the window structure, False otherwise.

`curses.isendwin()`

Return True if `endwin()` has been called (that is, the curses library has been deinitialized).

`curses.keyname(k)`

Return the name of the key numbered *k* as a bytes object. The name of a key generating printable ASCII character is the key's character. The name of a control-key combination is a two-byte bytes object consisting of a caret (b'^') followed by the corresponding printable ASCII character. The name of an alt-key combination (128–255) is a bytes object consisting of the prefix b'M-' followed by the name of the corresponding ASCII character.

`curses.killchar()`

Return the user's current line kill character as a one-byte bytes object. Under Unix operating systems this is a property of the controlling tty of the curses program, and is not set by the curses library itself.

`curses.longname()`

Return a bytes object containing the terminfo long name field describing the current terminal. The maximum length of a verbose description is 128 characters. It is defined only after the call to `initscr()`.

`curses.meta(flag)`

If *flag* is True, allow 8-bit characters to be input. If *flag* is False, allow only 7-bit chars.

`curses.mouseinterval(interval)`

Set the maximum time in milliseconds that can elapse between press and release events in order for them to be recognized as a click, and return the previous interval value. The default value is 200 msec, or one fifth of a second.

`curses.mousemask(mousemask)`

Set the mouse events to be reported, and return a tuple (*availmask*, *oldmask*). *availmask* indicates which

of the specified mouse events can be reported; on complete failure it returns 0. *oldmask* is the previous value of the given window's mouse event mask. If this function is never called, no mouse events are ever reported.

`curses.napms(ms)`

Sleep for *ms* milliseconds.

`curses.newpad(nlines, ncols)`

Create and return a pointer to a new pad data structure with the given number of lines and columns. Return a pad as a window object.

A pad is like a window, except that it is not restricted by the screen size, and is not necessarily associated with a particular part of the screen. Pads can be used when a large window is needed, and only a part of the window will be on the screen at one time. Automatic refreshes of pads (such as from scrolling or echoing of input) do not occur. The `refresh()` and `noutrefresh()` methods of a pad require 6 arguments to specify the part of the pad to be displayed and the location on the screen to be used for the display. The arguments are *pminrow*, *pmincol*, *sminrow*, *smincol*, *smaxrow*, *smaxcol*; the *p* arguments refer to the upper left corner of the pad region to be displayed and the *s* arguments define a clipping box on the screen within which the pad region is to be displayed.

`curses.newwin(nlines, ncols)`

`curses.newwin(nlines, ncols, begin_y, begin_x)`

Return a new [window](#), whose left-upper corner is at (*begin_y*, *begin_x*), and whose height/width is *nlines/ncols*.

By default, the window will extend from the specified position to the lower right corner of the screen.

`curses.nl()`

Enter newline mode. This mode translates the return key into newline on input, and translates newline into return and line-feed on output. Newline mode is initially on.

`curses.nocbreak()`

Leave cbreak mode. Return to normal “cooked” mode with line buffering.

`curses.noecho()`

Leave echo mode. Echoing of input characters is turned off.

`curses.nonl()`

Leave newline mode. Disable translation of return into newline on input, and disable low-level translation of newline into newline/return on output (but this does not change the behavior of `addch('\n')`, which always does the equivalent of return and line feed on the virtual screen). With translation off, curses can sometimes speed up vertical motion a little; also, it will be able to detect the return key on input.

`curses.noqiflush()`

When the `noqiflush()` routine is used, normal flush of input and output queues associated with the INTR, QUIT and SUSP characters will not be done. You may want to call `noqiflush()` in a signal handler if you want output to continue as though the interrupt had not occurred, after the handler exits.

`curses.noraw()`

Leave raw mode. Return to normal “cooked” mode with line buffering.

`curses.pair_content(pair_number)`

Return a tuple (*fg*, *bg*) containing the colors for the requested color pair. The value of *pair_number* must be between 1 and `COLOR_PAIRS - 1`.

`curses.pair_number(attr)`

Return the number of the color-pair set by the attribute value *attr*. `color_pair()` is the counterpart to this function.

`curses.putp(str)`

Equivalent to `tputs(str, 1, putchar)`; emit the value of a specified terminfo capability for the current terminal. Note that the output of `putp()` always goes to standard output.

`curses.qiflush([flag])`

If *flag* is `False`, the effect is the same as calling `noqiflush()`. If *flag* is `True`, or no argument is provided, the queues will be flushed when these control characters are read.

`curses.raw()`

Enter raw mode. In raw mode, normal line buffering and processing of interrupt, quit, suspend, and flow control keys are turned off; characters are presented to curses input functions one by one.

`curses.reset_prog_mode()`

Restore the terminal to “program” mode, as previously saved by `def_prog_mode()`.

`curses.reset_shell_mode()`

Restore the terminal to “shell” mode, as previously saved by `def_shell_mode()`.

`curses.resetty()`

Restore the state of the terminal modes to what it was at the last call to `savetty()`.

`curses.resize_term(nlines, ncols)`

Backend function used by `resizeterm()`, performing most of the work; when resizing the windows, `resize_term()` blank-fills the areas that are extended. The calling application should fill in these areas with appropriate data. The `resize_term()` function attempts to resize all windows. However, due to the calling convention of pads, it is not possible to resize these without additional interaction with the application.

`curses.resizeterm(nlines, ncols)`

Resize the standard and current windows to the specified dimensions, and adjusts other bookkeeping data used by the curses library that record the window dimensions (in particular the SIGWINCH handler).

`curses.savetty()`

Save the current state of the terminal modes in a buffer, usable by `resetty()`.

`curses.setsyx(y, x)`

Set the virtual screen cursor to *y*, *x*. If *y* and *x* are both `-1`, then `leaveok` is set `True`.

`curses.setupterm(term=None, fd=-1)`

Initialize the terminal. *term* is a string giving the terminal name, or `None`; if omitted or `None`, the value of the `TERM` environment variable will be used. *fd* is the file descriptor to which any initialization sequences will be sent; if not supplied or `-1`, the file descriptor for `sys.stdout` will be used.

`curses.start_color()`

Must be called if the programmer wants to use colors, and before any other color manipulation routine is called. It is good practice to call this routine right after `initscr()`.

`start_color()` initializes eight basic colors (black, red, green, yellow, blue, magenta, cyan, and white), and two global variables in the `curses` module, `COLORS` and `COLOR_PAIRS`, containing the maximum number of colors and color-pairs the terminal can support. It also restores the colors on the terminal to the values they had when the terminal was just turned on.

`curses.termattrs()`

Return a logical OR of all video attributes supported by the terminal. This information is useful when a curses program needs complete control over the appearance of the screen.

`curses.termname()`

Return the value of the environment variable `TERM`, as a bytes object, truncated to 14 characters.

`curses.tigetflag(capname)`

Return the value of the Boolean capability corresponding to the terminfo capability name *capname* as an integer. Return the value `-1` if *capname* is not a Boolean capability, or `0` if it is canceled or absent from the terminal description.

`curses.tigetnum(capname)`

Return the value of the numeric capability corresponding to the terminfo capability name *capname* as an integer.

Return the value `-2` if *capname* is not a numeric capability, or `-1` if it is canceled or absent from the terminal description.

`curses.tigetstr(capname)`

Return the value of the string capability corresponding to the terminfo capability name *capname* as a bytes object. Return `None` if *capname* is not a terminfo “string capability”, or is canceled or absent from the terminal description.

`curses.tparm(str[, ...])`

Instantiate the bytes object *str* with the supplied parameters, where *str* should be a parameterized string obtained from the terminfo database. E.g. `tparm(tigetstr("cup"), 5, 3)` could result in `b'\033[6;4H'`, the exact result depending on terminal type.

`curses.typeahead(fd)`

Specify that the file descriptor *fd* be used for typeahead checking. If *fd* is `-1`, then no typeahead checking is done.

The curses library does “line-breakout optimization” by looking for typeahead periodically while updating the screen. If input is found, and it is coming from a tty, the current update is postponed until `refresh` or `doupdate` is called again, allowing faster response to commands typed in advance. This function allows specifying a different file descriptor for typeahead checking.

`curses.unctrl(ch)`

Return a bytes object which is a printable representation of the character *ch*. Control characters are represented as a caret followed by the character, for example as `b'^C'`. Printing characters are left as they are.

`curses.ungetch(ch)`

Push *ch* so the next `getch()` will return it.

注解: Only one *ch* can be pushed before `getch()` is called.

`curses.update_lines_cols()`

Update `LINES` and `COLS`. Useful for detecting manual screen resize.

3.5 新版功能.

`curses.unget_wch(ch)`

Push *ch* so the next `get_wch()` will return it.

注解: Only one *ch* can be pushed before `get_wch()` is called.

3.3 新版功能.

`curses.ungetmouse(id, x, y, z, bstate)`

Push a `KEY_MOUSE` event onto the input queue, associating the given state data with it.

`curses.use_env(flag)`

If used, this function should be called before `initscr()` or `newterm` are called. When *flag* is `False`, the values of lines and columns specified in the terminfo database will be used, even if environment variables `LINES` and `COLUMNS` (used by default) are set, or if curses is running in a window (in which case default behavior would be to use the window size if `LINES` and `COLUMNS` are not set).

`curses.use_default_colors()`

Allow use of default values for colors on terminals supporting this feature. Use this to support transparency in your application. The default color is assigned to the color number `-1`. After calling this function, `init_pair(x, curses.COLOR_RED, -1)` initializes, for instance, color pair *x* to a red foreground color on the default background.

`curses.wrapper(func, ...)`

Initialize curses and call another callable object, *func*, which should be the rest of your curses-using application.

If the application raises an exception, this function will restore the terminal to a sane state before re-raising the exception and generating a traceback. The callable object *func* is then passed the main window `'stdscr'` as its first argument, followed by any other arguments passed to `wrapper()`. Before calling *func*, `wrapper()` turns on cbreak mode, turns off echo, enables the terminal keypad, and initializes colors if the terminal has color support. On exit (whether normally or by exception) it restores cooked mode, turns on echo, and disables the terminal keypad.

16.10.2 Window Objects

Window objects, as returned by `initscr()` and `newwin()` above, have the following methods and attributes:

`window.addch(ch[, attr])`
`window.addch(y, x, ch[, attr])`

Paint character *ch* at (*y*, *x*) with attributes *attr*, overwriting any character previously painter at that location. By default, the character position and attributes are the current settings for the window object.

注解: Writing outside the window, subwindow, or pad raises a `curses.error`. Attempting to write to the lower right corner of a window, subwindow, or pad will cause an exception to be raised after the character is printed.

`window.addnstr(str, n[, attr])`
`window.addnstr(y, x, str, n[, attr])`

Paint at most *n* characters of the character string *str* at (*y*, *x*) with attributes *attr*, overwriting anything previously on the display.

`window.addstr(str[, attr])`
`window.addstr(y, x, str[, attr])`

Paint the character string *str* at (*y*, *x*) with attributes *attr*, overwriting anything previously on the display.

注解: Writing outside the window, subwindow, or pad raises `curses.error`. Attempting to write to the lower right corner of a window, subwindow, or pad will cause an exception to be raised after the string is printed.

`window.attroff(attr)`

Remove attribute *attr* from the “background” set applied to all writes to the current window.

`window.atttron(attr)`

Add attribute *attr* from the “background” set applied to all writes to the current window.

`window.attrset(attr)`

Set the “background” set of attributes to *attr*. This set is initially 0 (no attributes).

`window.bkgd(ch[, attr])`

Set the background property of the window to the character *ch*, with attributes *attr*. The change is then applied to every character position in that window:

- The attribute of every character in the window is changed to the new background attribute.
- Wherever the former background character appears, it is changed to the new background character.

`window.bkgdset(ch[, attr])`

Set the window's background. A window's background consists of a character and any combination of attributes. The attribute part of the background is combined (OR'ed) with all non-blank characters that are written into the window. Both the character and attribute parts of the background are combined with the blank characters. The background becomes a property of the character and moves with the character through any scrolling and insert/delete line/character operations.

`window.border([ls[, rs[, ts[, bs[, tl[, tr[, bl[, br]]]]]]]])`

在窗口边缘绘制边框。每个参数指定用于边界特定部分的字符; 请参阅下表了解更多详情。

注解: A 0 value for any parameter will cause the default character to be used for that parameter. Keyword parameters can *not* be used. The defaults are listed in this table:

参数	描述	默认值
<i>ls</i>	左侧	ACS_VLINE
<i>rs</i>	右侧	ACS_VLINE
<i>ts</i>	顶部	ACS_HLINE
<i>bs</i>	底部	ACS_HLINE
<i>tl</i>	左上角	ACS_ULCORNER
<i>tr</i>	右上角	ACS_URCORNER
<i>bl</i>	左下角	ACS_LLCORNER
<i>br</i>	右下角	ACS_LRCORNER

`window.box([vertch, horch])`

Similar to `border()`, but both *ls* and *rs* are *vertch* and both *ts* and *bs* are *horch*. The default corner characters are always used by this function.

`window.chgat(attr)`

`window.chgat(num, attr)`

`window.chgat(y, x, attr)`

`window.chgat(y, x, num, attr)`

Set the attributes of *num* characters at the current cursor position, or at position (*y*, *x*) if supplied. If *num* is not given or is -1, the attribute will be set on all the characters to the end of the line. This function moves cursor to position (*y*, *x*) if supplied. The changed line will be touched using the `touchline()` method so that the contents will be redisplayed by the next window refresh.

`window.clear()`

Like `erase()`, but also cause the whole window to be repainted upon next call to `refresh()`.

`window.clearok(flag)`

If *flag* is True, the next call to `refresh()` will clear the window completely.

`window.clrtoebot()`

Erase from cursor to the end of the window: all lines below the cursor are deleted, and then the equivalent of `clrtoeol()` is performed.

`window.clrtoeol()`

Erase from cursor to the end of the line.

`window.cursyncup()`

Update the current cursor position of all the ancestors of the window to reflect the current cursor position of the window.

`window.delch([y, x])`

Delete any character at (*y*, *x*).

`window.deleteln()`

Delete the line under the cursor. All following lines are moved up by one line.

`window.derwin(begin_y, begin_x)`

`window.derwin(nlines, ncols, begin_y, begin_x)`

An abbreviation for “derive window”, `derwin()` is the same as calling `subwin()`, except that *begin_y* and

begin_x are relative to the origin of the window, rather than relative to the entire screen. Return a window object for the derived window.

`window.echochar(ch[, attr])`

Add character *ch* with attribute *attr*, and immediately call `refresh()` on the window.

`window.enclose(y, x)`

Test whether the given pair of screen-relative character-cell coordinates are enclosed by the given window, returning True or False. It is useful for determining what subset of the screen windows enclose the location of a mouse event.

`window.encoding`

Encoding used to encode method arguments (Unicode strings and characters). The encoding attribute is inherited from the parent window when a subwindow is created, for example with `window.subwin()`. By default, the locale encoding is used (see `locale.getpreferredencoding()`).

3.3 新版功能.

`window.erase()`

Clear the window.

`window.getbegyx()`

Return a tuple (*y*, *x*) of co-ordinates of upper-left corner.

`window.getbkgd()`

Return the given window's current background character/attribute pair.

`window.getch([y, x])`

Get a character. Note that the integer returned does *not* have to be in ASCII range: function keys, keypad keys and so on are represented by numbers higher than 255. In no-delay mode, return -1 if there is no input, otherwise wait until a key is pressed.

`window.get_wch([y, x])`

Get a wide character. Return a character for most keys, or an integer for function keys, keypad keys, and other special keys. In no-delay mode, raise an exception if there is no input.

3.3 新版功能.

`window.getkey([y, x])`

Get a character, returning a string instead of an integer, as `getch()` does. Function keys, keypad keys and other special keys return a multibyte string containing the key name. In no-delay mode, raise an exception if there is no input.

`window.getmaxyx()`

Return a tuple (*y*, *x*) of the height and width of the window.

`window.getparyx()`

Return the beginning coordinates of this window relative to its parent window as a tuple (*y*, *x*). Return (-1, -1) if this window has no parent.

`window.getstr()`

`window.getstr(n)`

`window.getstr(y, x)`

`window.getstr(y, x, n)`

Read a bytes object from the user, with primitive line editing capacity.

`window.getyx()`

Return a tuple (*y*, *x*) of current cursor position relative to the window's upper-left corner.

`window.hline(ch, n)`

`window.hline(y, x, ch, n)`

Display a horizontal line starting at (*y*, *x*) with length *n* consisting of the character *ch*.

`window.idcok(flag)`

If *flag* is `False`, `curses` no longer considers using the hardware insert/delete character feature of the terminal; if *flag* is `True`, use of character insertion and deletion is enabled. When `curses` is first initialized, use of character insert/delete is enabled by default.

`window.idlok(flag)`

If *flag* is `True`, `curses` will try and use hardware line editing facilities. Otherwise, line insertion/deletion are disabled.

`window.immedok(flag)`

If *flag* is `True`, any change in the window image automatically causes the window to be refreshed; you no longer have to call `refresh()` yourself. However, it may degrade performance considerably, due to repeated calls to `wrefresh`. This option is disabled by default.

`window.inch([y, x])`

Return the character at the given position in the window. The bottom 8 bits are the character proper, and upper bits are the attributes.

`window.insch(ch[, attr])`

`window.insch(y, x, ch[, attr])`

Paint character *ch* at (*y*, *x*) with attributes *attr*, moving the line from position *x* right by one character.

`window.insdelln(nlines)`

Insert *nlines* lines into the specified window above the current line. The *nlines* bottom lines are lost. For negative *nlines*, delete *nlines* lines starting with the one under the cursor, and move the remaining lines up. The bottom *nlines* lines are cleared. The current cursor position remains the same.

`window.insertln()`

Insert a blank line under the cursor. All following lines are moved down by one line.

`window.insnstr(str, n[, attr])`

`window.insnstr(y, x, str, n[, attr])`

Insert a character string (as many characters as will fit on the line) before the character under the cursor, up to *n* characters. If *n* is zero or negative, the entire string is inserted. All characters to the right of the cursor are shifted right, with the rightmost characters on the line being lost. The cursor position does not change (after moving to *y*, *x*, if specified).

`window.insstr(str[, attr])`

`window.insstr(y, x, str[, attr])`

Insert a character string (as many characters as will fit on the line) before the character under the cursor. All characters to the right of the cursor are shifted right, with the rightmost characters on the line being lost. The cursor position does not change (after moving to *y*, *x*, if specified).

`window.instr([n])`

`window.instr(y, x[, n])`

Return a bytes object of characters, extracted from the window starting at the current cursor position, or at *y*, *x* if specified. Attributes are stripped from the characters. If *n* is specified, `instr()` returns a string at most *n* characters long (exclusive of the trailing NUL).

`window.is_linetouched(line)`

Return `True` if the specified line was modified since the last call to `refresh()`; otherwise return `False`. Raise a `curses.error` exception if *line* is not valid for the given window.

`window.is_wintouched()`

Return `True` if the specified window was modified since the last call to `refresh()`; otherwise return `False`.

`window.keypad(flag)`

If *flag* is `True`, escape sequences generated by some keys (keypad, function keys) will be interpreted by `curses`. If *flag* is `False`, escape sequences will be left as is in the input stream.

`window.leaveok(flag)`

If *flag* is `True`, cursor is left where it is on update, instead of being at “cursor position.” This reduces cursor movement where possible. If possible the cursor will be made invisible.

If *flag* is `False`, cursor will always be at “cursor position” after an update.

`window.move(new_y, new_x)`

Move cursor to `(new_y, new_x)`.

`window.mvderwin(y, x)`

Move the window inside its parent window. The screen-relative parameters of the window are not changed. This routine is used to display different parts of the parent window at the same physical position on the screen.

`window.mvwin(new_y, new_x)`

Move the window so its upper-left corner is at `(new_y, new_x)`.

`window.nodelay(flag)`

If *flag* is `True`, `getch()` will be non-blocking.

`window.notimeout(flag)`

If *flag* is `True`, escape sequences will not be timed out.

If *flag* is `False`, after a few milliseconds, an escape sequence will not be interpreted, and will be left in the input stream as is.

`window.noutrefresh()`

Mark for refresh but wait. This function updates the data structure representing the desired state of the window, but does not force an update of the physical screen. To accomplish that, call `doupdate()`.

`window.overlay(destwin[, sminrow, smincol, dminrow, dmincol, dmaxrow, dmaxcol])`

Overlay the window on top of *destwin*. The windows need not be the same size, only the overlapping region is copied. This copy is non-destructive, which means that the current background character does not overwrite the old contents of *destwin*.

To get fine-grained control over the copied region, the second form of `overlay()` can be used. *sminrow* and *smincol* are the upper-left coordinates of the source window, and the other variables mark a rectangle in the destination window.

`window.overwrite(destwin[, sminrow, smincol, dminrow, dmincol, dmaxrow, dmaxcol])`

Overwrite the window on top of *destwin*. The windows need not be the same size, in which case only the overlapping region is copied. This copy is destructive, which means that the current background character overwrites the old contents of *destwin*.

To get fine-grained control over the copied region, the second form of `overwrite()` can be used. *sminrow* and *smincol* are the upper-left coordinates of the source window, the other variables mark a rectangle in the destination window.

`window.putwin(file)`

Write all data associated with the window into the provided file object. This information can be later retrieved using the `getwin()` function.

`window.redrawln(beg, num)`

Indicate that the *num* screen lines, starting at line *beg*, are corrupted and should be completely redrawn on the next `refresh()` call.

`window.redrawwin()`

Touch the entire window, causing it to be completely redrawn on the next `refresh()` call.

`window.refresh([pminrow, pmincol, sminrow, smincol, smaxrow, smaxcol])`

Update the display immediately (sync actual screen with previous drawing/deleting methods).

The 6 optional arguments can only be specified when the window is a pad created with `newpad()`. The additional parameters are needed to indicate what part of the pad and screen are involved. `pminrow` and `pmincol` specify the upper left-hand corner of the rectangle to be displayed in the pad. `sminrow`, `smincol`, `smaxrow`, and `smaxcol` specify the edges of the rectangle to be displayed on the screen. The lower right-hand corner of the rectangle to be displayed in the pad is calculated from the screen coordinates, since the rectangles must be the same size. Both rectangles must be entirely contained within their respective structures. Negative values of `pminrow`, `pmincol`, `sminrow`, or `smincol` are treated as if they were zero.

`window.resize(nlines, ncols)`

Reallocate storage for a curses window to adjust its dimensions to the specified values. If either dimension is larger than the current values, the window's data is filled with blanks that have the current background rendition (as set by `bkgdset()`) merged into them.

`window.scroll([lines=1])`

Scroll the screen or scrolling region upward by *lines* lines.

`window.scrollok(flag)`

Control what happens when the cursor of a window is moved off the edge of the window or scrolling region, either as a result of a newline action on the bottom line, or typing the last character of the last line. If *flag* is `False`, the cursor is left on the bottom line. If *flag* is `True`, the window is scrolled up one line. Note that in order to get the physical scrolling effect on the terminal, it is also necessary to call `idlok()`.

`window.setscrreg(top, bottom)`

Set the scrolling region from line *top* to line *bottom*. All scrolling actions will take place in this region.

`window.standend()`

Turn off the standout attribute. On some terminals this has the side effect of turning off all attributes.

`window.standout()`

Turn on attribute `A_STANDOUT`.

`window.subpad(begin_y, begin_x)`

`window.subpad(nlines, ncols, begin_y, begin_x)`

Return a sub-window, whose upper-left corner is at (*begin_y*, *begin_x*), and whose width/height is *ncols/nlines*.

`window.subwin(begin_y, begin_x)`

`window.subwin(nlines, ncols, begin_y, begin_x)`

Return a sub-window, whose upper-left corner is at (*begin_y*, *begin_x*), and whose width/height is *ncols/nlines*.

By default, the sub-window will extend from the specified position to the lower right corner of the window.

`window.syncdown()`

Touch each location in the window that has been touched in any of its ancestor windows. This routine is called by `refresh()`, so it should almost never be necessary to call it manually.

`window.syncok(flag)`

If *flag* is `True`, then `syncup()` is called automatically whenever there is a change in the window.

`window.syncup()`

Touch all locations in ancestors of the window that have been changed in the window.

`window.timeout(delay)`

Set blocking or non-blocking read behavior for the window. If *delay* is negative, blocking read is used (which will wait indefinitely for input). If *delay* is zero, then non-blocking read is used, and `getch()` will return `-1` if no input is waiting. If *delay* is positive, then `getch()` will block for *delay* milliseconds, and return `-1` if there is still no input at the end of that time.

`window.touchline (start, count[, changed])`

Pretend *count* lines have been changed, starting with line *start*. If *changed* is supplied, it specifies whether the affected lines are marked as having been changed (*changed*=True) or unchanged (*changed*=False).

`window.touchwin ()`

Pretend the whole window has been changed, for purposes of drawing optimizations.

`window.untouchwin ()`

Mark all lines in the window as unchanged since the last call to `refresh()`.

`window.vline (ch, n)`

`window.vline (y, x, ch, n)`

Display a vertical line starting at (*y*, *x*) with length *n* consisting of the character *ch*.

16.10.3 常量

The `curses` module defines the following data members:

`curses.ERR`

Some curses routines that return an integer, such as `getch()`, return `ERR` upon failure.

`curses.OK`

Some curses routines that return an integer, such as `napms()`, return `OK` upon success.

`curses.version`

A bytes object representing the current version of the module. Also available as `__version__`.

Some constants are available to specify character cell attributes. The exact constants available are system dependent.

属性	意义
<code>A_ALTCHARSET</code>	备用字符集模式
<code>A_BLINK</code>	闪烁模式
<code>A_BOLD</code>	粗体模式
<code>A_DIM</code>	暗淡模式
<code>A_INVIS</code>	不可见或空白模式
<code>A_NORMAL</code>	正常属性
<code>A_PROTECT</code>	保护模式
<code>A_REVERSE</code>	反转背景色和前景色
<code>A_STANDOUT</code>	突出模式
<code>A_UNDERLINE</code>	下划线模式
<code>A_HORIZONTAL</code>	水平突出显示
<code>A_LEFT</code>	左高亮
<code>A_LOW</code>	底部高亮
<code>A_RIGHT</code>	右高亮
<code>A_TOP</code>	顶部高亮
<code>A_VERTICAL</code>	垂直突出显示
<code>A_CHARTEXT</code>	用于提取字符的位掩码

有几个常量可用于提取某些方法返回的相应属性。

位掩码	意义
<code>A_ATTRIBUTES</code>	用于提取属性的位掩码
<code>A_CHARTEXT</code>	用于提取字符的位掩码
<code>A_COLOR</code>	用于提取颜色对字段信息的位掩码

键由名称以 KEY_ 开头的整数常量引用。确切的可用键取决于系统。

关键常数	(Windows 注册表的) 键 (Windows 注册表的) 键
KEY_MIN	最小键值
KEY_BREAK	中断键 (不可靠)
KEY_DOWN	向下箭头
KEY_UP	向上箭头
KEY_LEFT	向左箭头
KEY_RIGHT	向右箭头
KEY_HOME	Home key (upward+left arrow)
KEY_BACKSPACE	退格 (不可靠)
KEY_F0	Function keys. Up to 64 function keys are supported.
KEY_Fn	Value of function key <i>n</i>
KEY_DL	删除行
KEY_IL	插入行
KEY_DC	Delete character
KEY_IC	Insert char or enter insert mode
KEY_EIC	Exit insert char mode
KEY_CLEAR	Clear screen
KEY_EOS	Clear to end of screen
KEY_EOL	Clear to end of line
KEY_SF	Scroll 1 line forward
KEY_SR	Scroll 1 line backward (reverse)
KEY_NPAGE	下一页
KEY_PPAGE	上一页
KEY_STAB	Set tab
KEY_CTAB	Clear tab
KEY_CATAB	Clear all tabs
KEY_ENTER	Enter or send (unreliable)
KEY_SRESET	Soft (partial) reset (unreliable)
KEY_RESET	Reset or hard reset (unreliable)
KEY_PRINT	打印
KEY_LL	Home down or bottom (lower left)
KEY_A1	键盘的左上角
KEY_A3	键盘的右上角
KEY_B2	键盘的中心
KEY_C1	键盘左下方
KEY_C3	键盘右下方
KEY_BTAB	Back tab
KEY_BEG	Beg (beginning)
KEY_CANCEL	取消
KEY_CLOSE	关闭
KEY_COMMAND	Cmd (命令行)
KEY_COPY	复制
KEY_CREATE	创建
KEY_END	End
KEY_EXIT	退出
KEY_FIND	查找
KEY_HELP	帮助
KEY_MARK	标记
KEY_MESSAGE	消息

下页继续

表 1 - 续上页

关键常数	(Windows 注册表的) 键 (Windows 注册表的) 键
KEY_MOVE	移动
KEY_NEXT	下一个
KEY_OPEN	打开
KEY_OPTIONS	选项
KEY_PREVIOUS	Prev (previous)
KEY_REDO	重做
KEY_REFERENCE	Ref (reference)
KEY_REFRESH	刷新
KEY_REPLACE	替换
KEY_RESTART	重启
KEY_RESUME	恢复
KEY_SAVE	保存
KEY_SBEG	Shifted Beg (beginning)
KEY_SCANCEL	Shifted Cancel
KEY_SCOMMAND	Shifted Command
KEY_SCOPY	Shifted Copy
KEY_SCREATE	Shifted Create
KEY_SDC	Shifted Delete char
KEY_SDL	Shifted Delete line
KEY_SELECT	Select
KEY_SEND	Shifted End
KEY_SEOL	Shifted Clear line
KEY_SEXIT	Shifted Exit
KEY_SFIND	Shifted Find
KEY_SHELP	Shifted Help
KEY_SHOME	Shifted Home
KEY_SIC	Shifted Input
KEY_SLEFT	Shifted Left arrow
KEY_SMESSAGE	Shifted Message
KEY_SMOVE	Shifted Move
KEY_SNEXT	Shifted Next
KEY_SOPTIONS	Shifted Options
KEY_SPREVIOUS	Shifted Prev
KEY_SPRINT	Shifted Print
KEY_SREDO	Shifted Redo
KEY_SREPLACE	Shifted Replace
KEY_SRIGHT	Shifted Right arrow
KEY_SRSUME	Shifted Resume
KEY_SSAVE	Shifted Save
KEY_SSUSPEND	Shifted Suspend
KEY_SUNDO	Shifted Undo
KEY_SUSPEND	Suspend
KEY_UNDO	撤销操作
KEY_MOUSE	Mouse event has occurred
KEY_RESIZE	Terminal resize event
KEY_MAX	Maximum key value

在 VT100 及其软件仿真（例如 X 终端仿真器）上，通常至少有四个功能键（KEY_F1, KEY_F2, KEY_F3, KEY_F4）可用，并且箭头键以明显的方式映射到 KEY_UP, KEY_DOWN, KEY_LEFT 和 KEY_RIGHT。如果您的机器有一个 PC 键盘，可以安全地使用箭头键和十二个功能键（旧的 PC 键盘可能只有十个功能键）；此外，

以下键盘映射是标准的：

键帽	常数
Insert	KEY_IC
Delete	KEY_DC
Home	KEY_HOME
End	KEY_END
Page Up	KEY_PPAGE
Page Down	KEY_NPAGE

The following table lists characters from the alternate character set. These are inherited from the VT100 terminal, and will generally be available on software emulations such as X terminals. When there is no graphic available, curses falls back on a crude printable ASCII approximation.

注解：只有在调用 `initscr()` 之后才能使用它们

ACS 代码	意义
ACS_BBSS	右上角的别名
ACS_BLOCK	实心方块
ACS_BOARD	正方形
ACS_BSBS	水平线的别名
ACS_BSSB	左上角的别名
ACS_BSSS	顶部 T 型的别名
ACS_BTEE	底部 T 型
ACS_BULLET	bullet
ACS_CKBOARD	棋盘（点刻）
ACS_DARROW	向下箭头
ACS_DEGREE	等级符
ACS_DIAMOND	菱形
ACS_GEQUAL	大于或等于
ACS_HLINE	水平线
ACS_LANTERN	灯形符号
ACS_LARROW	向左箭头
ACS_LEQUAL	小于或等于
ACS_LLCORNER	左下角
ACS_LRCORNER	右下角
ACS_LTEE	左侧 T 型
ACS_NEQUAL	不等号
ACS_PI	字母 π
ACS_PLMINUS	正负号
ACS_PLUS	加号
ACS_RARROW	向右箭头
ACS_RTEE	右侧 T 型
ACS_S1	扫描线 1
ACS_S3	扫描线 3
ACS_S7	扫描线 7
ACS_S9	扫描线 9
ACS_SBBS	右下角的别名
ACS_SBSB	垂直线的别名

下页继续

表 2 - 续上页

ACS 代码	意义
ACS_SBSS	右侧 T 型的别名
ACS_SSBB	左下角的别名
ACS_SSBS	底部 T 型的别名
ACS_SSSB	左侧 T 型的别名
ACS_SSSS	alternate name for crossover or big plus
ACS_STERLING	英镑
ACS_TTEE	顶部 T 型
ACS_UARROW	向上箭头
ACS_ULCORNER	左上角
ACS_URCORNER	右上角
ACS_VLINE	垂线

下表列出了预定义的颜色：

常数	颜色
COLOR_BLACK	黑色
COLOR_BLUE	蓝色
COLOR_CYAN	青色（浅绿蓝色）
COLOR_GREEN	绿色
COLOR_MAGENTA	洋红色（紫红色）
COLOR_RED	红色
COLOR_WHITE	白色
COLOR_YELLOW	黄色

16.11 `curses.textpad` —Text input widget for curses programs

The `curses.textpad` module provides a `Textbox` class that handles elementary text editing in a curses window, supporting a set of keybindings resembling those of Emacs (thus, also of Netscape Navigator, BBedit 6.x, FrameMaker, and many other programs). The module also provides a rectangle-drawing function useful for framing text boxes or for other purposes.

The module `curses.textpad` defines the following function:

`curses.textpad.rectangle(win, uly, ulx, lry, lrx)`

Draw a rectangle. The first argument must be a window object; the remaining arguments are coordinates relative to that window. The second and third arguments are the y and x coordinates of the upper left hand corner of the rectangle to be drawn; the fourth and fifth arguments are the y and x coordinates of the lower right hand corner. The rectangle will be drawn using VT100/IBM PC forms characters on terminals that make this possible (including xterm and most other software terminal emulators). Otherwise it will be drawn with ASCII dashes, vertical bars, and plus signs.

16.11.1 文本框对象

You can instantiate a `Textbox` object as follows:

class `curses.textpad.Textbox(win)`

Return a textbox widget object. The `win` argument should be a curses `window` object in which the textbox is to be contained. The edit cursor of the textbox is initially located at the upper left hand corner of the containing window, with coordinates `(0, 0)`. The instance's `stripspaces` flag is initially on.

`Textbox` objects have the following methods:

edit ([*validator*])

This is the entry point you will normally use. It accepts editing keystrokes until one of the termination keystrokes is entered. If *validator* is supplied, it must be a function. It will be called for each keystroke entered with the keystroke as a parameter; command dispatch is done on the result. This method returns the window contents as a string; whether blanks in the window are included is affected by the *stripspaces* attribute.

do_command (*ch*)

处理单个按键命令。以下是支持的特殊按键：

按键	动作
Control-A	转到窗口的左边缘。
Control-B	光标向左，如果可能，包含前一行。
Control-D	删除光标下的字符。
Control-E	Go to right edge (stripspaces off) or end of line (stripspaces on).
Control-F	向右移动光标，适当时换行到下一行。
Control-G	终止，返回窗口内容。
Control-H	向后删除字符。
Control-J	如果窗口是 1 行则终止，否则插入换行符。
Control-K	如果行为空，则删除它，否则清除到行尾。
Control-L	刷新屏幕。
Control-N	光标向下；向下移动一行。
Control-O	在光标位置插入一个空行。
Control-P	光标向上；向上移动一行。

如果光标位于无法移动的边缘，则移动操作不执行任何操作。在可能的情况下，支持以下同义词：

常数	按键
KEY_LEFT	Control-B
KEY_RIGHT	Control-F
KEY_UP	Control-P
KEY_DOWN	Control-N
KEY_BACKSPACE	Control-h

All other keystrokes are treated as a command to insert the given character and move right (with line wrapping).

gather ()

Return the window contents as a string; whether blanks in the window are included is affected by the *stripspaces* member.

stripspaces

This attribute is a flag which controls the interpretation of blanks in the window. When it is on, trailing blanks on each line are ignored; any cursor motion that would land the cursor on a trailing blank goes to the end of that line instead, and trailing blanks are stripped when the window contents are gathered.

16.12 `curses.ascii` —Utilities for ASCII characters

The `curses.ascii` module supplies name constants for ASCII characters and functions to test membership in various ASCII character classes. The constants supplied are names for control characters as follows:

名称	意义
NUL	
SOH	Start of heading, console interrupt
STX	Start of text
ETX	End of text
EOT	End of transmission
ENQ	Enquiry, goes with ACK flow control
ACK	Acknowledgement
BEL	Bell
BS	Backspace
TAB	Tab
HT	Alias for TAB: “Horizontal tab”
LF	Line feed
NL	Alias for LF: “New line”
VT	Vertical tab
FF	Form feed
CR	Carriage return
SO	Shift-out, begin alternate character set
SI	Shift-in, resume default character set
DLE	Data-link escape
DC1	XON, for flow control
DC2	Device control 2, block-mode flow control
DC3	XOFF, for flow control
DC4	Device control 4
NAK	Negative acknowledgement
SYN	Synchronous idle
ETB	End transmission block
CAN	取消
EM	End of medium
SUB	Substitute
ESC	Escape
FS	File separator
GS	Group separator
RS	Record separator, block-mode terminator
US	Unit separator
SP	Space
DEL	Delete

Note that many of these have little practical significance in modern usage. The mnemonics derive from teleprinter conventions that predate digital computers.

The module supplies the following functions, patterned on those in the standard C library:

`curses.ascii.isalnum(c)`

Checks for an ASCII alphanumeric character; it is equivalent to `isalpha(c)` or `isdigit(c)`.

`curses.ascii.isalpha(c)`

Checks for an ASCII alphabetic character; it is equivalent to `isupper(c)` or `islower(c)`.

`curses.ascii.isascii(c)`

Checks for a character value that fits in the 7-bit ASCII set.

`curses.ascii.isblank(c)`

Checks for an ASCII whitespace character; space or horizontal tab.

`curses.ascii.iscntrl(c)`

Checks for an ASCII control character (in the range 0x00 to 0x1f or 0x7f).

`curses.ascii.isdigit(c)`

Checks for an ASCII decimal digit, '0' through '9'. This is equivalent to `c in string.digits`.

`curses.ascii.isgraph(c)`

Checks for ASCII any printable character except space.

`curses.ascii.islower(c)`

Checks for an ASCII lower-case character.

`curses.ascii.isprint(c)`

Checks for any ASCII printable character including space.

`curses.ascii.ispunct(c)`

Checks for any printable ASCII character which is not a space or an alphanumeric character.

`curses.ascii.isspace(c)`

Checks for ASCII white-space characters; space, line feed, carriage return, form feed, horizontal tab, vertical tab.

`curses.ascii.isupper(c)`

Checks for an ASCII uppercase letter.

`curses.ascii.isxdigit(c)`

Checks for an ASCII hexadecimal digit. This is equivalent to `c in string.hexdigits`.

`curses.ascii.isctrl(c)`

Checks for an ASCII control character (ordinal values 0 to 31).

`curses.ascii.ismeta(c)`

Checks for a non-ASCII character (ordinal values 0x80 and above).

These functions accept either integers or single-character strings; when the argument is a string, it is first converted using the built-in function `ord()`.

Note that all these functions check ordinal bit values derived from the character of the string you pass in; they do not actually know anything about the host machine's character encoding.

The following two functions take either a single-character string or integer byte value; they return a value of the same type.

`curses.ascii.ascii(c)`

Return the ASCII value corresponding to the low 7 bits of *c*.

`curses.ascii.ctrl(c)`

Return the control character corresponding to the given character (the character bit value is bitwise-anded with 0x1f).

`curses.ascii.alt(c)`

Return the 8-bit character corresponding to the given ASCII character (the character bit value is bitwise-ored with 0x80).

The following function takes either a single-character string or integer value; it returns a string.

`curses.ascii.unctrl(c)`

Return a string representation of the ASCII character *c*. If *c* is printable, this string is the character itself. If the character is a control character (0x00–0x1f) the string consists of a caret ('^') followed by the corresponding uppercase letter. If the character is an ASCII delete (0x7f) the string is '^?'. If the character has its meta bit (0x80) set, the meta bit is stripped, the preceding rules applied, and '!' prepended to the result.

`curses.ascii.controlnames`

A 33-element string array that contains the ASCII mnemonics for the thirty-two ASCII control characters from 0 (NUL) to 0x1f (US), in order, plus the mnemonic SP for the space character.

16.13 `curses.panel` —A panel stack extension for `curses`

Panels are windows with the added feature of depth, so they can be stacked on top of each other, and only the visible portions of each window will be displayed. Panels can be added, moved up or down in the stack, and removed.

16.13.1 函数

The module `curses.panel` defines the following functions:

`curses.panel.bottom_panel()`

Returns the bottom panel in the panel stack.

`curses.panel.new_panel(win)`

Returns a panel object, associating it with the given window *win*. Be aware that you need to keep the returned panel object referenced explicitly. If you don't, the panel object is garbage collected and removed from the panel stack.

`curses.panel.top_panel()`

Returns the top panel in the panel stack.

`curses.panel.update_panels()`

Updates the virtual screen after changes in the panel stack. This does not call `curses.doupdate()`, so you'll have to do this yourself.

16.13.2 Panel Objects

Panel objects, as returned by `new_panel()` above, are windows with a stacking order. There's always a window associated with a panel which determines the content, while the panel methods are responsible for the window's depth in the panel stack.

Panel objects have the following methods:

`Panel.above()`

Returns the panel above the current panel.

`Panel.below()`

Returns the panel below the current panel.

`Panel.bottom()`

Push the panel to the bottom of the stack.

`Panel.hidden()`

Returns True if the panel is hidden (not visible), False otherwise.

`Panel.hide()`

Hide the panel. This does not delete the object, it just makes the window on screen invisible.

`Panel.move(y, x)`

Move the panel to the screen coordinates (y, x).

`Panel.replace(win)`

Change the window associated with the panel to the window *win*.

`Panel.set_userptr(obj)`

Set the panel's user pointer to *obj*. This is used to associate an arbitrary piece of data with the panel, and can be any Python object.

`Panel.show()`

Display the panel (which might have been hidden).

`Panel.top()`

Push panel to the top of the stack.

`Panel.userptr()`

Returns the user pointer for the panel. This might be any Python object.

`Panel.window()`

Returns the window object associated with the panel.

16.14 platform — 获取底层平台的标识数据

示例代码: [Lib/platform.py](#)

注解: Specific platforms listed alphabetically, with Linux included in the Unix section.

16.14.1 跨平台

`platform.architecture(executable=sys.executable, bits="", linkage="")`

Queries the given executable (defaults to the Python interpreter binary) for various architecture information.

Returns a tuple (*bits*, *linkage*) which contain information about the bit architecture and the linkage format used for the executable. Both values are returned as strings.

Values that cannot be determined are returned as given by the parameter presets. If *bits* is given as `''`, the `sizeof(pointer)` (or `sizeof(long)` on Python version < 1.5.2) is used as indicator for the supported pointer size.

The function relies on the system's `file` command to do the actual work. This is available on most if not all Unix platforms and some non-Unix platforms and then only if the executable points to the Python interpreter. Reasonable defaults are used when the above needs are not met.

注解: On Mac OS X (and perhaps other platforms), executable files may be universal files containing multiple architectures.

To get at the “64-bitness” of the current interpreter, it is more reliable to query the `sys.maxsize` attribute:

```
is_64bits = sys.maxsize > 2**32
```

`platform.machine()`

Returns the machine type, e.g. `'i386'`. An empty string is returned if the value cannot be determined.

`platform.node()`

Returns the computer's network name (may not be fully qualified!). An empty string is returned if the value cannot be determined.

`platform.platform(aliased=0, terse=0)`

Returns a single string identifying the underlying platform with as much useful information as possible.

The output is intended to be *human readable* rather than machine parseable. It may look different on different platforms and this is intended.

If *aliased* is true, the function will use aliases for various platforms that report system names which differ from their common names, for example SunOS will be reported as Solaris. The `system_alias()` function is used to implement this.

Setting *terse* to true causes the function to return only the absolute minimum information needed to identify the platform.

`platform.processor()`

Returns the (real) processor name, e.g. 'amd64'.

An empty string is returned if the value cannot be determined. Note that many platforms do not provide this information or simply return the same value as for `machine()`. NetBSD does this.

`platform.python_build()`

Returns a tuple (buildno, builddate) stating the Python build number and date as strings.

`platform.python_compiler()`

Returns a string identifying the compiler used for compiling Python.

`platform.python_branch()`

Returns a string identifying the Python implementation SCM branch.

`platform.python_implementation()`

Returns a string identifying the Python implementation. Possible return values are: 'CPython', 'IronPython', 'Jython', 'PyPy'.

`platform.python_revision()`

Returns a string identifying the Python implementation SCM revision.

`platform.python_version()`

Returns the Python version as string 'major.minor.patchlevel'.

Note that unlike the Python `sys.version`, the returned value will always include the patchlevel (it defaults to 0).

`platform.python_version_tuple()`

Returns the Python version as tuple (major, minor, patchlevel) of strings.

Note that unlike the Python `sys.version`, the returned value will always include the patchlevel (it defaults to '0').

`platform.release()`

Returns the system's release, e.g. '2.2.0' or 'NT'. An empty string is returned if the value cannot be determined.

`platform.system()`

返回系统/OS 的名字, e.g. 'Linux', 'Windows', 或者 'Java'。如果这个值不能被确定则会返回一个空字符串。

`platform.system_alias(system, release, version)`

Returns (system, release, version) aliased to common marketing names used for some systems. It also does some reordering of the information in some cases where it would otherwise cause confusion.

`platform.version()`

Returns the system's release version, e.g. '#3 on degas'. An empty string is returned if the value cannot be determined.

`platform.uname()`

Fairly portable `uname` interface. Returns a `namedtuple()` containing six attributes: `system`, `node`, `release`, `version`, `machine`, and `processor`.

Note that this adds a sixth attribute (`processor`) not present in the `os.uname()` result. Also, the attribute names are different for the first two attributes; `os.uname()` names them `sysname` and `nodename`.

Entries which cannot be determined are set to `' '`.

在 3.3 版更改: 将结果从元组改为命名元组。

16.14.2 Java 平台

`platform.java_ver(release="", vendor="", vminfo=("", "", ""), osinfo=("", "", ""))`

Jython 的版本接口

Returns a tuple (`release`, `vendor`, `vminfo`, `osinfo`) with `vminfo` being a tuple (`vm_name`, `vm_release`, `vm_vendor`) and `osinfo` being a tuple (`os_name`, `os_version`, `os_arch`). Values which cannot be determined are set to the defaults given as parameters (which all default to `' '`).

16.14.3 Windows 平台

`platform.win32_ver(release="", version="", csd="", ptype="")`

Get additional version information from the Windows Registry and return a tuple (`release`, `version`, `csd`, `ptype`) referring to OS release, version number, CSD level (service pack) and OS type (multi/single processor).

As a hint: `ptype` is `'Uniprocessor Free'` on single processor NT machines and `'Multiprocessor Free'` on multi processor machines. The `'Free'` refers to the OS version being free of debugging code. It could also state `'Checked'` which means the OS version uses debugging code, i.e. code that checks arguments, ranges, etc.

注解: This function works best with Mark Hammond's `win32all` package installed, but also on Python 2.3 and later (support for this was added in Python 2.6). It obviously only runs on Win32 compatible platforms.

Win95/98 specific

`platform.popen(cmd, mode='r', bufsize=-1)`

Portable `popen()` interface. Find a working `popen` implementation preferring `win32pipe.popen()`. On Windows NT, `win32pipe.popen()` should work; on Windows 9x it hangs due to bugs in the MS C library.

3.3 版后已移除: This function is obsolete. Use the `subprocess` module. Check especially the [Replacing Older Functions with the subprocess Module](#) section.

16.14.4 Mac OS 平台

`platform.mac_ver(release="", versioninfo=("", "", ""), machine="")`

Get Mac OS version information and return it as tuple (`release`, `versioninfo`, `machine`) with `versioninfo` being a tuple (`version`, `dev_stage`, `non_release_version`).

Entries which cannot be determined are set to `' '`. All tuple entries are strings.

16.14.5 Unix Platforms

`platform.dist` (*distname*=", *version*", *id*", *supported_dists*=('SuSE', 'debian', 'redhat', 'mandrake', ...))
这是 `linux_distribution()` 的另外一个名字。

Deprecated since version 3.5, will be removed in version 3.8: See alternative like the `distro` package.

`platform.linux_distribution` (*distname*=", *version*", *id*", *supported_dists*=('SuSE', 'debian', 'redhat', 'mandrake', ...), *full_distribution_name*=1)

Tries to determine the name of the Linux OS distribution name.

supported_dists may be given to define the set of Linux distributions to look for. It defaults to a list of currently supported Linux distributions identified by their release file name.

If *full_distribution_name* is true (default), the full distribution read from the OS is returned. Otherwise the short name taken from *supported_dists* is used.

Returns a tuple (*distname*, *version*, *id*) which defaults to the args given as parameters. *id* is the item in parentheses after the version number. It is usually the version codename.

Deprecated since version 3.5, will be removed in version 3.8: See alternative like the `distro` package.

`platform.libc_ver` (*executable*=`sys.executable`, *lib*", *version*", *chunksize*=16384)

Tries to determine the libc version against which the file *executable* (defaults to the Python interpreter) is linked. Returns a tuple of strings (*lib*, *version*) which default to the given parameters in case the lookup fails.

Note that this function has intimate knowledge of how different libc versions add symbols to the executable is probably only usable for executables compiled using `gcc`.

The file is read and scanned in chunks of *chunksize* bytes.

16.15 `errno` —Standard `errno` system symbols

This module makes available standard `errno` system symbols. The value of each symbol is the corresponding integer value. The names and descriptions are borrowed from `linux/include/errno.h`, which should be pretty all-inclusive.

`errno.errorcode`

Dictionary providing a mapping from the `errno` value to the string name in the underlying system. For instance, `errno.errorcode[errno.EPERM]` maps to `'EPERM'`.

To translate a numeric error code to an error message, use `os.strerror()`.

Of the following list, symbols that are not used on the current platform are not defined by the module. The specific list of defined symbols is available as `errno.errorcode.keys()`. Symbols available can include:

`errno.EPERM`

Operation not permitted

`errno.ENOENT`

No such file or directory

`errno.ESRCH`

No such process

`errno.EINTR`

Interrupted system call.

参见:

This error is mapped to the exception *InterruptedError*.

errno.EIO
I/O error

errno.ENXIO
No such device or address

errno.E2BIG
Arg list too long

errno.ENOEXEC
Exec format error

errno.EBADF
Bad file number

errno.ECHILD
No child processes

errno.EAGAIN
Try again

errno.ENOMEM
Out of memory

errno.EACCES
Permission denied

errno.EFAULT
Bad address

errno.ENOTBLK
Block device required

errno.EBUSY
Device or resource busy

errno.EEXIST
File exists

errno.EXDEV
Cross-device link

errno.ENODEV
No such device

errno.ENOTDIR
Not a directory

errno.EISDIR
Is a directory

errno.EINVAL
Invalid argument

errno.ENFILE
File table overflow

errno.EMFILE
Too many open files

errno.ENOTTY
Not a typewriter

`errno.ETXTBSY`
Text file busy

`errno.EFBIG`
File too large

`errno.ENOSPC`
No space left on device

`errno.ESPIPE`
Illegal seek

`errno.EROFS`
Read-only file system

`errno.EMLINK`
Too many links

`errno.EPIPE`
Broken pipe

`errno.EDOM`
Math argument out of domain of func

`errno.ERANGE`
Math result not representable

`errno.EDEADLK`
Resource deadlock would occur

`errno.ENAMETOOLONG`
File name too long

`errno.ENOLCK`
No record locks available

`errno.ENOSYS`
Function not implemented

`errno.ENOTEMPTY`
Directory not empty

`errno.ELOOP`
Too many symbolic links encountered

`errno.EWOULDBLOCK`
Operation would block

`errno.ENOMSG`
No message of desired type

`errno.EIDRM`
Identifier removed

`errno.ECHRNG`
Channel number out of range

`errno.EL2NSYNC`
Level 2 not synchronized

`errno.EL3HLT`
Level 3 halted

`errno.EL3RST`
Level 3 reset

`errno.ELNRNG`
Link number out of range

`errno.EUNATCH`
Protocol driver not attached

`errno.ENOCSI`
No CSI structure available

`errno.EL2HLT`
Level 2 halted

`errno.EBADE`
Invalid exchange

`errno.EBADR`
Invalid request descriptor

`errno.EXFULL`
Exchange full

`errno.ENOANO`
No anode

`errno.EBADRQC`
Invalid request code

`errno.EBADSLT`
Invalid slot

`errno.EDEADLOCK`
File locking deadlock error

`errno.EBFONT`
Bad font file format

`errno.ENOSTR`
Device not a stream

`errno.ENODATA`
No data available

`errno.ETIME`
Timer expired

`errno.ENOSR`
Out of streams resources

`errno.ENONET`
Machine is not on the network

`errno.ENOPKG`
Package not installed

`errno.EREMOTE`
Object is remote

`errno.ENOLINK`
Link has been severed

`errno.EADV`
Advertise error

`errno.ESRMNT`
Srmount error

`errno.ECOMM`
Communication error on send

`errno.EPROTO`
Protocol error

`errno.EMULTIHOP`
Multihop attempted

`errno.EDOTDOT`
RFS specific error

`errno.EBADMSG`
Not a data message

`errno.EOVERFLOW`
Value too large for defined data type

`errno.ENOTUNIQ`
Name not unique on network

`errno.EBADFD`
File descriptor in bad state

`errno.EREMCHG`
Remote address changed

`errno.ELIBACC`
Can not access a needed shared library

`errno.ELIBBAD`
Accessing a corrupted shared library

`errno.ELIBSCN`
.lib section in a.out corrupted

`errno.ELIBMAX`
Attempting to link in too many shared libraries

`errno.ELIBEXEC`
Cannot exec a shared library directly

`errno.EILSEQ`
Illegal byte sequence

`errno.ERESTART`
Interrupted system call should be restarted

`errno.ESTRPIPE`
Streams pipe error

`errno.EUSERS`
Too many users

`errno.ENOTSOCK`
Socket operation on non-socket

`errno.EDESTADDRREQ`
Destination address required

`errno.EMSGSIZE`
Message too long

`errno.EPROTOTYPE`
Protocol wrong type for socket

`errno.ENOPROTOOPT`
Protocol not available

`errno.EPROTONOSUPPORT`
Protocol not supported

`errno.ESOCKTNOSUPPORT`
Socket type not supported

`errno.EOPNOTSUPP`
Operation not supported on transport endpoint

`errno.EPFNOSUPPORT`
Protocol family not supported

`errno.EAFNOSUPPORT`
Address family not supported by protocol

`errno.EADDRINUSE`
Address already in use

`errno.EADDRNOTAVAIL`
Cannot assign requested address

`errno.ENETDOWN`
Network is down

`errno.ENETUNREACH`
Network is unreachable

`errno.ENETRESET`
Network dropped connection because of reset

`errno.ECONNABORTED`
Software caused connection abort

`errno.ECONNRESET`
Connection reset by peer

`errno.ENOBUFS`
No buffer space available

`errno.EISCONN`
Transport endpoint is already connected

`errno.ENOTCONN`
Transport endpoint is not connected

`errno.ESHUTDOWN`
Cannot send after transport endpoint shutdown

`errno.ETOOMANYREFS`
Too many references: cannot splice

```

errno.ETIMEDOUT
    Connection timed out

errno.ECONNREFUSED
    Connection refused

errno.EHOSTDOWN
    Host is down

errno.EHOSTUNREACH
    No route to host

errno.EALREADY
    Operation already in progress

errno.EINPROGRESS
    Operation now in progress

errno.ESTALE
    Stale NFS file handle

errno.EUCLEAN
    Structure needs cleaning

errno.ENOTNAM
    Not a XENIX named type file

errno.ENAVAIL
    No XENIX semaphores available

errno.EISNAM
    Is a named type file

errno.EREMOTEIO
    Remote I/O error

errno.EDQUOT
    Quota exceeded

```

16.16 ctypes —Python 的外部函数库

`ctypes` 是 Python 的外部函数库。它提供了与 C 兼容的数据类型，并允许调用 DLL 或共享库中的函数。可使用该模块以纯 Python 形式对这些库进行封装。

16.16.1 ctypes 教程

注意：在本教程中的示例代码使用 `doctest` 进行过测试，保证其正确运行。由于有些代码在 Linux, Windows 或 Mac OS X 下的表现不同，这些代码会在 `doctest` 中包含相关的指令注解。

注意：部分示例代码引用了 `ctypes c_int` 类型。在 `sizeof(long) == sizeof(int)` 的平台上此类型是 `c_long` 的一个别名。所以，在程序输出 `c_long` 而不是你期望的 `c_int` 时不必感到迷惑——它们实际上是同一种类型。

载入动态链接库

`ctypes` 导出了 `cdll` 对象，在 Windows 系统中还导出了 `windll` 和 `oledll` 对象用于载入动态链接库。

通过操作这些对象的属性，你可以载入外部的动态链接库。`cdll` 载入按标准的 `cdecl` 调用协议导出的函数，而 `windll` 导入的库按 `stdcall` 调用协议调用其中的函数。`oledll` 也按 `stdcall` 调用协议调用其中的函数，并假定该函数返回的是 Windows HRESULT 错误代码，并当函数调用失败时，自动根据该代码甩出一个 `OSError` 异常。

在 3.3 版更改：原来在 Windows 下甩出的异常类型 `WindowsError` 现在是 `OSError` 的一个别名。

这是一些 Windows 下的例子。注意：`msvcrt` 是微软 C 标准库，包含了大部分 C 标准函数，这些函数都是以 `cdecl` 调用协议进行调用的。

```
>>> from ctypes import *
>>> print(windll.kernel32)
<WinDLL 'kernel32', handle ... at ...>
>>> print(cdll.msvcrt)
<CDLL 'msvcrt', handle ... at ...>
>>> libc = cdll.msvcrt
>>>
```

Windows 会自动添加通常的 `.dll` 文件扩展名。

注解：通过 `cdll.msvcrt` 调用的标准 C 函数，可能会导致调用一个过时的，与当前 Python 所不兼容的函数。因此，请尽量使用标准的 Python 函数，而不要使用 `msvcrt` 模块。

在 Linux 下，必须使用包含文件扩展名的文件名来导入共享库。因此不能简单使用对象属性的方式来导入库。因此，你可以使用方法 `LoadLibrary()`，或构造 `CDLL` 对象来导入库。

```
>>> cdll.LoadLibrary("libc.so.6")
<CDLL 'libc.so.6', handle ... at ...>
>>> libc = CDLL("libc.so.6")
>>> libc
<CDLL 'libc.so.6', handle ... at ...>
>>>
```

操作导入的动态链接库中的函数

通过操作 `dll` 对象的属性来操作这些函数。

```
>>> from ctypes import *
>>> libc.printf
<_FuncPtr object at 0x...>
>>> print(windll.kernel32.GetModuleHandleA)
<_FuncPtr object at 0x...>
>>> print(windll.kernel32.MyOwnFunction)
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
  File "ctypes.py", line 239, in __getattr__
    func = _StdcallFuncPtr(name, self)
AttributeError: function 'MyOwnFunction' not found
>>>
```

注意：Win32 系统的动态库，比如 `kernel32` 和 `user32`，通常会同时导出同一个函数的 ANSI 版本和 UNICODE 版本。UNICODE 版本通常会在名字最后以 `w` 结尾，而 ANSI 版本的则以 `A` 结尾。win32 的

GetModuleHandle 函数会根据一个模块名返回一个 模块句柄，该函数暨同时包含这样的两个版本的原型函数，并通过宏 UNICODE 是否定义，来决定宏 GetModuleHandle 导出的是哪个具体函数。

```
/* ANSI version */
HMODULE GetModuleHandleA(LPCSTR lpModuleName);
/* UNICODE version */
HMODULE GetModuleHandleW(LPCWSTR lpModuleName);
```

windll 不会通过这样的魔法手段来帮你决定选择哪一种函数，你必须显式的调用 GetModuleHandleA 或 GetModuleHandleW，并分别使用字节对象或字符串对象作参数。

有时候，dlls 的导出的函数名不符合 Python 的标识符规范，比如 "??2@YAPAXI@Z"。此时，你必须使用 `getattr()` 方法来获得该函数。

```
>>> getattr(cdll.msvcrt, "??2@YAPAXI@Z")
<_FuncPtr object at 0x...>
>>>
```

Windows 下，有些 dll 导出的函数没有函数名，而是通过其顺序号调用。对此类函数，你也可以通过 dll 对象的数值索引来操作这些函数。

```
>>> cdll.kernel32[1]
<_FuncPtr object at 0x...>
>>> cdll.kernel32[0]
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
  File "ctypes.py", line 310, in __getitem__
    func = _StdcallFuncPtr(name, self)
AttributeError: function ordinal 0 not found
>>>
```

调用函数

你可以貌似是调用其它 Python 函数那样直接调用这些函数。在这个例子中，我们调用了 `time()` 函数，该函数返回一个系统时间戳（从 Unix 时间起点到现在的秒数），而“`GetModuleHandleA()`”函数返回一个 win32 模块句柄。

此函数中调用的两个函数都使用了空指针（用 `None` 作为空指针）：

```
>>> print(libc.time(None))
1150640792
>>> print(hex(windll.kernel32.GetModuleHandleA(None)))
0x1d000000
>>>
```

注解：调用该函数，若 `ctypes` 发现传入的参数个数不符，则会甩出一个异常 `ValueError`。但该行为并不可靠。它在 3.6.2 中被废弃，会在 3.7 中彻底移除。

如果你用 `cdecl` 调用方式调用 `stdcall` 约定的函数，则会甩出一个异常 `ValueError`。反之亦然。

```
>>> cdll.kernel32.GetModuleHandleA(None)
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
ValueError: Procedure probably called with not enough arguments (4 bytes missing)
```

(下页继续)

(续上页)

```
>>>
>>> windll.msvcrt.printf(b"spam")
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
ValueError: Procedure probably called with too many arguments (4 bytes in excess)
>>>
```

你必须阅读这些库的头文件或说明文档来确定它们的调用协议。

在 Windows 中, `ctypes` 使用 win32 结构化异常处理来防止由于在调用函数时使用非法参数导致的程序崩溃。

```
>>> windll.kernel32.GetModuleHandleA(32)
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
OSError: exception: access violation reading 0x00000020
>>>
```

然而, 总有许多办法, 通过调用 `ctypes` 使得 Python 程序崩溃。因此, 你必须小心使用。 `faulthandler` 模块可以用于帮助诊断程序崩溃的原因。(比如由于错误的 C 库函数调用导致的段错误)。

None, 整型, 字节对象和 (UNICODE) 字符串是仅有的可以直接作为函数参数使用的四种 Python 本地数据类型。None 作为 C 的空指针 (NULL), 字节和字符串类型作为一个指向其保存数据的内存块指针 (char * 或 wchar_t *)。Python 的整型则作为平台默认的 C 的 int 类型, 他们的数值被截断以适应 C 类型的整型长度。

在我们开始调用函数前, 我们必须先了解作为函数参数的 `ctypes` 数据类型。

基础数据类型

`ctypes` 定义了一些和 C 兼容的基本数据类型:

ctypes 类型	C 类型	Python 数据类型
<code>c_bool</code>	<code>_Bool</code>	<code>bool</code> (1)
<code>c_char</code>	<code>char</code>	单字符字节对象
<code>c_wchar</code>	<code>wchar_t</code>	单字符字符串
<code>c_byte</code>	<code>char</code>	<code>int</code>
<code>c_ubyte</code>	<code>unsigned char</code>	<code>int</code>
<code>c_short</code>	<code>short</code>	<code>int</code>
<code>c_ushort</code>	<code>unsigned short</code>	<code>int</code>
<code>c_int</code>	<code>int</code>	<code>int</code>
<code>c_uint</code>	<code>unsigned int</code>	<code>int</code>
<code>c_long</code>	<code>long</code>	<code>int</code>
<code>c_ulong</code>	<code>unsigned long</code>	<code>int</code>
<code>c_longlong</code>	<code>__int64</code> 或 <code>long long</code>	<code>int</code>
<code>c_ulonglong</code>	<code>unsigned __int64</code> 或 <code>unsigned long long</code>	<code>int</code>
<code>c_size_t</code>	<code>size_t</code>	<code>int</code>
<code>c_ssize_t</code>	<code>ssize_t</code> 或 <code>Py_ssize_t</code>	<code>int</code>
<code>c_float</code>	<code>float</code>	<code>float</code>
<code>c_double</code>	<code>double</code>	<code>float</code>
<code>c_longdouble</code>	<code>long double</code>	<code>float</code>
<code>c_char_p</code>	<code>char *</code> (NUL terminated)	字节串对象或 <code>None</code>
<code>c_wchar_p</code>	<code>wchar_t *</code> (NUL terminated)	字符串或 <code>None</code>
<code>c_void_p</code>	<code>void *</code>	<code>int</code> 或 <code>None</code>

(1) 构造函数接受任何具有真值的对象。

所有这些类型都可以通过使用正确类型和值的可选初始值调用它们来创建:

```
>>> c_int()
c_long(0)
>>> c_wchar_p("Hello, World")
c_wchar_p(140018365411392)
>>> c_ushort(-3)
c_ushort(65533)
>>>
```

由于这些类型是可变的，它们的值也可以在以后更改:

```
>>> i = c_int(42)
>>> print(i)
c_long(42)
>>> print(i.value)
42
>>> i.value = -99
>>> print(i.value)
-99
>>>
```

当给指针类型的对象 `c_char_p`, `c_wchar_p` 和 `c_void_p` 等赋值时，将改变它们所指向的内存地址，而不是它们所指向的内存区域的内容 (这是理所当然的，因为 Python 的 `bytes` 对象是不可变的):

```
>>> s = "Hello, World"
>>> c_s = c_wchar_p(s)
>>> print(c_s)
c_wchar_p(139966785747344)
>>> print(c_s.value)
```

(下页继续)

(续上页)

```

Hello World
>>> c_s.value = "Hi, there"
>>> print(c_s)           # the memory location has changed
c_wchar_p(139966783348904)
>>> print(c_s.value)
Hi, there
>>> print(s)             # first object is unchanged
Hello, World
>>>

```

但你要注意不能将它们传递给会改变指针所指内存的函数。如果你需要可改变的内存块，`ctypes` 提供了 `create_string_buffer()` 函数，它提供多种方式创建这种内存块。当前的内存块内容可以通过 `raw` 属性存取，如果你希望将它作为 NUL 结束的字符串，请使用 `value` 属性：

```

>>> from ctypes import *
>>> p = create_string_buffer(3)           # create a 3 byte buffer, initialized to_
↳ NUL bytes
>>> print(sizeof(p), repr(p.raw))
3 b'\x00\x00\x00'
>>> p = create_string_buffer(b"Hello")    # create a buffer containing a NUL_
↳ terminated string
>>> print(sizeof(p), repr(p.raw))
6 b'Hello\x00'
>>> print(repr(p.value))
b'Hello'
>>> p = create_string_buffer(b"Hello", 10) # create a 10 byte buffer
>>> print(sizeof(p), repr(p.raw))
10 b'Hello\x00\x00\x00\x00\x00'
>>> p.value = b"Hi"
>>> print(sizeof(p), repr(p.raw))
10 b'Hi\x00lo\x00\x00\x00\x00'
>>>

```

`create_string_buffer()` 函数替代以前的 `ctypes` 版本中的 `c_buffer()` 函数（仍然可当作别名使用）和 `c_string()` 函数。`create_unicode_buffer()` 函数创建包含 `unicode` 字符的可变内存块，与之对应的 C 语言类型是 `wchar_t`。

调用函数，继续

注意 `printf` 将打印到真正标准输出设备，而 `*` 不是 `* sys.stdout`，因此这些实例只能在控制台提示符下工作，而不能在 `IDLE` 或 `PythonWin` 中运行。

```

>>> printf = libc.printf
>>> printf(b"Hello, %s\n", b"World!")
Hello, World!
14
>>> printf(b"Hello, %S\n", "World!")
Hello, World!
14
>>> printf(b"%d bottles of beer\n", 42)
42 bottles of beer
19
>>> printf(b"%f bottles of beer\n", 42.5)
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>

```

(下页继续)

(续上页)

```
ArgumentError: argument 2: exceptions.TypeError: Don't know how to convert parameter 2
>>>
```

正如前面所提到过的, 除了整数、字符串以及字节串之外, 所有的 Python 类型都必须使用它们对应的 *ctypes* 类型包装, 才能够被正确地转换为所需的 C 语言类型。

```
>>> printf(b"An int %d, a double %f\n", 1234, c_double(3.14))
An int 1234, a double 3.140000
31
>>>
```

使用自定义的数据类型调用函数

You can also customize *ctypes* argument conversion to allow instances of your own classes be used as function arguments. *ctypes* looks for an `_as_parameter_` attribute and uses this as the function argument. Of course, it must be one of integer, string, or bytes:

```
>>> class Bottles:
...     def __init__(self, number):
...         self._as_parameter_ = number
...
>>> bottles = Bottles(42)
>>> printf(b"%d bottles of beer\n", bottles)
42 bottles of beer
19
>>>
```

If you don't want to store the instance's data in the `_as_parameter_` instance variable, you could define a *property* which makes the attribute available on request.

Specifying the required argument types (function prototypes)

It is possible to specify the required argument types of functions exported from DLLs by setting the `argtypes` attribute.

`argtypes` must be a sequence of C data types (the `printf` function is probably not a good example here, because it takes a variable number and different types of parameters depending on the format string, on the other hand this is quite handy to experiment with this feature):

```
>>> printf.argtypes = [c_char_p, c_char_p, c_int, c_double]
>>> printf(b"String '%s', Int %d, Double %f\n", b"Hi", 10, 2.2)
String 'Hi', Int 10, Double 2.200000
37
>>>
```

Specifying a format protects against incompatible argument types (just as a prototype for a C function), and tries to convert the arguments to valid types:

```
>>> printf(b"%d %d %d", 1, 2, 3)
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
ArgumentError: argument 2: exceptions.TypeError: wrong type
>>> printf(b"%s %d %f\n", b"X", 2, 3)
X 2 3.000000
```

(下页继续)

(续上页)

```
13
>>>
```

If you have defined your own classes which you pass to function calls, you have to implement a `from_param()` class method for them to be able to use them in the `argtypes` sequence. The `from_param()` class method receives the Python object passed to the function call, it should do a typecheck or whatever is needed to make sure this object is acceptable, and then return the object itself, its `_as_parameter_` attribute, or whatever you want to pass as the C function argument in this case. Again, the result should be an integer, string, bytes, a `ctypes` instance, or an object with an `_as_parameter_` attribute.

Return types

By default functions are assumed to return the C `int` type. Other return types can be specified by setting the `restype` attribute of the function object.

Here is a more advanced example, it uses the `strchr` function, which expects a string pointer and a char, and returns a pointer to a string:

```
>>> strchr = libc.strchr
>>> strchr(b"abcdef", ord("d"))
8059983
>>> strchr.restype = c_char_p      # c_char_p is a pointer to a string
>>> strchr(b"abcdef", ord("d"))
b'def'
>>> print(strchr(b"abcdef", ord("x")))
None
>>>
```

If you want to avoid the `ord("x")` calls above, you can set the `argtypes` attribute, and the second argument will be converted from a single character Python bytes object into a C char:

```
>>> strchr.restype = c_char_p
>>> strchr.argtypes = [c_char_p, c_char]
>>> strchr(b"abcdef", b"d")
'def'
>>> strchr(b"abcdef", b"def")
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
ArgumentError: argument 2: exceptions.TypeError: one character string expected
>>> print(strchr(b"abcdef", b"x"))
None
>>> strchr(b"abcdef", b"d")
'def'
>>>
```

You can also use a callable Python object (a function or a class for example) as the `restype` attribute, if the foreign function returns an integer. The callable will be called with the *integer* the C function returns, and the result of this call will be used as the result of your function call. This is useful to check for error return values and automatically raise an exception:

```
>>> GetModuleHandle = windll.kernel32.GetModuleHandleA
>>> def ValidHandle(value):
...     if value == 0:
...         raise WinError()
...     return value
```

(下页继续)

(续上页)

```

...
>>>
>>> GetModuleHandle.restype = ValidHandle
>>> GetModuleHandle(None)
486539264
>>> GetModuleHandle("something silly")
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
  File "<stdin>", line 3, in ValidHandle
OSError: [Errno 126] The specified module could not be found.
>>>

```

WinError is a function which will call Windows `FormatMessage()` api to get the string representation of an error code, and *returns* an exception. WinError takes an optional error code parameter, if no one is used, it calls `GetLastError()` to retrieve it.

Please note that a much more powerful error checking mechanism is available through the `errcheck` attribute; see the reference manual for details.

Passing pointers (or: passing parameters by reference)

Sometimes a C api function expects a *pointer* to a data type as parameter, probably to write into the corresponding location, or if the data is too large to be passed by value. This is also known as *passing parameters by reference*.

`ctypes` exports the `byref()` function which is used to pass parameters by reference. The same effect can be achieved with the `pointer()` function, although `pointer()` does a lot more work since it constructs a real pointer object, so it is faster to use `byref()` if you don't need the pointer object in Python itself:

```

>>> i = c_int()
>>> f = c_float()
>>> s = create_string_buffer(b'\000' * 32)
>>> print(i.value, f.value, repr(s.value))
0 0.0 b''
>>> libc.sscanf(b"1 3.14 Hello", b"%d %f %s",
...             byref(i), byref(f), s)
3
>>> print(i.value, f.value, repr(s.value))
1 3.1400001049 b'Hello'
>>>

```

Structures and unions

Structures and unions must derive from the `Structure` and `Union` base classes which are defined in the `ctypes` module. Each subclass must define a `_fields_` attribute. `_fields_` must be a list of 2-tuples, containing a *field name* and a *field type*.

The field type must be a `ctypes` type like `c_int`, or any other derived `ctypes` type: structure, union, array, pointer.

Here is a simple example of a POINT structure, which contains two integers named *x* and *y*, and also shows how to initialize a structure in the constructor:

```

>>> from ctypes import *
>>> class POINT(Structure):
...     _fields_ = [("x", c_int),

```

(下页继续)

(续上页)

```

...         ("y", c_int)]
...
>>> point = POINT(10, 20)
>>> print(point.x, point.y)
10 20
>>> point = POINT(y=5)
>>> print(point.x, point.y)
0 5
>>> POINT(1, 2, 3)
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
ValueError: too many initializers
>>>

```

You can, however, build much more complicated structures. A structure can itself contain other structures by using a structure as a field type.

Here is a RECT structure which contains two POINTs named *upperleft* and *lowerright*:

```

>>> class RECT(Structure):
...     _fields_ = [("upperleft", POINT),
...                 ("lowerright", POINT)]
...
>>> rc = RECT(point)
>>> print(rc.upperleft.x, rc.upperleft.y)
0 5
>>> print(rc.lowerright.x, rc.lowerright.y)
0 0
>>>

```

Nested structures can also be initialized in the constructor in several ways:

```

>>> r = RECT(POINT(1, 2), POINT(3, 4))
>>> r = RECT((1, 2), (3, 4))

```

Field *descriptors* can be retrieved from the *class*, they are useful for debugging because they can provide useful information:

```

>>> print(POINT.x)
<Field type=c_long, ofs=0, size=4>
>>> print(POINT.y)
<Field type=c_long, ofs=4, size=4>
>>>

```

警告: *ctypes* does not support passing unions or structures with bit-fields to functions by value. While this may work on 32-bit x86, it's not guaranteed by the library to work in the general case. Unions and structures with bit-fields should always be passed to functions by pointer.

Structure/union alignment and byte order

By default, Structure and Union fields are aligned in the same way the C compiler does it. It is possible to override this behavior by specifying a `_pack_` class attribute in the subclass definition. This must be set to a positive integer and specifies the maximum alignment for the fields. This is what `#pragma pack(n)` also does in MSVC.

`ctypes` uses the native byte order for Structures and Unions. To build structures with non-native byte order, you can use one of the `BigEndianStructure`, `LittleEndianStructure`, `BigEndianUnion`, and `LittleEndianUnion` base classes. These classes cannot contain pointer fields.

Bit fields in structures and unions

It is possible to create structures and unions containing bit fields. Bit fields are only possible for integer fields, the bit width is specified as the third item in the `_fields_` tuples:

```
>>> class Int(Structure):
...     _fields_ = [("first_16", c_int, 16),
...                 ("second_16", c_int, 16)]
...
>>> print(Int.first_16)
<Field type=c_long, ofs=0:0, bits=16>
>>> print(Int.second_16)
<Field type=c_long, ofs=0:16, bits=16>
>>>
```

Arrays

Arrays are sequences, containing a fixed number of instances of the same type.

The recommended way to create array types is by multiplying a data type with a positive integer:

```
TenPointsArrayType = POINT * 10
```

Here is an example of a somewhat artificial data type, a structure containing 4 POINTs among other stuff:

```
>>> from ctypes import *
>>> class POINT(Structure):
...     _fields_ = ("x", c_int), ("y", c_int)
...
>>> class MyStruct(Structure):
...     _fields_ = [("a", c_int),
...                 ("b", c_float),
...                 ("point_array", POINT * 4)]
>>>
>>> print(len(MyStruct().point_array))
4
>>>
```

Instances are created in the usual way, by calling the class:

```
arr = TenPointsArrayType()
for pt in arr:
    print(pt.x, pt.y)
```

The above code print a series of 0 0 lines, because the array contents is initialized to zeros.

Initializers of the correct type can also be specified:

```
>>> from ctypes import *
>>> TenIntegers = c_int * 10
>>> ii = TenIntegers(1, 2, 3, 4, 5, 6, 7, 8, 9, 10)
```

(下页继续)

(续上页)

```
>>> print(ii)
<c_long_Array_10 object at 0x...>
>>> for i in ii: print(i, end=" ")
...
1 2 3 4 5 6 7 8 9 10
>>>
```

Pointers

Pointer instances are created by calling the `pointer()` function on a `ctypes` type:

```
>>> from ctypes import *
>>> i = c_int(42)
>>> pi = pointer(i)
>>>
```

Pointer instances have a `contents` attribute which returns the object to which the pointer points, the `i` object above:

```
>>> pi.contents
c_long(42)
>>>
```

Note that `ctypes` does not have OOR (original object return), it constructs a new, equivalent object each time you retrieve an attribute:

```
>>> pi.contents is i
False
>>> pi.contents is pi.contents
False
>>>
```

Assigning another `c_int` instance to the pointer's `contents` attribute would cause the pointer to point to the memory location where this is stored:

```
>>> i = c_int(99)
>>> pi.contents = i
>>> pi.contents
c_long(99)
>>>
```

Pointer instances can also be indexed with integers:

```
>>> pi[0]
99
>>>
```

Assigning to an integer index changes the pointed to value:

```
>>> print(i)
c_long(99)
>>> pi[0] = 22
>>> print(i)
c_long(22)
>>>
```

It is also possible to use indexes different from 0, but you must know what you're doing, just as in C: You can access or change arbitrary memory locations. Generally you only use this feature if you receive a pointer from a C function, and you *know* that the pointer actually points to an array instead of a single item.

Behind the scenes, the `pointer()` function does more than simply create pointer instances, it has to create pointer *types* first. This is done with the `POINTER()` function, which accepts any *ctypes* type, and returns a new type:

```
>>> PI = POINTER(c_int)
>>> PI
<class 'ctypes.LP_c_long'>
>>> PI(42)
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
TypeError: expected c_long instead of int
>>> PI(c_int(42))
<ctypes.LP_c_long object at 0x...>
>>>
```

Calling the pointer type without an argument creates a NULL pointer. NULL pointers have a `False` boolean value:

```
>>> null_ptr = POINTER(c_int)()
>>> print(bool(null_ptr))
False
>>>
```

ctypes checks for NULL when dereferencing pointers (but dereferencing invalid non-NULL pointers would crash Python):

```
>>> null_ptr[0]
Traceback (most recent call last):
  ....
ValueError: NULL pointer access
>>>

>>> null_ptr[0] = 1234
Traceback (most recent call last):
  ....
ValueError: NULL pointer access
>>>
```

Type conversions

Usually, *ctypes* does strict type checking. This means, if you have `POINTER(c_int)` in the `argtypes` list of a function or as the type of a member field in a structure definition, only instances of exactly the same type are accepted. There are some exceptions to this rule, where *ctypes* accepts other objects. For example, you can pass compatible array instances instead of pointer types. So, for `POINTER(c_int)`, *ctypes* accepts an array of `c_int`:

```
>>> class Bar(Structure):
...     _fields_ = [("count", c_int), ("values", POINTER(c_int))]
...
>>> bar = Bar()
>>> bar.values = (c_int * 3)(1, 2, 3)
>>> bar.count = 3
>>> for i in range(bar.count):
...     print(bar.values[i])
...
>>>
```

(下页继续)

(续上页)

```
1
2
3
>>>
```

In addition, if a function argument is explicitly declared to be a pointer type (such as `POINTER(c_int)`) in `argtypes`, an object of the pointed type (`c_int` in this case) can be passed to the function. `ctypes` will apply the required `byref()` conversion in this case automatically.

To set a `POINTER` type field to `NULL`, you can assign `None`:

```
>>> bar.values = None
>>>
```

Sometimes you have instances of incompatible types. In C, you can cast one type into another type. `ctypes` provides a `cast()` function which can be used in the same way. The `Bar` structure defined above accepts `POINTER(c_int)` pointers or `c_int` arrays for its `values` field, but not instances of other types:

```
>>> bar.values = (c_byte * 4)()
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
TypeError: incompatible types, c_byte_Array_4 instance instead of LP_c_long instance
>>>
```

For these cases, the `cast()` function is handy.

The `cast()` function can be used to cast a `ctypes` instance into a pointer to a different `ctypes` data type. `cast()` takes two parameters, a `ctypes` object that is or can be converted to a pointer of some kind, and a `ctypes` pointer type. It returns an instance of the second argument, which references the same memory block as the first argument:

```
>>> a = (c_byte * 4)()
>>> cast(a, POINTER(c_int))
<ctypes.LP_c_long object at ...>
>>>
```

So, `cast()` can be used to assign to the `values` field of `Bar` the structure:

```
>>> bar = Bar()
>>> bar.values = cast((c_byte * 4)(), POINTER(c_int))
>>> print(bar.values[0])
0
>>>
```

Incomplete Types

Incomplete Types are structures, unions or arrays whose members are not yet specified. In C, they are specified by forward declarations, which are defined later:

```
struct cell; /* forward declaration */

struct cell {
    char *name;
    struct cell *next;
};
```

The straightforward translation into `ctypes` code would be this, but it does not work:

```
>>> class cell(Structure):
...     _fields_ = [("name", c_char_p),
...                 ("next", POINTER(cell))]
...
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
  File "<stdin>", line 2, in cell
NameError: name 'cell' is not defined
>>>
```

because the new class `cell` is not available in the class statement itself. In `ctypes`, we can define the `cell` class and set the `_fields_` attribute later, after the class statement:

```
>>> from ctypes import *
>>> class cell(Structure):
...     pass
...
>>> cell._fields_ = [("name", c_char_p),
...                  ("next", POINTER(cell))]
>>>
```

Lets try it. We create two instances of `cell`, and let them point to each other, and finally follow the pointer chain a few times:

```
>>> c1 = cell()
>>> c1.name = "foo"
>>> c2 = cell()
>>> c2.name = "bar"
>>> c1.next = pointer(c2)
>>> c2.next = pointer(c1)
>>> p = c1
>>> for i in range(8):
...     print(p.name, end=" ")
...     p = p.next[0]
...
foo bar foo bar foo bar foo bar
>>>
```

Callback functions

`ctypes` allows creating C callable function pointers from Python callables. These are sometimes called *callback functions*.

First, you must create a class for the callback function. The class knows the calling convention, the return type, and the number and types of arguments this function will receive.

The `CFUNCTYPE()` factory function creates types for callback functions using the `cdecl` calling convention. On Windows, the `WINFUNCTYPE()` factory function creates types for callback functions using the `stdcall` calling convention.

Both of these factory functions are called with the result type as first argument, and the callback functions expected argument types as the remaining arguments.

I will present an example here which uses the standard C library's `qsort()` function, that is used to sort items with the help of a callback function. `qsort()` will be used to sort an array of integers:

```
>>> IntArray5 = c_int * 5
>>> ia = IntArray5(5, 1, 7, 33, 99)
>>> qsort = libc.qsort
>>> qsort.restype = None
>>>
```

`qsort()` must be called with a pointer to the data to sort, the number of items in the data array, the size of one item, and a pointer to the comparison function, the callback. The callback will then be called with two pointers to items, and it must return a negative integer if the first item is smaller than the second, a zero if they are equal, and a positive integer otherwise.

So our callback function receives pointers to integers, and must return an integer. First we create the `type` for the callback function:

```
>>> CMPFUNC = CFUNCTYPE(c_int, POINTER(c_int), POINTER(c_int))
>>>
```

To get started, here is a simple callback that shows the values it gets passed:

```
>>> def py_cmp_func(a, b):
...     print("py_cmp_func", a[0], b[0])
...     return 0
...
>>> cmp_func = CMPFUNC(py_cmp_func)
>>>
```

The result:

```
>>> qsort(ia, len(ia), sizeof(c_int), cmp_func)
py_cmp_func 5 1
py_cmp_func 33 99
py_cmp_func 7 33
py_cmp_func 5 7
py_cmp_func 1 7
>>>
```

Now we can actually compare the two items and return a useful result:

```
>>> def py_cmp_func(a, b):
...     print("py_cmp_func", a[0], b[0])
...     return a[0] - b[0]
...
>>>
>>> qsort(ia, len(ia), sizeof(c_int), CMPFUNC(py_cmp_func))
py_cmp_func 5 1
py_cmp_func 33 99
py_cmp_func 7 33
py_cmp_func 1 7
py_cmp_func 5 7
>>>
```

As we can easily check, our array is sorted now:

```
>>> for i in ia: print(i, end=" ")
...
1 5 7 33 99
>>>
```

The function factories can be used as decorator factories, so we may as well write:

```
>>> @CFUNCTYPE(c_int, POINTER(c_int), POINTER(c_int))
... def py_cmp_func(a, b):
...     print("py_cmp_func", a[0], b[0])
...     return a[0] - b[0]
...
>>> qsort(ia, len(ia), sizeof(c_int), py_cmp_func)
py_cmp_func 5 1
py_cmp_func 33 99
py_cmp_func 7 33
py_cmp_func 1 7
py_cmp_func 5 7
>>>
```

注解: Make sure you keep references to `CFUNCTYPE()` objects as long as they are used from C code. `ctypes` doesn't, and if you don't, they may be garbage collected, crashing your program when a callback is made.

Also, note that if the callback function is called in a thread created outside of Python's control (e.g. by the foreign code that calls the callback), `ctypes` creates a new dummy Python thread on every invocation. This behavior is correct for most purposes, but it means that values stored with `threading.local` will *not* survive across different callbacks, even when those calls are made from the same C thread.

Accessing values exported from dlls

Some shared libraries not only export functions, they also export variables. An example in the Python library itself is the `Py_OptimizeFlag`, an integer set to 0, 1, or 2, depending on the `-O` or `-OO` flag given on startup.

`ctypes` can access values like this with the `in_dll()` class methods of the type. `pythonapi` is a predefined symbol giving access to the Python C api:

```
>>> opt_flag = c_int.in_dll(pythonapi, "Py_OptimizeFlag")
>>> print(opt_flag)
c_long(0)
>>>
```

If the interpreter would have been started with `-O`, the sample would have printed `c_long(1)`, or `c_long(2)` if `-OO` would have been specified.

An extended example which also demonstrates the use of pointers accesses the `PyImport_FrozenModules` pointer exported by Python.

Quoting the docs for that value:

This pointer is initialized to point to an array of struct `_frozen` records, terminated by one whose members are all `NULL` or zero. When a frozen module is imported, it is searched in this table. Third-party code could play tricks with this to provide a dynamically created collection of frozen modules.

So manipulating this pointer could even prove useful. To restrict the example size, we show only how this table can be read with `ctypes`:

```
>>> from ctypes import *
>>>
>>> class struct_frozen(Structure):
...     _fields_ = [("name", c_char_p),
...                 ("code", POINTER(c_ubyte)),
```

(下页继续)

(续上页)

```
...         ("size", c_int)]
...
>>>
```

We have defined the struct `_frozen` data type, so we can get the pointer to the table:

```
>>> FrozenTable = POINTER(struct_frozen)
>>> table = FrozenTable.in_dll(pythonsapi, "PyImport_FrozenModules")
>>>
```

Since `table` is a pointer to the array of `struct_frozen` records, we can iterate over it, but we just have to make sure that our loop terminates, because pointers have no size. Sooner or later it would probably crash with an access violation or whatever, so it's better to break out of the loop when we hit the NULL entry:

```
>>> for item in table:
...     if item.name is None:
...         break
...     print(item.name.decode("ascii"), item.size)
...
_frozen_importlib 31764
_frozen_importlib_external 41499
__hello__ 161
__phello__ -161
__phello__.spam 161
>>>
```

The fact that standard Python has a frozen module and a frozen package (indicated by the negative size member) is not well known, it is only used for testing. Try it out with `import __hello__` for example.

Surprises

There are some edges in `ctypes` where you might expect something other than what actually happens.

Consider the following example:

```
>>> from ctypes import *
>>> class POINT(Structure):
...     _fields_ = ("x", c_int), ("y", c_int)
...
>>> class RECT(Structure):
...     _fields_ = ("a", POINT), ("b", POINT)
...
>>> p1 = POINT(1, 2)
>>> p2 = POINT(3, 4)
>>> rc = RECT(p1, p2)
>>> print(rc.a.x, rc.a.y, rc.b.x, rc.b.y)
1 2 3 4
>>> # now swap the two points
>>> rc.a, rc.b = rc.b, rc.a
>>> print(rc.a.x, rc.a.y, rc.b.x, rc.b.y)
3 4 3 4
>>>
```

Hm. We certainly expected the last statement to print `3 4 1 2`. What happened? Here are the steps of the `rc.a, rc.b = rc.b, rc.a` line above:

```
>>> temp0, temp1 = rc.b, rc.a
>>> rc.a = temp0
>>> rc.b = temp1
>>>
```

Note that `temp0` and `temp1` are objects still using the internal buffer of the `rc` object above. So executing `rc.a = temp0` copies the buffer contents of `temp0` into `rc`'s buffer. This, in turn, changes the contents of `temp1`. So, the last assignment `rc.b = temp1`, doesn't have the expected effect.

Keep in mind that retrieving sub-objects from Structure, Unions, and Arrays doesn't *copy* the sub-object, instead it retrieves a wrapper object accessing the root-object's underlying buffer.

Another example that may behave different from what one would expect is this:

```
>>> s = c_char_p()
>>> s.value = "abc def ghi"
>>> s.value
'abc def ghi'
>>> s.value is s.value
False
>>>
```

Why is it printing `False`? `ctypes` instances are objects containing a memory block plus some *descriptors* accessing the contents of the memory. Storing a Python object in the memory block does not store the object itself, instead the contents of the object is stored. Accessing the contents again constructs a new Python object each time!

Variable-sized data types

`ctypes` provides some support for variable-sized arrays and structures.

The `resize()` function can be used to resize the memory buffer of an existing `ctypes` object. The function takes the object as first argument, and the requested size in bytes as the second argument. The memory block cannot be made smaller than the natural memory block specified by the objects type, a `ValueError` is raised if this is tried:

```
>>> short_array = (c_short * 4)()
>>> print(sizeof(short_array))
8
>>> resize(short_array, 4)
Traceback (most recent call last):
...
ValueError: minimum size is 8
>>> resize(short_array, 32)
>>> sizeof(short_array)
32
>>> sizeof(type(short_array))
8
>>>
```

This is nice and fine, but how would one access the additional elements contained in this array? Since the type still only knows about 4 elements, we get errors accessing other elements:

```
>>> short_array[:]
[0, 0, 0, 0]
>>> short_array[7]
Traceback (most recent call last):
...
```

(下页继续)

(续上页)

```

IndexError: invalid index
>>>

```

Another way to use variable-sized data types with `ctypes` is to use the dynamic nature of Python, and (re-)define the data type after the required size is already known, on a case by case basis.

16.16.2 ctypes reference

Finding shared libraries

When programming in a compiled language, shared libraries are accessed when compiling/linking a program, and when the program is run.

The purpose of the `find_library()` function is to locate a library in a way similar to what the compiler or runtime loader does (on platforms with several versions of a shared library the most recent should be loaded), while the `ctypes` library loaders act like when a program is run, and call the runtime loader directly.

The `ctypes.util` module provides a function which can help to determine the library to load.

`ctypes.util.find_library(name)`

Try to find a library and return a pathname. *name* is the library name without any prefix like *lib*, suffix like *.so*, *.dylib* or version number (this is the form used for the *posix* linker option *-l*). If no library can be found, returns *None*.

The exact functionality is system dependent.

On Linux, `find_library()` tries to run external programs (`/sbin/ldconfig`, `gcc`, `objdump` and `ld`) to find the library file. It returns the filename of the library file.

在 3.6 版更改: On Linux, the value of the environment variable `LD_LIBRARY_PATH` is used when searching for libraries, if a library cannot be found by any other means.

Here are some examples:

```

>>> from ctypes.util import find_library
>>> find_library("m")
'libm.so.6'
>>> find_library("c")
'libc.so.6'
>>> find_library("bz2")
'libbz2.so.1.0'
>>>

```

On OS X, `find_library()` tries several predefined naming schemes and paths to locate the library, and returns a full pathname if successful:

```

>>> from ctypes.util import find_library
>>> find_library("c")
'/usr/lib/libc.dylib'
>>> find_library("m")
'/usr/lib/libm.dylib'
>>> find_library("bz2")
'/usr/lib/libbz2.dylib'
>>> find_library("AGL")
'/System/Library/Frameworks/AGL.framework/AGL'
>>>

```

On Windows, `find_library()` searches along the system search path, and returns the full pathname, but since there is no predefined naming scheme a call like `find_library("c")` will fail and return `None`.

If wrapping a shared library with `ctypes`, it *may* be better to determine the shared library name at development time, and hardcode that into the wrapper module instead of using `find_library()` to locate the library at runtime.

Loading shared libraries

There are several ways to load shared libraries into the Python process. One way is to instantiate one of the following classes:

class `ctypes.CDLL` (*name*, *mode*=`DEFAULT_MODE`, *handle*=`None`, *use_errno*=`False`,
 use_last_error=`False`)
 Instances of this class represent loaded shared libraries. Functions in these libraries use the standard C calling convention, and are assumed to return `int`.

class `ctypes.OleDLL` (*name*, *mode*=`DEFAULT_MODE`, *handle*=`None`, *use_errno*=`False`,
 use_last_error=`False`)
 Windows only: Instances of this class represent loaded shared libraries, functions in these libraries use the `stdcall` calling convention, and are assumed to return the windows specific `HRESULT` code. `HRESULT` values contain information specifying whether the function call failed or succeeded, together with additional error code. If the return value signals a failure, an `OSError` is automatically raised.

在 3.3 版更改: `WindowsError` used to be raised.

class `ctypes.WinDLL` (*name*, *mode*=`DEFAULT_MODE`, *handle*=`None`, *use_errno*=`False`,
 use_last_error=`False`)
 Windows only: Instances of this class represent loaded shared libraries, functions in these libraries use the `stdcall` calling convention, and are assumed to return `int` by default.

On Windows CE only the standard calling convention is used, for convenience the `WinDLL` and `OleDLL` use the standard calling convention on this platform.

The Python *global interpreter lock* is released before calling any function exported by these libraries, and reacquired afterwards.

class `ctypes.PyDLL` (*name*, *mode*=`DEFAULT_MODE`, *handle*=`None`)
 Instances of this class behave like `CDLL` instances, except that the Python GIL is *not* released during the function call, and after the function execution the Python error flag is checked. If the error flag is set, a Python exception is raised.

Thus, this is only useful to call Python C api functions directly.

All these classes can be instantiated by calling them with at least one argument, the pathname of the shared library. If you have an existing handle to an already loaded shared library, it can be passed as the `handle` named parameter, otherwise the underlying platforms `dlopen` or `LoadLibrary` function is used to load the library into the process, and to get a handle to it.

The *mode* parameter can be used to specify how the library is loaded. For details, consult the `dlopen(3)` manpage. On Windows, *mode* is ignored. On posix systems, `RTLD_NOW` is always added, and is not configurable.

The *use_errno* parameter, when set to true, enables a ctypes mechanism that allows accessing the system `errno` error number in a safe way. `ctypes` maintains a thread-local copy of the systems `errno` variable; if you call foreign functions created with `use_errno=True` then the `errno` value before the function call is swapped with the ctypes private copy, the same happens immediately after the function call.

The function `ctypes.get_errno()` returns the value of the ctypes private copy, and the function `ctypes.set_errno()` changes the ctypes private copy to a new value and returns the former value.

The `use_last_error` parameter, when set to true, enables the same mechanism for the Windows error code which is managed by the `GetLastError()` and `SetLastError()` Windows API functions; `ctypes.get_last_error()` and `ctypes.set_last_error()` are used to request and change the ctypes private copy of the windows error code.

`ctypes.RTLD_GLOBAL`

Flag to use as `mode` parameter. On platforms where this flag is not available, it is defined as the integer zero.

`ctypes.RTLD_LOCAL`

Flag to use as `mode` parameter. On platforms where this is not available, it is the same as `RTLD_GLOBAL`.

`ctypes.DEFAULT_MODE`

The default mode which is used to load shared libraries. On OSX 10.3, this is `RTLD_GLOBAL`, otherwise it is the same as `RTLD_LOCAL`.

Instances of these classes have no public methods. Functions exported by the shared library can be accessed as attributes or by index. Please note that accessing the function through an attribute caches the result and therefore accessing it repeatedly returns the same object each time. On the other hand, accessing it through an index returns a new object each time:

```
>>> libc.time == libc.time
True
>>> libc['time'] == libc['time']
False
```

The following public attributes are available, their name starts with an underscore to not clash with exported function names:

`PyDLL._handle`

The system handle used to access the library.

`PyDLL._name`

The name of the library passed in the constructor.

Shared libraries can also be loaded by using one of the prefabricated objects, which are instances of the `LibraryLoader` class, either by calling the `LoadLibrary()` method, or by retrieving the library as attribute of the loader instance.

class `ctypes.LibraryLoader` (*dlltype*)

Class which loads shared libraries. *dlltype* should be one of the `CDLL`, `PyDLL`, `WinDLL`, or `OleDLL` types.

`__getattr__()` has special behavior: It allows loading a shared library by accessing it as attribute of a library loader instance. The result is cached, so repeated attribute accesses return the same library each time.

LoadLibrary (*name*)

Load a shared library into the process and return it. This method always returns a new instance of the library.

These prefabricated library loaders are available:

`ctypes.cdll`

Creates `CDLL` instances.

`ctypes.windll`

仅 Windows 中: 创建 `WinDLL` 实例。

`ctypes.oledll`

仅 Windows 中: 创建 `OleDLL` 实例。

`ctypes.pydll`

创建 `PyDLL` 实例。

For accessing the C Python api directly, a ready-to-use Python shared library object is available:

ctypes.pythonapi

An instance of *PyDLL* that exposes Python C API functions as attributes. Note that all these functions are assumed to return C `int`, which is of course not always the truth, so you have to assign the correct `restype` attribute to use these functions.

Foreign functions

As explained in the previous section, foreign functions can be accessed as attributes of loaded shared libraries. The function objects created in this way by default accept any number of arguments, accept any ctypes data instances as arguments, and return the default result type specified by the library loader. They are instances of a private class:

class ctypes._FuncPtr

Base class for C callable foreign functions.

Instances of foreign functions are also C compatible data types; they represent C function pointers.

This behavior can be customized by assigning to special attributes of the foreign function object.

restype

Assign a ctypes type to specify the result type of the foreign function. Use `None` for `void`, a function not returning anything.

It is possible to assign a callable Python object that is not a ctypes type, in this case the function is assumed to return a C `int`, and the callable will be called with this integer, allowing further processing or error checking. Using this is deprecated, for more flexible post processing or error checking use a ctypes data type as `restype` and assign a callable to the `errcheck` attribute.

argtypes

Assign a tuple of ctypes types to specify the argument types that the function accepts. Functions using the `stdcall` calling convention can only be called with the same number of arguments as the length of this tuple; functions using the C calling convention accept additional, unspecified arguments as well.

When a foreign function is called, each actual argument is passed to the `from_param()` class method of the items in the `argtypes` tuple, this method allows adapting the actual argument to an object that the foreign function accepts. For example, a `c_char_p` item in the `argtypes` tuple will convert a string passed as argument into a bytes object using ctypes conversion rules.

New: It is now possible to put items in `argtypes` which are not ctypes types, but each item must have a `from_param()` method which returns a value usable as argument (integer, string, ctypes instance). This allows defining adapters that can adapt custom objects as function parameters.

errcheck

Assign a Python function or another callable to this attribute. The callable will be called with three or more arguments:

callable (*result, func, arguments*)

result is what the foreign function returns, as specified by the `restype` attribute.

func is the foreign function object itself, this allows reusing the same callable object to check or post process the results of several functions.

arguments is a tuple containing the parameters originally passed to the function call, this allows specializing the behavior on the arguments used.

The object that this function returns will be returned from the foreign function call, but it can also check the result value and raise an exception if the foreign function call failed.

exception ctypes.ArgumentError

This exception is raised when a foreign function call cannot convert one of the passed arguments.

Function prototypes

Foreign functions can also be created by instantiating function prototypes. Function prototypes are similar to function prototypes in C; they describe a function (return type, argument types, calling convention) without defining an implementation. The factory functions must be called with the desired result type and the argument types of the function, and can be used as decorator factories, and as such, be applied to functions through the `@wrapper` syntax. See [Callback functions](#) for examples.

`ctypes.CFUNCTYPE (restype, *argtypes, use_errno=False, use_last_error=False)`

The returned function prototype creates functions that use the standard C calling convention. The function will release the GIL during the call. If `use_errno` is set to true, the ctypes private copy of the system `errno` variable is exchanged with the real `errno` value before and after the call; `use_last_error` does the same for the Windows error code.

`ctypes.WINFUNCTYPE (restype, *argtypes, use_errno=False, use_last_error=False)`

Windows only: The returned function prototype creates functions that use the `stdcall` calling convention, except on Windows CE where `WINFUNCTYPE ()` is the same as `CFUNCTYPE ()`. The function will release the GIL during the call. `use_errno` and `use_last_error` have the same meaning as above.

`ctypes.PYFUNCTYPE (restype, *argtypes)`

The returned function prototype creates functions that use the Python calling convention. The function will *not* release the GIL during the call.

Function prototypes created by these factory functions can be instantiated in different ways, depending on the type and number of the parameters in the call:

prototype (*address*)

Returns a foreign function at the specified address which must be an integer.

prototype (*callable*)

Create a C callable function (a callback function) from a Python *callable*.

prototype (*func_spec* [, *paramflags*])

Returns a foreign function exported by a shared library. *func_spec* must be a 2-tuple (*name_or_ordinal*, *library*). The first item is the name of the exported function as string, or the ordinal of the exported function as small integer. The second item is the shared library instance.

prototype (*vtbl_index*, *name* [, *paramflags* [, *iid*]])

Returns a foreign function that will call a COM method. *vtbl_index* is the index into the virtual function table, a small non-negative integer. *name* is name of the COM method. *iid* is an optional pointer to the interface identifier which is used in extended error reporting.

COM methods use a special calling convention: They require a pointer to the COM interface as first argument, in addition to those parameters that are specified in the *argtypes* tuple.

The optional *paramflags* parameter creates foreign function wrappers with much more functionality than the features described above.

paramflags must be a tuple of the same length as *argtypes*.

Each item in this tuple contains further information about a parameter, it must be a tuple containing one, two, or three items.

The first item is an integer containing a combination of direction flags for the parameter:

- 1 Specifies an input parameter to the function.
- 2 Output parameter. The foreign function fills in a value.
- 4 Input parameter which defaults to the integer zero.

The optional second item is the parameter name as string. If this is specified, the foreign function can be called with named parameters.

The optional third item is the default value for this parameter.

This example demonstrates how to wrap the Windows `MessageBoxW` function so that it supports default parameters and named arguments. The C declaration from the windows header file is this:

```
WINUSERAPI int WINAPI
MessageBoxW(
    HWND hWnd,
    LPCWSTR lpText,
    LPCWSTR lpCaption,
    UINT uType);
```

Here is the wrapping with `ctypes`:

```
>>> from ctypes import c_int, WINFUNCTYPE, windll
>>> from ctypes.wintypes import HWND, LPCWSTR, UINT
>>> prototype = WINFUNCTYPE(c_int, HWND, LPCWSTR, LPCWSTR, UINT)
>>> paramflags = (1, "hwnd", 0), (1, "text", "Hi"), (1, "caption", "Hello from ctypes"), (1, "flags", 0)
>>> MessageBox = prototype(("MessageBoxW", windll.user32), paramflags)
```

The `MessageBox` foreign function can now be called in these ways:

```
>>> MessageBox()
>>> MessageBox(text="Spam, spam, spam")
>>> MessageBox(flags=2, text="foo bar")
```

A second example demonstrates output parameters. The win32 `GetWindowRect` function retrieves the dimensions of a specified window by copying them into `RECT` structure that the caller has to supply. Here is the C declaration:

```
WINUSERAPI BOOL WINAPI
GetWindowRect(
    HWND hWnd,
    LPRECT lpRect);
```

Here is the wrapping with `ctypes`:

```
>>> from ctypes import POINTER, WINFUNCTYPE, windll, WinError
>>> from ctypes.wintypes import BOOL, HWND, RECT
>>> prototype = WINFUNCTYPE(BOOL, HWND, POINTER(RECT))
>>> paramflags = (1, "hwnd"), (2, "lprect")
>>> GetWindowRect = prototype(("GetWindowRect", windll.user32), paramflags)
>>>
```

Functions with output parameters will automatically return the output parameter value if there is a single one, or a tuple containing the output parameter values when there are more than one, so the `GetWindowRect` function now returns a `RECT` instance, when called.

Output parameters can be combined with the `errcheck` protocol to do further output processing and error checking. The win32 `GetWindowRect` api function returns a `BOOL` to signal success or failure, so this function could do the error checking, and raises an exception when the api call failed:

```
>>> def errcheck(result, func, args):
...     if not result:
...         raise WinError()
```

(下页继续)

(续上页)

```

...     return args
...
>>> GetWindowRect.errcheck = errcheck
>>>

```

If the `errcheck` function returns the argument tuple it receives unchanged, `ctypes` continues the normal processing it does on the output parameters. If you want to return a tuple of window coordinates instead of a `RECT` instance, you can retrieve the fields in the function and return them instead, the normal processing will no longer take place:

```

>>> def errcheck(result, func, args):
...     if not result:
...         raise WinError()
...     rc = args[1]
...     return rc.left, rc.top, rc.bottom, rc.right
...
>>> GetWindowRect.errcheck = errcheck
>>>

```

Utility functions

`ctypes.addressof(obj)`

Returns the address of the memory buffer as integer. *obj* must be an instance of a `ctypes` type.

`ctypes.alignment(obj_or_type)`

Returns the alignment requirements of a `ctypes` type. *obj_or_type* must be a `ctypes` type or instance.

`ctypes.byref(obj[, offset])`

Returns a light-weight pointer to *obj*, which must be an instance of a `ctypes` type. *offset* defaults to zero, and must be an integer that will be added to the internal pointer value.

`byref(obj, offset)` corresponds to this C code:

```
((char *)&obj) + offset)
```

The returned object can only be used as a foreign function call parameter. It behaves similar to `pointer(obj)`, but the construction is a lot faster.

`ctypes.cast(obj, type)`

This function is similar to the cast operator in C. It returns a new instance of *type* which points to the same memory block as *obj*. *type* must be a pointer type, and *obj* must be an object that can be interpreted as a pointer.

`ctypes.create_string_buffer(init_or_size, size=None)`

This function creates a mutable character buffer. The returned object is a `ctypes` array of `c_char`.

init_or_size must be an integer which specifies the size of the array, or a bytes object which will be used to initialize the array items.

If a bytes object is specified as first argument, the buffer is made one item larger than its length so that the last element in the array is a NUL termination character. An integer can be passed as second argument which allows specifying the size of the array if the length of the bytes should not be used.

`ctypes.create_unicode_buffer(init_or_size, size=None)`

This function creates a mutable unicode character buffer. The returned object is a `ctypes` array of `c_wchar`.

init_or_size must be an integer which specifies the size of the array, or a string which will be used to initialize the array items.

If a string is specified as first argument, the buffer is made one item larger than the length of the string so that the last element in the array is a NUL termination character. An integer can be passed as second argument which allows specifying the size of the array if the length of the string should not be used.

`ctypes.DllCanUnloadNow()`

Windows only: This function is a hook which allows implementing in-process COM servers with ctypes. It is called from the `DllCanUnloadNow` function that the `_ctypes` extension dll exports.

`ctypes.DllGetClassObject()`

Windows only: This function is a hook which allows implementing in-process COM servers with ctypes. It is called from the `DllGetClassObject` function that the `_ctypes` extension dll exports.

`ctypes.util.find_library(name)`

Try to find a library and return a pathname. *name* is the library name without any prefix like `lib`, suffix like `.so`, `.dylib` or version number (this is the form used for the posix linker option `-l`). If no library can be found, returns `None`.

The exact functionality is system dependent.

`ctypes.util.find_msvcrt()`

Windows only: return the filename of the VC runtime library used by Python, and by the extension modules. If the name of the library cannot be determined, `None` is returned.

If you need to free memory, for example, allocated by an extension module with a call to the `free(void *)`, it is important that you use the function in the same library that allocated the memory.

`ctypes.FormatError([code])`

Windows only: Returns a textual description of the error code *code*. If no error code is specified, the last error code is used by calling the Windows api function `GetLastError`.

`ctypes.GetLastError()`

Windows only: Returns the last error code set by Windows in the calling thread. This function calls the Windows `GetLastError()` function directly, it does not return the ctypes-private copy of the error code.

`ctypes.get_errno()`

Returns the current value of the ctypes-private copy of the system `errno` variable in the calling thread.

`ctypes.get_last_error()`

Windows only: returns the current value of the ctypes-private copy of the system `LastError` variable in the calling thread.

`ctypes.memmove(dst, src, count)`

Same as the standard C `memmove` library function: copies *count* bytes from *src* to *dst*. *dst* and *src* must be integers or ctypes instances that can be converted to pointers.

`ctypes.memset(dst, c, count)`

Same as the standard C `memset` library function: fills the memory block at address *dst* with *count* bytes of value *c*. *dst* must be an integer specifying an address, or a ctypes instance.

`ctypes.POINTER(type)`

This factory function creates and returns a new ctypes pointer type. Pointer types are cached and reused internally, so calling this function repeatedly is cheap. *type* must be a ctypes type.

`ctypes.pointer(obj)`

This function creates a new pointer instance, pointing to *obj*. The returned object is of the type `POINTER(type(obj))`.

Note: If you just want to pass a pointer to an object to a foreign function call, you should use `byref(obj)` which is much faster.

`ctypes.resize(obj, size)`

This function resizes the internal memory buffer of *obj*, which must be an instance of a ctypes type. It is not possible

to make the buffer smaller than the native size of the objects type, as given by `sizeof(type(obj))`, but it is possible to enlarge the buffer.

`ctypes.set_errno(value)`

Set the current value of the ctypes-private copy of the system `errno` variable in the calling thread to *value* and return the previous value.

`ctypes.set_last_error(value)`

Windows only: set the current value of the ctypes-private copy of the system `LastError` variable in the calling thread to *value* and return the previous value.

`ctypes.sizeof(obj_or_type)`

Returns the size in bytes of a ctypes type or instance memory buffer. Does the same as the C `sizeof` operator.

`ctypes.string_at(address, size=-1)`

This function returns the C string starting at memory address *address* as a bytes object. If *size* is specified, it is used as size, otherwise the string is assumed to be zero-terminated.

`ctypes.WinError(code=None, descr=None)`

Windows only: this function is probably the worst-named thing in ctypes. It creates an instance of `OSError`. If *code* is not specified, `GetLastError` is called to determine the error code. If *descr* is not specified, `FormatError()` is called to get a textual description of the error.

在 3.3 版更改: An instance of `WindowsError` used to be created.

`ctypes.wstring_at(address, size=-1)`

This function returns the wide character string starting at memory address *address* as a string. If *size* is specified, it is used as the number of characters of the string, otherwise the string is assumed to be zero-terminated.

Data types

class `ctypes._CData`

This non-public class is the common base class of all ctypes data types. Among other things, all ctypes type instances contain a memory block that hold C compatible data; the address of the memory block is returned by the `addressof()` helper function. Another instance variable is exposed as `_objects`; this contains other Python objects that need to be kept alive in case the memory block contains pointers.

Common methods of ctypes data types, these are all class methods (to be exact, they are methods of the *metaclass*):

from_buffer (*source*[, *offset*])

This method returns a ctypes instance that shares the buffer of the *source* object. The *source* object must support the writeable buffer interface. The optional *offset* parameter specifies an offset into the source buffer in bytes; the default is zero. If the source buffer is not large enough a `ValueError` is raised.

from_buffer_copy (*source*[, *offset*])

This method creates a ctypes instance, copying the buffer from the *source* object buffer which must be readable. The optional *offset* parameter specifies an offset into the source buffer in bytes; the default is zero. If the source buffer is not large enough a `ValueError` is raised.

from_address (*address*)

This method returns a ctypes type instance using the memory specified by *address* which must be an integer.

from_param (*obj*)

This method adapts *obj* to a ctypes type. It is called with the actual object used in a foreign function call when the type is present in the foreign function's `argtypes` tuple; it must return an object that can be used as a function call parameter.

All ctypes data types have a default implementation of this classmethod that normally returns *obj* if that is an instance of the type. Some types accept other objects as well.

in_dll (*library*, *name*)

This method returns a ctypes type instance exported by a shared library. *name* is the name of the symbol that exports the data, *library* is the loaded shared library.

Common instance variables of ctypes data types:

_b_base_

Sometimes ctypes data instances do not own the memory block they contain, instead they share part of the memory block of a base object. The `_b_base_` read-only member is the root ctypes object that owns the memory block.

_b_needsfree_

This read-only variable is true when the ctypes data instance has allocated the memory block itself, false otherwise.

_objects

This member is either `None` or a dictionary containing Python objects that need to be kept alive so that the memory block contents is kept valid. This object is only exposed for debugging; never modify the contents of this dictionary.

基础数据类型

class ctypes._SimpleCData

This non-public class is the base class of all fundamental ctypes data types. It is mentioned here because it contains the common attributes of the fundamental ctypes data types. `_SimpleCData` is a subclass of `_CData`, so it inherits their methods and attributes. ctypes data types that are not and do not contain pointers can now be pickled.

Instances have a single attribute:

value

This attribute contains the actual value of the instance. For integer and pointer types, it is an integer, for character types, it is a single character bytes object or string, for character pointer types it is a Python bytes object or string.

When the `value` attribute is retrieved from a ctypes instance, usually a new object is returned each time. `ctypes` does *not* implement original object return, always a new object is constructed. The same is true for all other ctypes object instances.

Fundamental data types, when returned as foreign function call results, or, for example, by retrieving structure field members or array items, are transparently converted to native Python types. In other words, if a foreign function has a `restype` of `c_char_p`, you will always receive a Python bytes object, *not* a `c_char_p` instance.

Subclasses of fundamental data types do *not* inherit this behavior. So, if a foreign functions `restype` is a subclass of `c_void_p`, you will receive an instance of this subclass from the function call. Of course, you can get the value of the pointer by accessing the `value` attribute.

These are the fundamental ctypes data types:

class ctypes.c_byte

Represents the C `signed char` datatype, and interprets the value as small integer. The constructor accepts an optional integer initializer; no overflow checking is done.

class ctypes.c_char

Represents the C `char` datatype, and interprets the value as a single character. The constructor accepts an optional string initializer, the length of the string must be exactly one character.

class ctypes.c_char_p

Represents the C `char *` datatype when it points to a zero-terminated string. For a general character pointer that may also point to binary data, `POINTER(c_char)` must be used. The constructor accepts an integer address, or a bytes object.

class ctypes.c_double

Represents the C double datatype. The constructor accepts an optional float initializer.

class ctypes.c_longdouble

Represents the C long double datatype. The constructor accepts an optional float initializer. On platforms where `sizeof(long double) == sizeof(double)` it is an alias to `c_double`.

class ctypes.c_float

Represents the C float datatype. The constructor accepts an optional float initializer.

class ctypes.c_int

Represents the C signed int datatype. The constructor accepts an optional integer initializer; no overflow checking is done. On platforms where `sizeof(int) == sizeof(long)` it is an alias to `c_long`.

class ctypes.c_int8

Represents the C 8-bit signed int datatype. Usually an alias for `c_byte`.

class ctypes.c_int16

Represents the C 16-bit signed int datatype. Usually an alias for `c_short`.

class ctypes.c_int32

Represents the C 32-bit signed int datatype. Usually an alias for `c_int`.

class ctypes.c_int64

Represents the C 64-bit signed int datatype. Usually an alias for `c_longlong`.

class ctypes.c_long

Represents the C signed long datatype. The constructor accepts an optional integer initializer; no overflow checking is done.

class ctypes.c_longlong

Represents the C signed long long datatype. The constructor accepts an optional integer initializer; no overflow checking is done.

class ctypes.c_short

Represents the C signed short datatype. The constructor accepts an optional integer initializer; no overflow checking is done.

class ctypes.c_size_t

Represents the C size_t datatype.

class ctypes.c_ssize_t

Represents the C ssize_t datatype.

3.2 新版功能.

class ctypes.c_ubyte

Represents the C unsigned char datatype, it interprets the value as small integer. The constructor accepts an optional integer initializer; no overflow checking is done.

class ctypes.c_uint

Represents the C unsigned int datatype. The constructor accepts an optional integer initializer; no overflow checking is done. On platforms where `sizeof(int) == sizeof(long)` it is an alias for `c_ulong`.

class ctypes.c_uint8

Represents the C 8-bit unsigned int datatype. Usually an alias for `c_ubyte`.

class ctypes.c_uint16

Represents the C 16-bit unsigned int datatype. Usually an alias for `c_ushort`.

class ctypes.c_uint32

Represents the C 32-bit unsigned int datatype. Usually an alias for `c_uint`.

class `ctypes.c_uint64`

Represents the C 64-bit unsigned int datatype. Usually an alias for `c_ulonglong`.

class `ctypes.c_ulong`

Represents the C unsigned long datatype. The constructor accepts an optional integer initializer; no overflow checking is done.

class `ctypes.c_ulonglong`

Represents the C unsigned long long datatype. The constructor accepts an optional integer initializer; no overflow checking is done.

class `ctypes.c_ushort`

Represents the C unsigned short datatype. The constructor accepts an optional integer initializer; no overflow checking is done.

class `ctypes.c_void_p`

Represents the C void * type. The value is represented as integer. The constructor accepts an optional integer initializer.

class `ctypes.c_wchar`

Represents the C wchar_t datatype, and interprets the value as a single character unicode string. The constructor accepts an optional string initializer, the length of the string must be exactly one character.

class `ctypes.c_wchar_p`

Represents the C wchar_t * datatype, which must be a pointer to a zero-terminated wide character string. The constructor accepts an integer address, or a string.

class `ctypes.c_bool`

Represent the C bool datatype (more accurately, _Bool from C99). Its value can be True or False, and the constructor accepts any object that has a truth value.

class `ctypes.HRESULT`

Windows only: Represents a HRESULT value, which contains success or error information for a function or method call.

class `ctypes.py_object`

Represents the C PyObject * datatype. Calling this without an argument creates a NULL PyObject * pointer.

The `ctypes.wintypes` module provides quite some other Windows specific data types, for example `HWND`, `WPARAM`, or `DWORD`. Some useful structures like `MSG` or `RECT` are also defined.

Structured data types

class `ctypes.Union(*args, **kw)`

Abstract base class for unions in native byte order.

class `ctypes.BigEndianStructure(*args, **kw)`

Abstract base class for structures in *big endian* byte order.

class `ctypes.LittleEndianStructure(*args, **kw)`

Abstract base class for structures in *little endian* byte order.

Structures with non-native byte order cannot contain pointer type fields, or any other data types containing pointer type fields.

class `ctypes.Structure(*args, **kw)`

Abstract base class for structures in *native* byte order.

Concrete structure and union types must be created by subclassing one of these types, and at least define a `__fields__` class variable. `ctypes` will create *descriptors* which allow reading and writing the fields by direct attribute accesses. These are the

`__fields__`

A sequence defining the structure fields. The items must be 2-tuples or 3-tuples. The first item is the name of the field, the second item specifies the type of the field; it can be any `ctypes` data type.

For integer type fields like `c_int`, a third optional item can be given. It must be a small positive integer defining the bit width of the field.

Field names must be unique within one structure or union. This is not checked, only one field can be accessed when names are repeated.

It is possible to define the `__fields__` class variable *after* the class statement that defines the Structure subclass, this allows creating data types that directly or indirectly reference themselves:

```
class List(Structure):
    pass
List.__fields__ = [("pNext", POINTER(List)),
                  ...
                  ]
```

The `__fields__` class variable must, however, be defined before the type is first used (an instance is created, `sizeof()` is called on it, and so on). Later assignments to the `__fields__` class variable will raise an `AttributeError`.

It is possible to defined sub-subclasses of structure types, they inherit the fields of the base class plus the `__fields__` defined in the sub-subclass, if any.

`__pack__`

An optional small integer that allows overriding the alignment of structure fields in the instance. `__pack__` must already be defined when `__fields__` is assigned, otherwise it will have no effect.

`__anonymous__`

An optional sequence that lists the names of unnamed (anonymous) fields. `__anonymous__` must be already defined when `__fields__` is assigned, otherwise it will have no effect.

The fields listed in this variable must be structure or union type fields. `ctypes` will create descriptors in the structure type that allows accessing the nested fields directly, without the need to create the structure or union field.

Here is an example type (Windows):

```
class _U(Union):
    __fields__ = [("lptdesc", POINTER(TYPEDESC)),
                  ("lpadesc", POINTER(ARRAYDESC)),
                  ("hreftype", HREFTYPE)]

class TYPEDESC(Structure):
    __anonymous__ = ("u",)
    __fields__ = [("u", _U),
                  ("vt", VARTYPE)]
```

The `TYPEDESC` structure describes a COM data type, the `vt` field specifies which one of the union fields is valid. Since the `u` field is defined as anonymous field, it is now possible to access the members directly off the `TYPEDESC` instance. `td.lptdesc` and `td.u.lptdesc` are equivalent, but the former is faster since it does not need to create a temporary union instance:

```
td = TYPEDESC()
td.vt = VT_PTR
td.lptdesc = POINTER(some_type)
td.u.lptdesc = POINTER(some_type)
```

It is possible to defined sub-subclasses of structures, they inherit the fields of the base class. If the subclass definition has a separate `__fields__` variable, the fields specified in this are appended to the fields of the base class.

Structure and union constructors accept both positional and keyword arguments. Positional arguments are used to initialize member fields in the same order as they are appear in `__fields__`. Keyword arguments in the constructor are interpreted as attribute assignments, so they will initialize `__fields__` with the same name, or create new attributes for names not present in `__fields__`.

Arrays and pointers

class `ctypes.Array(*args)`

Abstract base class for arrays.

The recommended way to create concrete array types is by multiplying any `ctypes` data type with a positive integer. Alternatively, you can subclass this type and define `__length__` and `__type__` class variables. Array elements can be read and written using standard subscript and slice accesses; for slice reads, the resulting object is *not* itself an `Array`.

__length__

A positive integer specifying the number of elements in the array. Out-of-range subscripts result in an `IndexError`. Will be returned by `len()`.

__type__

Specifies the type of each element in the array.

Array subclass constructors accept positional arguments, used to initialize the elements in order.

class `ctypes._Pointer`

Private, abstract base class for pointers.

Concrete pointer types are created by calling `POINTER()` with the type that will be pointed to; this is done automatically by `pointer()`.

If a pointer points to an array, its elements can be read and written using standard subscript and slice accesses. Pointer objects have no size, so `len()` will raise `TypeError`. Negative subscripts will read from the memory *before* the pointer (as in C), and out-of-range subscripts will probably crash with an access violation (if you're lucky).

__type__

Specifies the type pointed to.

contents

Returns the object to which to pointer points. Assigning to this attribute changes the pointer to point to the assigned object.

本章中描述的模块支持并发执行代码。适当的工具选择取决于要执行的任务（CPU 密集型或 IO 密集型）和偏好的开发风格（事件驱动的协作式多任务或抢占式多任务处理）。这是一个概述：

17.1 threading — 基于线程的并行

源代码: [Lib/threading.py](#)

这个模块在较低级的模块 `_thread` 基础上建立较高级的线程接口。参见: `queue` 模块。

The `dummy_threading` module is provided for situations where `threading` cannot be used because `_thread` is missing.

注解：虽然他们没有在下面列出，这个模块仍然支持 Python 2.x 系列的这个模块下以 camelCase（驼峰法）命名的方法和函数。

这个模块定义了以下函数：

`threading.active_count()`

返回当前存活的线程类 `Thread` 对象。返回的计数等于 `enumerate()` 返回的列表长度。

`threading.current_thread()`

返回当前对应调用者的控制线程的 `Thread` 对象。如果调用者的控制线程不是利用 `threading` 创建，会返回一个功能受限的虚拟线程对象。

`threading.get_ident()`

返回当前线程的“线程标识符”。它是一个非零的整数。它的值没有直接含义，主要是用作 magic cookie，比如作为含有线程相关数据的字典的索引。线程标识符可能会在线程退出，新线程创建时被复用。

3.3 新版功能.

`threading.enumerate()`

以列表形式返回当前所有存活的 `Thread` 对象。该列表包含守护线程, `current_thread()` 创建的虚拟线程对象和主线程。它不包含已终结的线程和尚未开始的线程。

`threading.main_thread()`

返回主 `Thread` 对象。一般情况下, 主线程是 Python 解释器开始时创建的线程。

3.4 新版功能.

`threading.settrace(func)`

为所有 `threading` 模块开始的线程设置追踪函数。在每个线程的 `run()` 方法被调用前, `func` 会被传递给 `sys.settrace()`。

`threading.setprofile(func)`

为所有 `threading` 模块开始的线程设置性能测试函数。在每个线程的 `run()` 方法被调用前, `func` 会被传递给 `sys.setprofile()`。

`threading.stack_size([size])`

Return the thread stack size used when creating new threads. The optional *size* argument specifies the stack size to be used for subsequently created threads, and must be 0 (use platform or configured default) or a positive integer value of at least 32,768 (32 KiB). If *size* is not specified, 0 is used. If changing the thread stack size is unsupported, a `RuntimeError` is raised. If the specified stack size is invalid, a `ValueError` is raised and the stack size is unmodified. 32 KiB is currently the minimum supported stack size value to guarantee sufficient stack space for the interpreter itself. Note that some platforms may have particular restrictions on values for the stack size, such as requiring a minimum stack size > 32 KiB or requiring allocation in multiples of the system memory page size - platform documentation should be referred to for more information (4 KiB pages are common; using multiples of 4096 for the stack size is the suggested approach in the absence of more specific information). Availability: Windows, systems with POSIX threads.

这个模块同时定义了以下常量:

`threading.TIMEOUT_MAX`

阻塞函数 (`Lock.acquire()`, `RLock.acquire()`, `Condition.wait()`, ...) 中形参 *timeout* 允许的最大值。传入超过这个值的 *timeout* 会抛出 `OverflowError` 异常。

3.2 新版功能.

这个模块定义了许多类, 详见以下部分。

该模块的设计基于 Java 的线程模型。但是, 在 Java 里面, 锁和条件变量是每个对象的基础特性, 而在 Python 里面, 这些被独立成了单独的对象。Python 的 `Thread` 类只是 Java 的 `Thread` 类的一个子集; 目前还没有优先级, 没有线程组, 线程还不能被销毁、停止、暂停、恢复或中断。Java 的 `Thread` 类的静态方法在实现时会映射为模块级函数。

下列描述的方法都是自动执行的。

17.1.1 线程本地数据

线程本地数据是特定线程的数据。管理线程本地数据, 只需要创建一个 `local` (或者一个子类型) 的实例并在实例中储存属性:

```
mydata = threading.local()
mydata.x = 1
```

在不同的线程中, 实例的值会不同。

class `threading.local`

一个代表线程本地数据的类。

更多相关细节和大量示例, 参见 `_threading_local` 模块的文档。

17.1.2 线程对象

The `Thread` class represents an activity that is run in a separate thread of control. There are two ways to specify the activity: by passing a callable object to the constructor, or by overriding the `run()` method in a subclass. No other methods (except for the constructor) should be overridden in a subclass. In other words, *only* override the `__init__()` and `run()` methods of this class.

当线程对象一但被创建，其活动一定会因调用线程的 `start()` 方法开始。这会在独立的控制线程调用 `run()` 方法。

一旦线程活动开始，该线程会被认为是‘存活的’。当它的 `run()` 方法终结了（不管是正常的还是抛出未被处理的异常），就不是‘存活的’。`is_alive()` 方法用于检查线程是否存活。

其他线程可以调用一个线程的 `join()` 方法。这会阻塞调用该方法的线程，直到被调用 `join()` 方法的线程终结。

线程有名字。名字可以传递给构造函数，也可以通过 `name` 属性读取或者修改。

一个线程可以被标记成一个“守护线程”。这个标志的意义是，只有守护线程都终结，整个 Python 程序才会退出。初始值继承于创建线程。这个标志可以通过 `daemon` 特征属性或者 守护构造函数参数来设置。

注解：守护线程在程序关闭时会突然关闭。他们的资源（例如已经打开的文档，数据库事务等等）可能没有被正确释放。如果你想你的线程正常停止，设置他们成为非守护模式并且使用合适的信号机制，例如：`Event`。

有个“主线程”对象；这对应 Python 程序里面初始的控制线程。它不是一个守护线程。

“虚拟线程对象”是可以被创建的。这些是对应于“外部线程”的线程对象，它们是在线程模块外部启动的控制线程，例如直接来自 C 代码。虚拟线程对象功能受限；他们总是被认为是存活的和守护模式，不能被 `join()`。因为无法检测外来线程的终结，它们永远不会被删除。

class `threading.Thread` (`group=None`, `target=None`, `name=None`, `args=()`, `kwargs={}`, `*`, `daemon=None`)

调用这个构造函数时，必需带有关键字参数。参数如下：

`group` 应该为 `None`；为了日后扩展 `ThreadGroup` 类实现而保留。

`target` 是用于 `run()` 方法调用的可调用对象。默认是 `None`，表示不需要调用任何方法。

`name` 是线程名称。默认情况下，由“Thread-*N*”格式构成一个唯一的名称，其中 *N* 是小的十进制数。

`args` 是用于调用目标函数的参数元组。默认是 `()`。

`kwargs` 是用于调用目标函数的关键字参数字典。默认是 `{}`。

如果不是 `None`，线程将被显式的设置为 守护模式，不管该线程是否是守护模式。如果是 `None`（默认值），线程将继承当前线程的守护模式属性。

如果子类型重载了构造函数，它一定要确保在做任何事前，先发起调用基类构造器 (`Thread.__init__()`)。

在 3.3 版更改：加入 `daemon` 参数。

start()

开始线程活动。

它在一个线程里最多只能被调用一次。它安排对象的 `run()` 方法在一个独立的控制进程中调用。

如果同一个线程对象中调用这个方法次数大于一次，会抛出 `RuntimeError`。

run()

代表线程活动的方法。

You may override this method in a subclass. The standard `run()` method invokes the callable object passed to the object's constructor as the *target* argument, if any, with sequential and keyword arguments taken from the *args* and *kwargs* arguments, respectively.

join (timeout=None)

等待，直到线程终结。这会阻塞调用这个方法的线程，直到被调用`join()`的线程终结 – 不管是正常终结还是抛出未处理异常 – 或者直到发生超时，超时选项是可选的。

当 *timeout* 参数存在而且不是 `None` 时，它应该是一个用于指定操作超时的以秒为单位的浮点数（或者分数）。因为`join()`总是返回 `None`，所以你一定要在`join()`后调用`is_alive()`才能判断是否发生超时 – 如果线程仍然存活，则`join()`超时。

当 *timeout* 参数不存在或者是 `None`，这个操作会阻塞直到线程终结。

一个线程可以被`join()`很多次。

如果尝试加入当前线程会导致死锁，`join()`会引起`RuntimeError`异常。如果尝试`join()`一个尚未开始的线程，也会抛出相同的异常。

name

只用于识别的字符串。它没有语义。多个线程可以赋予相同的名称。初始名称由构造函数设置。

getName ()

setName ()

旧的`name`取值/设值 API；直接当做特征属性使用它。

ident

这个线程的‘线程标识符’，如果线程尚未开始则为 `None`。这是个非零整数。参见`get_ident()`函数。当一个线程退出而另外一个线程被创建，线程标识符会被复用。即使线程退出后，仍可得到标识符。

is_alive ()

返回线程是否存活。

当`run()`方法刚开始直到`run()`方法刚结束，这个方法返回 `True`。模块函数`enumerate()`返回包含所有存活线程的列表。

daemon

一个表示这个线程是 (`True`) 否 (`False`) 守护线程的布尔值。一定要在调用`start()`前设置好，否则会抛出`RuntimeError`。初始值继承于创建线程；主线程不是守护线程，因此主线程创建的所有线程默认都是`daemon=False`。

当没有存活的非守护线程时，整个 Python 程序才会退出。

isDaemon ()

setDaemon ()

旧的`name`取值/设值 API；建议直接当做特征属性使用它。

CPython implementation detail: CPython 下，因为`Global Interpreter Lock`，一个时刻只有一个线程可以执行 Python 代码（尽管如此，某些性能导向的库可能会克服这个限制）。如果你想让你的应用更好的利用多核计算机的计算性能，推荐你使用`multiprocessing`或者`concurrent.futures.ProcessPoolExecutor`。但是如果你想同时运行多个 I/O 绑定任务，线程仍然是一个合适的模型。

17.1.3 锁对象

原始锁是一个在锁定时不属于特定线程的同步基元组件。在 Python 中，它是能用的最低级的同步基元组件，由`_thread`扩展模块直接实现。

原始锁处于“锁定”或者“非锁定”两种状态之一。它被创建时为非锁定状态。它有两个基本方法，`acquire()`和`release()`。当状态为非锁定时，`acquire()`将状态改为锁定并立即返回。当状态是锁定时，`acquire()`

将阻塞至其他线程调用 `release()` 将其改为非锁定状态，然后 `acquire()` 调用重置其为锁定状态并返回。`release()` 只在锁定状态下调用；它将状态改为非锁定并立即返回。如果尝试释放一个非锁定的锁，则会引发 `RuntimeError` 异常。

锁同样支持上下文管理协议。

当多个线程在 `acquire()` 等待状态转变为未锁定被阻塞，然后 `release()` 重置状态为未锁定时，只有一个线程能继续执行；至于哪个等待线程继续执行没有定义，并且会根据实现而不同。

所有方法都是自动执行的。

class `threading.Lock`

实现原始锁对象的类。一旦一个线程获得一个锁，会阻塞随后尝试获得锁的线程，直到它被释放；任何线程都可以释放它。

需要注意的是 `Lock` 其实是一个工厂函数，返回平台支持的具体锁类中最有效的版本的实例。

acquire (*blocking=True, timeout=-1*)

可以阻塞或非阻塞地获得锁。

当调用时参数 *blocking* 设置为 `True`（缺省值），阻塞直到锁被释放，然后将锁锁定并返回 `True`。

在参数 *blocking* 被设置为 `False` 的情况下调用，将不会发生阻塞。如果调用时 *blocking* 设为 `True` 会阻塞，并立即返回 `False`；否则，将锁锁定并返回 `True`。

当浮点型 *timeout* 参数被设置为正值调用时，只要无法获得锁，将最多阻塞 *timeout* 设定的秒数。*timeout* 参数被设置为 `-1` 时将无限等待。当 *blocking* 为 `false` 时，*timeout* 指定的值将被忽略。

如果成功获得锁，则返回 `True`，否则返回 `False`（例如发生 超时的时候）。

在 3.2 版更改：新的 *timeout* 形参。

在 3.2 版更改：现在如果底层线程实现支持，则可以通过 POSIX 上的信号中断锁的获取。

release ()

释放一个锁。这个方法可以在任何线程中调用，不单指获得锁的线程。

当锁被锁定，将它重置为未锁定，并返回。如果其他线程正在等待这个锁解锁而被阻塞，只允许其中一个允许。

在未锁定的锁调用时，会引发 `RuntimeError` 异常。

没有返回值。

17.1.4 递归锁对象

重入锁是一个可以被同一个线程多次获取的同步基元组件。在内部，它在基元锁的锁定/非锁定状态上附加了“所属线程”和“递归等级”的概念。在锁定状态下，某些线程拥有锁；在非锁定状态下，没有线程拥有它。

若要锁定锁，线程调用其 `acquire()` 方法；一旦线程拥有了锁，方法将返回。若要解锁，线程调用 `release()` 方法。`acquire()/release()` 对可以嵌套；只有最终 `release()`（最外面一对的 `release()`）将锁解开，才能让其他线程继续处理 `acquire()` 阻塞。

递归锁也支持上下文管理协议。

class `threading.RLock`

此类实现了重入锁对象。重入锁必须由获取它的线程释放。一旦线程获得了重入锁，同一个线程再次获取它将不阻塞；线程必须在每次获取它时释放一次。

需要注意的是 `RLock` 其实是一个工厂函数，返回平台支持的具体递归锁类中最有效的版本的实例。

acquire (*blocking=True, timeout=-1*)

可以阻塞或非阻塞地获得锁。

当无参数调用时：如果这个线程已经拥有锁，递归级别增加一，并立即返回。否则，如果其他线程拥有该锁，则阻塞至该锁解锁。一旦锁被解锁（不属于任何线程），则抢夺所有权，设置递归等级为一，并返回。如果多个线程被阻塞，等待锁被解锁，一次只有一个线程能抢到锁的所有权。在这种情况下，没有返回值。

当调用时参数 *blocking* 设置为 `True`，和没带参数调用一样做同样的事，然后返回 `True`。

当 *blocking* 参数设置为 `false` 的情况下调用，不进行阻塞。如果一个无参数的调用已经阻塞，立即返回 `false`；否则，执行和无参数调用一样的操作，并返回 `true`。

当浮点数 *timeout* 参数被设置为正值调用时，只要无法获得锁，将最多阻塞 *timeout* 设定的秒数。如果锁被获取返回 `true`，如果超时返回 `false`。

在 3.2 版更改：新的 *timeout* 形参。

`release()`

释放锁，自减递归等级。如果减到零，则将锁重置为非锁定状态（不被任何线程拥有），并且，如果其他线程正被阻塞着等待锁被解锁，则仅允许其中一个线程继续。如果自减后，递归等级仍然不是零，则锁保持锁定，仍由调用线程拥有。

只有当前线程拥有锁才能调用这个方法。如果锁被释放后调用这个方法，会引起 `RuntimeError` 异常。

没有返回值。

17.1.5 条件对象

条件变量总是与某种类型的锁对象相关联，锁对象可以通过传入获得，或者在缺省的情况下自动创建。当多个条件变量需要共享同一个锁时，传入一个锁很有用。锁是条件对象的一部分，你不必单独地跟踪它。

条件变量服从上下文管理协议：使用 `with` 语句会在它包围的代码块内获取关联的锁。`acquire()` 和 `release()` 方法也能调用关联锁的相关方法。

其它方法必须在持有关联的锁的情况下调用。`wait()` 方法释放锁，然后阻塞直到其它线程调用 `notify()` 方法或 `notify_all()` 方法唤醒它。一旦被唤醒，`wait()` 方法重新获取锁并返回。它也可以指定超时时间。

The `notify()` method wakes up one of the threads waiting for the condition variable, if any are waiting. The `notify_all()` method wakes up all threads waiting for the condition variable.

注意：`notify()` 方法和 `notify_all()` 方法并不会释放锁，这意味着被唤醒的线程不会立即从它们的 `wait()` 方法调用中返回，而是会在调用了 `notify()` 方法或 `notify_all()` 方法的线程最终放弃了锁的所有权后返回。

使用条件变量的典型编程风格是将锁用于同步某些共享状态的权限，那些对状态的某些特定改变感兴趣的线程，它们重复调用 `wait()` 方法，直到看到所期望的改变发生；而对于修改状态的线程，它们将当前状态改变为可能是等待者所期待的新状态后，调用 `notify()` 方法或者 `notify_all()` 方法。例如，下面的代码是一个通用的无限缓冲区容量的生产者-消费者情形：

```
# Consume one item
with cv:
    while not an_item_is_available():
        cv.wait()
    get_an_available_item()

# Produce one item
with cv:
    make_an_item_available()
    cv.notify()
```

使用 `while` 循环检查所要求的条件成立与否是有必要的，因为 `wait()` 方法可能要经过不确定长度的时间后才会返回，而此时导致 `notify()` 方法调用的那个条件可能已经不再成立。这是多线程编程所固有的问题。`wait_for()` 方法可自动化条件检查，并简化超时计算。

```
# Consume an item
with cv:
    cv.wait_for(an_item_is_available)
    get_an_available_item()
```

选择 `notify()` 还是 `notify_all()`，取决于一次状态改变是只能被一个还是能被多个等待线程所用。例如在一个典型的生产者-消费者情形中，添加一个项目到缓冲区只需唤醒一个消费者线程。

class `threading.Condition` (`lock=None`)

实现条件变量对象的类。一个条件变量对象允许一个或多个线程在被其它线程所通知之前进行等待。

如果给出了非 `None` 的 `lock` 参数，则它必须为 `Lock` 或者 `RLock` 对象，并且它将被用作底层锁。否则，将会创建新的 `RLock` 对象，并将其用作底层锁。

在 3.3 版更改：从工厂函数变为类。

acquire (*args)

请求底层锁。此方法调用底层锁的相应方法，返回值是底层锁相应方法的返回值。

release ()

释放底层锁。此方法调用底层锁的相应方法。没有返回值。

wait (`timeout=None`)

等待直到被通知或发生超时。如果线程在调用此方法时没有获得锁，将会引发 `RuntimeError` 异常。

这个方法释放底层锁，然后阻塞，直到在另外一个线程中调用同一个条件变量的 `notify()` 或 `notify_all()` 唤醒它，或者直到可选的超时发生。一旦被唤醒或者超时，它重新获得锁并返回。

当提供了 `timeout` 参数且不是 `None` 时，它应该是一个浮点数，代表操作的超时时间，以秒为单位（可以为小数）。

当底层锁是个 `RLock`，不会使用它的 `release()` 方法释放锁，因为当它被递归多次获取时，实际上可能无法解锁。相反，使用了 `RLock` 类的内部接口，即使多次递归获取它也能解锁它。然后，在重新获取锁时，使用另一个内部接口来恢复递归级别。

返回 `True`，除非提供的 `timeout` 过期，这种情况下返回 `False`。

在 3.2 版更改：很明显，方法总是返回 `None`。

wait_for (`predicate`, `timeout=None`)

等待，直到条件计算为真。`predicate` 应该是一个可调用对象而且它的返回值可被解释为一个布尔值。可以提供 `timeout` 参数给出最大等待时间。

这个实用方法会重复地调用 `wait()` 直到满足判断式或者发生超时。返回值是判断式最后一个返回值，而且如果方法发生超时会返回 `False`。

忽略超时功能，调用此方法大致相当于编写：

```
while not predicate():
    cv.wait()
```

因此，规则同样适用于 `wait()`：锁必须在被调用时保持获取，并在返回时重新获取。随着锁定执行判断式。

3.2 新版功能。

notify(*n*=1)

默认唤醒一个等待这个条件的线程。如果调用线程在没有获得锁的情况下调用这个方法，会引发 `RuntimeError` 异常。

这个方法唤醒最多 *n* 个正在等待这个条件变量的线程；如果没有线程在等待，这是一个空操作。

当前实现中，如果至少有 *n* 个线程正在等待，准确唤醒 *n* 个线程。但是依赖这个行为并不安全。未来，优化的实现有时会唤醒超过 *n* 个线程。

注意：被唤醒的线程实际上不会返回它调用的 `wait()`，直到它可以重新获得锁。因为 `notify()` 不会释放锁，只有它的调用者应该这样做。

notify_all()

唤醒所有正在等待这个条件的线程。这个方法行为与 `notify()` 相似，但并不只唤醒单一线程，而是唤醒所有等待线程。如果调用线程在调用这个方法时没有获得锁，会引发 `RuntimeError` 异常。

17.1.6 信号量对象

这是计算机科学史上最古老的同步原语之一，早期的荷兰科学家 Edsger W. Dijkstra 发明了它。（他使用名称 `P()` 和 `V()` 而不是 `acquire()` 和 `release()`）。

一个信号量管理一个内部计数器，该计数器因 `acquire()` 方法的调用而递减，因 `release()` 方法的调用而递增。计数器的值永远不会小于零；当 `acquire()` 方法发现计数器为零时，将会阻塞，直到其它线程调用 `release()` 方法。

信号量对象也支持上下文管理协议。

class threading.Semaphore(*value*=1)

该类实现信号量对象。信号量对象管理一个原子性的计数器，代表 `release()` 方法的调用次数减去 `acquire()` 的调用次数再加上一个初始值。如果需要，`acquire()` 方法将会阻塞直到可以返回而不会使得计数器变成负数。在没有显式给出 *value* 的值时，默认为 1。

可选参数 *value* 赋予内部计数器初始值，默认值为 1。如果 *value* 被赋予小于 0 的值，将会引发 `ValueError` 异常。

在 3.3 版更改：从工厂函数变为类。

acquire(*blocking*=True, *timeout*=None)

获取一个信号量。

在不带参数的情况下调用时：

- 如果在进入时，内部计数器的值大于 0，将其减 1 并立即返回 `true`。
- 如果在进入时，内部计数器的值为 0，将会阻塞直到被 `release()` 方法的调用唤醒。一旦被唤醒（并且计数器值大于 0），将计数器值减少 1 并返回 `true`。线程会被每次 `release()` 方法的调用唤醒。线程被唤醒的次序是不确定的。

在参数 *blocking* 被设置为 `false` 的情况下调用，将不会发生阻塞。如果不带参数的调用会发生阻塞的话，带参数的调用在相同情况下将会立即返回 `false`。否则，执行和不带参数的调用一样的操作并返回 `true`。

当参数 *timeout* 不为 `None` 时，将最多阻塞 *timeout* 秒。如果未能在时间间隔内成功获取信号量，将返回 `false`，否则返回 `true`。

在 3.2 版更改：新的 *timeout* 形参。

release()

释放一个信号量，将内部计数器的值增加 1。当计数器原先的值为 0 且有其它线程正在等待它再次大于 0 时，唤醒正在等待的线程。

class `threading.BoundedSemaphore` (*value=1*)

该类实现有界信号量。有界信号量通过检查以确保它当前的值不会超过初始值。如果超过了初始值，将会引发 `ValueError` 异常。在大多情况下，信号量用于保护数量有限的资源。如果信号量被释放的次数过多，则表明出现了错误。没有指定时，*value* 的值默认为 1。

在 3.3 版更改: 从工厂函数变为类。

Semaphore 例子

信号量通常用于保护数量有限的资源，例如数据库服务器。在资源数量固定的任何情况下，都应该使用有界信号量。在生成任何工作线程前，应该在主线程中初始化信号量。

```
maxconnections = 5
# ...
pool_sema = BoundedSemaphore(value=maxconnections)
```

工作线程生成后，当需要连接服务器时，这些线程将调用信号量的 `acquire` 和 `release` 方法：

```
with pool_sema:
    conn = connectdb()
    try:
        # ... use connection ...
    finally:
        conn.close()
```

使用有界信号量能减少这种编程错误：信号量的释放次数多于其请求次数。

17.1.7 事件对象

这是线程之间通信的最简单机制之一：一个线程发出事件信号，而其他线程等待该信号。

一个事件对象管理一个内部标志，调用 `set()` 方法可将其设置为 `true`，调用 `clear()` 方法可将其设置为 `false`，调用 `wait()` 方法将进入阻塞直到标志为 `true`。

class `threading.Event`

实现事件对象的类。事件对象管理一个内部标志，调用 `set()` 方法可将其设置为 `true`。调用 `clear()` 方法可将其设置为 `false`。调用 `wait()` 方法将进入阻塞直到标志为 `true`。这个标志初始时为 `false`。

在 3.3 版更改: 从工厂函数变为类。

is_set()

当且仅当内部标志为 `true` 时返回 `true`。

set()

将内部标志设置为 `true`。所有正在等待这个事件的线程将被唤醒。当标志为 `true` 时，调用 `wait()` 方法的线程不会被阻塞。

clear()

将内部标志设置为 `false`。之后调用 `wait()` 方法的线程将会被阻塞，直到调用 `set()` 方法将内部标志再次设置为 `true`。

wait (*timeout=None*)

阻塞线程直到内部变量为 `true`。如果调用时内部标志为 `true`，将立即返回。否则将阻塞线程，直到调用 `set()` 方法将标志设置为 `true` 或者发生可选的超时。

当提供了 `timeout` 参数且不是 `None` 时，它应该是一个浮点数，代表操作的超时时间，以秒为单位（可以为小数）。

当内部标志在调用 `wait` 进入阻塞后被设置为 `true`，或者调用 `wait` 时已经被设置为 `true` 时，方法返回 `true`。也就是说，除非设定了超时且发生了超时的情况下将会返回 `false`，其他情况该方法都将返回 `True`。

在 3.1 版更改：很明显，方法总是返回 `None`。

17.1.8 定时器对象

此类表示一个操作应该在等待一定的时间之后运行—相当于一个定时器。`Timer` 类是 `Thread` 类的子类，因此可以像一个自定义线程一样工作。

与线程一样，通过调用 `start()` 方法启动定时器。而 `cancel()` 方法可以停止计时器（在计时结束前），定时器在执行其操作之前等待的时间间隔可能与用户指定的时间间隔不完全相同。

例如：

```
def hello():
    print("hello, world")

t = Timer(30.0, hello)
t.start()  # after 30 seconds, "hello, world" will be printed
```

class `threading.Timer` (*interval*, *function*, *args=None*, *kwargs=None*)

创建一个定时器，在经过 *interval* 秒的间隔事件后，将会用参数 *args* 和关键字参数 *kwargs* 调用 *function*。如果 *args* 为 `None`（默认值），则会使用一个空列表。如果 *kwargs* 为 `None`（默认值），则会使用一个空字典。

在 3.3 版更改：从工厂函数变为类。

cancel()

停止定时器并取消执行计时器将要执行的操作。仅当计时器仍处于等待状态时有效。

17.1.9 栅栏对象

3.2 新版功能。

栅栏类提供一个简单的同步原语，用于应对固定数量的线程需要彼此相互等待的情况。线程调用 `wait()` 方法后将阻塞，直到所有线程都调用了 `wait()` 方法。此时所有线程将被同时释放。

栅栏对象可以被多次使用，但进程的数量不能改变。

这是一个使用简便的方法实现客户端进程与服务端进程同步的例子：

```
b = Barrier(2, timeout=5)

def server():
    start_server()
    b.wait()
    while True:
        connection = accept_connection()
        process_server_connection(connection)

def client():
    b.wait()
    while True:
        connection = make_connection()
        process_client_connection(connection)
```

class `threading.Barrier` (*parties*, *action=None*, *timeout=None*)

创建一个需要 *parties* 个线程的栅栏对象。如果提供了可调用的 *action* 参数，它会在所有线程被释放时在其中一个线程中自动调用。*timeout* 是默认的超时时间，如果没有在 `wait()` 方法中指定超时时间的话。

wait (*timeout=None*)

冲出栅栏。当栅栏中所有线程都已经调用了这个函数，它们将同时被释放。如果提供了 *timeout* 参数，这里的 *timeout* 参数优先于创建栅栏对象时提供的 *timeout* 参数。

函数返回值是一个整数，取值范围在 0 到 *parties* - 1，在每个线程中的返回值不相同。可用于从所有线程中选择唯一的一个线程执行一些特别的工作。例如：

```
i = barrier.wait()
if i == 0:
    # Only one thread needs to print this
    print("passed the barrier")
```

如果创建栅栏对象时在构造函数中提供了 *action* 参数，它将在其中一个线程释放前被调用。如果此调用引发了异常，栅栏对象将进入损坏态。

如果发生了超时，栅栏对象将进入破损态。

如果栅栏对象进入破损态，或重置栅栏时仍有线程等待释放，将会引发 `BrokenBarrierError` 异常。

reset ()

重置栅栏为默认的初始态。如果栅栏中仍有线程等待释放，这些线程将会收到 `BrokenBarrierError` 异常。

注意使用此函数时，如果有某些线程状态未知，则可能需其它的同步来确保线程已被释放。如果栅栏进入了破损态，最好废弃它并新建一个栅栏。

abort ()

使栅栏进入破损态。这将导致所有已经调用和未来调用的 `wait()` 方法中引发 `BrokenBarrierError` 异常。使用这个方法的一种情况是需要中止程序以避免死锁。

更好的方式是：创建栅栏时提供一个合理的超时时间，来自动避免某个线程出错。

parties

冲出栅栏所需要的线程数量。

n_waiting

当前时刻正在栅栏中阻塞的线程数量。

broken

一个布尔值，值为 `True` 表明栅栏为破损态。

exception `threading.BrokenBarrierError`

异常类，是 `RuntimeError` 异常的子类，在 `Barrier` 对象重置时仍有线程阻塞时和对象进入破损态时被引发。

17.1.10 Using locks, conditions, and semaphores in the with statement

这个模块提供的带有 `acquire()` 和 `release()` 方法的对象，可以被用作 `with` 语句的上下文管理器。当进入语句块时 `acquire()` 方法会被调用，退出语句块时 `release()` 会被调用。因此，以下片段：

```
with some_lock:
    # do something...
```

相当于：

```
some_lock.acquire()
try:
    # do something...
finally:
    some_lock.release()
```

现在`Lock`、`RLock`、`Condition`、`Semaphore`和`BoundedSemaphore`对象可以用作`with`语句的上下文管理器。

17.2 multiprocessing — 基于进程的并行

源代码 `Lib/multiprocessing/`

17.2.1 概述

`multiprocessing` 是一个用与`threading`模块相似 API 的支持产生进程的包。`multiprocessing`包同时提供本地和远程并发，使用子进程代替线程，有效避免`Global Interpreter Lock`带来的影响。因此，`multiprocessing`模块允许程序员充分利用机器上的多个核心。Unix 和 Windows 上都可以运行。

`multiprocessing`模块还引入了在`threading`模块中没有类似物的 API。这方面的一个主要例子是`Pool`对象，它提供了一种方便的方法，可以跨多个输入值并行化函数的执行，跨进程分配输入数据（数据并行）。以下示例演示了在模块中定义此类函数的常见做法，以便子进程可以成功导入该模块。这个数据并行的基本例子使用`Pool`，

```
from multiprocessing import Pool

def f(x):
    return x*x

if __name__ == '__main__':
    with Pool(5) as p:
        print(p.map(f, [1, 2, 3]))
```

将打印到标准输出

```
[1, 4, 9]
```

Process 类

在`multiprocessing`中，通过创建一个`Process`对象然后调用它的`start()`方法来生成进程。`Process`和`threading.Thread` API 相同。一个简单的多进程序例是：

```
from multiprocessing import Process

def f(name):
    print('hello', name)

if __name__ == '__main__':
    p = Process(target=f, args=('bob',))
    p.start()
    p.join()
```

要显示所涉及的所有进程 ID，这是一个扩展示例：

```
from multiprocessing import Process
import os

def info(title):
    print(title)
    print('module name:', __name__)
    print('parent process:', os.getppid())
    print('process id:', os.getpid())

def f(name):
    info('function f')
    print('hello', name)

if __name__ == '__main__':
    info('main line')
    p = Process(target=f, args=('bob',))
    p.start()
    p.join()
```

为了解释为什么 `if __name__ == '__main__':` 部分是必需的，请参见 *Programming guidelines*。

上下文和启动方法

根据不同的平台，*multiprocessing* 支持三种启动进程的方法。这些启动方法有

spawn 父进程启动一个新的 Python 解释器进程。子进程只会继承那些运行进程对象的 *run()* 方法所需的资源。特别是父进程中非必须的文件描述符和句柄不会被继承。相对于使用 *fork* 或者 *forkserver*，使用这个方法启动进程相当慢。

可在 Unix 和 Windows 上使用。Windows 上的默认设置。

fork 父进程使用 *os.fork()* 来产生 Python 解释器分叉。子进程在开始时实际上与父进程相同。父进程的所有资源都由子进程继承。请注意，安全分叉多线程进程是棘手的。

只存在于 Unix。Unix 中的默认值。

forkserver 程序启动并选择 *forkserver* 启动方法时，将启动服务器进程。从那时起，每当需要一个新进程时，父进程就会连接到服务器并请求它分叉一个新进程。分叉服务器进程是单线程的，因此使用 *os.fork()* 是安全的。没有不必要的资源被继承。

可在 Unix 平台上使用，支持通过 Unix 管道传递文件描述符。

在 3.4 版更改：*spawn* 在所有 unix 平台上添加，并且为一些 unix 平台添加了 *forkserver*。子进程不再继承 Windows 上的所有上级进程可继承的句柄。

在 Unix 上使用 *spawn* 或 *forkserver* 启动方法也将启动一个信号量跟踪器进程，该进程跟踪由程序进程创建的未链接的命名信号量。当所有进程退出时，信号量跟踪器取消链接任何剩余的信号量。通常不应该有，但如果一个进程被信号杀死，可能会有一些“泄露”的信号量。（取消链接命名的信号量是一个严重的问题，因为系统只允许有限的数量，并且在下次重新启动之前它们不会自动取消链接。）

要选择一个启动方法，你应该在主模块的 `if __name__ == '__main__':` 子句中调用 *set_start_method()*。例如：

```
import multiprocessing as mp

def foo(q):
    q.put('hello')
```

(下页继续)

(续上页)

```

if __name__ == '__main__':
    mp.set_start_method('spawn')
    q = mp.Queue()
    p = mp.Process(target=foo, args=(q,))
    p.start()
    print(q.get())
    p.join()

```

在程序中 `set_start_method()` 不应该被多次调用。

或者，你可以使用 `get_context()` 来获取上下文对象。上下文对象与多处理模块具有相同的 API，并允许在同一程序中使用多个启动方法。：

```

import multiprocessing as mp

def foo(q):
    q.put('hello')

if __name__ == '__main__':
    ctx = mp.get_context('spawn')
    q = ctx.Queue()
    p = ctx.Process(target=foo, args=(q,))
    p.start()
    print(q.get())
    p.join()

```

请注意，与一个上下文相关的对象可能与不同上下文的进程不兼容。特别是，使用 `fork` 上下文创建的锁不能传递给使用 `spawn` 或 `forkserver` 启动方法启动的进程。

想要使用特定启动方法的库应该使用 `get_context()` 以避免干扰库用户的选择。

在进程之间交换对象

`multiprocessing` 支持进程之间的两种通信通道：

队列

`Queue` 类是一个近似 `queue.Queue` 的克隆。例如：

```

from multiprocessing import Process, Queue

def f(q):
    q.put([42, None, 'hello'])

if __name__ == '__main__':
    q = Queue()
    p = Process(target=f, args=(q,))
    p.start()
    print(q.get())    # prints "[42, None, 'hello']"
    p.join()

```

队列是线程和进程安全的。

管道

`Pipe()` 函数返回一个由管道连接的连接对象，默认情况下是双工（双向）。例如：

```

from multiprocessing import Process, Pipe

def f(conn):
    conn.send([42, None, 'hello'])
    conn.close()

if __name__ == '__main__':
    parent_conn, child_conn = Pipe()
    p = Process(target=f, args=(child_conn,))
    p.start()
    print(parent_conn.recv())    # prints "[42, None, 'hello']"
    p.join()

```

返回的两个连接对象`Pipe()`表示管道的两端。每个连接对象都有`send()`和`recv()`方法（相互之间的）。请注意，如果两个进程（或线程）同时尝试读取或写入管道的同一端，则管道中的数据可能会损坏。当然，同时使用管道的不同端的进程不存在损坏的风险。

进程之间的同步

`multiprocessing` 包含来自`threading`的所有同步基本体的等价物。例如，可以使用锁来确保一次只有一个进程打印到标准输出：

```

from multiprocessing import Process, Lock

def f(l, i):
    l.acquire()
    try:
        print('hello world', i)
    finally:
        l.release()

if __name__ == '__main__':
    lock = Lock()

    for num in range(10):
        Process(target=f, args=(lock, num)).start()

```

不使用来自不同进程的锁输出容易产生混淆。

在进程之间共享状态

如上所述，在进行并发编程时，通常最好尽量避免使用共享状态。使用多个进程时尤其如此。

但是，如果你真的需要使用一些共享数据，那么`multiprocessing`提供了两种方法。

共享内存

可以使用`Value`或`Array`将数据存储在共享内存映射中。例如，以下代码：

```

from multiprocessing import Process, Value, Array

def f(n, a):
    n.value = 3.1415927
    for i in range(len(a)):
        a[i] = -a[i]

```

(下页继续)

(续上页)

```

if __name__ == '__main__':
    num = Value('d', 0.0)
    arr = Array('i', range(10))

    p = Process(target=f, args=(num, arr))
    p.start()
    p.join()

    print(num.value)
    print(arr[:])

```

将打印

```

3.1415927
[0, -1, -2, -3, -4, -5, -6, -7, -8, -9]

```

创建 `num` 和 `arr` 时使用的 `'d'` 和 `'i'` 参数是 `array` 模块使用的类型的 `typecode`：`'d'` 表示双精度浮点数，`'i'` 表示有符号整数。这些共享对象将是进程和线程安全的。

为了更灵活地使用共享内存，可以使用 `multiprocessing.sharedctypes` 模块，该模块支持创建从共享内存分配的任意 `ctypes` 对象。

服务器进程

由 `Manager()` 返回的管理器对象控制一个服务器进程，该进程保存 Python 对象并允许其他进程使用代理操作它们。

`Manager()` 返回的管理器支持类型：`list`、`dict`、`Namespace`、`Lock`、`RLock`、`Semaphore`、`BoundedSemaphore`、`Condition`、`Event`、`Barrier`、`Queue`、`Value` 和 `Array`。例如

```

from multiprocessing import Process, Manager

def f(d, l):
    d[1] = '1'
    d['2'] = 2
    d[0.25] = None
    l.reverse()

if __name__ == '__main__':
    with Manager() as manager:
        d = manager.dict()
        l = manager.list(range(10))

        p = Process(target=f, args=(d, l))
        p.start()
        p.join()

        print(d)
        print(l)

```

将打印

```

{0.25: None, 1: '1', '2': 2}
[9, 8, 7, 6, 5, 4, 3, 2, 1, 0]

```

服务器进程管理器比使用共享内存对象更灵活，因为它们可以支持任意对象类型。此外，单个管理器可以通过网络由不同计算机上的进程共享。但是，它们比使用共享内存慢。

使用工作进程

`Pool` 类表示一个工作进程池。它具有允许以几种不同方式将任务分配到工作进程的方法。

例如:

```
from multiprocessing import Pool, TimeoutError
import time
import os

def f(x):
    return x*x

if __name__ == '__main__':
    # start 4 worker processes
    with Pool(processes=4) as pool:

        # print "[0, 1, 4, ..., 81]"
        print(pool.map(f, range(10)))

        # print same numbers in arbitrary order
        for i in pool.imap_unordered(f, range(10)):
            print(i)

        # evaluate "f(20)" asynchronously
        res = pool.apply_async(f, (20,)) # runs in *only* one process
        print(res.get(timeout=1))        # prints "400"

        # evaluate "os.getpid()" asynchronously
        res = pool.apply_async(os.getpid, ()) # runs in *only* one process
        print(res.get(timeout=1))           # prints the PID of that process

        # launching multiple evaluations asynchronously *may* use more processes
        multiple_results = [pool.apply_async(os.getpid, ()) for i in range(4)]
        print([res.get(timeout=1) for res in multiple_results])

        # make a single worker sleep for 10 secs
        res = pool.apply_async(time.sleep, (10,))
        try:
            print(res.get(timeout=1))
        except TimeoutError:
            print("We lacked patience and got a multiprocessing.TimeoutError")

        print("For the moment, the pool remains available for more work")

    # exiting the 'with'-block has stopped the pool
    print("Now the pool is closed and no longer available")
```

请注意，池的方法只能由创建它的进程使用。

注解：该软件包中的功能要求子项可以导入 `__main__` 模块。这包含在 *Programming guidelines* 中，但值得指出。这意味着一些示例，例如 `multiprocessing.pool.Pool` 示例在交互式解释器中不起作用。例如：

```
>>> from multiprocessing import Pool
>>> p = Pool(5)
>>> def f(x):
...     return x*x
```

(下页继续)

(续上页)

```

...
>>> p.map(f, [1,2,3])
Process PoolWorker-1:
Process PoolWorker-2:
Process PoolWorker-3:
Traceback (most recent call last):
AttributeError: 'module' object has no attribute 'f'
AttributeError: 'module' object has no attribute 'f'
AttributeError: 'module' object has no attribute 'f'

```

(如果你尝试这个，它实际上会以半随机的方式输出三个完整的回溯，然后你可能不得不以某种方式停止主进程。)

17.2.2 参考

multiprocessing 包大部分复制了 *threading* 模块的 API。

Process 和异常

class multiprocessing.**Process** (*group=None, target=None, name=None, args=(), kwargs={}, *, daemon=None*)

进程对象表示在单独进程中运行的活动。*Process* 类等价于 *threading.Thread*。

应始终使用关键字参数调用构造函数。*group* 应该始终是 *None*；它仅用于兼容 *threading.Thread*。*target* 是由 *run()* 方法调用的可调用对象。它默认为 *None*，意味着什么都没有被调用。*name* 是进程名称（有关详细信息，请参阅 *name*）。*args* 是目标调用的参数元组。*kwargs* 是目标调用的关键字参数字典。如果提供，则键参数 *daemon* 将进程 *daemon* 标志设置为 *True* 或 *False*。如果是 *None*（默认值），则该标志将从创建的进程继承。

默认情况下，不会将任何参数传递给 *target*。

如果子类重写构造函数，它必须确保它在对进程执行任何其他操作之前调用基类构造函数 (*Process.__init__()*)。

在 3.3 版更改：加入 *daemon* 参数。

run()

表示进程活动的方法。

你可以在子类中重载此方法。标准 *run()* 方法调用传递给对象构造函数的可调用对象作为目标参数（如果有），分别从 *args* 和 *kwargs* 参数中获取顺序和关键字参数。

start()

启动进程活动。

每个进程对象最多只能调用一次。它安排对象的 *run()* 方法在一个单独的进程中调用。

join([*timeout*])

如果可选参数 *timeout* 是 *None*（默认值），则该方法将阻塞，直到调用 *join()* 方法的进程终止。如果 *timeout* 是一个正数，它最多会阻塞 *timeout* 秒。请注意，如果进程终止或方法超时，则该方法返回 *None*。检查进程的 *exitcode* 以确定它是否终止。

一个进程可以合并多次。

进程无法并入自身，因为这会导致死锁。尝试在启动进程之前合并进程是错误的。

name

进程的名称。该名称是一个字符串，仅用于识别目的。它没有语义。可以为多个进程指定相同的名称。

初始名称由构造器设定。如果没有为构造器提供显式名称，则会构造一个形式为 ‘Process-N₁:N₂:...:N_k’ 的名称，其中每个 N_k 是其父亲的第 N 个孩子。

is_alive()

返回进程是否还活着。

粗略地说，从 `start()` 方法返回到子进程终止之前，进程对象仍处于活动状态。

daemon

进程的守护标志，一个布尔值。这必须在 `start()` 被调用之前设置。

初始值继承自创建进程。

当进程退出时，它会尝试终止其所有守护进程子进程。

请注意，不允许守护进程创建子进程。否则，守护进程会在子进程退出时终止其子进程。另外，这些 **不是** Unix 守护进程或服务，它们是正常进程，如果非守护进程已经退出，它们将被终止（并且不被合并）。

除了 `threading.Thread` API，`Process` 对象还支持以下属性和方法：

pid

返回进程 ID。在生成该进程之前，这将是 `None`。

exitcode

的退出进程出代码。如果进程尚未终止，这将是 `None`。负值 `-N` 表示孩子被信号 `N` 终止。

authkey

进程的身份验证密钥（字节字符串）。

当 `multiprocessing` 初始化时，主进程使用 `os.urandom()` 分配一个随机字符串。

当创建 `Process` 对象时，它将继承其父进程的身份验证密钥，尽管可以通过将 `authkey` 设置为另一个字节字符串来更改。

参见 *Authentication keys*。

sentinel

系统对象的数字句柄，当进程结束时将变为 “ready”。

如果要使用 `multiprocessing.connection.wait()` 一次等待多个事件，可以使用此值。否则调用 `join()` 更简单。

在 Windows 上，这是一个操作系统句柄，可以与 `WaitForSingleObject` 和 `WaitForMultipleObjects` 系列 API 调用一起使用。在 Unix 上，这是一个文件描述符，可以使用来自 `select` 模块的原语。

3.3 新版功能。

terminate()

终止进程。在 Unix 上，这是使用 `SIGTERM` 信号完成的；在 Windows 上使用 `TerminateProcess()`。请注意，不会执行退出处理程序和 `finally` 子句等。

请注意，进程的后代进程将不会被终止——它们将简单地变成孤立的。

警告：如果在关联进程使用管道或队列时使用此方法，则管道或队列可能会损坏，并可能无法被其他进程使用。类似地，如果进程已获得锁或信号量等，则终止它可能导致其他进程死锁。

注意`start()`、`join()`、`is_alive()`、`terminate()`和`exitcode`方法只能由创建进程对象的进程调用。

Process 一些方法的示例用法：

```
>>> import multiprocessing, time, signal
>>> p = multiprocessing.Process(target=time.sleep, args=(1000,))
>>> print(p, p.is_alive())
<Process(Process-1, initial)> False
>>> p.start()
>>> print(p, p.is_alive())
<Process(Process-1, started)> True
>>> p.terminate()
>>> time.sleep(0.1)
>>> print(p, p.is_alive())
<Process(Process-1, stopped[SIGTERM])> False
>>> p.exitcode == -signal.SIGTERM
True
```

exception `multiprocessing.ProcessError`

所有*multiprocessing*异常的基类。

exception `multiprocessing.BufferTooShort`

当提供的缓冲区对象太小而无法读取消息时，`Connection.recv_bytes_into()`引发的异常。

如果 `e` 是一个 *BufferTooShort* 实例，那么 `e.args[0]` 将把消息作为字节字符串给出。

exception `multiprocessing.AuthenticationError`

出现身份验证错误时引发。

exception `multiprocessing.TimeoutError`

有超时的方法超时引发。

管道和队列

使用多进程时，一般使用消息机制实现进程间通信，尽可能避免使用同步原语，例如锁。

消息机制包含：*Pipe()*（可以用于在两个进程间传递消息），以及队列（能够在多个生产者和消费者之间通信）。

The *Queue*, *SimpleQueue* and *JoinableQueue* types are multi-producer, multi-consumer FIFO queues modelled on the *queue.Queue* class in the standard library. They differ in that *Queue* lacks the *task_done()* and *join()* methods introduced into Python 2.5's *queue.Queue* class.

如果你使用了 *JoinableQueue*，那么你**必须**对每个已经移出队列的任务调用 *JoinableQueue.task_done()*。不然的话用于统计未完成任务的信号量最终会溢出并抛出异常。

另外还可以通过使用一个管理器对象创建一个共享队列，详见 *Managers*。

注解： *multiprocessing* 使用了普通的 *queue.Empty* 和 *queue.Full* 异常去表示超时。你需要从 *queue* 中导入它们，因为它们并不在 *multiprocessing* 的命名空间中。

注解： 当一个对象被放入一个队列中时，这个对象首先会被一个后台线程用 *pickle* 序列化，并将序列化后的数据通过一个底层管道的管道传递到队列中。这种做法会有点让人惊讶，但一般不会出现什么问题。如果它们确实妨碍了你，你可以使用一个由管理器 *manager* 创建的队列替换它。

- (1) 将一个对象放入一个空队列后，可能需要极小的延迟，队列的方法`empty()` 才会返回`False`。而`get_nowait()` 可以不抛出`queue.Empty` 直接返回。
- (2) 如果有多个进程同时将对象放入队列，那么在队列的另一端接受到的对象可能是无序的。但是由同一个进程放入的多个对象的顺序在另一端输出时总是一样的。

警告： 如果一个进程通过调用`Process.terminate()` 或`os.kill()` 在尝试使用`Queue` 期间被终止了，那么队列中的数据很可能被破坏。这可能导致其他进程在尝试使用该队列时遇到异常。

警告： 正如刚才提到的，如果一个子进程将一些对象放进队列中 (并且它没有用`JoinableQueue.cancel_join_thread` 方法)，那么这个进程在所有缓冲区的对象被刷新进管道之前，是不会终止的。

这意味着，除非你确定所有放入队列中的对象都已经被消费了，否则如果你试图等待这个进程，你可能会陷入死锁中。相似地，如果该子进程不是后台进程，那么父进程可能在试图等待所有非后台进程退出时挂起。

注意用管理器创建的队列不存在这个问题，详见[Programming guidelines](#)。

该示例 展示了如何使用队列实现进程间通信。

`multiprocessing.Pipe([duplex])`

返回一对 `Connection` 对象 `` (conn1, conn2)`, 分别表示管道的两端。

如果 `duplex` 被置为 `True` (默认值)，那么该管道是双向的。如果 `duplex` 被置为 `False`，那么该管道是单向的，即 `conn1` 只能用于接收消息，而 `conn2` 仅能用于发送消息。

`class multiprocessing.Queue([maxsize])`

返回一个使用一个管道和少量锁和信号量实现的共享队列实例。当一个进程将一个对象放进队列中时，一个写入线程会启动并将对象从缓冲区写入管道中。

一旦超时，将抛出标准库`queue` 模块中常见的异常`queue.Empty` 和`queue.Full`。

除了`task_done()` 和`join()` 之外，`Queue` 实现了标准库类`queue.Queue` 中所有的方法。

`qsize()`

返回队列的大致长度。由于多线程或者多进程的上下文，这个数字是不可靠的。

注意，在 `Unix` 平台上，例如 `Mac OS X`，这个方法可能会抛出`NotImplementedError` 异常，因为该平台没有实现 `sem_getvalue()`。

`empty()`

如果队列是空的，返回 `True`，反之返回 `False`。由于多线程或多进程的环境，该状态是不可靠的。

`full()`

如果队列是满的，返回 `True`，反之返回 `False`。由于多线程或多进程的环境，该状态是不可靠的。

`put(obj[, block[, timeout]])`

将 `obj` 放入队列。如果可选参数 `block` 是 `True` (默认值) 而且 `timeout` 是 `None` (默认值)，将会阻塞当前进程，直到有空的缓冲槽。如果 `timeout` 是正数，将会在阻塞了最多 `timeout` 秒之后还是没有可用的缓冲槽时抛出`queue.Full` 异常。反之 (`block` 是 `False` 时)，仅当有可用缓冲槽时才放入对象，否则抛出`queue.Full` 异常 (在这种情形下 `timeout` 参数会被忽略)。

`put_nowait(obj)`

相当于 `put(obj, False)`。

get ([*block*, *timeout*]])

从队列中取出并返回对象。如果可选参数 *block* 是 `True` (默认值) 而且 *timeout* 是 `None` (默认值), 将会阻塞当前进程, 直到队列中出现可用的对象。如果 *timeout* 是正数, 将会在阻塞了最多 *timeout* 秒之后还是没有可用的对象时抛出 `queue.Empty` 异常。反之 (*block* 是 `False` 时), 仅当有可用对象能够取出时返回, 否则抛出 `queue.Empty` 异常 (在这种情形下 *timeout* 参数会被忽略)。

get_nowait ()

相当于 `get(False)`。

`multiprocessing.Queue` 类有一些在 `queue.Queue` 类中没有出现的方法。这些方法在大多数情形下并不是必须的。

close ()

指示当前进程将不会再往队列中放入对象。一旦所有缓冲区中的数据被写入管道之后, 后台的线程会退出。这个方法在队列被 `gc` 回收时会自动调用。

join_thread ()

等待后台线程。这个方法仅在调用了 `close()` 方法之后可用。这会阻塞当前进程, 直到后台线程退出, 确保所有缓冲区中的数据都被写入管道中。

默认情况下, 如果一个不是队列创建者的进程试图退出, 它会尝试等待这个队列的后台线程。这个进程可以使用 `cancel_join_thread()` 让 `join_thread()` 方法什么都不做直接跳过。

cancel_join_thread ()

防止 `join_thread()` 方法阻塞当前进程。具体而言, 这防止进程退出时自动等待后台线程退出。详见 `join_thread()`。

可能这个方法称为 “`allow_exit_without_flush()`” 会更好。这有可能会产生正在排队进入队列的数据丢失, 大多数情况下你不需要用到这个方法, 仅当你不关心底层管道中可能丢失的数据, 只是希望进程能够马上退出时使用。

注解: 该类的功能依赖于宿主操作系统具有可用的共享信号量实现。否则该类将被禁用, 任何试图实例化一个 `Queue` 对象的操作都会抛出 `ImportError` 异常, 更多信息详见 [bpo-3770](#)。后续说明的任何专用队列对象亦如此。

class `multiprocessing.SimpleQueue`

这是一个简化的 `Queue` 类的实现, 很像带锁的 `Pipe`。

empty ()

如果队列为空返回 `True`, 否则返回 `False`。

get ()

从队列中移出并返回一个对象。

put (*item*)

将 *item* 放入队列。

class `multiprocessing.JoinableQueue` ([*maxsize*])

`JoinableQueue` 类是 `Queue` 的子类, 额外添加了 `task_done()` 和 `join()` 方法。

task_done ()

指出之前进入队列的任务已经完成。由队列的消费者进程使用。对于每次调用 `get()` 获取的任务, 执行完成后调用 `task_done()` 告诉队列该任务已经处理完成。

如果 `join()` 方法正在阻塞之中, 该方法会在所有对象都被处理完的时候返回 (即对之前使用 `put()` 放进队列中的所有对象都已经返回了对应的 `task_done()`)。

如果被调用的次数多于放入队列中的项目数量, 将引发 `ValueError` 异常。

join()

阻塞至队列中所有的元素都被接收和处理完毕。

当条目添加到队列的时候，未完成任务的计数就会增加。每当消费者进程调用`task_done()`表示这个条目已经被回收，该条目所有工作已经完成，未完成计数就会减少。当未完成计数降到零的时候，`join()`阻塞被解除。

杂项**multiprocessing.active_children()**

返回当前进程存活的子进程的列表。

调用该方法有“等待”已经结束的进程的副作用。

multiprocessing.cpu_count()

返回系统的 CPU 数量。

该数量不同于当前进程可以使用的 CPU 数量。可用的 CPU 数量可以由 `len(os.sched_getaffinity(0))` 方法获得。

可能引发`NotImplementedError`。

参见：

`os.cpu_count()`

multiprocessing.current_process()

返回与当前进程相对应的`Process`对象。

和`threading.current_thread()`相同。

multiprocessing.freeze_support()

为使用了`multiprocessing`的程序，提供冻结以产生 Windows 可执行文件的支持。(在 `py2exe`, `PyInstaller` 和 `cx_Freeze` 上测试通过)

需要在 main 模块的 `if __name__ == '__main__':` 该行之后马上调用该函数。例如：

```
from multiprocessing import Process, freeze_support

def f():
    print('hello world!')

if __name__ == '__main__':
    freeze_support()
    Process(target=f).start()
```

如果没有调用 `freeze_support()` 在尝试运行被冻结的可执行文件时会抛出`RuntimeError`异常。

对 `freeze_support()` 的调用在非 Windows 平台上是无效的。如果该模块在 Windows 平台的 Python 解释器中正常运行(该程序没有被冻结)，调用“`freeze_support()`”也是无效的。

multiprocessing.get_all_start_methods()

返回支持的启动方法的列表，该列表的首项即为默认选项。可能的启动方法有 `'fork'`, `'spawn'` 和 `'forkserver'`。在 Windows 中，只有 `'spawn'` 是可用的。Unix 平台总是支持 `'fork'` 和 `'spawn'`，且 `'fork'` 是默认值。

3.4 新版功能。

multiprocessing.get_context(method=None)

返回一个 Context 对象。该对象具有和`multiprocessing`模块相同的 API。

如果 *method* 设置成 `None` 那么将返回默认上下文对象。否则 *method* 应该是 `'fork'`, `'spawn'`, `'forkserver'`。如果指定的启动方法不存在, 将抛出 `ValueError` 异常。

3.4 新版功能.

`multiprocessing.get_start_method(allow_none=False)`

返回启动进程时使用的启动方法名。

如果启动方法已经固定, 并且 *allow_none* 被设置成 `False`, 那么启动方法将被固定为默认的启动方法, 并且返回其方法名。如果启动方法没有设定, 并且 *allow_none* 被设置成 `True`, 那么将返回 `None`。

返回值将为 `'fork'`, `'spawn'`, `'forkserver'` 或者 `None`。`'fork'` 是 Unix 的默认值, `'spawn'` 是 Windows 的默认值。

3.4 新版功能.

`multiprocessing.set_executable()`

设置在启动子进程时使用的 Python 解释器路径。(默认使用 `sys.executable`) 嵌入式编程人员可能需要这样做:

```
set_executable(os.path.join(sys.exec_prefix, 'pythonw.exe'))
```

以使他们可以创建子进程。

在 3.4 版更改: 现在在 Unix 平台上使用 `'spawn'` 启动方法时支持调用该方法。

`multiprocessing.set_start_method(method)`

设置启动子进程的方法。*method* 可以是 `'fork'`, `'spawn'` 或者 `'forkserver'`。

注意这最多只能调用一次, 并且需要藏在 `main` 模块中, 由 `if __name__ == '__main__':` 保护着。

3.4 新版功能.

注解: `multiprocessing` 并没有包含类似 `threading.active_count()`, `threading.enumerate()`, `threading.settrace()`, `threading.setprofile()`, `threading.Timer`, 或者 `threading.local` 的方法和类。

连接 (Connection) 对象

`Connection` 对象允许收发可以序列化的对象或字符串。它们可以看作面向消息的连接套接字。

通常使用 `Pipe` 创建 `Connection` 对象。详见: [Listeners and Clients](#).

class `multiprocessing.connection.Connection`

send(obj)

将一个对象发送到连接的另一端, 可以用 `recv()` 读取。

The object must be picklable. Very large pickles (approximately 32 MB+, though it depends on the OS) may raise a `ValueError` exception.

recv()

返回一个由另一端使用 `send()` 发送的对象。该方法会一直阻塞直到接收到对象。如果对端关闭了连接或者没有东西可接收, 将抛出 `EOFError` 异常。

fileno()

返回由连接对象使用的描述符或者句柄。

close()

关闭连接对象。

当连接对象被垃圾回收时会自动调用。

poll([timeout])

返回连接对象中是否有可以读取的数据。

如果未指定 *timeout*，此方法会马上返回。如果 *timeout* 是一个数字，则指定了最大阻塞的秒数。如果 *timeout* 是 *None*，那么将一直等待，不会超时。

注意通过使用 `multiprocessing.connection.wait()` 可以一次轮询多个连接对象。

send_bytes(buffer[, offset[, size]])

从一个 *bytes-like object*（字节类对象）对象中取出字节数组并作为一条完整消息发送。

If *offset* is given then data is read from that position in *buffer*. If *size* is given then that many bytes will be read from *buffer*. Very large buffers (approximately 32 MB+, though it depends on the OS) may raise a *ValueError* exception

recv_bytes([maxlength])

以字符串形式返回一条从连接对象另一端发送过来的字节数据。此方法在接收到数据前将一直阻塞。如果连接对象被对端关闭或者没有数据可读取，将抛出 *EOFError* 异常。

如果给定了 *maxlength* 并且消息短于 *maxlength* 那么将抛出 *OSError* 并且该连接对象将不再可读。

在 3.3 版更改: 曾经该函数抛出 *IOError*，现在这是 *OSError* 的别名。

recv_bytes_into(buffer[, offset])

将一条完整的字节数据消息读入 *buffer* 中并返回消息的字节数。此方法在接收到数据前将一直阻塞。如果连接对象被对端关闭或者没有数据可读取，将抛出 *EOFError* 异常。

buffer must be a writable *bytes-like object*. If *offset* is given then the message will be written into the buffer from that position. Offset must be a non-negative integer less than the length of *buffer* (in bytes).

如果缓冲区太小，则将引发 *BufferTooShort* 异常，并且完整的消息将会存放在异常实例 *e* 的 *e.args[0]* 中。

在 3.3 版更改: 现在连接对象自身可以通过 `Connection.send()` 和 `Connection.recv()` 在进程之间传递。

3.3 新版功能: 连接对象现已支持上下文管理协议 – 参见 `see 上下文管理器类型`。`__enter__()` 返回连接对象，`__exit__()` 会调用 `close()`。

例如:

```
>>> from multiprocessing import Pipe
>>> a, b = Pipe()
>>> a.send([1, 'hello', None])
>>> b.recv()
[1, 'hello', None]
>>> b.send_bytes(b'thank you')
>>> a.recv_bytes()
b'thank you'
>>> import array
>>> arr1 = array.array('i', range(5))
>>> arr2 = array.array('i', [0] * 10)
>>> a.send_bytes(arr1)
>>> count = b.recv_bytes_into(arr2)
>>> assert count == len(arr1) * arr1.itemsize
>>> arr2
array('i', [0, 1, 2, 3, 4, 0, 0, 0, 0, 0])
```

警告： The `Connection.recv()` method automatically unpickles the data it receives, which can be a security risk unless you can trust the process which sent the message.

因此，除非连接对象是由 `Pipe()` 产生的，在通过一些认证手段之前你应该只使用 `recv()` 和 `send()` 方法。参考 *Authentication keys*。

警告： 如果一个进程在试图读写管道时被终止了，那么管道中的数据很可能是不完整的，因为此时可能无法确定消息的边界。

同步原语

通常来说同步愿意在多进程环境中并不像它们在多线程环境中那么必要。参考 *threading* 模块的文档。

注意可以使用管理器对象创建同步原语，参考 *Managers*。

class multiprocessing.Barrier(*parties*[, *action*[, *timeout*]])

类似 *threading.Barrier* 的栅栏对象。

3.3 新版功能。

class multiprocessing.BoundedSemaphore(*value*)

非常类似 *threading.BoundedSemaphore* 的有界信号量对象。

一个小小的不同在于，它的 `acquire` 方法的第一个参数名是和 *Lock.acquire()* 一样的 *block*。

注解： 在 Mac OS X 平台上，该对象于 *Semaphore* 不同在于 `sem_getvalue()` 方法并没有在该平台上实现。

class multiprocessing.Condition(*lock*)

条件变量：*threading.Condition* 的别名。

指定的 *lock* 参数应该是 *multiprocessing* 模块中的 *Lock* 或者 *RLock* 对象。

在 3.3 版更改：新增了 `wait_for()` 方法。

class multiprocessing.Event

A clone of *threading.Event*.

class multiprocessing.Lock

原始锁（非递归锁）对象，类似于 *threading.Lock*。一旦一个进程或者线程拿到了锁，后续的任何其他进程或线程的其他请求都会被阻塞直到锁被释放。任何进程或线程都可以释放锁。除非另有说明，否则 *multiprocessing.Lock* 用于进程或者线程的概念和行为都和 *threading.Lock* 一致。

注意 *Lock* 实际上是一个工厂函数。它返回由默认上下文初始化的 *multiprocessing.synchronize.Lock* 对象。

Lock supports the *context manager* protocol and thus may be used in `with` statements.

acquire (*block=True*, *timeout=None*)

获得锁，阻塞或非阻塞的。

如果 *block* 参数被设为 `True`（默认值），对该方法的调用在锁处于释放状态之前都会阻塞，然后将锁设置为锁住状态并返回 `True`。需要注意的是第一个参数名与 *threading.Lock.acquire()* 的不同。

如果 *block* 参数被设置成 `False`，方法的调用将不会阻塞。如果锁当前处于锁住状态，将返回 `False`；否则将锁设置成锁住状态，并返回 `True`。

When invoked with a positive, floating-point value for *timeout*, block for at most the number of seconds specified by *timeout* as long as the lock can not be acquired. Invocations with a negative value for *timeout* are equivalent to a *timeout* of zero. Invocations with a *timeout* value of `None` (the default) set the timeout period to infinite. Note that the treatment of negative or `None` values for *timeout* differs from the implemented behavior in `threading.Lock.acquire()`. The *timeout* argument has no practical implications if the *block* argument is set to `False` and is thus ignored. Returns `True` if the lock has been acquired or `False` if the timeout period has elapsed.

release()

Release a lock. This can be called from any process or thread, not only the process or thread which originally acquired the lock.

Behavior is the same as in `threading.Lock.release()` except that when invoked on an unlocked lock, a `ValueError` is raised.

class multiprocessing.RLock

A recursive lock object: a close analog of `threading.RLock`. A recursive lock must be released by the process or thread that acquired it. Once a process or thread has acquired a recursive lock, the same process or thread may acquire it again without blocking; that process or thread must release it once for each time it has been acquired.

Note that `RLock` is actually a factory function which returns an instance of `multiprocessing.synchronize.RLock` initialized with a default context.

`RLock` supports the *context manager* protocol and thus may be used in `with` statements.

acquire (block=True, timeout=None)

获得锁，阻塞或非阻塞的。

When invoked with the *block* argument set to `True`, block until the lock is in an unlocked state (not owned by any process or thread) unless the lock is already owned by the current process or thread. The current process or thread then takes ownership of the lock (if it does not already have ownership) and the recursion level inside the lock increments by one, resulting in a return value of `True`. Note that there are several differences in this first argument's behavior compared to the implementation of `threading.RLock.acquire()`, starting with the name of the argument itself.

When invoked with the *block* argument set to `False`, do not block. If the lock has already been acquired (and thus is owned) by another process or thread, the current process or thread does not take ownership and the recursion level within the lock is not changed, resulting in a return value of `False`. If the lock is in an unlocked state, the current process or thread takes ownership and the recursion level is incremented, resulting in a return value of `True`.

Use and behaviors of the *timeout* argument are the same as in `Lock.acquire()`. Note that some of these behaviors of *timeout* differ from the implemented behaviors in `threading.RLock.acquire()`.

release()

Release a lock, decrementing the recursion level. If after the decrement the recursion level is zero, reset the lock to unlocked (not owned by any process or thread) and if any other processes or threads are blocked waiting for the lock to become unlocked, allow exactly one of them to proceed. If after the decrement the recursion level is still nonzero, the lock remains locked and owned by the calling process or thread.

Only call this method when the calling process or thread owns the lock. An `AssertionError` is raised if this method is called by a process or thread other than the owner or if the lock is in an unlocked (un-owned) state. Note that the type of exception raised in this situation differs from the implemented behavior in `threading.RLock.release()`.

class multiprocessing.Semaphore ([value])

A semaphore object: a close analog of `threading.Semaphore`.

一个小小的不同在于，它的 `acquire` 方法的第一个参数名是和 `Lock.acquire()` 一样的 `block`。

注解： On Mac OS X, `sem_timedwait` is unsupported, so calling `acquire()` with a timeout will emulate that function's behavior using a sleeping loop.

注解： If the SIGINT signal generated by Ctrl-C arrives while the main thread is blocked by a call to `BoundedSemaphore.acquire()`, `Lock.acquire()`, `RLock.acquire()`, `Semaphore.acquire()`, `Condition.acquire()` or `Condition.wait()` then the call will be immediately interrupted and `KeyboardInterrupt` will be raised.

This differs from the behaviour of `threading` where SIGINT will be ignored while the equivalent blocking calls are in progress.

注解： Some of this package's functionality requires a functioning shared semaphore implementation on the host operating system. Without one, the `multiprocessing.synchronize` module will be disabled, and attempts to import it will result in an `ImportError`. See [bpo-3770](#) for additional information.

Shared ctypes Objects

It is possible to create shared objects using shared memory which can be inherited by child processes.

`multiprocessing.Value` (*typecode_or_type*, **args*, *lock=True*)

Return a `ctypes` object allocated from shared memory. By default the return value is actually a synchronized wrapper for the object. The object itself can be accessed via the `value` attribute of a `Value`.

typecode_or_type determines the type of the returned object: it is either a `ctypes` type or a one character typecode of the kind used by the `array` module. **args* is passed on to the constructor for the type.

If *lock* is `True` (the default) then a new recursive lock object is created to synchronize access to the value. If *lock* is a `Lock` or `RLock` object then that will be used to synchronize access to the value. If *lock* is `False` then access to the returned object will not be automatically protected by a lock, so it will not necessarily be “process-safe”.

Operations like `+=` which involve a read and write are not atomic. So if, for instance, you want to atomically increment a shared value it is insufficient to just do

```
counter.value += 1
```

Assuming the associated lock is recursive (which it is by default) you can instead do

```
with counter.get_lock():
    counter.value += 1
```

Note that *lock* is a keyword-only argument.

`multiprocessing.Array` (*typecode_or_type*, *size_or_initializer*, *, *lock=True*)

Return a `ctypes` array allocated from shared memory. By default the return value is actually a synchronized wrapper for the array.

typecode_or_type determines the type of the elements of the returned array: it is either a `ctypes` type or a one character typecode of the kind used by the `array` module. If *size_or_initializer* is an integer, then it determines the length of the array, and the array will be initially zeroed. Otherwise, *size_or_initializer* is a sequence which is used to initialize the array and whose length determines the length of the array.

If *lock* is `True` (the default) then a new lock object is created to synchronize access to the value. If *lock* is a `Lock` or `RLock` object then that will be used to synchronize access to the value. If *lock* is `False` then access to the returned object will not be automatically protected by a lock, so it will not necessarily be “process-safe”.

Note that *lock* is a keyword only argument.

Note that an array of `ctypes.c_char` has *value* and *raw* attributes which allow one to use it to store and retrieve strings.

The `multiprocessing.sharedctypes` module

The `multiprocessing.sharedctypes` module provides functions for allocating `ctypes` objects from shared memory which can be inherited by child processes.

注解： Although it is possible to store a pointer in shared memory remember that this will refer to a location in the address space of a specific process. However, the pointer is quite likely to be invalid in the context of a second process and trying to dereference the pointer from the second process may cause a crash.

`multiprocessing.sharedctypes.RawArray` (*typecode_or_type*, *size_or_initializer*)

Return a `ctypes` array allocated from shared memory.

typecode_or_type determines the type of the elements of the returned array: it is either a `ctypes` type or a one character typecode of the kind used by the `array` module. If *size_or_initializer* is an integer then it determines the length of the array, and the array will be initially zeroed. Otherwise *size_or_initializer* is a sequence which is used to initialize the array and whose length determines the length of the array.

Note that setting and getting an element is potentially non-atomic – use `Array()` instead to make sure that access is automatically synchronized using a lock.

`multiprocessing.sharedctypes.RawValue` (*typecode_or_type*, **args*)

Return a `ctypes` object allocated from shared memory.

typecode_or_type determines the type of the returned object: it is either a `ctypes` type or a one character typecode of the kind used by the `array` module. **args* is passed on to the constructor for the type.

Note that setting and getting the value is potentially non-atomic – use `Value()` instead to make sure that access is automatically synchronized using a lock.

Note that an array of `ctypes.c_char` has *value* and *raw* attributes which allow one to use it to store and retrieve strings – see documentation for `ctypes`.

`multiprocessing.sharedctypes.Array` (*typecode_or_type*, *size_or_initializer*, ***, *lock=True*)

The same as `RawArray()` except that depending on the value of *lock* a process-safe synchronization wrapper may be returned instead of a raw `ctypes` array.

If *lock* is `True` (the default) then a new lock object is created to synchronize access to the value. If *lock* is a `Lock` or `RLock` object then that will be used to synchronize access to the value. If *lock* is `False` then access to the returned object will not be automatically protected by a lock, so it will not necessarily be “process-safe”.

Note that *lock* is a keyword-only argument.

`multiprocessing.sharedctypes.Value` (*typecode_or_type*, **args*, *lock=True*)

The same as `RawValue()` except that depending on the value of *lock* a process-safe synchronization wrapper may be returned instead of a raw `ctypes` object.

If *lock* is `True` (the default) then a new lock object is created to synchronize access to the value. If *lock* is a `Lock` or `RLock` object then that will be used to synchronize access to the value. If *lock* is `False` then access to the returned object will not be automatically protected by a lock, so it will not necessarily be “process-safe”.

Note that *lock* is a keyword-only argument.

`multiprocessing.sharedctypes.copy(obj)`

Return a ctypes object allocated from shared memory which is a copy of the ctypes object *obj*.

`multiprocessing.sharedctypes.synchronized(obj[, lock])`

Return a process-safe wrapper object for a ctypes object which uses *lock* to synchronize access. If *lock* is None (the default) then a `multiprocessing.RLock` object is created automatically.

A synchronized wrapper will have two methods in addition to those of the object it wraps: `get_obj()` returns the wrapped object and `get_lock()` returns the lock object used for synchronization.

Note that accessing the ctypes object through the wrapper can be a lot slower than accessing the raw ctypes object.

在 3.5 版更改: Synchronized objects support the *context manager* protocol.

The table below compares the syntax for creating shared ctypes objects from shared memory with the normal ctypes syntax. (In the table `MyStruct` is some subclass of `ctypes.Structure`.)

ctypes	使用类型的共享 ctypes	使用 typecode 的共享 ctypes
<code>c_double(2.4)</code>	<code>RawValue(c_double, 2.4)</code>	<code>RawValue('d', 2.4)</code>
<code>MyStruct(4, 6)</code>	<code>RawValue(MyStruct, 4, 6)</code>	
<code>(c_short * 7)()</code>	<code>RawArray(c_short, 7)</code>	<code>RawArray('h', 7)</code>
<code>(c_int * 3)(9, 2, 8)</code>	<code>RawArray(c_int, (9, 2, 8))</code>	<code>RawArray('i', (9, 2, 8))</code>

Below is an example where a number of ctypes objects are modified by a child process:

```
from multiprocessing import Process, Lock
from multiprocessing.sharedctypes import Value, Array
from ctypes import Structure, c_double

class Point(Structure):
    _fields_ = [('x', c_double), ('y', c_double)]

def modify(n, x, s, A):
    n.value **= 2
    x.value **= 2
    s.value = s.value.upper()
    for a in A:
        a.x **= 2
        a.y **= 2

if __name__ == '__main__':
    lock = Lock()

    n = Value('i', 7)
    x = Value(c_double, 1.0/3.0, lock=False)
    s = Array('c', b'hello world', lock=lock)
    A = Array(Point, [(1.875, -6.25), (-5.75, 2.0), (2.375, 9.5)], lock=lock)

    p = Process(target=modify, args=(n, x, s, A))
    p.start()
    p.join()

    print(n.value)
    print(x.value)
    print(s.value)
    print([(a.x, a.y) for a in A])
```

The results printed are

```
49
0.11111111111111111
HELLO WORLD
[(3.515625, 39.0625), (33.0625, 4.0), (5.640625, 90.25)]
```

Managers

Managers provide a way to create data which can be shared between different processes, including sharing over a network between processes running on different machines. A manager object controls a server process which manages *shared objects*. Other processes can access the shared objects by using proxies.

`multiprocessing.Manager()`

Returns a started *SyncManager* object which can be used for sharing objects between processes. The returned manager object corresponds to a spawned child process and has methods which will create shared objects and return corresponding proxies.

Manager processes will be shutdown as soon as they are garbage collected or their parent process exits. The manager classes are defined in the `multiprocessing.managers` module:

class `multiprocessing.managers.BaseManager([address[, authkey]])`

Create a BaseManager object.

Once created one should call `start()` or `get_server().serve_forever()` to ensure that the manager object refers to a started manager process.

address is the address on which the manager process listens for new connections. If *address* is `None` then an arbitrary one is chosen.

authkey is the authentication key which will be used to check the validity of incoming connections to the server process. If *authkey* is `None` then `current_process().authkey` is used. Otherwise *authkey* is used and it must be a byte string.

start `([initializer[, initargs]])`

Start a subprocess to start the manager. If *initializer* is not `None` then the subprocess will call `initializer(*initargs)` when it starts.

get_server `()`

Returns a *Server* object which represents the actual server under the control of the Manager. The *Server* object supports the `serve_forever()` method:

```
>>> from multiprocessing.managers import BaseManager
>>> manager = BaseManager(address=(' ', 50000), authkey=b'abc')
>>> server = manager.get_server()
>>> server.serve_forever()
```

Server additionally has an *address* attribute.

connect `()`

Connect a local manager object to a remote manager process:

```
>>> from multiprocessing.managers import BaseManager
>>> m = BaseManager(address=('127.0.0.1', 5000), authkey=b'abc')
>>> m.connect()
```

shutdown `()`

Stop the process used by the manager. This is only available if `start()` has been used to start the server process.

它可以被多次调用。

register (*typeid*[, *callable*[, *proxytype*[, *exposed*[, *method_to_typeid*[, *create_method*]]]]])

A classmethod which can be used for registering a type or callable with the manager class.

typeid is a “type identifier” which is used to identify a particular type of shared object. This must be a string.

callable is a callable used for creating objects for this type identifier. If a manager instance will be connected to the server using the `connect()` method, or if the *create_method* argument is `False` then this can be left as `None`.

proxytype is a subclass of `BaseProxy` which is used to create proxies for shared objects with this *typeid*. If `None` then a proxy class is created automatically.

exposed is used to specify a sequence of method names which proxies for this *typeid* should be allowed to access using `BaseProxy._callmethod()`. (If *exposed* is `None` then `proxytype._exposed_` is used instead if it exists.) In the case where no exposed list is specified, all “public methods” of the shared object will be accessible. (Here a “public method” means any attribute which has a `__call__()` method and whose name does not begin with `'_'`.)

method_to_typeid is a mapping used to specify the return type of those exposed methods which should return a proxy. It maps method names to *typeid* strings. (If *method_to_typeid* is `None` then `proxytype._method_to_typeid_` is used instead if it exists.) If a method’s name is not a key of this mapping or if the mapping is `None` then the object returned by the method will be copied by value.

create_method determines whether a method should be created with name *typeid* which can be used to tell the server process to create a new shared object and return a proxy for it. By default it is `True`.

`BaseManager` instances also have one read-only property:

address

管理器所用的地址。

在 3.3 版更改: Manager objects support the context management protocol – see [上下文管理器类型](#). `__enter__()` starts the server process (if it has not already started) and then returns the manager object. `__exit__()` calls `shutdown()`.

In previous versions `__enter__()` did not start the manager’s server process if it was not already started.

class `multiprocessing.managers.SyncManager`

A subclass of `BaseManager` which can be used for the synchronization of processes. Objects of this type are returned by `multiprocessing.Manager()`.

Its methods create and return 代理对象 for a number of commonly used data types to be synchronized across processes. This notably includes shared lists and dictionaries.

Barrier (*parties*[, *action*[, *timeout*]])

Create a shared `threading.Barrier` object and return a proxy for it.

3.3 新版功能.

BoundedSemaphore ([*value*])

Create a shared `threading.BoundedSemaphore` object and return a proxy for it.

Condition ([*lock*])

Create a shared `threading.Condition` object and return a proxy for it.

If *lock* is supplied then it should be a proxy for a `threading.Lock` or `threading.RLock` object.

在 3.3 版更改: 新增了 `wait_for()` 方法。

Event ()

Create a shared `threading.Event` object and return a proxy for it.

Lock()Create a shared `threading.Lock` object and return a proxy for it.**Namespace()**Create a shared `Namespace` object and return a proxy for it.**Queue** ([*maxsize*])Create a shared `queue.Queue` object and return a proxy for it.**RLock()**Create a shared `threading.RLock` object and return a proxy for it.**Semaphore** ([*value*])Create a shared `threading.Semaphore` object and return a proxy for it.**Array** (*typecode*, *sequence*)

Create an array and return a proxy for it.

Value (*typecode*, *value*)Create an object with a writable `value` attribute and return a proxy for it.**dict** ()**dict** (*mapping*)**dict** (*sequence*)Create a shared `dict` object and return a proxy for it.**list** ()**list** (*sequence*)Create a shared `list` object and return a proxy for it.

在 3.6 版更改: Shared objects are capable of being nested. For example, a shared container object such as a shared list can contain other shared objects which will all be managed and synchronized by the `SyncManager`.

class multiprocessing.managers.NamespaceA type that can register with `SyncManager`.

A namespace object has no public methods, but does have writable attributes. Its representation shows the values of its attributes.

However, when using a proxy for a namespace object, an attribute beginning with `'_'` will be an attribute of the proxy and not an attribute of the referent:

```
>>> manager = multiprocessing.Manager()
>>> Global = manager.Namespace()
>>> Global.x = 10
>>> Global.y = 'hello'
>>> Global._z = 12.3      # this is an attribute of the proxy
>>> print(Global)
Namespace(x=10, y='hello')
```

Customized managers

To create one's own manager, one creates a subclass of `BaseManager` and uses the `register()` classmethod to register new types or callables with the manager class. For example:

```
from multiprocessing.managers import BaseManager

class MathsClass:
    def add(self, x, y):
```

(下页继续)

(续上页)

```

        return x + y
    def mul(self, x, y):
        return x * y

class MyManager(BaseManager):
    pass

MyManager.register('Maths', MathsClass)

if __name__ == '__main__':
    with MyManager() as manager:
        maths = manager.Maths()
        print(maths.add(4, 3))           # prints 7
        print(maths.mul(7, 8))         # prints 56

```

Using a remote manager

It is possible to run a manager server on one machine and have clients use it from other machines (assuming that the firewalls involved allow it).

Running the following commands creates a server for a single shared queue which remote clients can access:

```

>>> from multiprocessing.managers import BaseManager
>>> from queue import Queue
>>> queue = Queue()
>>> class QueueManager(BaseManager): pass
>>> QueueManager.register('get_queue', callable=lambda: queue)
>>> m = QueueManager(address=('', 50000), authkey=b'abracadabra')
>>> s = m.get_server()
>>> s.serve_forever()

```

One client can access the server as follows:

```

>>> from multiprocessing.managers import BaseManager
>>> class QueueManager(BaseManager): pass
>>> QueueManager.register('get_queue')
>>> m = QueueManager(address=('foo.bar.org', 50000), authkey=b'abracadabra')
>>> m.connect()
>>> queue = m.get_queue()
>>> queue.put('hello')

```

Another client can also use it:

```

>>> from multiprocessing.managers import BaseManager
>>> class QueueManager(BaseManager): pass
>>> QueueManager.register('get_queue')
>>> m = QueueManager(address=('foo.bar.org', 50000), authkey=b'abracadabra')
>>> m.connect()
>>> queue = m.get_queue()
>>> queue.get()
'hello'

```

Local processes can also access that queue, using the code from above on the client to access it remotely:

```

>>> from multiprocessing import Process, Queue
>>> from multiprocessing.managers import BaseManager
>>> class Worker(Process):
...     def __init__(self, q):
...         self.q = q
...         super(Worker, self).__init__()
...     def run(self):
...         self.q.put('local hello')
...
>>> queue = Queue()
>>> w = Worker(queue)
>>> w.start()
>>> class QueueManager(BaseManager): pass
...
>>> QueueManager.register('get_queue', callable=lambda: queue)
>>> m = QueueManager(address=('', 50000), authkey=b'abracadabra')
>>> s = m.get_server()
>>> s.serve_forever()

```

代理对象

A proxy is an object which *refers* to a shared object which lives (presumably) in a different process. The shared object is said to be the *referent* of the proxy. Multiple proxy objects may have the same referent.

A proxy object has methods which invoke corresponding methods of its referent (although not every method of the referent will necessarily be available through the proxy). In this way, a proxy can be used just like its referent can:

```

>>> from multiprocessing import Manager
>>> manager = Manager()
>>> l = manager.list([i*i for i in range(10)])
>>> print(l)
[0, 1, 4, 9, 16, 25, 36, 49, 64, 81]
>>> print(repr(l))
<ListProxy object, typeid 'list' at 0x...>
>>> l[4]
16
>>> l[2:5]
[4, 9, 16]

```

Notice that applying `str()` to a proxy will return the representation of the referent, whereas applying `repr()` will return the representation of the proxy.

An important feature of proxy objects is that they are picklable so they can be passed between processes. As such, a referent can contain 代理对象. This permits nesting of these managed lists, dicts, and other 代理对象:

```

>>> a = manager.list()
>>> b = manager.list()
>>> a.append(b)           # referent of a now contains referent of b
>>> print(a, b)
[<ListProxy object, typeid 'list' at ...>] []
>>> b.append('hello')
>>> print(a[0], b)
['hello'] ['hello']

```

Similarly, dict and list proxies may be nested inside one another:

```
>>> l_outer = manager.list([ manager.dict() for i in range(2) ])
>>> d_first_inner = l_outer[0]
>>> d_first_inner['a'] = 1
>>> d_first_inner['b'] = 2
>>> l_outer[1]['c'] = 3
>>> l_outer[1]['z'] = 26
>>> print(l_outer[0])
{'a': 1, 'b': 2}
>>> print(l_outer[1])
{'c': 3, 'z': 26}
```

If standard (non-proxy) `list` or `dict` objects are contained in a referent, modifications to those mutable values will not be propagated through the manager because the proxy has no way of knowing when the values contained within are modified. However, storing a value in a container proxy (which triggers a `__setitem__` on the proxy object) does propagate through the manager and so to effectively modify such an item, one could re-assign the modified value to the container proxy:

```
# create a list proxy and append a mutable object (a dictionary)
lproxy = manager.list()
lproxy.append({})
# now mutate the dictionary
d = lproxy[0]
d['a'] = 1
d['b'] = 2
# at this point, the changes to d are not yet synced, but by
# updating the dictionary, the proxy is notified of the change
lproxy[0] = d
```

This approach is perhaps less convenient than employing nested 代理对象 for most use cases but also demonstrates a level of control over the synchronization.

注解: The proxy types in `multiprocessing` do nothing to support comparisons by value. So, for instance, we have:

```
>>> manager.list([1,2,3]) == [1,2,3]
False
```

One should just use a copy of the referent instead when making comparisons.

class `multiprocessing.managers.BaseProxy`

Proxy objects are instances of subclasses of `BaseProxy`.

`__callmethod` (`methodname` [, `args` [, `kwds`]])

Call and return the result of a method of the proxy's referent.

If proxy is a proxy whose referent is `obj` then the expression

```
proxy.__callmethod(methodname, args, kwds)
```

will evaluate the expression

```
getattr(obj, methodname)(*args, **kwds)
```

in the manager's process.

The returned value will be a copy of the result of the call or a proxy to a new shared object – see documentation for the `method_to_typeid` argument of `BaseManager.register()`.

If an exception is raised by the call, then is re-raised by `_callmethod()`. If some other exception is raised in the manager's process then this is converted into a `RemoteError` exception and is raised by `_callmethod()`.

Note in particular that an exception will be raised if *methodname* has not been *exposed*.

An example of the usage of `_callmethod()`:

```
>>> l = manager.list(range(10))
>>> l._callmethod('__len__')
10
>>> l._callmethod('__getitem__', (slice(2, 7),)) # equivalent to l[2:7]
[2, 3, 4, 5, 6]
>>> l._callmethod('__getitem__', (20,))          # equivalent to l[20]
Traceback (most recent call last):
...
IndexError: list index out of range
```

`_getvalue()`

Return a copy of the referent.

If the referent is unpicklable then this will raise an exception.

`__repr__()`

Return a representation of the proxy object.

`__str__()`

Return the representation of the referent.

Cleanup

A proxy object uses a weakref callback so that when it gets garbage collected it deregisters itself from the manager which owns its referent.

A shared object gets deleted from the manager process when there are no longer any proxies referring to it.

进程池

One can create a pool of processes which will carry out tasks submitted to it with the `Pool` class.

```
class multiprocessing.pool.Pool([processes[, initializer[, initargs[, maxtasksperchild[, context]]]])
```

A process pool object which controls a pool of worker processes to which jobs can be submitted. It supports asynchronous results with timeouts and callbacks and has a parallel map implementation.

processes is the number of worker processes to use. If *processes* is `None` then the number returned by `os.cpu_count()` is used.

If *initializer* is not `None` then each worker process will call `initializer(*initargs)` when it starts.

maxtasksperchild is the number of tasks a worker process can complete before it will exit and be replaced with a fresh worker process, to enable unused resources to be freed. The default *maxtasksperchild* is `None`, which means worker processes will live as long as the pool.

context can be used to specify the context used for starting the worker processes. Usually a pool is created using the function `multiprocessing.Pool()` or the `Pool()` method of a context object. In both cases *context* is set appropriately.

Note that the methods of the pool object should only be called by the process which created the pool.

3.2 新版功能: *maxtasksperchild*

3.4 新版功能: *context*

注解: Worker processes within a *Pool* typically live for the complete duration of the Pool's work queue. A frequent pattern found in other systems (such as Apache, mod_wsgi, etc) to free resources held by workers is to allow a worker within a pool to complete only a set amount of work before being exiting, being cleaned up and a new process spawned to replace the old one. The *maxtasksperchild* argument to the *Pool* exposes this ability to the end user.

apply (*func*[, *args*[, *kws*]])

Call *func* with arguments *args* and keyword arguments *kws*. It blocks until the result is ready. Given this blocks, *apply_async()* is better suited for performing work in parallel. Additionally, *func* is only executed in one of the workers of the pool.

apply_async (*func*[, *args*[, *kws*[, *callback*[, *error_callback*]]]])

A variant of the *apply()* method which returns a result object.

If *callback* is specified then it should be a callable which accepts a single argument. When the result becomes ready *callback* is applied to it, that is unless the call failed, in which case the *error_callback* is applied instead.

If *error_callback* is specified then it should be a callable which accepts a single argument. If the target function fails, then the *error_callback* is called with the exception instance.

Callbacks should complete immediately since otherwise the thread which handles the results will get blocked.

map (*func*, *iterable*[, *chunksize*])

A parallel equivalent of the *map()* built-in function (it supports only one *iterable* argument though). It blocks until the result is ready.

This method chops the iterable into a number of chunks which it submits to the process pool as separate tasks. The (approximate) size of these chunks can be specified by setting *chunksize* to a positive integer.

map_async (*func*, *iterable*[, *chunksize*[, *callback*[, *error_callback*]]])

A variant of the *map()* method which returns a result object.

If *callback* is specified then it should be a callable which accepts a single argument. When the result becomes ready *callback* is applied to it, that is unless the call failed, in which case the *error_callback* is applied instead.

If *error_callback* is specified then it should be a callable which accepts a single argument. If the target function fails, then the *error_callback* is called with the exception instance.

Callbacks should complete immediately since otherwise the thread which handles the results will get blocked.

imap (*func*, *iterable*[, *chunksize*])

A lazier version of *map()*.

The *chunksize* argument is the same as the one used by the *map()* method. For very long iterables using a large value for *chunksize* can make the job complete **much** faster than using the default value of 1.

Also if *chunksize* is 1 then the *next()* method of the iterator returned by the *imap()* method has an optional *timeout* parameter: *next(timeout)* will raise *multiprocessing.TimeoutError* if the result cannot be returned within *timeout* seconds.

imap_unordered (*func*, *iterable*[, *chunksize*])

The same as *imap()* except that the ordering of the results from the returned iterator should be considered arbitrary. (Only when there is only one worker process is the order guaranteed to be “correct”.)

starmap (*func*, *iterable*[, *chunksize*])

Like *map()* except that the elements of the *iterable* are expected to be iterables that are unpacked as arguments.

Hence an *iterable* of `[(1, 2), (3, 4)]` results in `[func(1, 2), func(3, 4)]`.

3.3 新版功能.

starmap_async(*func*, *iterable*`[, chunksize``[, callback``[, error_callback]``]])`

A combination of `starmap()` and `map_async()` that iterates over *iterable* of iterables and calls *func* with the iterables unpacked. Returns a result object.

3.3 新版功能.

close()

Prevents any more tasks from being submitted to the pool. Once all the tasks have been completed the worker processes will exit.

terminate()

Stops the worker processes immediately without completing outstanding work. When the pool object is garbage collected `terminate()` will be called immediately.

join()

Wait for the worker processes to exit. One must call `close()` or `terminate()` before using `join()`.

3.3 新版功能: Pool objects now support the context management protocol – see 上下文管理器类型. `__enter__()` returns the pool object, and `__exit__()` calls `terminate()`.

class multiprocessing.pool.**AsyncResult**

The class of the result returned by `Pool.apply_async()` and `Pool.map_async()`.

get(`[timeout]`)

Return the result when it arrives. If *timeout* is not None and the result does not arrive within *timeout* seconds then `multiprocessing.TimeoutError` is raised. If the remote call raised an exception then that exception will be reraised by `get()`.

wait(`[timeout]`)

Wait until the result is available or until *timeout* seconds pass.

ready()

Return whether the call has completed.

successful()

Return whether the call completed without raising an exception. Will raise `AssertionError` if the result is not ready.

The following example demonstrates the use of a pool:

```
from multiprocessing import Pool
import time

def f(x):
    return x*x

if __name__ == '__main__':
    with Pool(processes=4) as pool:          # start 4 worker processes
        result = pool.apply_async(f, (10,)) # evaluate "f(10)" asynchronously in a
        ↪ single process
        print(result.get(timeout=1))        # prints "100" unless your computer is
        ↪ *very* slow

        print(pool.map(f, range(10)))      # prints "[0, 1, 4, ..., 81]"

        it = pool.imap(f, range(10))
        print(next(it))                    # prints "0"
```

(下页继续)

(续上页)

```

print(next(it))                # prints "1"
print(it.next(timeout=1))      # prints "4" unless your computer is
↪ *very* slow

result = pool.apply_async(time.sleep, (10,))
print(result.get(timeout=1))    # raises multiprocessing.TimeoutError

```

Listeners and Clients

Usually message passing between processes is done using queues or by using *Connection* objects returned by *Pipe()*.

However, the *multiprocessing.connection* module allows some extra flexibility. It basically gives a high level message oriented API for dealing with sockets or Windows named pipes. It also has support for *digest authentication* using the *hmac* module, and for polling multiple connections at the same time.

multiprocessing.connection.deliver_challenge (*connection*, *authkey*)

Send a randomly generated message to the other end of the connection and wait for a reply.

If the reply matches the digest of the message using *authkey* as the key then a welcome message is sent to the other end of the connection. Otherwise *AuthenticationError* is raised.

multiprocessing.connection.answer_challenge (*connection*, *authkey*)

Receive a message, calculate the digest of the message using *authkey* as the key, and then send the digest back.

If a welcome message is not received, then *AuthenticationError* is raised.

multiprocessing.connection.Client (*address*[, *family*[, *authkey*]])

Attempt to set up a connection to the listener which is using address *address*, returning a *Connection*.

The type of the connection is determined by *family* argument, but this can generally be omitted since it can usually be inferred from the format of *address*. (See *Address Formats*)

If *authkey* is given and not None, it should be a byte string and will be used as the secret key for an HMAC-based authentication challenge. No authentication is done if *authkey* is None. *AuthenticationError* is raised if authentication fails. See *Authentication keys*.

class *multiprocessing.connection.Listener* ([*address*[, *family*[, *backlog*[, *authkey*]]]])

A wrapper for a bound socket or Windows named pipe which is 'listening' for connections.

address is the address to be used by the bound socket or named pipe of the listener object.

注解: If an address of '0.0.0.0' is used, the address will not be a connectable end point on Windows. If you require a connectable end-point, you should use '127.0.0.1'.

family is the type of socket (or named pipe) to use. This can be one of the strings 'AF_INET' (for a TCP socket), 'AF_UNIX' (for a Unix domain socket) or 'AF_PIPE' (for a Windows named pipe). Of these only the first is guaranteed to be available. If *family* is None then the family is inferred from the format of *address*. If *address* is also None then a default is chosen. This default is the family which is assumed to be the fastest available. See *Address Formats*. Note that if *family* is 'AF_UNIX' and *address* is None then the socket will be created in a private temporary directory created using *tempfile.mkstemp()*.

If the listener object uses a socket then *backlog* (1 by default) is passed to the *listen()* method of the socket once it has been bound.

If *authkey* is given and not None, it should be a byte string and will be used as the secret key for an HMAC-based authentication challenge. No authentication is done if *authkey* is None. *AuthenticationError* is raised if authentication fails. See *Authentication keys*.

accept()

Accept a connection on the bound socket or named pipe of the listener object and return a *Connection* object. If authentication is attempted and fails, then *AuthenticationError* is raised.

close()

Close the bound socket or named pipe of the listener object. This is called automatically when the listener is garbage collected. However it is advisable to call it explicitly.

Listener objects have the following read-only properties:

address

The address which is being used by the Listener object.

last_accepted

The address from which the last accepted connection came. If this is unavailable then it is *None*.

3.3 新版功能: Listener objects now support the context management protocol – see 上下文管理器类型. `__enter__()` returns the listener object, and `__exit__()` calls `close()`.

`multiprocessing.connection.wait(object_list, timeout=None)`

Wait till an object in *object_list* is ready. Returns the list of those objects in *object_list* which are ready. If *timeout* is a float then the call blocks for at most that many seconds. If *timeout* is *None* then it will block for an unlimited period. A negative timeout is equivalent to a zero timeout.

For both Unix and Windows, an object can appear in *object_list* if it is

- a readable *Connection* object;
- a connected and readable *socket.socket* object; or
- the *sentinel* attribute of a *Process* object.

A connection or socket object is ready when there is data available to be read from it, or the other end has been closed.

Unix: `wait(object_list, timeout)` almost equivalent `select.select(object_list, [], [], timeout)`. The difference is that, if `select.select()` is interrupted by a signal, it can raise *OSError* with an error number of *EINTR*, whereas `wait()` will not.

Windows: An item in *object_list* must either be an integer handle which is waitable (according to the definition used by the documentation of the Win32 function `WaitForMultipleObjects()`) or it can be an object with a `fileno()` method which returns a socket handle or pipe handle. (Note that pipe handles and socket handles are **not** waitable handles.)

3.3 新版功能.

Examples

The following server code creates a listener which uses 'secret password' as an authentication key. It then waits for a connection and sends some data to the client:

```
from multiprocessing.connection import Listener
from array import array

address = ('localhost', 6000)      # family is deduced to be 'AF_INET'

with Listener(address, authkey=b'secret password') as listener:
    with listener.accept() as conn:
        print('connection accepted from', listener.last_accepted)

        conn.send([2.25, None, 'junk', float])
```

(下页继续)

(续上页)

```
conn.send_bytes(b'hello')

conn.send_bytes(array('i', [42, 1729]))
```

The following code connects to the server and receives some data from the server:

```
from multiprocessing.connection import Client
from array import array

address = ('localhost', 6000)

with Client(address, authkey=b'secret password') as conn:
    print(conn.recv())           # => [2.25, None, 'junk', float]

    print(conn.recv_bytes())     # => 'hello'

    arr = array('i', [0, 0, 0, 0, 0])
    print(conn.recv_bytes_into(arr)) # => 8
    print(arr)                   # => array('i', [42, 1729, 0, 0, 0])
```

The following code uses `wait()` to wait for messages from multiple processes at once:

```
import time, random
from multiprocessing import Process, Pipe, current_process
from multiprocessing.connection import wait

def foo(w):
    for i in range(10):
        w.send((i, current_process().name))
    w.close()

if __name__ == '__main__':
    readers = []

    for i in range(4):
        r, w = Pipe(duplex=False)
        readers.append(r)
        p = Process(target=foo, args=(w,))
        p.start()
        # We close the writable end of the pipe now to be sure that
        # p is the only process which owns a handle for it. This
        # ensures that when p closes its handle for the writable end,
        # wait() will promptly report the readable end as being ready.
        w.close()

    while readers:
        for r in wait(readers):
            try:
                msg = r.recv()
            except EOFError:
                readers.remove(r)
            else:
                print(msg)
```

Address Formats

- An 'AF_INET' address is a tuple of the form (hostname, port) where *hostname* is a string and *port* is an integer.
- An 'AF_UNIX' address is a string representing a filename on the filesystem.
- An 'AF_PIPE' address is a string of the form `r'\.\pipe{PipeName}'`. To use `Client()` to connect to a named pipe on a remote computer called *ServerName* one should use an address of the form `r'\ServerName\pipe{PipeName}'` instead.

Note that any string beginning with two backslashes is assumed by default to be an 'AF_PIPE' address rather than an 'AF_UNIX' address.

Authentication keys

When one uses `Connection.recv`, the data received is automatically unpickled. Unfortunately unpickling data from an untrusted source is a security risk. Therefore `Listener` and `Client()` use the `hmac` module to provide digest authentication.

An authentication key is a byte string which can be thought of as a password: once a connection is established both ends will demand proof that the other knows the authentication key. (Demonstrating that both ends are using the same key does **not** involve sending the key over the connection.)

If authentication is requested but no authentication key is specified then the return value of `current_process().authkey` is used (see `Process`). This value will be automatically inherited by any `Process` object that the current process creates. This means that (by default) all processes of a multi-process program will share a single authentication key which can be used when setting up connections between themselves.

Suitable authentication keys can also be generated by using `os.urandom()`.

日志

Some support for logging is available. Note, however, that the `logging` package does not use process shared locks so it is possible (depending on the handler type) for messages from different processes to get mixed up.

`multiprocessing.get_logger()`

Returns the logger used by `multiprocessing`. If necessary, a new one will be created.

When first created the logger has level `logging.NOTSET` and no default handler. Messages sent to this logger will not by default propagate to the root logger.

Note that on Windows child processes will only inherit the level of the parent process' s logger – any other customization of the logger will not be inherited.

`multiprocessing.log_to_stderr()`

This function performs a call to `get_logger()` but in addition to returning the logger created by `get_logger`, it adds a handler which sends output to `sys.stderr` using format `'[(levelname)s/%(processName)s] %(message)s'`.

Below is an example session with logging turned on:

```
>>> import multiprocessing, logging
>>> logger = multiprocessing.log_to_stderr()
>>> logger.setLevel(logging.INFO)
>>> logger.warning('doomed')
[WARNING/MainProcess] doomed
```

(下页继续)

(续上页)

```
>>> m = multiprocessing.Manager()
[INFO/SyncManager-...] child process calling self.run()
[INFO/SyncManager-...] created temp directory /.../pymp-...
[INFO/SyncManager-...] manager serving at '/.../listener-...'
>>> del m
[INFO/MainProcess] sending shutdown message to manager
[INFO/SyncManager-...] manager exiting with exitcode 0
```

For a full table of logging levels, see the *logging* module.

The `multiprocessing.dummy` module

multiprocessing.dummy replicates the API of *multiprocessing* but is no more than a wrapper around the *threading* module.

17.2.3 Programming guidelines

There are certain guidelines and idioms which should be adhered to when using *multiprocessing*.

All start methods

The following applies to all start methods.

Avoid shared state

As far as possible one should try to avoid shifting large amounts of data between processes.

It is probably best to stick to using queues or pipes for communication between processes rather than using the lower level synchronization primitives.

Picklability

Ensure that the arguments to the methods of proxies are picklable.

Thread safety of proxies

Do not use a proxy object from more than one thread unless you protect it with a lock.

(There is never a problem with different processes using the *same* proxy.)

Joining zombie processes

On Unix when a process finishes but has not been joined it becomes a zombie. There should never be very many because each time a new process starts (or *active_children()* is called) all completed processes which have not yet been joined will be joined. Also calling a finished process' s *Process.is_alive* will join the process. Even so it is probably good practice to explicitly join all the processes that you start.

Better to inherit than pickle/unpickle

When using the *spawn* or *forkserver* start methods many types from *multiprocessing* need to be picklable so that child processes can use them. However, one should generally avoid sending shared objects to other processes using pipes or queues. Instead you should arrange the program so that a process which needs access to a shared resource created elsewhere can inherit it from an ancestor process.

Avoid terminating processes

Using the `Process.terminate` method to stop a process is liable to cause any shared resources (such as locks, semaphores, pipes and queues) currently being used by the process to become broken or unavailable to other processes.

Therefore it is probably best to only consider using `Process.terminate` on processes which never use any shared resources.

Joining processes that use queues

Bear in mind that a process that has put items in a queue will wait before terminating until all the buffered items are fed by the “feeder” thread to the underlying pipe. (The child process can call the `Queue.cancel_join_thread` method of the queue to avoid this behaviour.)

This means that whenever you use a queue you need to make sure that all items which have been put on the queue will eventually be removed before the process is joined. Otherwise you cannot be sure that processes which have put items on the queue will terminate. Remember also that non-daemonic processes will be joined automatically.

An example which will deadlock is the following:

```
from multiprocessing import Process, Queue

def f(q):
    q.put('X' * 1000000)

if __name__ == '__main__':
    queue = Queue()
    p = Process(target=f, args=(queue,))
    p.start()
    p.join()                # this deadlocks
    obj = queue.get()
```

A fix here would be to swap the last two lines (or simply remove the `p.join()` line).

Explicitly pass resources to child processes

On Unix using the *fork* start method, a child process can make use of a shared resource created in a parent process using a global resource. However, it is better to pass the object as an argument to the constructor for the child process.

Apart from making the code (potentially) compatible with Windows and the other start methods this also ensures that as long as the child process is still alive the object will not be garbage collected in the parent process. This might be important if some resource is freed when the object is garbage collected in the parent process.

所以对于实例：

```
from multiprocessing import Process, Lock

def f():
    ... do something using "lock" ...

if __name__ == '__main__':
    lock = Lock()
    for i in range(10):
        Process(target=f).start()
```

应当重写成这样：

```
from multiprocessing import Process, Lock

def f(l):
    ... do something using "l" ...

if __name__ == '__main__':
    lock = Lock()
    for i in range(10):
        Process(target=f, args=(lock,)).start()
```

Beware of replacing `sys.stdin` with a “file like object”

`multiprocessing` originally unconditionally called:

```
os.close(sys.stdin.fileno())
```

in the `multiprocessing.Process._bootstrap()` method —this resulted in issues with processes-in-processes. This has been changed to:

```
sys.stdin.close()
sys.stdin = open(os.open(os.devnull, os.O_RDONLY), closefd=False)
```

Which solves the fundamental issue of processes colliding with each other resulting in a bad file descriptor error, but introduces a potential danger to applications which replace `sys.stdin()` with a “file-like object” with output buffering. This danger is that if multiple processes call `close()` on this file-like object, it could result in the same data being flushed to the object multiple times, resulting in corruption.

If you write a file-like object and implement your own caching, you can make it fork-safe by storing the pid whenever you append to the cache, and discarding the cache when the pid changes. For example:

```
@property
def cache(self):
    pid = os.getpid()
    if pid != self._pid:
        self._pid = pid
        self._cache = []
    return self._cache
```

For more information, see [bpo-5155](#), [bpo-5313](#) and [bpo-5331](#)

The `spawn` and `forkserver` start methods

There are a few extra restriction which don’t apply to the `fork` start method.

More picklability

Ensure that all arguments to `Process.__init__()` are picklable. Also, if you subclass `Process` then make sure that instances will be picklable when the `Process.start` method is called.

Global variables

Bear in mind that if code run in a child process tries to access a global variable, then the value it sees (if any) may not be the same as the value in the parent process at the time that `Process.start` was called.

However, global variables which are just module level constants cause no problems.

Safe importing of main module

Make sure that the main module can be safely imported by a new Python interpreter without causing unintended side effects (such as starting a new process).

For example, using the *spawn* or *forkserver* start method running the following module would fail with a *RuntimeError*:

```
from multiprocessing import Process

def foo():
    print('hello')

p = Process(target=foo)
p.start()
```

Instead one should protect the “entry point” of the program by using `if __name__ == '__main__':` as follows:

```
from multiprocessing import Process, freeze_support, set_start_method

def foo():
    print('hello')

if __name__ == '__main__':
    freeze_support()
    set_start_method('spawn')
    p = Process(target=foo)
    p.start()
```

(The `freeze_support()` line can be omitted if the program will be run normally instead of frozen.)

This allows the newly spawned Python interpreter to safely import the module and then run the module's `foo()` function.

Similar restrictions apply if a pool or manager is created in the main module.

17.2.4 示例

Demonstration of how to create and use customized managers and proxies:

```
from multiprocessing import freeze_support
from multiprocessing.managers import BaseManager, BaseProxy
import operator

##

class Foo:
    def f(self):
        print('you called Foo.f()')
    def g(self):
        print('you called Foo.g()')
    def _h(self):
        print('you called Foo._h()')

# A simple generator function
def baz():
    for i in range(10):
        yield i*i
```

(下页继续)

(续上页)

```

# Proxy type for generator objects
class GeneratorProxy(BaseProxy):
    _exposed_ = ['__next__']
    def __iter__(self):
        return self
    def __next__(self):
        return self._callmethod('__next__')

# Function to return the operator module
def get_operator_module():
    return operator

##

class MyManager(BaseManager):
    pass

# register the Foo class; make `f()` and `g()` accessible via proxy
MyManager.register('Foo1', Foo)

# register the Foo class; make `g()` and `_h()` accessible via proxy
MyManager.register('Foo2', Foo, exposed=('g', '_h'))

# register the generator function baz; use `GeneratorProxy` to make proxies
MyManager.register('baz', baz, proxytype=GeneratorProxy)

# register get_operator_module(); make public functions accessible via proxy
MyManager.register('operator', get_operator_module)

##

def test():
    manager = MyManager()
    manager.start()

    print('-' * 20)

    f1 = manager.Foo1()
    f1.f()
    f1.g()
    assert not hasattr(f1, '_h')
    assert sorted(f1._exposed_) == sorted(['f', 'g'])

    print('-' * 20)

    f2 = manager.Foo2()
    f2.g()
    f2._h()
    assert not hasattr(f2, 'f')
    assert sorted(f2._exposed_) == sorted(['g', '_h'])

    print('-' * 20)

    it = manager.baz()
    for i in it:

```

(下页继续)

(续上页)

```

        print('<%d>' % i, end=' ')
    print()

    print('-' * 20)

    op = manager.operator()
    print('op.add(23, 45) =', op.add(23, 45))
    print('op.pow(2, 94) =', op.pow(2, 94))
    print('op._exposed_ =', op._exposed_)

##

if __name__ == '__main__':
    freeze_support()
    test()

```

Using *Pool*:

```

import multiprocessing
import time
import random
import sys

#
# Functions used by test code
#

def calculate(func, args):
    result = func(*args)
    return '%s says that %s%s = %s' % (
        multiprocessing.current_process().name,
        func.__name__, args, result
    )

def calculatestar(args):
    return calculate(*args)

def mul(a, b):
    time.sleep(0.5 * random.random())
    return a * b

def plus(a, b):
    time.sleep(0.5 * random.random())
    return a + b

def f(x):
    return 1.0 / (x - 5.0)

def pow3(x):
    return x ** 3

def noop(x):
    pass

#
# Test code

```

(下页继续)

(续上页)

```

#
def test():
    PROCESSES = 4
    print('Creating pool with %d processes\n' % PROCESSES)

    with multiprocessing.Pool(PROCESSES) as pool:
        #
        # Tests
        #

        TASKS = [(mul, (i, 7)) for i in range(10)] + \
                [(plus, (i, 8)) for i in range(10)]

        results = [pool.apply_async(calculate, t) for t in TASKS]
        imap_it = pool.imap(calculatestar, TASKS)
        imap_unordered_it = pool.imap_unordered(calculatestar, TASKS)

        print('Ordered results using pool.apply_async():')
        for r in results:
            print('\t', r.get())
        print()

        print('Ordered results using pool.imap():')
        for x in imap_it:
            print('\t', x)
        print()

        print('Unordered results using pool.imap_unordered():')
        for x in imap_unordered_it:
            print('\t', x)
        print()

        print('Ordered results using pool.map() --- will block till complete:')
        for x in pool.map(calculatestar, TASKS):
            print('\t', x)
        print()

        #
        # Test error handling
        #

        print('Testing error handling:')

        try:
            print(pool.apply(f, (5,)))
        except ZeroDivisionError:
            print('\tGot ZeroDivisionError as expected from pool.apply()')
        else:
            raise AssertionError('expected ZeroDivisionError')

        try:
            print(pool.map(f, list(range(10))))
        except ZeroDivisionError:
            print('\tGot ZeroDivisionError as expected from pool.map()')
        else:

```

(下页继续)

(续上页)

```

        raise AssertionError('expected ZeroDivisionError')

    try:
        print(list(pool.imap(f, list(range(10)))))
    except ZeroDivisionError:
        print('\tGot ZeroDivisionError as expected from list(pool.imap())')
    else:
        raise AssertionError('expected ZeroDivisionError')

it = pool.imap(f, list(range(10)))
for i in range(10):
    try:
        x = next(it)
    except ZeroDivisionError:
        if i == 5:
            pass
        except StopIteration:
            break
    else:
        if i == 5:
            raise AssertionError('expected ZeroDivisionError')

assert i == 9
print('\tGot ZeroDivisionError as expected from IMapIterator.next()')
print()

#
# Testing timeouts
#

print('Testing ApplyResult.get() with timeout:', end=' ')
res = pool.apply_async(calculate, TASKS[0])
while 1:
    sys.stdout.flush()
    try:
        sys.stdout.write('\n\t%s' % res.get(0.02))
        break
    except multiprocessing.TimeoutError:
        sys.stdout.write('.')

print()
print()

print('Testing IMapIterator.next() with timeout:', end=' ')
it = pool.imap(calculatestar, TASKS)
while 1:
    sys.stdout.flush()
    try:
        sys.stdout.write('\n\t%s' % it.next(0.02))
    except StopIteration:
        break
    except multiprocessing.TimeoutError:
        sys.stdout.write('.')

print()
print()

```

(下页继续)

(续上页)

```
if __name__ == '__main__':
    multiprocessing.freeze_support()
    test()
```

An example showing how to use queues to feed tasks to a collection of worker processes and collect the results:

```
import time
import random

from multiprocessing import Process, Queue, current_process, freeze_support

#
# Function run by worker processes
#

def worker(input, output):
    for func, args in iter(input.get, 'STOP'):
        result = calculate(func, args)
        output.put(result)

#
# Function used to calculate result
#

def calculate(func, args):
    result = func(*args)
    return '%s says that %s%s = %s' % \
        (current_process().name, func.__name__, args, result)

#
# Functions referenced by tasks
#

def mul(a, b):
    time.sleep(0.5*random.random())
    return a * b

def plus(a, b):
    time.sleep(0.5*random.random())
    return a + b

#
#
#

def test():
    NUMBER_OF_PROCESSES = 4
    TASKS1 = [(mul, (i, 7)) for i in range(20)]
    TASKS2 = [(plus, (i, 8)) for i in range(10)]

    # Create queues
    task_queue = Queue()
    done_queue = Queue()

    # Submit tasks
    for task in TASKS1:
```

(下页继续)

(续上页)

```

        task_queue.put(task)

    # Start worker processes
    for i in range(NUMBER_OF_PROCESSES):
        Process(target=worker, args=(task_queue, done_queue)).start()

    # Get and print results
    print('Unordered results:')
    for i in range(len(TASKS1)):
        print('\t', done_queue.get())

    # Add more tasks using `put()`
    for task in TASKS2:
        task_queue.put(task)

    # Get and print some more results
    for i in range(len(TASKS2)):
        print('\t', done_queue.get())

    # Tell child processes to stop
    for i in range(NUMBER_OF_PROCESSES):
        task_queue.put('STOP')

if __name__ == '__main__':
    freeze_support()
    test()

```

17.3 concurrent 包

目前，此包中只有一个模块：

- `concurrent.futures` —— 启动并行任务

17.4 concurrent.futures — 启动并行任务

3.2 新版功能.

源码: `Lib/concurrent/futures/thread.py` 和 `Lib/concurrent/futures/process.py`

`concurrent.futures` 模块提供异步执行回调高层接口。

异步执行可以由 `ThreadPoolExecutor` 使用线程或由 `ProcessPoolExecutor` 使用单独的进程来实现。两者都是实现抽象类 `Executor` 定义的接口。

17.4.1 执行器对象

class `concurrent.futures.Executor`

抽象类提供异步执行调用方法。要通过它的子类调用，而不是直接调用。

submit (*fn*, **args*, ***kwargs*)

调度可调用对象 *fn*，以 *fn*(**args* ***kwargs*) 方式执行并返回 *Future* 对象代表可调用对象的执行。：

```
with ThreadPoolExecutor(max_workers=1) as executor:
    future = executor.submit(pow, 323, 1235)
    print(future.result())
```

map (*func*, **iterables*, *timeout=None*, *chunksize=1*)

类似于 *map(func, *iterables)* 除去：

- 应立即收集 *iterables* 不要延迟再收集；
- *func* 是异步执行的且对 *func* 的调用可以并发执行。

如果 *__next__()* 已被调用且返回的结果在对 *Executor.map()* 的原始调用经过 *timeout* 秒后还不可用，则已返回的迭代器将引发 *concurrent.futures.TimeoutError*。*timeout* 可以为 *int* 或 *float* 类型。如果 *timeout* 未指定或为 *None*，则不限制等待时间。

如果 *func* 调用引发一个异常，当从迭代器中取回它的值时这个异常将被引发。

使用 *ProcessPoolExecutor* 时，这个方法会将 *iterables* 分割任务块并作为独立的的任务并提交到执行池中。这些块的大概数量可以由 *chunksize* 指定正整数设置。对很长的迭代器来说，使用大的 *chunksize* 值比默认值 1 能显著地提高性能。*chunksize* 对 *ThreadPoolExecutor* 没有效果。

在 3.5 版更改：加入 *chunksize* 参数。

shutdown (*wait=True*)

当待执行的期程完成执行后向执行者发送信号，它就会释放正在使用的任何资源。调用 *Executor.submit()* 和 *Executor.submit()* 会在关闭后触发 *RuntimeError*。

如果 *wait* 为 *True* 则此方法只有在所有待执行的期程完成执行且释放已分配的资源后才会返回。如果 *wait* 为 *False*，方法立即返回，所有待执行的期程完成执行后会释放已分配的资源。不管 *wait* 的值是什么，整个 Python 程序将等到所有待执行的期程完成执行后才退出。

如果使用 *with* 语句，你就可以避免显式调用这个方法，它将会停止 *Executor* (就好像 *Executor.shutdown()* 调用时 *wait* 设为 *True* 一样等待)：

```
import shutil
with ThreadPoolExecutor(max_workers=4) as e:
    e.submit(shutil.copy, 'src1.txt', 'dest1.txt')
    e.submit(shutil.copy, 'src2.txt', 'dest2.txt')
    e.submit(shutil.copy, 'src3.txt', 'dest3.txt')
    e.submit(shutil.copy, 'src4.txt', 'dest4.txt')
```

17.4.2 ThreadPoolExecutor

ThreadPoolExecutor 是 *Executor* 的子类，它使用线程池来异步执行调用。

当回调已关联了一个 *Future* 然后再等待另一个 *Future* 的结果时就会发生死锁情况。例如：

```
import time
def wait_on_b():
    time.sleep(5)
    print(b.result()) # b will never complete because it is waiting on a.
    return 5

def wait_on_a():
```

(下页继续)

(续上页)

```

time.sleep(5)
print(a.result()) # a will never complete because it is waiting on b.
return 6

executor = ThreadPoolExecutor(max_workers=2)
a = executor.submit(wait_on_b)
b = executor.submit(wait_on_a)

```

And:

```

def wait_on_future():
    f = executor.submit(pow, 5, 2)
    # This will never complete because there is only one worker thread and
    # it is executing this function.
    print(f.result())

executor = ThreadPoolExecutor(max_workers=1)
executor.submit(wait_on_future)

```

class `concurrent.futures.ThreadPoolExecutor` (*max_workers=None, thread_name_prefix=""*)
Executor 的一个子类，使用最多 *max_workers* 个线程的线程池来异步执行调用。

在 3.5 版更改：如果 *max_workers* 为 `None` 或没有指定，将默认为机器处理器的个数，假如 *ThreadPoolExecutor* 则重于 I/O 操作而不是 CPU 运算，那么可以乘以 5，同时工作线程的数量可以比 *ProcessPoolExecutor* 的数量高。

3.6 新版功能：添加 *thread_name_prefix* 参数允许用户控制由线程池创建的 *threading.Thread* 工作线程名称以方便调试。

ThreadPoolExecutor 例子

```

import concurrent.futures
import urllib.request

URLS = ['http://www.foxnews.com/',
        'http://www.cnn.com/',
        'http://europe.wsj.com/',
        'http://www.bbc.co.uk/',
        'http://some-made-up-domain.com/']

# Retrieve a single page and report the URL and contents
def load_url(url, timeout):
    with urllib.request.urlopen(url, timeout=timeout) as conn:
        return conn.read()

# We can use a with statement to ensure threads are cleaned up promptly
with concurrent.futures.ThreadPoolExecutor(max_workers=5) as executor:
    # Start the load operations and mark each future with its URL
    future_to_url = {executor.submit(load_url, url, 60): url for url in URLS}
    for future in concurrent.futures.as_completed(future_to_url):
        url = future_to_url[future]
        try:
            data = future.result()
        except Exception as exc:

```

(下页继续)

(续上页)

```

        print('%r generated an exception: %s' % (url, exc))
    else:
        print('%r page is %d bytes' % (url, len(data)))

```

17.4.3 ProcessPoolExecutor

ProcessPoolExecutor 是 *Executor* 的子类，它使用进程池来实现异步执行调用。*ProcessPoolExecutor* 使用 *multiprocessing* 回避 *Global Interpreter Lock* 但也意味着只可以处理和返回可序列化的对象。

`__main__` 模块必须可以被工作者子进程导入。这意味着 *ProcessPoolExecutor* 不可以工作在交互式解释器中。

从提交给 *ProcessPoolExecutor* 的回调中调用 *Executor* 或 *Future* 方法会导致死锁。

class `concurrent.futures.ProcessPoolExecutor` (*max_workers=None*)

An *Executor* subclass that executes calls asynchronously using a pool of at most *max_workers* processes. If *max_workers* is `None` or not given, it will default to the number of processors on the machine. If *max_workers* is lower or equal to 0, then a *ValueError* will be raised.

在 3.3 版更改: 如果其中一个工作进程被突然终止, `BrokenProcessPool` 就会马上触发。可预计的行为没有定义, 但执行器上的操作或它的期程会被冻结或死锁。

ProcessPoolExecutor 例子

```

import concurrent.futures
import math

PRIMES = [
    112272535095293,
    112582705942171,
    112272535095293,
    115280095190773,
    115797848077099,
    1099726899285419]

def is_prime(n):
    if n % 2 == 0:
        return False

    sqrt_n = int(math.floor(math.sqrt(n)))
    for i in range(3, sqrt_n + 1, 2):
        if n % i == 0:
            return False
    return True

def main():
    with concurrent.futures.ProcessPoolExecutor() as executor:
        for number, prime in zip(PRIMES, executor.map(is_prime, PRIMES)):
            print('%d is prime: %s' % (number, prime))

if __name__ == '__main__':
    main()

```

17.4.4 期程对象

`Future` 类将可调对象封装为异步执行。`Future` 实例由 `Executor.submit()` 创建。

class `concurrent.futures.Future`

将可调对象封装为异步执行。`Future` 实例由 `Executor.submit()` 创建，除非测试，不应直接创建。

cancel()

Attempt to cancel the call. If the call is currently being executed and cannot be cancelled then the method will return `False`, otherwise the call will be cancelled and the method will return `True`.

cancelled()

如果调用成功取消返回 `True`。

running()

如果调用正在执行而且不能被取消那么返回 “`True`”。

done()

如果调用已被取消或正常结束那么返回 `True`。

result(timeout=None)

返回调用返回的值。如果调用还没完成那么这个方法将等待 `timeout` 秒。如果在 `timeout` 秒内没有执行完成，`concurrent.futures.TimeoutError` 将会被触发。`timeout` 可以是整数或浮点数。如果 `timeout` 没有指定或为 `None`，那么等待时间就没有限制。

如果 `future` 在完成前被取消则 `CancelledError` 将被触发。

如果调用引发了一个异常，这个方法也会引发同样的异常。

exception(timeout=None)

返回由调用引发的异常。如果调用还没完成那么这个方法将等待 `timeout` 秒。如果在 `timeout` 秒内没有执行完成，`concurrent.futures.TimeoutError` 将会被触发。`timeout` 可以是整数或浮点数。如果 `timeout` 没有指定或为 `None`，那么等待时间就没有限制。

如果 `future` 在完成前被取消则 `CancelledError` 将被触发。

如果调用正常完成那么返回 `None`。

add_done_callback(fn)

附加可调 `fn` 到期程。当期程被取消或完成运行时，将会调用 `fn`，而这个期程将作为它唯一的参数。

加入的可调对象总被属于添加它们的进程中的线程按加入的顺序调用。如果可调对象引发一个 `Exception` 子类，它会被记录下来并被忽略掉。如果可调对象引发一个 `BaseException` 子类，这个行为没有定义。

如果期程已经完成或已取消，`fn` 会被立即调用。

下面这些 `Future` 方法用于单元测试和 `Executor` 实现。

set_running_or_notify_cancel()

这个方法只可以在执行关联 `Future` 工作之前由 `Executor` 实现调用或由单测试调用。

如果这个方法返回 `False` 那么 `Future` 已被取消，即 `Future.cancel()` 已被调用并返回 `True`。等待 `Future` 完成(即通过 `as_completed()` 或 `wait()`) 的线程将被唤醒。

如果这个方法返回 `True` 那么 `Future` 不会被取消并已将它变为正在运行状态，也就是说调用 `Future.running()` 时将返回 `True`。

这个方法只可以被调用一次并且不能在调用 `Future.set_result()` 或 `Future.set_exception()` 之后再调用。

set_result (*result*)

设置将`Future` 关联工作的结果给 *result* 。

这个方法只可以由`Executor` 实现和单元测试使用。

set_exception (*exception*)

设置`Future` 关联工作的结果给`Exception exception` 。

这个方法只可以由`Executor` 实现和单元测试使用。

17.4.5 模块函数

`concurrent.futures.wait` (*fs*, *timeout=None*, *return_when=ALL_COMPLETED*)

Wait for the `Future` instances (possibly created by different `Executor` instances) given by *fs* to complete. Returns a named 2-tuple of sets. The first set, named `done`, contains the futures that completed (finished or were cancelled) before the wait completed. The second set, named `not_done`, contains uncompleted futures.

timeout 可以用来控制返回前最大的等待秒数。*timeout* 可以为 `int` 或 `float` 类型。如果 *timeout* 未指定或为 `None` , 则不限制等待时间。

return_when 指定此函数应在何时返回。它必须为以下常数之一:

常数	描述
<code>FIRST_COMPLETED</code>	函数将在任意可等待对象结束或取消时返回。
<code>FIRST_EXCEPTION</code>	函数将在任意可等待对象因引发异常而结束时返回。当没有引发任何异常时它就相当于 <code>ALL_COMPLETED</code> 。
<code>ALL_COMPLETED</code>	函数将在所有可等待对象结束或取消时返回。

`concurrent.futures.as_completed` (*fs*, *timeout=None*)

Returns an iterator over the `Future` instances (possibly created by different `Executor` instances) given by *fs* that yields futures as they complete (finished or were cancelled). Any futures given by *fs* that are duplicated will be returned once. Any futures that completed before `as_completed()` is called will be yielded first. The returned iterator raises a `concurrent.futures.TimeoutError` if `__next__()` is called and the result isn't available after *timeout* seconds from the original call to `as_completed()`. *timeout* can be an `int` or `float`. If *timeout* is not specified or `None`, there is no limit to the wait time.

参见:

PEP 3148 – futures - 异步执行指令。 该提案描述了 Python 标准库中包含的这个特性。

17.4.6 Exception 类

exception `concurrent.futures.CancelledError`

future 被取消时会触发。

exception `concurrent.futures.TimeoutError`

future 运算超出给定的超时数值时触发。

exception `concurrent.futures.process.BrokenProcessPool`

Derived from `RuntimeError`, this exception class is raised when one of the workers of a `ProcessPoolExecutor` has terminated in a non-clean fashion (for example, if it was killed from the outside).

3.3 新版功能.

17.5 subprocess — 子进程管理

源代码: [Lib/subprocess.py](#)

`subprocess` 模块允许你生成新的进程，连接它们的输入、输出、错误管道，并且获取它们的返回码。此模块打算代替一些老旧的模块与功能：

```
os.system
os.spawn*
```

在下面的段落中，你可以找到关于 `subprocess` 模块如何代替这些模块和功能的相关信息。

参见：

PEP 324 – 提出 `subprocess` 模块的 PEP

17.5.1 使用 `subprocess` 模块

推荐的调用子进程的方式是在任何它支持的用例中使用 `run()` 函数。对于更进阶的用例，也可以使用底层的 `Popen` 接口。

`run()` 函数是在 Python 3.5 被添加的；如果你需要与旧版本保持兼容，查看 [Older high-level API](#) 段落。

`subprocess.run(args, *, stdin=None, input=None, stdout=None, stderr=None, shell=False, cwd=None, timeout=None, check=False, encoding=None, errors=None, env=None)`
运行被 `arg` 描述的指令。等待指令完成，然后返回一个 `CompletedProcess` 示例。

The arguments shown above are merely the most common ones, described below in [常用参数](#) (hence the use of keyword-only notation in the abbreviated signature). The full function signature is largely the same as that of the `Popen` constructor - apart from `timeout`, `input` and `check`, all the arguments to this function are passed through to that interface.

This does not capture stdout or stderr by default. To do so, pass `PIPE` for the `stdout` and/or `stderr` arguments.

`timeout` 参数将被传递给 `Popen.communicate()`。如果发生超时，子进程将被杀死并等待。`TimeoutExpired` 异常将在子进程中断后被抛出。

The `input` argument is passed to `Popen.communicate()` and thus to the subprocess's stdin. If used it must be a byte sequence, or a string if `encoding` or `errors` is specified or `universal_newlines` is true. When used, the internal `Popen` object is automatically created with `stdin=PIPE`, and the `stdin` argument may not be used as well.

如果 `check` 设为 `True`，并且进程以非零状态码退出，一个 `CalledProcessError` 异常将被抛出。这个异常的属性将设置为参数，退出码，以及标准输出和标准错误，如果被捕获到。

If `encoding` or `errors` are specified, or `universal_newlines` is true, file objects for stdin, stdout and stderr are opened in text mode using the specified `encoding` and `errors` or the `io.TextIOWrapper` default. Otherwise, file objects are opened in binary mode.

如果 `env` 不是 `None`，它必须是一个字典，为新的进程设置环境变量；它用于替换继承的当前进程的环境的默认行为。它将直接被传递给 `Popen`。

例如：

```
>>> subprocess.run(["ls", "-l"]) # doesn't capture output
CompletedProcess(args=['ls', '-l'], returncode=0)

>>> subprocess.run("exit 1", shell=True, check=True)
```

(下页继续)

(续上页)

```
Traceback (most recent call last):
...
subprocess.CalledProcessError: Command 'exit 1' returned non-zero exit status 1

>>> subprocess.run(["ls", "-l", "/dev/null"], stdout=subprocess.PIPE)
CompletedProcess(args=['ls', '-l', '/dev/null'], returncode=0,
stdout=b'crw-rw-rw- 1 root root 1, 3 Jan 23 16:23 /dev/null\n')
```

3.5 新版功能.

在 3.6 版更改: 添加了 *encoding* 和 *errors* 形参.

class `subprocess.CompletedProcess`
`run()` 的返回值, 代表一个进程已经结束.

args
 被用作启动进程的参数. 可能是一个列表或字符串.

returncode
 子进程的退出状态码. 通常来说, 一个为 0 的退出码表示进程运行正常.
 一个负值 $-N$ 表示子进程被信号 N 中断 (仅 POSIX).

stdout
 Captured stdout from the child process. A bytes sequence, or a string if `run()` was called with an encoding or errors. None if stdout was not captured.

如果你通过 `stderr=subprocess.STDOUT` 运行, 标准输入和标准错误将被组合在一起, 并且 `stderr`` 将为 None.

stderr
 Captured stderr from the child process. A bytes sequence, or a string if `run()` was called with an encoding or errors. None if stderr was not captured.

check_returncode()
 如果 `returncode` 非零, 抛出 `CalledProcessError`.

3.5 新版功能.

`subprocess.DEVNULL`
 可被 `Popen` 的 `stdin`, `stdout` 或者 `stderr` 参数使用的特殊值, 表示使用特殊文件 `os.devnull`.

3.3 新版功能.

`subprocess.PIPE`
 可被 `Popen` 的 `stdin`, `stdout` 或者 `stderr` 参数使用的特殊值, 表示打开标准流的管道. 常用于 `Popen.communicate()`.

`subprocess.STDOUT`
 可被 `Popen` 的 `stdin`, `stdout` 或者 `stderr` 参数使用的特殊值, 表示标准错误与标准输出使用同一句柄。

exception `subprocess.SubprocessError`
 此模块的其他异常的基类。

3.3 新版功能.

exception `subprocess.TimeoutExpired`
`SubprocessError` 的子类, 等待子进程的过程中发生超时时被抛出。

cmd
 用于创建子进程的指令。

timeout

超时秒数。

output

子进程的输出，如果被`run()`或`check_output()`捕获。否则为`None`。

stdout

对`output`的别名，对应的有`stderr`。

stderr

子进程的标准错误输出，如果被`run()`捕获。否则为`None`。

3.3 新版功能.

在 3.5 版更改: 添加了`stdout`和`stderr`属性。

exception subprocess.CalledProcessError

`SubprocessError` 的子类，当一个被`check_call()`或`check_output()`函数运行的子进程返回了非零退出码时被抛出。

returncode

子进程的退出状态。如果程序由一个信号终止，这将会被设为一个负的信号码。

cmd

用于创建子进程的指令。

output

子进程的输出，如果被`run()`或`check_output()`捕获。否则为`None`。

stdout

对`output`的别名，对应的有`stderr`。

stderr

子进程的标准错误输出，如果被`run()`捕获。否则为`None`。

在 3.5 版更改: 添加了`stdout`和`stderr`属性。

常用参数

为了支持丰富的使用案例，`Popen`的构造函数（以及方便的函数）接受大量可选的参数。对于大多数典型的应用例，许多参数可以被安全地留以它们的默认值。通常需要的参数有：

`args` 被所有调用需要，应当为一个字符串，或者一个程序参数序列。提供一个参数序列通常更好，它可以更小心地使用参数中的转义字符以及引用（例如允许文件名中的空格）。如果传递一个简单的字符串，则`shell`参数必须为`True`（见下文）或者该字符串中将被运行的程序名必须用简单的命名而不指定任何参数。

`stdin`、`stdout`和`stderr`分别指定了执行的程序的标准输入、输出和标准错误文件句柄。合法的值有`PIPE`、`DEVNULL`、一个现存的文件描述符（一个正整数）、一个现存的文件对象以及`None`。`PIPE`表示应该新建一个对子进程的管道。`DEVNULL`表示使用特殊的文件`os.devnull`。当使用默认设置`None`时，将不会进行重定向，子进程的文件流将继承自父进程。另外，`stderr`可能为`STDOUT`，表示来自于子进程的标准错误数据应该被`stdout`相同的句柄捕获。

If `encoding` or `errors` are specified, or `universal_newlines` is true, the file objects `stdin`, `stdout` and `stderr` will be opened in text mode using the `encoding` and `errors` specified in the call or the defaults for `io.TextIOWrapper`.

当构造函数的`newline`参数为`None`时。对于`stdin`，输入的换行符`'\n'`将被转换为默认的换行符`os.linesep`。对于`stdout`和`stderr`，所有输出的换行符都被转换为`'\n'`。更多信息，查看`io.TextIOWrapper`类的文档。

如果文本模式未被使用, `stdin`, `stdout` 和 `stderr` 将会以二进制流模式打开。没有编码与换行符转换发生。

3.6 新版功能: 添加了 `encoding` 和 `errors` 形参。

注解: 文件对象 `Popen.stdin`、`Popen.stdout` 和 `Popen.stderr` 的换行符属性不会被 `Popen.communicate()` 方法更新。

如果 `shell` 设为 `True`, 则使用 `shell` 执行指定的指令。如果您主要使用 Python 增强的控制流 (它比大多数系统 `shell` 提供的强大), 并且仍然希望方便地使用其他 `shell` 功能, 如 `shell` 管道、文件通配符、环境变量展开以及 `~` 展开到用户家目录, 这将非常有用。但是, 注意 Python 自己也实现了许多类似 `shell` 的特性 (例如 `glob`, `fnmatch`, `os.walk()`, `os.path.expandvars()`, `os.path.expanduser()` 和 `shutil`)。

在 3.3 版更改: 当 `universal_newline` 被设为 `True`, 则类使用 `locale.getpreferredencoding(False)` 编码来代替 `locale.getpreferredencoding()`。关于它们的区别的更多信息, 见 `io.TextIOWrapper`。

注解: 在使用 `shell=True` 之前, 请阅读 [Security Considerations](#) 段落。

这些选项以及所有其他选项在 `Popen` 构造函数文档中有更详细的描述。

Popen 构造函数

此模块的底层的进程创建与管理由 `Popen` 类处理。它提供了很大的灵活性, 因此开发者能够处理未被便利函数覆盖的不常见用例。

```
class subprocess.Popen(args, bufsize=-1, executable=None, stdin=None, stdout=None, stderr=None,
                        preexec_fn=None, close_fds=True, shell=False, cwd=None, env=None, univer-
                        sal_newlines=False, startupinfo=None, creationflags=0, restore_signals=True,
                        start_new_session=False, pass_fds=(), *, encoding=None, errors=None)
```

在一个新的进程中执行子程序。在 POSIX, 此类使用类似于 `os.execvp()` 的行为来执行子程序。在 Windows, 此类使用了 `Windows CreateProcess()` 函数。 `Popen` 的参数如下:

`args` 应当是一个程序的参数列表或者一个简单的字符串。默认情况下, 如果 `args` 是一个序列, 将运行的程序是此序列的第一项。如果 `args` 是一个字符串, 解释是平台相关的, 如下所述。有关默认行为的其他差异, 见 `shell` 和 `executable` 参数。除非另有说明, 推荐将 `args` 作为序列传递。

在 POSIX, 如果 `args` 是一个字符串, 此字符串被作为将被执行的程序的命名或路径解释。但是, 只有在不传递任何参数给程序的情况下才能这么做。

注解: `shlex.split()` 在确定正确 `args` 的正确标记化时非常有用, 尤其是在复杂情况下:

```
>>> import shlex, subprocess
>>> command_line = input()
/bin/vikings -input eggs.txt -output "spam spam.txt" -cmd "echo '$MONEY'"
>>> args = shlex.split(command_line)
>>> print(args)
['/bin/vikings', '-input', 'eggs.txt', '-output', 'spam spam.txt', '-cmd', "echo '
↪$MONEY'"]
>>> p = subprocess.Popen(args) # Success!
```

特别注意，由 `shell` 中的空格分隔的选项（例如 `-input`）和参数（例如 `eggs.txt`）位于分开的列表元素中，而在需要时使用引号或反斜杠转义的参数在 `shell`（例如包含空格的文件名或上面显示的 `echo` 命令）是单独的列表元素。

在 Windows，如果 `args` 是一个序列，他将通过一个在 *Converting an argument sequence to a string on Windows* 描述的方式被转换为一个字符串。这是因为底层的 `CreateProcess()` 只处理字符串。

参数 `shell`（默认为 `False`）指定是否使用 `shell` 执行程序。如果 `shell` 为 `True`，更推荐将 `args` 作为字符串传递而非序列。

在 POSIX，当 `shell=True`，`shell` 默认为 `/bin/sh`。如果 `args` 是一个字符串，此字符串指定将通过 `shell` 执行的命令。这意味着字符串的格式必须和在命令提示符中所输入的完全相同。这包括，例如，引号和反斜杠转义包含空格的文件名。如果 `args` 是一个序列，第一项指定了命令，另外的项目将作为传递给 `shell`（而非命令）的参数对待。也就是说，`Popen` 等同于：

```
Popen(['/bin/sh', '-c', args[0], args[1], ...])
```

在 Windows，使用 `shell=True`，环境变量 `COMSPEC` 指定了默认 `shell`。在 Windows 你唯一需要指定 `shell=True` 的情况是你想要执行内置在 `shell` 中的命令（例如 `dir` 或者 `copy`）。在运行一个批处理文件或者基于控制台的可执行文件时，不需要 `shell=True`。

注解： 在使用 `shell=True` 之前，请阅读 *Security Considerations* 段落。

`bufsize` 将在 `open()` 函数创建了 `stdin/stdout/stderr` 管道文件对象时作为对应的参数供应：

- 0 表示不使用缓冲区（读取与写入是一个系统调用并且可以返回短内容）
- 1 表示行缓冲（只有 `universal_newlines=True` 时才有用，例如，在文本模式中）
- 任何其他正值表示使用一个约为对应大小的缓冲区
- 负的 `bufsize`（默认）表示使用系统默认的 `io.DEFAULT_BUFFER_SIZE`。

在 3.3.1 版更改：`bufsize` 现在默认为 -1 来启用缓冲，以符合大多数代码所期望的行为。在 Python 3.2.4 和 3.3.1 之前的版本中，它错误地将默认值设为了 0，这是无缓冲的并且允许短读取。这是无意的，并且与大多数代码所期望的 Python 2 的行为不一致。

`executable` 参数指定一个要执行的替换程序。这很少需要。当 `shell=True`，`executable` 替换 `args` 指定运行的程序。但是，原始的 `args` 仍然被传递给程序。大多数程序将被 `args` 指定的程序作为命令名对待，这可以与实际运行的程序不同。在 POSIX，`args` 名作为实际调用程序中可执行文件的显示名称，例如 `ps`。如果 `shell=True`，在 POSIX，`executable` 参数指定用于替换默认 `shell /bin/sh` 的 `shell`。

`stdin`，`stdout` 和 `stderr` 分别指定被运行的程序的标准输入、输出和标准错误的文件句柄。合法的值有 `PIPE`，`DEVNULL`，一个存在的文件描述符（一个正整数），一个存在的文件对象以及 `None`。`PIPE` 表示应创建一个新的对子进程的管道。`DEVNULL` 表示使用特殊的 `os.devnull` 文件。使用默认的 `None`，则不进行成定向；子进程的文件流将继承自父进程。另外，`stderr` 可设为 `STDOUT`，表示应用程序的标准错误数据应和标准输出一同捕获。

如果 `preexec_fn` 被设为一个可调对象，此对象将在子进程刚创建时被调用。（仅 POSIX）

警告： `preexec_fn` 形参在应用程序中存在多线程时是不安全的。子进程在调用前可能死锁。如果你必须使用它，保持警惕！最小化你调用的库的数量。

注解： 如果你需要修改子进程环境，使用 `env` 形参而非在 `preexec_fn` 中进行。`start_new_session` 形参可以

代替之前常用的 `preexec_fn` 来在子进程中调用 `os.setsid()`。

If `close_fds` is true, all file descriptors except 0, 1 and 2 will be closed before the child process is executed. (POSIX only). The default varies by platform: Always true on POSIX. On Windows it is true when `stdin/stdout/stderr` are `None`, false otherwise. On Windows, if `close_fds` is true then no handles will be inherited by the child process. Note that on Windows, you cannot set `close_fds` to true and also redirect the standard handles by setting `stdin`, `stdout` or `stderr`.

在 3.2 版更改: `close_fds` 的默认值已经从 `False` 修改为上述值。

`pass_fds` 是一个可选的在父子进程间保持打开的文件描述符序列。提供任何 `pass_fds` 将强制 `close_fds` 为 `True`。(仅 POSIX)

3.2 新版功能: 加入了 `pass_fds` 形参。

如果 `cwd` 不为 `None`, 此函数在执行子进程前改变当前工作目录为 `cwd`。`cwd` 可以为一个 `str` 和 `path-like` 对象。特别要注意, 当可执行文件的路径为相对路径时, 此函数按相对于 `cwd` 的路径来寻找 `executable` (或者 `args` 的第一项)。

在 3.6 版更改: `cwd` 形参接受一个 `path-like object`。

如果 `restore_signals` 为 true (默认值), 则 Python 设置为 `SIG_IGN` 的所有信号将在 `exec` 之前的子进程中恢复为 `SIG_DFL`。目前, 这包括 `SIGPIPE`, `SIGXFZ` 和 `SIGXFSZ` 信号。(仅 POSIX)

在 3.2 版更改: `restore_signals` 被加入。

如果 `start_new_session` 为 true, 则 `setsid()` 系统调用将在子进程执行之前被执行。(仅 POSIX)

在 3.2 版更改: `start_new_session` 被添加。

如果 `env` 不为 `None`, 则必须为一个为新进程定义了环境变量的字典; 这些用于替换继承的当前进程环境的默认行为。

注解: 如果指定, `env` 必须提供所有被子进程需求的变量。在 Windows, 为了运行一个 `side-by-side assembly`, 指定的 `env` 必须包含一个有效的 `SystemRoot`。

If `encoding` or `errors` are specified, the file objects `stdin`, `stdout` and `stderr` are opened in text mode with the specified encoding and `errors`, as described above in 常用参数. If `universal_newlines` is True, they are opened in text mode with default encoding. Otherwise, they are opened as binary streams.

3.6 新版功能: `encoding` 和 `errors` 被添加。

If given, `startupinfo` will be a `STARTUPINFO` object, which is passed to the underlying `CreateProcess` function. `creationflags`, if given, can be `CREATE_NEW_CONSOLE` or `CREATE_NEW_PROCESS_GROUP`. (Windows only)

`Popen` 对象支持通过 `with` 语句作为上下文管理器, 在退出时关闭文件描述符并等待进程:

```
with Popen(["ifconfig"], stdout=PIPE) as proc:
    log.write(proc.stdout.read())
```

在 3.2 版更改: 添加了上下文管理器支持。

在 3.6 版更改: 现在, 如果 `Popen` 析构时子进程仍然在运行, 则析构器会发送一个 `ResourceWarning` 警告。

异常

Exceptions raised in the child process, before the new program has started to execute, will be re-raised in the parent. Additionally, the exception object will have one extra attribute called `child_traceback`, which is a string containing traceback information from the child's point of view.

最常见的被抛出异常是 `OSError`。例如, 当尝试执行一个不存在的文件时就会发生。应用程序需要为 `OSError` 异常做好准备。

如果 `Popen` 调用时有无效的参数, 则一个 `ValueError` 将被抛出。

`check_all()` 与 `check_output()` 在调用的进程返回非零退出码时将抛出 `CalledProcessError`。

所有接受 `timeout` 形参的函数与方法, 例如 `call()` 和 `Popen.communicate()` 将会在进程退出前超时到期时抛出 `TimeoutExpired`。

此模块中定义的异常都继承自 `SubprocessError`。

3.3 新版功能: 基类 `SubprocessError` 被添加。

17.5.2 安全考量

不像一些其他的 `popen` 功能, 此实现绝不会隐式调用一个系统 `shell`。这意味着任何字符, 包括 `shell` 元字符, 可以安全地被传递给子进程。如果 `shell` 被明确地调用, 通过 `shell=True` 设置, 则确保所有空白字符和元字符被恰当地包裹在引号内以避免 `shell` 注入 漏洞就由应用程序负责了。

当使用 `shell=True`, `shlex.quote()` 函数可以作为在将被用于构造 `shell` 指令的字符串中转义空白字符以及 `shell` 元字符的方案。

17.5.3 Popen 对象

`Popen` 类的实例拥有以下方法:

`Popen.poll()`

检查子进程是否已被终止。设置并返回 `returncode` 属性。否则返回 `None`。

`Popen.wait(timeout=None)`

等待子进程被终止。设置并返回 `returncode` 属性。

如果进程在 `timeout` 秒后未中断, 抛出一个 `TimeoutExpired` 异常, 可以安全地捕获此异常并重新等待。

注解: 当 `stdout=PIPE` 或者 `stderr=PIPE` 并且子进程产生了足以阻塞 OS 管道缓冲区接收更多数据的输出到管道时, 将会发生死锁。当使用管道时用 `Popen.communicate()` 来规避它。

注解: 此函数使用了一个 `busy loop` (非阻塞调用以及短睡眠) 实现。使用 `asyncio` 模块进行异步等待: 参阅 `asyncio.create_subprocess_exec`。

在 3.3 版更改: `timeout` 被添加

3.4 版后已移除: Do not use the `endtime` parameter. It is was unintentionally exposed in 3.3 but was left undocumented as it was intended to be private for internal use. Use `timeout` instead.

`Popen.communicate(input=None, timeout=None)`

与进程交互：向 `stdin` 传输数据。从 `stdout` 和 `stderr` 读取数据，直到文件结束符。等待进程终止。可选的 `input` 参数应当未被传输给子进程的数据，如果没有数据应被传输给子进程则为 `None`。如果流以文本模式打开，`input` 必须为字符串。否则，它必须为字节。

`communicate()` 返回一个 `(stdout_data, stderr_data)` 元组。如果文件以文本模式打开则为字符串；否则字节。

注意如果你想要向进程的 `stdin` 传输数据，你需要通过 `stdin=PIPE` 创建此 `Popen` 对象。类似的，要从结果元组获取任何非 `None` 值，你同样需要设置 `stdout=PIPE` 或者 `stderr=PIPE`。

如果进程在 `timeout` 秒后未终止，一个 `TimeoutExpired` 异常将被抛出。捕获此异常并重新等待将不会丢失任何输出。

如果超时到期，子进程不会被杀死，所以为了正确清理一个行为良好的应用程序应该杀死子进程并完成通讯。

```
proc = subprocess.Popen(...)
try:
    outs, errs = proc.communicate(timeout=15)
except TimeoutExpired:
    proc.kill()
    outs, errs = proc.communicate()
```

注解： 内存里数据读取是缓冲的，所以如果数据尺寸过大或无限，不要使用此方法。

在 3.3 版更改： `timeout` 被添加

`Popen.send_signal(signal)`

将信号 `signal` 发送给子进程。

注解： 在 Windows，`SIGTERM` 是一个 `terminate()` 的别名。`CTRL_C_EVENT` 和 `CTRL_BREAK_EVENT` 可以被发送给以包含 `CREATE_NEW_PROCESS` 的 `creationflags` 形参启动的进程。

`Popen.terminate()`

停止子进程。在 Posix 操作系统上，此方法发送 `SIGTERM`。在 Windows，调用 Win32 API 函数 `TerminateProcess()` 来停止子进程。

`Popen.kill()`

杀死子进程。在 Posix 操作系统上，此函数给子进程发送 `SIGKILL` 信号。在 Windows 上，`kill()` 是 `terminate()` 的别名。

以下属性也是可用的：

`Popen.args`

`args` 参数传递给 `Popen` — 一个程序参数的序列或者一个简单字符串。

3.3 新版功能。

`Popen.stdin`

如果 `stdin` 参数为 `PIPE`，此属性是一个类似 `open()` 返回的可写的流对象。如果 `encoding` 或 `errors` 参数被指定或者 `universal_newlines` 参数为 `True`，则此流是一个文本流，否则是字节流。如果 `stdin` 参数非 `PIPE`，此属性为 `None`。

`Popen.stdout`

如果 `stdout` 参数是 `PIPE`，此属性是一个类似 `open()` 返回的可读流。从流中读取子进程提供的输出。

如果 *encoding* 或 *errors* 参数被指定或者 *universal_newlines* 参数为 `True`，此流为文本流，否则为字节流。如果 *stdout* 参数非 *PIPE*，此属性为 `None`。

`Popen.stderr`

如果 *stderr* 参数是 *PIPE*，此属性是一个类似 *open()* 返回的可读流。从流中读取子进程提供的输出。如果 *encoding* 或 *errors* 参数被指定或者 *universal_newlines* 参数为 `True`，此流为文本流，否则为字节流。如果 *stderr* 参数非 *PIPE*，此属性为 `None`。

警告： 使用 *communicate()* 而非 *.stdin.write*、*.stdout.read* 或者 *.stderr.read* 来避免由于任意其他 OS 管道缓冲区被子进程填满阻塞而导致的死锁。

`Popen.pid`

子进程的进程号。

注意如果你设置了 *shell* 参数为 `True`，则这是生成的子 *shell* 的进程号。

`Popen.returncode`

此进程的退出码，由 *poll()* 和 *wait()* 设置（以及直接由 *communicate()* 设置）。一个 `None` 值表示此进程仍未结束。

一个负值 `-N` 表示子进程被信号 `N` 中断（仅 POSIX）。

17.5.4 Windows Popen 助手

STARTUPINFO 类和以下常数仅在 Windows 有效。

`class subprocess.STARTUPINFO`

Partial support of the Windows *STARTUPINFO* structure is used for *Popen* creation.

`dwFlags`

A bit field that determines whether certain *STARTUPINFO* attributes are used when the process creates a window.

```
si = subprocess.STARTUPINFO()
si.dwFlags = subprocess.STARTF_USESTDHANDLES | subprocess.STARTF_USESHOWWINDOW
```

`hStdInput`

If *dwFlags* specifies *STARTF_USESTDHANDLES*, this attribute is the standard input handle for the process. If *STARTF_USESTDHANDLES* is not specified, the default for standard input is the keyboard buffer.

`hStdOutput`

If *dwFlags* specifies *STARTF_USESTDHANDLES*, this attribute is the standard output handle for the process. Otherwise, this attribute is ignored and the default for standard output is the console window's buffer.

`hStdError`

If *dwFlags* specifies *STARTF_USESTDHANDLES*, this attribute is the standard error handle for the process. Otherwise, this attribute is ignored and the default for standard error is the console window's buffer.

`wShowWindow`

If *dwFlags* specifies *STARTF_USESHOWWINDOW*, this attribute can be any of the values that can be specified in the *nCmdShow* parameter for the *ShowWindow* function, except for *SW_SHOWDEFAULT*. Otherwise, this attribute is ignored.

SW_HIDE is provided for this attribute. It is used when *Popen* is called with *shell=True*.

Constants

The `subprocess` module exposes the following constants.

`subprocess.STD_INPUT_HANDLE`

The standard input device. Initially, this is the console input buffer, `CONIN$`.

`subprocess.STD_OUTPUT_HANDLE`

The standard output device. Initially, this is the active console screen buffer, `CONOUT$`.

`subprocess.STD_ERROR_HANDLE`

The standard error device. Initially, this is the active console screen buffer, `CONOUT$`.

`subprocess.SW_HIDE`

Hides the window. Another window will be activated.

`subprocess.STARTF_USESTDHANDLES`

Specifies that the `STARTUPINFO.hStdInput`, `STARTUPINFO.hStdOutput`, and `STARTUPINFO.hStdError` attributes contain additional information.

`subprocess.STARTF_USESHOWWINDOW`

Specifies that the `STARTUPINFO.wShowWindow` attribute contains additional information.

`subprocess.CREATE_NEW_CONSOLE`

The new process has a new console, instead of inheriting its parent's console (the default).

`subprocess.CREATE_NEW_PROCESS_GROUP`

A `Popen` `creationflags` parameter to specify that a new process group will be created. This flag is necessary for using `os.kill()` on the subprocess.

This flag is ignored if `CREATE_NEW_CONSOLE` is specified.

17.5.5 Older high-level API

Prior to Python 3.5, these three functions comprised the high level API to `subprocess`. You can now use `run()` in many cases, but lots of existing code calls these functions.

`subprocess.call(args, *, stdin=None, stdout=None, stderr=None, shell=False, cwd=None, timeout=None)`

Run the command described by `args`. Wait for command to complete, then return the `returncode` attribute.

这相当于:

```
run(...).returncode
```

(except that the `input` and `check` parameters are not supported)

The arguments shown above are merely the most common ones. The full function signature is largely the same as that of the `Popen` constructor - this function passes all supplied arguments other than `timeout` directly through to that interface.

注解: Do not use `stdout=PIPE` or `stderr=PIPE` with this function. The child process will block if it generates enough output to a pipe to fill up the OS pipe buffer as the pipes are not being read from.

在 3.3 版更改: `timeout` 被添加

`subprocess.check_call(args, *, stdin=None, stdout=None, stderr=None, shell=False, cwd=None, timeout=None)`

Run command with arguments. Wait for command to complete. If the return code was zero then return, oth-

otherwise raise `CalledProcessError`. The `CalledProcessError` object will have the return code in the `returncode` attribute.

这相当于:

```
run(..., check=True)
```

(except that the `input` parameter is not supported)

The arguments shown above are merely the most common ones. The full function signature is largely the same as that of the `Popen` constructor - this function passes all supplied arguments other than `timeout` directly through to that interface.

注解: Do not use `stdout=PIPE` or `stderr=PIPE` with this function. The child process will block if it generates enough output to a pipe to fill up the OS pipe buffer as the pipes are not being read from.

在 3.3 版更改: `timeout` 被添加

`subprocess.check_output` (*args*, *, *stdin=None*, *stderr=None*, *shell=False*, *cwd=None*, *encoding=None*, *errors=None*, *universal_newlines=False*, *timeout=None*)
Run command with arguments and return its output.

If the return code was non-zero it raises a `CalledProcessError`. The `CalledProcessError` object will have the return code in the `returncode` attribute and any output in the `output` attribute.

这相当于:

```
run(..., check=True, stdout=PIPE).stdout
```

The arguments shown above are merely the most common ones. The full function signature is largely the same as that of `run()` - most arguments are passed directly through to that interface. However, explicitly passing `input=None` to inherit the parent's standard input file handle is not supported.

By default, this function will return the data as encoded bytes. The actual encoding of the output data may depend on the command being invoked, so the decoding to text will often need to be handled at the application level.

This behaviour may be overridden by setting `universal_newlines` to `True` as described above in 常用参数.

To also capture standard error in the result, use `stderr=subprocess.STDOUT`:

```
>>> subprocess.check_output (
...     "ls non_existent_file; exit 0",
...     stderr=subprocess.STDOUT,
...     shell=True)
'ls: non_existent_file: No such file or directory\n'
```

3.1 新版功能.

在 3.3 版更改: `timeout` 被添加

在 3.4 版更改: Support for the `input` keyword argument was added.

在 3.6 版更改: `encoding` and `errors` were added. See `run()` for details.

17.5.6 Replacing Older Functions with the `subprocess` Module

In this section, “a becomes b” means that b can be used as a replacement for a.

注解： All “a” functions in this section fail (more or less) silently if the executed program cannot be found; the “b” replacements raise `OSError` instead.

In addition, the replacements using `check_output()` will fail with a `CalledProcessError` if the requested operation produces a non-zero return code. The output is still available as the `output` attribute of the raised exception.

In the following examples, we assume that the relevant functions have already been imported from the `subprocess` module.

Replacing `/bin/sh` shell backquote

```
output=`mycmd myarg`
```

becomes:

```
output = check_output(["mycmd", "myarg"])
```

Replacing shell pipeline

```
output=`dmesg | grep hda`
```

becomes:

```
p1 = Popen(["dmesg"], stdout=PIPE)
p2 = Popen(["grep", "hda"], stdin=p1.stdout, stdout=PIPE)
p1.stdout.close() # Allow p1 to receive a SIGPIPE if p2 exits.
output = p2.communicate()[0]
```

The `p1.stdout.close()` call after starting the `p2` is important in order for `p1` to receive a `SIGPIPE` if `p2` exits before `p1`.

Alternatively, for trusted input, the shell’s own pipeline support may still be used directly:

```
output=`dmesg | grep hda`
```

becomes:

```
output=check_output("dmesg | grep hda", shell=True)
```

Replacing `os.system()`

```
sts = os.system("mycmd" + " myarg")
# becomes
sts = call("mycmd" + " myarg", shell=True)
```

注释:

- Calling the program through the shell is usually not required.

A more realistic example would look like this:

```

try:
    retcode = call("mycmd" + " myarg", shell=True)
    if retcode < 0:
        print("Child was terminated by signal", -retcode, file=sys.stderr)
    else:
        print("Child returned", retcode, file=sys.stderr)
except OSError as e:
    print("Execution failed:", e, file=sys.stderr)

```

Replacing the `os.spawn` family

`P_NOWAIT` example:

```

pid = os.spawnlp(os.P_NOWAIT, "/bin/mycmd", "mycmd", "myarg")
==>
pid = Popen(["/bin/mycmd", "myarg"]).pid

```

`P_WAIT` example:

```

retcode = os.spawnlp(os.P_WAIT, "/bin/mycmd", "mycmd", "myarg")
==>
retcode = call(["/bin/mycmd", "myarg"])

```

Vector example:

```

os.spawnvp(os.P_NOWAIT, path, args)
==>
Popen([path] + args[1:])

```

Environment example:

```

os.spawnlpe(os.P_NOWAIT, "/bin/mycmd", "mycmd", "myarg", env)
==>
Popen(["/bin/mycmd", "myarg"], env={"PATH": "/usr/bin"})

```

Replacing `os.popen()`, `os.popen2()`, `os.popen3()`

```

(child_stdin, child_stdout) = os.popen2(cmd, mode, bufsize)
==>
p = Popen(cmd, shell=True, bufsize=bufsize,
          stdin=PIPE, stdout=PIPE, close_fds=True)
(child_stdin, child_stdout) = (p.stdin, p.stdout)

```

```

(child_stdin,
 child_stdout,
 child_stderr) = os.popen3(cmd, mode, bufsize)
==>
p = Popen(cmd, shell=True, bufsize=bufsize,
          stdin=PIPE, stdout=PIPE, stderr=PIPE, close_fds=True)
(child_stdin,
 child_stdout,
 child_stderr) = (p.stdin, p.stdout, p.stderr)

```

```
(child_stdin, child_stdout_and_stderr) = os.popen4(cmd, mode, bufsize)
==>
p = Popen(cmd, shell=True, bufsize=bufsize,
          stdin=PIPE, stdout=PIPE, stderr=STDOUT, close_fds=True)
(child_stdin, child_stdout_and_stderr) = (p.stdin, p.stdout)
```

Return code handling translates as follows:

```
pipe = os.popen(cmd, 'w')
...
rc = pipe.close()
if rc is not None and rc >> 8:
    print("There were some errors")
==>
process = Popen(cmd, stdin=PIPE)
...
process.stdin.close()
if process.wait() != 0:
    print("There were some errors")
```

Replacing functions from the `popen2` module

注解: If the `cmd` argument to `popen2` functions is a string, the command is executed through `/bin/sh`. If it is a list, the command is directly executed.

```
(child_stdout, child_stdin) = popen2.popen2("somestring", bufsize, mode)
==>
p = Popen("somestring", shell=True, bufsize=bufsize,
          stdin=PIPE, stdout=PIPE, close_fds=True)
(child_stdout, child_stdin) = (p.stdout, p.stdin)
```

```
(child_stdout, child_stdin) = popen2.popen2(["mycmd", "myarg"], bufsize, mode)
==>
p = Popen(["mycmd", "myarg"], bufsize=bufsize,
          stdin=PIPE, stdout=PIPE, close_fds=True)
(child_stdout, child_stdin) = (p.stdout, p.stdin)
```

`popen2.Popen3` and `popen2.Popen4` basically work as `subprocess.Popen`, except that:

- `Popen` raises an exception if the execution fails.
- the `capturestderr` argument is replaced with the `stderr` argument.
- `stdin=PIPE` and `stdout=PIPE` must be specified.
- `popen2` closes all file descriptors by default, but you have to specify `close_fds=True` with `Popen` to guarantee this behavior on all platforms or past Python versions.

17.5.7 Legacy Shell Invocation Functions

This module also provides the following legacy functions from the `2.x` `commands` module. These operations implicitly invoke the system shell and none of the guarantees described above regarding security and exception handling consistency are valid for these functions.

`subprocess.getstatusoutput(cmd)`

Return (exitcode, output) of executing *cmd* in a shell.

Execute the string *cmd* in a shell with `Popen.check_output()` and return a 2-tuple (exitcode, output). The locale encoding is used; see the notes on [常用参数](#) for more details.

A trailing newline is stripped from the output. The exit code for the command can be interpreted as the return code of `subprocess`. Example:

```
>>> subprocess.getstatusoutput('ls /bin/ls')
(0, '/bin/ls')
>>> subprocess.getstatusoutput('cat /bin/junk')
(1, 'cat: /bin/junk: No such file or directory')
>>> subprocess.getstatusoutput('/bin/junk')
(127, 'sh: /bin/junk: not found')
>>> subprocess.getstatusoutput('/bin/kill $$')
(-15, '')
```

Availability: POSIX & Windows

在 3.3.4 版更改: Windows support was added.

The function now returns (exitcode, output) instead of (status, output) as it did in Python 3.3.3 and earlier. exitcode has the same value as [returncode](#).

`subprocess.getoutput(cmd)`

Return output (stdout and stderr) of executing *cmd* in a shell.

Like [getstatusoutput\(\)](#), except the exit code is ignored and the return value is a string containing the command's output. Example:

```
>>> subprocess.getoutput('ls /bin/ls')
'/bin/ls'
```

Availability: POSIX & Windows

在 3.3.4 版更改: Windows support added

17.5.8 注释

Converting an argument sequence to a string on Windows

On Windows, an *args* sequence is converted to a string that can be parsed using the following rules (which correspond to the rules used by the MS C runtime):

1. Arguments are delimited by white space, which is either a space or a tab.
2. A string surrounded by double quotation marks is interpreted as a single argument, regardless of white space contained within. A quoted string can be embedded in an argument.
3. A double quotation mark preceded by a backslash is interpreted as a literal double quotation mark.
4. Backslashes are interpreted literally, unless they immediately precede a double quotation mark.
5. If backslashes immediately precede a double quotation mark, every pair of backslashes is interpreted as a literal backslash. If the number of backslashes is odd, the last backslash escapes the next double quotation mark as described in rule 3.

参见:

[shlex](#) Module which provides function to parse and escape command lines.

17.6 sched — 事件调度器

源码: [Lib/sched.py](#)

`sched` 模块定义了一个实现通用事件调度程序的类:

class `sched.scheduler` (*timefunc=time.monotonic, delayfunc=time.sleep*)

The `scheduler` class defines a generic interface to scheduling events. It needs two functions to actually deal with the “outside world” — *timefunc* should be callable without arguments, and return a number (the “time”, in any units whatsoever). If `time.monotonic` is not available, the *timefunc* default is `time.time` instead. The *delayfunc* function should be callable with one argument, compatible with the output of *timefunc*, and should delay that many time units. *delayfunc* will also be called with the argument 0 after each event is run to allow other threads an opportunity to run in multi-threaded applications.

在 3.3 版更改: *timefunc* 和 *delayfunc* 参数是可选的。

在 3.3 版更改: `scheduler` 类可以安全的在多线程环境中使用。

示例:

```
>>> import sched, time
>>> s = sched.scheduler(time.time, time.sleep)
>>> def print_time(a='default'):
...     print("From print_time", time.time(), a)
...
>>> def print_some_times():
...     print(time.time())
...     s.enter(10, 1, print_time)
...     s.enter(5, 2, print_time, argument=('positional',))
...     s.enter(5, 1, print_time, kwargs={'a': 'keyword'})
...     s.run()
...     print(time.time())
...
>>> print_some_times()
930343690.257
From print_time 930343695.274 positional
From print_time 930343695.275 keyword
From print_time 930343700.273 default
930343700.276
```

17.6.1 调度器对象

`scheduler` 实例拥有以下方法和属性:

`scheduler.enterabs` (*time, priority, action, argument=(), kwargs={}*)

安排一个新事件。*time* 参数应该有一个数字类型兼容的返回值, 与传递给构造函数的 *timefunc* 函数的返回值兼容。计划在相同 *time* 的事件将按其 *priority* 的顺序执行。数字越小表示优先级越高。

执行事件意为执行 `action(*argument, **kwargs)`。*argument* 是包含有 *action* 的位置参数的序列。*kwargs* 是包含 *action* 的关键字参数的字典。

返回值是一个事件, 可用于以后取消事件 (参见 `cancel()`)。

在 3.3 版更改: *argument* 参数是可选的。

3.3 新版功能: 添加了 *kwargs* 形参。

`scheduler.enter(delay, priority, action, argument=(), kwargs={})`

安排延后 *delay* 时间单位的事件。除了相对时间，其他参数、效果和返回值与 `enterabs()` 的相同。

在 3.3 版更改: *argument* 参数是可选的。

3.3 新版功能: 添加了 *kwargs* 形参。

`scheduler.cancel(event)`

从队列中删除事件。如果 *event* 不是当前队列中的事件，则此方法将引发 `ValueError`。

`scheduler.empty()`

如果事件队列为空，则返回真值。

`scheduler.run(blocking=True)`

运行所有预定事件。此方法将等待（使用传递给构造函数的 `delayfunc()` 函数）进行下一个事件，然后执行它，依此类推，直到没有更多的计划事件。

如果 *blocking* 为 `false`，则执行由于最快到期（如果有）的预定事件，然后在调度程序中返回下一个预定调用的截止时间（如果有）。

action 或 *delayfunc* 都可以引发异常。在任何一种情况下，调度程序都将保持一致状态并传播异常。如果 *action* 引发异常，则在将来调用 `run()` 时不会尝试该事件。

如果一系列事件的运行时间比下一个事件之前的可用时间长，那么调度程序将完全落后。不会发生任何事件；调用代码负责取消不再相关的事件。

3.3 新版功能: 添加了 *blocking* 形参。

`scheduler.queue`

只读属性按照将要运行的顺序返回即将发生的事件列表。每个事件都显示为 *named tuple*，包含以下字段: *time*、*priority*、*action*、*argument*、*kwargs*。

17.7 queue — 一个同步的队列类

源代码: `Lib/queue.py`

`queue` 模块实现多生产者，多消费者队列。当信息必须安全的在多线程之间交换时，它在线程编程中是特别有用的。此模块中的 `Queue` 类实现了所有锁定需求的语义。它依赖于 Python 支持的线程可用性；请参阅 `threading` 模块。

模块实现了三种类型的队列，它们的区别仅仅是条目取回的顺序。在 FIFO 队列中，先添加的任务先取回。在 LIFO 队列中，最近被添加的条目先取回（操作类似一个堆栈）。优先级队列中，条目将保持排序（使用 `heapq` 模块）并且最小值的条目第一个返回。

Internally, the module uses locks to temporarily block competing threads; however, it is not designed to handle reentrancy within a thread.

`queue` 模块定义了下列类和异常：

class `queue.Queue(maxsize=0)`

Constructor for a FIFO queue. *maxsize* is an integer that sets the upperbound limit on the number of items that can be placed in the queue. Insertion will block once this size has been reached, until queue items are consumed. If *maxsize* is less than or equal to zero, the queue size is infinite.

class `queue.LifoQueue(maxsize=0)`

LIFO 队列构造函数。*maxsize* 是个整数，用于设置可以放入队列中的项目数的上限。当达到这个大小的时候，插入操作将阻塞至队列中的项目被消费掉。如果 *maxsize* 小于等于零，队列尺寸为无限大。

class queue.PriorityQueue (maxsize=0)

优先级队列构造函数。*maxsize* 是个整数，用于设置可以放入队列中的项目数的上限。当达到这个大小的时候，插入操作将阻塞至队列中的项目被消费掉。如果 *maxsize* 小于等于零，队列尺寸为无限大。

最小值先被取出（最小值条目是由 `sorted(list(entries))[0]` 返回的条目）。条目的典型模式是一个以下形式的元组：(priority_number, data)。

exception queue.Empty

对空的 *Queue* 对象，调用非阻塞的 *get()* (or *get_nowait()*) 时，引发的异常。

exception queue.Full

对满的 *Queue* 对象，调用非阻塞的 *put()* (or *put_nowait()*) 时，引发的异常。

17.7.1 Queue 对象

队列对象 (*Queue*, *LifoQueue*, 或者 *PriorityQueue*) 提供下列描述的公共方法。

Queue.qsize()

返回队列的大致大小。注意，*qsize()* > 0 不保证后续的 *get()* 不被阻塞，*qsize()* < *maxsize* 也不保证 *put()* 不被阻塞。

Queue.empty()

如果队列为空，返回 True，否则返回 False。如果 *empty()* 返回 True，不保证后续调用的 *put()* 不被阻塞。类似的，如果 *empty()* 返回 False，也不保证后续调用的 *get()* 不被阻塞。

Queue.full()

如果队列是满的返回 True，否则返回 False。如果 *full()* 返回 True 不保证后续调用的 *get()* 不被阻塞。类似的，如果 *full()* 返回 False 也不保证后续调用的 *put()* 不被阻塞。

Queue.put (item, block=True, timeout=None)

将 *item* 放入队列。如果可选参数 *block* 是 true 并且 *timeout* 是 None (默认)，则在必要时阻塞至有空闲插槽可用。如果 *timeout* 是个正数，将最多阻塞 *timeout* 秒，如果在这段时间没有可用的空闲插槽，将引发 *Full* 异常。反之 (*block* 是 false)，如果空闲插槽立即可用，则把 *item* 放入队列，否则引发 *Full* 异常 (在这种情况下，*timeout* 将被忽略)。

Queue.put_nowait (item)

相当于 *put* (item, False)。

Queue.get (block=True, timeout=None)

从队列中移除并返回一个项目。如果可选参数 *block* 是 true 并且 *timeout* 是 None (默认值)，则在必要时阻塞至项目可得到。如果 *timeout* 是个正数，将最多阻塞 *timeout* 秒，如果在这段时间内项目不能得到，将引发 *Empty* 异常。反之 (*block* 是 false)，如果一个项目立即可得到，则返回一个项目，否则引发 *Empty* 异常 (这种情况下，*timeout* 将被忽略)。

Queue.get_nowait ()

相当于 *get* (False)。

提供了两个方法，用于支持跟踪排队的任务是否被守护的消费者线程完整的处理。

Queue.task_done ()

表示前面排队的任务已经被完成。被队列的消费者线程使用。每个 *get()* 被用于获取一个任务，后续调用 *task_done()* 告诉队列，该任务的处理已经完成。

如果 *join()* 当前正在阻塞，在所有条目都被处理后，将解除阻塞 (意味着每个 *put()* 进队列的条目的 *task_done()* 都被收到)。

如果被调用的次数多于放入队列中的项目数量，将引发 *ValueError* 异常。

Queue.join ()

阻塞至队列中所有的元素都被接收和处理完毕。

当条目添加到队列的时候，未完成任务的计数就会增加。每当消费者线程调用`task_done()`表示这个条目已经被回收，该条目所有工作已经完成，未完成计数就会减少。当未完成计数降到零的时候，`join()`阻塞被解除。

如何等待排队的任务被完成的示例：

```
def worker():
    while True:
        item = q.get()
        if item is None:
            break
        do_work(item)
        q.task_done()

q = queue.Queue()
threads = []
for i in range(num_worker_threads):
    t = threading.Thread(target=worker)
    t.start()
    threads.append(t)

for item in source():
    q.put(item)

# block until all tasks are done
q.join()

# stop workers
for i in range(num_worker_threads):
    q.put(None)
for t in threads:
    t.join()
```

参见：

类`multiprocessing.Queue` 一个用于多进程上下文的队列类（而不是多线程）。

`collections.deque` 是无界队列的替代实现，具有快速原子的`append()`和`popleft()`操作，不需要锁定。

以下是上述某些服务的支持模块：

17.8 dummy_threading —可直接替代 threading 模块。

源代码： [Lib/dummy_threading.py](#)

This module provides a duplicate interface to the `threading` module. It is meant to be imported when the `_thread` module is not provided on a platform.

Suggested usage is:

```
try:
    import threading
except ImportError:
    import dummy_threading as threading
```

如果线程需要阻塞等待另一个线程被创建的话，可能会造成死锁，这通常是由于阻塞 I/O 引起的。这种场景下请不要使用这个模块。

17.9 `_thread` — 底层多线程 API

该模块提供了操作多个线程（也被称为 轻量级进程或 任务）的底层原语——多个控制线程共享全局数据空间。为了处理同步问题，也提供了简单的锁机制（也称为 互斥锁或 二进制信号）。`threading` 模块基于该模块提供了更易用的高级多线程 API。

The module is optional. It is supported on Windows, Linux, SGI IRIX, Solaris 2.x, as well as on systems that have a POSIX thread (a.k.a. “pthread”) implementation. For systems lacking the `_thread` module, the `_dummy_thread` module is available. It duplicates this module’s interface and can be used as a drop-in replacement.

It defines the following constants and functions:

exception `_thread.error`

发生线程相关错误时抛出。

在 3.3 版更改：现在是内建异常 `RuntimeError` 的别名。

`_thread.LockType`

锁对象的类型。

`_thread.start_new_thread(function, args[, kwargs])`

启动一个线程，并返回其标识符。线程会用 `args` 作为参数（必须是元组）执行 `function` 函数。可选的 `kwargs` 参数使用字典来指定有名参数。当函数返回时，线程会静默退出，当函数由于未处理的异常而中止时，会打印一条堆栈追踪信息，然后该线程会退出（但其他线程还是会继续运行）。

`_thread.interrupt_main()`

Raise a `KeyboardInterrupt` exception in the main thread. A subthread can use this function to interrupt the main thread.

`_thread.exit()`

抛出 `SystemExit` 异常。如果没有捕获的话，这个异常会使线程退出。

`_thread.allocate_lock()`

返回一个新的锁对象。锁中的方法在后面描述。初始情况下锁处于解锁状态。

`_thread.get_ident()`

返回当前线程的“线程描述符”。它是一个非零的整型数。它的值没有什么含义，主要是作为 magic cookie 使用，比如作为含有线程相关数据的字典的索引。线程描述符可能会在线程退出，新线程创建时复用。

`_thread.stack_size([size])`

Return the thread stack size used when creating new threads. The optional `size` argument specifies the stack size to be used for subsequently created threads, and must be 0 (use platform or configured default) or a positive integer value of at least 32,768 (32 KiB). If `size` is not specified, 0 is used. If changing the thread stack size is unsupported, a `RuntimeError` is raised. If the specified stack size is invalid, a `ValueError` is raised and the stack size is unmodified. 32 KiB is currently the minimum supported stack size value to guarantee sufficient stack space for the interpreter itself. Note that some platforms may have particular restrictions on values for the stack size, such as requiring a minimum stack size > 32 KiB or requiring allocation in multiples of the system memory page size - platform documentation should be referred to for more information (4 KiB pages are common; using multiples of 4096 for the stack size is the suggested approach in the absence of more specific information). Availability: Windows, systems with POSIX threads.

_thread.TIMEOUT_MAX

Lock.acquire() 方法中 *timeout* 参数允许的最大值。传入超过这个值的 *timeout* 会抛出 *OverflowError* 异常。

3.2 新版功能。

锁对象有以下方法：

lock.acquire(waitflag=1, timeout=-1)

没有任何可选参数时，该方法无条件申请获得锁，有必要的话会等待其他线程释放锁（同时只有一个线程能获得锁——这正是锁存在的原因）。

如果传入了整型参数 *waitflag*，具体的行为取决于传入的值：如果是 0 的话，只会在能够立刻获取到锁时才获取，不会等待，如果是非零的话，会像之前提到的一样，无条件获取锁。

如果传入正浮点数参数 *timeout*，相当于指定了返回之前等待得最大秒数。如果传入负的 *timeout*，相当于无限期等待。如果 *waitflag* 是 0 的话，不能指定 *timeout*。

如果成功获取到锁会返回 True，否则返回 False。

在 3.2 版更改：*timeout* 形参是新增的。

在 3.2 版更改：现在获取锁的操作可以被 POSIX 信号中断。

lock.release()

释放锁。锁必须已经被获取过，但不一定是同一个线程获取的。

lock.locked()

返回锁的状态：如果已被某个线程获取，返回 True，否则返回 False。

除了这些方法之外，锁对象也可以通过 with 语句使用，例如：

```
import _thread

a_lock = _thread.allocate_lock()

with a_lock:
    print("a_lock is locked while this executes")
```

注意事项：

- 线程与中断奇怪地交互：*KeyboardInterrupt* 异常可能会被任意一个线程捕获。（如果 *signal* 模块可用的话，中断总是会进入主线程。）
- 调用 *sys.exit()* 或是抛出 *SystemExit* 异常等效于调用 *_thread.exit()*。
- 不可能中断锁的 *acquire()* 方法——*KeyboardInterrupt* 一场会在锁获取到之后发生。
- 当主线程退出时，由系统决定其他线程是否存活。在大多数系统中，这些线程会直接被杀掉，不会执行 *try...finally* 语句，也不会执行对象析构函数。
- 当主线程退出时，不会进行正常的清理工作（除非使用了 *try...finally* 语句），标准 I/O 文件也不会刷新。

17.10 _dummy_thread — _thread 的替代模块

源代码： [Lib/_dummy_thread.py](#)

This module provides a duplicate interface to the `_thread` module. It is meant to be imported when the `_thread` module is not provided on a platform.

Suggested usage is:

```
try:
    import _thread
except ImportError:
    import _dummy_thread as _thread
```

如果线程需要阻塞等待另一个线程被创建的话，可能会造成死锁，这通常是由于阻塞 I/O 引起的。这种场景下请不要使用这个模块。

Interprocess Communication and Networking

The modules described in this chapter provide mechanisms for different processes to communicate.

某些模块仅适用于同一台机器上的两个进程，例如 *signal* 和 *mmap*。其他模块支持两个或多个进程可用于跨机器通信的网络协议。

本章中描述的模块列表是：

18.1 socket — 底层网络接口

源代码: [Lib/socket.py](#)

这个模块提供了访问 BSD* 套接字 * 的接口。在所有现代 Unix 系统、Windows、macOS 和其他一些平台上可用。

注解： 一些行为可能因平台不同而异，因为调用的是操作系统的套接字 API。

这个 Python 接口是用 Python 的面向对象风格对 Unix 系统调用和套接字库接口的直译：函数 *socket()* 返回一个套接字对象，其方法是对各种套接字系统调用的实现。形参类型一般与 C 接口相比更高级：例如在 Python 文件 *read()* 和 *write()* 操作中，接收操作的缓冲区分配是自动的，发送操作的缓冲区长度是隐式的。

参见：

模块 *socketserver* 用于简化网络服务端编写的类。

模块 *ssl* 套接字对象的 TLS/SSL 封装。

18.1.1 套接字协议族

Depending on the system and the build options, various socket families are supported by this module.

The address format required by a particular socket object is automatically selected based on the address family specified when the socket object was created. Socket addresses are represented as follows:

- The address of an `AF_UNIX` socket bound to a file system node is represented as a string, using the file system encoding and the 'surrogateescape' error handler (see [PEP 383](#)). An address in Linux's abstract namespace is returned as a *bytes-like object* with an initial null byte; note that sockets in this namespace can communicate with normal file system sockets, so programs intended to run on Linux may need to deal with both types of address. A string or bytes-like object can be used for either type of address when passing it as an argument.

在 3.3 版更改: Previously, `AF_UNIX` socket paths were assumed to use UTF-8 encoding.

在 3.5 版更改: Writable *bytes-like object* is now accepted.

- A pair (`host`, `port`) is used for the `AF_INET` address family, where `host` is a string representing either a hostname in Internet domain notation like 'daring.cwi.nl' or an IPv4 address like '100.50.200.5', and `port` is an integer.
 - For IPv4 addresses, two special forms are accepted instead of a host address: '' represents `INADDR_ANY`, which is used to bind to all interfaces, and the string '<broadcast>' represents `INADDR_BROADCAST`. This behavior is not compatible with IPv6, therefore, you may want to avoid these if you intend to support IPv6 with your Python programs.
- For `AF_INET6` address family, a four-tuple (`host`, `port`, `flowinfo`, `scopeid`) is used, where `flowinfo` and `scopeid` represent the `sin6_flowinfo` and `sin6_scope_id` members in `struct sockaddr_in6` in C. For `socket` module methods, `flowinfo` and `scopeid` can be omitted just for backward compatibility. Note, however, omission of `scopeid` can cause problems in manipulating scoped IPv6 addresses.
- `AF_NETLINK` sockets are represented as pairs (`pid`, `groups`).
- Linux-only support for TIPC is available using the `AF_TIPC` address family. TIPC is an open, non-IP based networked protocol designed for use in clustered computer environments. Addresses are represented by a tuple, and the fields depend on the address type. The general tuple form is (`addr_type`, `v1`, `v2`, `v3` [, `scope`]), where:
 - `addr_type` is one of `TIPC_ADDR_NAMESEQ`, `TIPC_ADDR_NAME`, or `TIPC_ADDR_ID`.
 - `scope` is one of `TIPC_ZONE_SCOPE`, `TIPC_CLUSTER_SCOPE`, and `TIPC_NODE_SCOPE`.
 - If `addr_type` is `TIPC_ADDR_NAME`, then `v1` is the server type, `v2` is the port identifier, and `v3` should be 0.
If `addr_type` is `TIPC_ADDR_NAMESEQ`, then `v1` is the server type, `v2` is the lower port number, and `v3` is the upper port number.
If `addr_type` is `TIPC_ADDR_ID`, then `v1` is the node, `v2` is the reference, and `v3` should be set to 0.
- A tuple (`interface`,) is used for the `AF_CAN` address family, where `interface` is a string representing a network interface name like 'can0'. The network interface name '' can be used to receive packets from all network interfaces of this family.
- A string or a tuple (`id`, `unit`) is used for the `SYSPROTO_CONTROL` protocol of the `PF_SYSTEM` family. The string is the name of a kernel control using a dynamically-assigned ID. The tuple can be used if ID and unit number of the kernel control are known or if a registered ID is used.

3.3 新版功能.

- `AF_BLUETOOTH` supports the following protocols and address formats:
 - `BTPROTO_L2CAP` accepts (`bdaddr`, `psm`) where `bdaddr` is the Bluetooth address as a string and `psm` is an integer.

- BTPROTO_RFCOMM accepts (bdaddr, channel) where bdaddr is the Bluetooth address as a string and channel is an integer.
- BTPROTO_HCI accepts (device_id,) where device_id is either an integer or a string with the Bluetooth address of the interface. (This depends on your OS; NetBSD and DragonFlyBSD expect a Bluetooth address while everything else expects an integer.)

在 3.2 版更改: NetBSD and DragonFlyBSD support added.

- BTPROTO_SCO accepts bdaddr where bdaddr is a *bytes* object containing the Bluetooth address in a string format. (ex. b'12:23:34:45:56:67') This protocol is not supported under FreeBSD.
- *AF_ALG* is a Linux-only socket based interface to Kernel cryptography. An algorithm socket is configured with a tuple of two to four elements (type, name [, feat [, mask]]), where:
 - type is the algorithm type as string, e.g. aead, hash, skcipher or rng.
 - name is the algorithm name and operation mode as string, e.g. sha256, hmac (sha256), cbc (aes) or drbg_nopr_ctr_aes256.
 - feat and mask are unsigned 32bit integers.

Availability Linux 2.6.38, some algorithm types require more recent Kernels.

3.6 新版功能.

- *AF_PACKET* is a low-level interface directly to network devices. The packets are represented by the tuple (ifname, proto[, pkttype[, hatype[, addr]]) where:
 - ifname - String specifying the device name.
 - proto - An in network-byte-order integer specifying the Ethernet protocol number.
 - pkttype - Optional integer specifying the packet type:
 - * PACKET_HOST (the default) - Packet addressed to the local host.
 - * PACKET_BROADCAST - Physical-layer broadcast packet.
 - * PACKET_MULTICAST - Packet sent to a physical-layer multicast address.
 - * PACKET_OTHERHOST - Packet to some other host that has been caught by a device driver in promiscuous mode.
 - * PACKET_OUTGOING - Packet originating from the local host that is looped back to a packet socket.
 - hatype - Optional integer specifying the ARP hardware address type.
 - addr - Optional bytes-like object specifying the hardware physical address, whose interpretation depends on the device.

If you use a hostname in the *host* portion of IPv4/v6 socket address, the program may show a nondeterministic behavior, as Python uses the first address returned from the DNS resolution. The socket address will be resolved differently into an actual IPv4/v6 address, depending on the results from DNS resolution and/or the host configuration. For deterministic behavior use a numeric address in *host* portion.

All errors raise exceptions. The normal exceptions for invalid argument types and out-of-memory conditions can be raised; starting from Python 3.3, errors related to socket or address semantics raise *OSError* or one of its subclasses (they used to raise *socket.error*).

Non-blocking mode is supported through *setblocking()*. A generalization of this based on timeouts is supported through *settimeout()*.

18.1.2 Module contents

The module `socket` exports the following elements.

异常

exception `socket.error`

A deprecated alias of `OSError`.

在 3.3 版更改: Following [PEP 3151](#), this class was made an alias of `OSError`.

exception `socket.herror`

A subclass of `OSError`, this exception is raised for address-related errors, i.e. for functions that use `h_errno` in the POSIX C API, including `gethostbyname_ex()` and `gethostbyaddr()`. The accompanying value is a pair (`h_errno`, `string`) representing an error returned by a library call. `h_errno` is a numeric value, while `string` represents the description of `h_errno`, as returned by the `hstrerror()` C function.

在 3.3 版更改: This class was made a subclass of `OSError`.

exception `socket.gaierror`

A subclass of `OSError`, this exception is raised for address-related errors by `getaddrinfo()` and `getnameinfo()`. The accompanying value is a pair (`error`, `string`) representing an error returned by a library call. `string` represents the description of `error`, as returned by the `gai_strerror()` C function. The numeric `error` value will match one of the `EAI_*` constants defined in this module.

在 3.3 版更改: This class was made a subclass of `OSError`.

exception `socket.timeout`

A subclass of `OSError`, this exception is raised when a timeout occurs on a socket which has had timeouts enabled via a prior call to `settimeout()` (or implicitly through `setdefaulttimeout()`). The accompanying value is a string whose value is currently always “timed out”.

在 3.3 版更改: This class was made a subclass of `OSError`.

常数

The `AF_*` and `SOCK_*` constants are now `AddressFamily` and `SocketKind` `IntEnum` collections.

3.4 新版功能.

`socket.AF_UNIX`

`socket.AF_INET`

`socket.AF_INET6`

These constants represent the address (and protocol) families, used for the first argument to `socket()`. If the `AF_UNIX` constant is not defined then this protocol is unsupported. More constants may be available depending on the system.

`socket.SOCK_STREAM`

`socket.SOCK_DGRAM`

`socket.SOCK_RAW`

`socket.SOCK_RDM`

`socket.SOCK_SEQPACKET`

These constants represent the socket types, used for the second argument to `socket()`. More constants may be available depending on the system. (Only `SOCK_STREAM` and `SOCK_DGRAM` appear to be generally useful.)

`socket.SOCK_CLOEXEC`

socket.SOCK_NONBLOCK

These two constants, if defined, can be combined with the socket types and allow you to set some flags atomically (thus avoiding possible race conditions and the need for separate calls).

参见:

[Secure File Descriptor Handling](#) for a more thorough explanation.

Availability: Linux >= 2.6.27.

3.2 新版功能.

SO_***socket.SOMAXCONN****MSG_*****SOL_*****SCM_*****IPPROTO_*****IPPORT_*****INADDR_*****IP_*****IPV6_*****EAI_*****AI_*****NI_*****TCP_***

Many constants of these forms, documented in the Unix documentation on sockets and/or the IP protocol, are also defined in the socket module. They are generally used in arguments to the `setsockopt()` and `getsockopt()` methods of socket objects. In most cases, only those symbols that are defined in the Unix header files are defined; for a few symbols, default values are provided.

在 3.6 版更改: `SO_DOMAIN`, `SO_PROTOCOL`, `SO_PEERSEC`, `SO_PASSSEC`, `TCP_USER_TIMEOUT`, `TCP_CONGESTION` were added.

在 3.6.5 版更改: On Windows, `TCP_FASTOPEN`, `TCP_KEEPCNT` appear if run-time Windows supports.

socket.AF_CAN**socket.PF_CAN****SOL_CAN_*****CAN_***

Many constants of these forms, documented in the Linux documentation, are also defined in the socket module.

Availability: Linux >= 2.6.25.

3.3 新版功能.

socket.CAN_BCM**CAN_BCM_***

`CAN_BCM`, in the CAN protocol family, is the broadcast manager (BCM) protocol. Broadcast manager constants, documented in the Linux documentation, are also defined in the socket module.

Availability: Linux >= 2.6.25.

3.4 新版功能.

socket.CAN_RAW_FD_FRAMES

Enables CAN FD support in a `CAN_RAW` socket. This is disabled by default. This allows your application to send both CAN and CAN FD frames; however, you one must accept both CAN and CAN FD frames when reading from the socket.

This constant is documented in the Linux documentation.

Availability: Linux >= 3.6.

3.5 新版功能.

`socket.AF_PACKET`

`socket.PF_PACKET`

PACKET_*

Many constants of these forms, documented in the Linux documentation, are also defined in the socket module.

Availability: Linux >= 2.2.

`socket.AF_RDS`

`socket.PF_RDS`

`socket.SOL_RDS`

RDS_*

Many constants of these forms, documented in the Linux documentation, are also defined in the socket module.

Availability: Linux >= 2.6.30.

3.3 新版功能.

`socket.SIO_RCVALL`

`socket.SIO_KEEPA_LIVE_VALS`

`socket.SIO_LOOPBACK_FAST_PATH`

RCVALL_*

Constants for Windows' `WSAIoctl()`. The constants are used as arguments to the `ioctl()` method of socket objects.

在 3.6 版更改: `SIO_LOOPBACK_FAST_PATH` was added.

TIPC_*

TIPC related constants, matching the ones exported by the C socket API. See the TIPC documentation for more information.

`socket.AF_ALG`

`socket.SOL_ALG`

ALG_*

Constants for Linux Kernel cryptography.

Availability: Linux >= 2.6.38.

3.6 新版功能.

`socket.AF_LINK`

Availability: BSD, OSX.

3.4 新版功能.

`socket.has_ipv6`

This constant contains a boolean value which indicates if IPv6 is supported on this platform.

`socket.BDADDR_ANY`

`socket.BDADDR_LOCAL`

These are string constants containing Bluetooth addresses with special meanings. For example, `BDADDR_ANY` can be used to indicate any address when specifying the binding socket with `BTPROTO_RFCOMM`.

`socket.HCI_FILTER`

`socket.HCI_TIME_STAMP`

`socket.HCI_DATA_DIR`

For use with `BTPROTO_HCI`. `HCI_FILTER` is not available for NetBSD or DragonFlyBSD. `HCI_TIME_STAMP` and `HCI_DATA_DIR` are not available for FreeBSD, NetBSD, or DragonFlyBSD.

函数

Creating sockets

The following functions all create *socket objects*.

`socket.socket (family=AF_INET, type=SOCK_STREAM, proto=0, fileno=None)`

Create a new socket using the given address family, socket type and protocol number. The address family should be `AF_INET` (the default), `AF_INET6`, `AF_UNIX`, `AF_CAN`, `AF_PACKET`, or `AF_RDS`. The socket type should be `SOCK_STREAM` (the default), `SOCK_DGRAM`, `SOCK_RAW` or perhaps one of the other `SOCK_` constants. The protocol number is usually zero and may be omitted or in the case where the address family is `AF_CAN` the protocol should be one of `CAN_RAW` or `CAN_BCM`. If `fileno` is specified, the other arguments are ignored, causing the socket with the specified file descriptor to return. Unlike `socket.fromfd()`, `fileno` will return the same socket and not a duplicate. This may help close a detached socket using `socket.close()`.

The newly created socket is *non-inheritable*.

在 3.3 版更改: The `AF_CAN` family was added. The `AF_RDS` family was added.

在 3.4 版更改: The `CAN_BCM` protocol was added.

在 3.4 版更改: The returned socket is now non-inheritable.

`socket.socketpair ([family[, type[, proto]]])`

Build a pair of connected socket objects using the given address family, socket type, and protocol number. Address family, socket type, and protocol number are as for the `socket()` function above. The default family is `AF_UNIX` if defined on the platform; otherwise, the default is `AF_INET`.

The newly created sockets are *non-inheritable*.

在 3.2 版更改: The returned socket objects now support the whole socket API, rather than a subset.

在 3.4 版更改: The returned sockets are now non-inheritable.

在 3.5 版更改: Windows support added.

`socket.create_connection (address[, timeout[, source_address]])`

Connect to a TCP service listening on the Internet *address* (a 2-tuple (host, port)), and return the socket object. This is a higher-level function than `socket.connect()`: if *host* is a non-numeric hostname, it will try to resolve it for both `AF_INET` and `AF_INET6`, and then try to connect to all possible addresses in turn until a connection succeeds. This makes it easy to write clients that are compatible to both IPv4 and IPv6.

Passing the optional *timeout* parameter will set the timeout on the socket instance before attempting to connect. If no *timeout* is supplied, the global default timeout setting returned by `getdefaulttimeout()` is used.

If supplied, *source_address* must be a 2-tuple (host, port) for the socket to bind to as its source address before connecting. If host or port are `''` or 0 respectively the OS default behavior will be used.

在 3.2 版更改: 添加了 *source_address*。

`socket.fromfd (fd, family, type, proto=0)`

Duplicate the file descriptor *fd* (an integer as returned by a file object's `fileno()` method) and build a socket object from the result. Address family, socket type and protocol number are as for the `socket()` function above. The file descriptor should refer to a socket, but this is not checked —subsequent operations on the object may fail if the file descriptor is invalid. This function is rarely needed, but can be used to get or set socket options on a socket passed to a program as standard input or output (such as a server started by the Unix `inetd` daemon). The socket is assumed to be in blocking mode.

The newly created socket is *non-inheritable*.

在 3.4 版更改: The returned socket is now non-inheritable.

`socket.fromshare(data)`

Instantiate a socket from data obtained from the `socket.share()` method. The socket is assumed to be in blocking mode.

Availability: Windows.

3.3 新版功能.

`socket.SocketType`

This is a Python type object that represents the socket object type. It is the same as `type(socket(...))`.

其他功能

The `socket` module also offers various network-related services:

`socket.getaddrinfo(host, port, family=0, type=0, proto=0, flags=0)`

Translate the *host/port* argument into a sequence of 5-tuples that contain all the necessary arguments for creating a socket connected to that service. *host* is a domain name, a string representation of an IPv4/v6 address or `None`. *port* is a string service name such as `'http'`, a numeric port number or `None`. By passing `None` as the value of *host* and *port*, you can pass `NULL` to the underlying C API.

The *family*, *type* and *proto* arguments can be optionally specified in order to narrow the list of addresses returned. Passing zero as a value for each of these arguments selects the full range of results. The *flags* argument can be one or several of the `AI_*` constants, and will influence how results are computed and returned. For example, `AI_NUMERICHOST` will disable domain name resolution and will raise an error if *host* is a domain name.

The function returns a list of 5-tuples with the following structure:

```
(family, type, proto, canonname, sockaddr)
```

In these tuples, *family*, *type*, *proto* are all integers and are meant to be passed to the `socket()` function. *canonname* will be a string representing the canonical name of the *host* if `AI_CANONNAME` is part of the *flags* argument; else *canonname* will be empty. *sockaddr* is a tuple describing a socket address, whose format depends on the returned *family* (a (address, port) 2-tuple for `AF_INET`, a (address, port, flow info, scope id) 4-tuple for `AF_INET6`), and is meant to be passed to the `socket.connect()` method.

The following example fetches address information for a hypothetical TCP connection to `example.org` on port 80 (results may differ on your system if IPv6 isn't enabled):

```
>>> socket.getaddrinfo("example.org", 80, proto=socket.IPPROTO_TCP)
[(<AddressFamily.AF_INET6: 10>, <SocketType.SOCK_STREAM: 1>,
  6, '', ('2606:2800:220:1:248:1893:25c8:1946', 80, 0, 0)),
 (<AddressFamily.AF_INET: 2>, <SocketType.SOCK_STREAM: 1>,
  6, '', ('93.184.216.34', 80))]
```

在 3.2 版更改: parameters can now be passed using keyword arguments.

`socket.getfqdn([name])`

Return a fully qualified domain name for *name*. If *name* is omitted or empty, it is interpreted as the local host. To find the fully qualified name, the hostname returned by `gethostbyaddr()` is checked, followed by aliases for the host, if available. The first name which includes a period is selected. In case no fully qualified domain name is available, the hostname as returned by `gethostname()` is returned.

`socket.gethostbyname(hostname)`

Translate a host name to IPv4 address format. The IPv4 address is returned as a string, such as `'100.50.200.5'`. If the host name is an IPv4 address itself it is returned unchanged. See `gethostbyname_ex()` for a more complete interface. `gethostbyname()` does not support IPv6 name resolution, and `getaddrinfo()` should be used instead for IPv4/v6 dual stack support.

`socket.gethostbyname_ex(hostname)`

Translate a host name to IPv4 address format, extended interface. Return a triple (`hostname`, `aliaslist`, `ipaddrlist`) where `hostname` is the primary host name responding to the given `ip_address`, `aliaslist` is a (possibly empty) list of alternative host names for the same address, and `ipaddrlist` is a list of IPv4 addresses for the same interface on the same host (often but not always a single address). `gethostbyname_ex()` does not support IPv6 name resolution, and `getaddrinfo()` should be used instead for IPv4/v6 dual stack support.

`socket.gethostname()`

Return a string containing the hostname of the machine where the Python interpreter is currently executing.

Note: `gethostname()` doesn't always return the fully qualified domain name; use `getfqdn()` for that.

`socket.gethostbyaddr(ip_address)`

Return a triple (`hostname`, `aliaslist`, `ipaddrlist`) where `hostname` is the primary host name responding to the given `ip_address`, `aliaslist` is a (possibly empty) list of alternative host names for the same address, and `ipaddrlist` is a list of IPv4/v6 addresses for the same interface on the same host (most likely containing only a single address). To find the fully qualified domain name, use the function `getfqdn()`. `gethostbyaddr()` supports both IPv4 and IPv6.

`socket.getnameinfo(sockaddr, flags)`

Translate a socket address `sockaddr` into a 2-tuple (`host`, `port`). Depending on the settings of `flags`, the result can contain a fully-qualified domain name or numeric address representation in `host`. Similarly, `port` can contain a string port name or a numeric port number.

`socket.getprotobyname(protocolname)`

Translate an Internet protocol name (for example, `'icmp'`) to a constant suitable for passing as the (optional) third argument to the `socket()` function. This is usually only needed for sockets opened in “raw” mode (`SOCK_RAW`); for the normal socket modes, the correct protocol is chosen automatically if the protocol is omitted or zero.

`socket.getservbyname(servicename[, protocolname])`

Translate an Internet service name and protocol name to a port number for that service. The optional protocol name, if given, should be `'tcp'` or `'udp'`, otherwise any protocol will match.

`socket.getservbyport(port[, protocolname])`

Translate an Internet port number and protocol name to a service name for that service. The optional protocol name, if given, should be `'tcp'` or `'udp'`, otherwise any protocol will match.

`socket.ntohl(x)`

Convert 32-bit positive integers from network to host byte order. On machines where the host byte order is the same as network byte order, this is a no-op; otherwise, it performs a 4-byte swap operation.

`socket.ntohs(x)`

Convert 16-bit positive integers from network to host byte order. On machines where the host byte order is the same as network byte order, this is a no-op; otherwise, it performs a 2-byte swap operation.

`socket.htonl(x)`

Convert 32-bit positive integers from host to network byte order. On machines where the host byte order is the same as network byte order, this is a no-op; otherwise, it performs a 4-byte swap operation.

`socket.htons(x)`

Convert 16-bit positive integers from host to network byte order. On machines where the host byte order is the same as network byte order, this is a no-op; otherwise, it performs a 2-byte swap operation.

`socket.inet_aton(ip_string)`

Convert an IPv4 address from dotted-quad string format (for example, `'123.45.67.89'`) to 32-bit packed binary format, as a bytes object four characters in length. This is useful when conversing with a program that uses the standard C library and needs objects of type `struct in_addr`, which is the C type for the 32-bit packed binary this function returns.

`inet_aton()` also accepts strings with less than three dots; see the Unix manual page `inet(3)` for details.

If the IPv4 address string passed to this function is invalid, `OSError` will be raised. Note that exactly what is valid depends on the underlying C implementation of `inet_aton()`.

`inet_aton()` does not support IPv6, and `inet_pton()` should be used instead for IPv4/v6 dual stack support.

`socket.inet_ntoa(packed_ip)`

Convert a 32-bit packed IPv4 address (a *bytes-like object* four bytes in length) to its standard dotted-quad string representation (for example, `'123.45.67.89'`). This is useful when conversing with a program that uses the standard C library and needs objects of type `struct in_addr`, which is the C type for the 32-bit packed binary data this function takes as an argument.

If the byte sequence passed to this function is not exactly 4 bytes in length, `OSError` will be raised. `inet_ntoa()` does not support IPv6, and `inet_ntop()` should be used instead for IPv4/v6 dual stack support.

在 3.5 版更改: Writable *bytes-like object* is now accepted.

`socket.inet_pton(address_family, ip_string)`

Convert an IP address from its family-specific string format to a packed, binary format. `inet_pton()` is useful when a library or network protocol calls for an object of type `struct in_addr` (similar to `inet_aton()`) or `struct in6_addr`.

Supported values for `address_family` are currently `AF_INET` and `AF_INET6`. If the IP address string `ip_string` is invalid, `OSError` will be raised. Note that exactly what is valid depends on both the value of `address_family` and the underlying implementation of `inet_pton()`.

Availability: Unix (maybe not all platforms), Windows.

在 3.4 版更改: Windows support added

`socket.inet_ntop(address_family, packed_ip)`

Convert a packed IP address (a *bytes-like object* of some number of bytes) to its standard, family-specific string representation (for example, `'7.10.0.5'` or `'5aef:2b::8'`). `inet_ntop()` is useful when a library or network protocol returns an object of type `struct in_addr` (similar to `inet_ntoa()`) or `struct in6_addr`.

Supported values for `address_family` are currently `AF_INET` and `AF_INET6`. If the bytes object `packed_ip` is not the correct length for the specified address family, `ValueError` will be raised. `OSError` is raised for errors from the call to `inet_ntop()`.

Availability: Unix (maybe not all platforms), Windows.

在 3.4 版更改: Windows support added

在 3.5 版更改: Writable *bytes-like object* is now accepted.

`socket.CMSG_LEN(length)`

Return the total length, without trailing padding, of an ancillary data item with associated data of the given `length`. This value can often be used as the buffer size for `recvmsg()` to receive a single item of ancillary data, but **RFC 3542** requires portable applications to use `CMSG_SPACE()` and thus include space for padding, even when the item will be the last in the buffer. Raises `OverflowError` if `length` is outside the permissible range of values.

Availability: most Unix platforms, possibly others.

3.3 新版功能.

`socket.CMSG_SPACE(length)`

Return the buffer size needed for `recvmsg()` to receive an ancillary data item with associated data of the given `length`, along with any trailing padding. The buffer space needed to receive multiple items is the sum of the `CMSG_SPACE()` values for their associated data lengths. Raises `OverflowError` if `length` is outside the permissible range of values.

Note that some systems might support ancillary data without providing this function. Also note that setting the buffer size using the results of this function may not precisely limit the amount of ancillary data that can be received, since additional data may be able to fit into the padding area.

Availability: most Unix platforms, possibly others.

3.3 新版功能.

`socket.getdefaulttimeout()`

Return the default timeout in seconds (float) for new socket objects. A value of `None` indicates that new socket objects have no timeout. When the socket module is first imported, the default is `None`.

`socket.setdefaulttimeout(timeout)`

Set the default timeout in seconds (float) for new socket objects. When the socket module is first imported, the default is `None`. See `settimeout()` for possible values and their respective meanings.

`socket.sethostname(name)`

Set the machine's hostname to *name*. This will raise an `OSError` if you don't have enough rights.

Availability: Unix.

3.3 新版功能.

`socket.if_nameindex()`

Return a list of network interface information (index int, name string) tuples. `OSError` if the system call fails.

Availability: Unix.

3.3 新版功能.

`socket.if_nameindex(if_name)`

Return a network interface index number corresponding to an interface name. `OSError` if no interface with the given name exists.

Availability: Unix.

3.3 新版功能.

`socket.if_indextoname(if_index)`

Return a network interface name corresponding to an interface index number. `OSError` if no interface with the given index exists.

Availability: Unix.

3.3 新版功能.

18.1.3 Socket Objects

Socket objects have the following methods. Except for `makefile()`, these correspond to Unix system calls applicable to sockets.

在 3.2 版更改: Support for the *context manager* protocol was added. Exiting the context manager is equivalent to calling `close()`.

`socket.accept()`

Accept a connection. The socket must be bound to an address and listening for connections. The return value is a pair (*conn*, *address*) where *conn* is a new socket object usable to send and receive data on the connection, and *address* is the address bound to the socket on the other end of the connection.

The newly created socket is *non-inheritable*.

在 3.4 版更改: The socket is now non-inheritable.

在 3.5 版更改: If the system call is interrupted and the signal handler does not raise an exception, the method now retries the system call instead of raising an `InterruptedError` exception (see [PEP 475](#) for the rationale).

`socket.bind(address)`

Bind the socket to *address*. The socket must not already be bound. (The format of *address* depends on the address family—see above.)

`socket.close()`

Mark the socket closed. The underlying system resource (e.g. a file descriptor) is also closed when all file objects from `makefile()` are closed. Once that happens, all future operations on the socket object will fail. The remote end will receive no more data (after queued data is flushed).

Sockets are automatically closed when they are garbage-collected, but it is recommended to `close()` them explicitly, or to use a `with` statement around them.

在 3.6 版更改: `OSError` is now raised if an error occurs when the underlying `close()` call is made.

注解: `close()` releases the resource associated with a connection but does not necessarily close the connection immediately. If you want to close the connection in a timely fashion, call `shutdown()` before `close()`.

`socket.connect(address)`

Connect to a remote socket at *address*. (The format of *address* depends on the address family—see above.)

If the connection is interrupted by a signal, the method waits until the connection completes, or raise a `socket.timeout` on timeout, if the signal handler doesn't raise an exception and the socket is blocking or has a timeout. For non-blocking sockets, the method raises an `InterruptedError` exception if the connection is interrupted by a signal (or the exception raised by the signal handler).

在 3.5 版更改: The method now waits until the connection completes instead of raising an `InterruptedError` exception if the connection is interrupted by a signal, the signal handler doesn't raise an exception and the socket is blocking or has a timeout (see the [PEP 475](#) for the rationale).

`socket.connect_ex(address)`

Like `connect(address)`, but return an error indicator instead of raising an exception for errors returned by the C-level `connect()` call (other problems, such as “host not found,” can still raise exceptions). The error indicator is 0 if the operation succeeded, otherwise the value of the `errno` variable. This is useful to support, for example, asynchronous connects.

`socket.detach()`

Put the socket object into closed state without actually closing the underlying file descriptor. The file descriptor is returned, and can be reused for other purposes.

3.2 新版功能.

`socket.dup()`

Duplicate the socket.

The newly created socket is *non-inheritable*.

在 3.4 版更改: The socket is now non-inheritable.

`socket.fileno()`

Return the socket's file descriptor (a small integer), or -1 on failure. This is useful with `select.select()`.

Under Windows the small integer returned by this method cannot be used where a file descriptor can be used (such as `os.fdopen()`). Unix does not have this limitation.

`socket.get_inheritable()`

Get the *inheritable flag* of the socket's file descriptor or socket's handle: `True` if the socket can be inherited in child processes, `False` if it cannot.

3.4 新版功能.

`socket.getpeername()`

Return the remote address to which the socket is connected. This is useful to find out the port number of a remote IPv4/v6 socket, for instance. (The format of the address returned depends on the address family —see above.) On some systems this function is not supported.

`socket.getsockname()`

Return the socket's own address. This is useful to find out the port number of an IPv4/v6 socket, for instance. (The format of the address returned depends on the address family —see above.)

`socket.getsockopt(level, optname[, buflen])`

Return the value of the given socket option (see the Unix man page *getsockopt(2)*). The needed symbolic constants (`SO_*` etc.) are defined in this module. If *buflen* is absent, an integer option is assumed and its integer value is returned by the function. If *buflen* is present, it specifies the maximum length of the buffer used to receive the option in, and this buffer is returned as a bytes object. It is up to the caller to decode the contents of the buffer (see the optional built-in module *struct* for a way to decode C structures encoded as byte strings).

`socket.gettimeout()`

Return the timeout in seconds (float) associated with socket operations, or `None` if no timeout is set. This reflects the last call to *setblocking()* or *settimeout()*.

`socket.ioctl(control, option)`**Platform Windows**

The *ioctl()* method is a limited interface to the `WSAIoctl` system interface. Please refer to the [Win32 documentation](#) for more information.

On other platforms, the generic *fcntl.fcntl()* and *fcntl.ioctl()* functions may be used; they accept a socket object as their first argument.

Currently only the following control codes are supported: `SIO_RCVALL`, `SIO_KEEPA_LIVE_VALS`, and `SIO_LOOPBACK_FAST_PATH`.

在 3.6 版更改: `SIO_LOOPBACK_FAST_PATH` was added.

`socket.listen([backlog])`

Enable a server to accept connections. If *backlog* is specified, it must be at least 0 (if it is lower, it is set to 0); it specifies the number of unaccepted connections that the system will allow before refusing new connections. If not specified, a default reasonable value is chosen.

在 3.5 版更改: The *backlog* parameter is now optional.

`socket.makefile(mode='r', buffering=None, *, encoding=None, errors=None, newline=None)`

Return a *file object* associated with the socket. The exact returned type depends on the arguments given to *makefile()*. These arguments are interpreted the same way as by the built-in *open()* function, except the only supported *mode* values are `'r'` (default), `'w'` and `'b'`.

The socket must be in blocking mode; it can have a timeout, but the file object's internal buffer may end up in an inconsistent state if a timeout occurs.

Closing the file object returned by *makefile()* won't close the original socket unless all other file objects have been closed and *socket.close()* has been called on the socket object.

注解: On Windows, the file-like object created by *makefile()* cannot be used where a file object with a file descriptor is expected, such as the stream arguments of *subprocess.Popen()*.

`socket.recv(buFSIZE[, flags])`

Receive data from the socket. The return value is a bytes object representing the data received. The maximum

amount of data to be received at once is specified by *bufsize*. See the Unix manual page *recv(2)* for the meaning of the optional argument *flags*; it defaults to zero.

注解: For best match with hardware and network realities, the value of *bufsize* should be a relatively small power of 2, for example, 4096.

在 3.5 版更改: If the system call is interrupted and the signal handler does not raise an exception, the method now retries the system call instead of raising an *InterruptedError* exception (see [PEP 475](#) for the rationale).

`socket.recvfrom(bufsize[, flags])`

Receive data from the socket. The return value is a pair (*bytes*, *address*) where *bytes* is a bytes object representing the data received and *address* is the address of the socket sending the data. See the Unix manual page *recv(2)* for the meaning of the optional argument *flags*; it defaults to zero. (The format of *address* depends on the address family —see above.)

在 3.5 版更改: If the system call is interrupted and the signal handler does not raise an exception, the method now retries the system call instead of raising an *InterruptedError* exception (see [PEP 475](#) for the rationale).

`socket.recvmsg(bufsize[, ancbufsize[, flags]])`

Receive normal data (up to *bufsize* bytes) and ancillary data from the socket. The *ancbufsize* argument sets the size in bytes of the internal buffer used to receive the ancillary data; it defaults to 0, meaning that no ancillary data will be received. Appropriate buffer sizes for ancillary data can be calculated using *MSG_SPACE()* or *MSG_LEN()*, and items which do not fit into the buffer might be truncated or discarded. The *flags* argument defaults to 0 and has the same meaning as for *recv()*.

The return value is a 4-tuple: (*data*, *ancdata*, *msg_flags*, *address*). The *data* item is a *bytes* object holding the non-ancillary data received. The *ancdata* item is a list of zero or more tuples (*cmsg_level*, *cmsg_type*, *cmsg_data*) representing the ancillary data (control messages) received: *cmsg_level* and *cmsg_type* are integers specifying the protocol level and protocol-specific type respectively, and *cmsg_data* is a *bytes* object holding the associated data. The *msg_flags* item is the bitwise OR of various flags indicating conditions on the received message; see your system documentation for details. If the receiving socket is unconnected, *address* is the address of the sending socket, if available; otherwise, its value is unspecified.

On some systems, *sendmsg()* and *recvmsg()* can be used to pass file descriptors between processes over an *AF_UNIX* socket. When this facility is used (it is often restricted to *SOCK_STREAM* sockets), *recvmsg()* will return, in its ancillary data, items of the form (*socket.SOL_SOCKET*, *socket.SCM_RIGHTS*, *fds*), where *fds* is a *bytes* object representing the new file descriptors as a binary array of the native C *int* type. If *recvmsg()* raises an exception after the system call returns, it will first attempt to close any file descriptors received via this mechanism.

Some systems do not indicate the truncated length of ancillary data items which have been only partially received. If an item appears to extend beyond the end of the buffer, *recvmsg()* will issue a *RuntimeWarning*, and will return the part of it which is inside the buffer provided it has not been truncated before the start of its associated data.

On systems which support the *SCM_RIGHTS* mechanism, the following function will receive up to *maxfds* file descriptors, returning the message data and a list containing the descriptors (while ignoring unexpected conditions such as unrelated control messages being received). See also *sendmsg()*.

```
import socket, array

def recv_fds(sock, msglen, maxfds):
    fds = array.array("i")    # Array of ints
    msg, ancdata, flags, addr = sock.recvmsg(msglen, socket.CMSG_LEN(maxfds * fds.
    ↪ itemsize))
    for cmsg_level, cmsg_type, cmsg_data in ancdata:
```

(下页继续)

(续上页)

```

    if (cmsg_level == socket.SOL_SOCKET and cmsg_type == socket.SCM_RIGHTS):
        # Append data, ignoring any truncated integers at the end.
        fds.fromstring(cmsg_data[:len(cmsg_data) - (len(cmsg_data) % fds.
↪itemsize)])
    return msg, list(fds)

```

Availability: most Unix platforms, possibly others.

3.3 新版功能.

在 3.5 版更改: If the system call is interrupted and the signal handler does not raise an exception, the method now retries the system call instead of raising an *InterruptedError* exception (see [PEP 475](#) for the rationale).

`socket.recvmsg_into(buffers[, ancbufsize[, flags]])`

Receive normal data and ancillary data from the socket, behaving as *recvmsg()* would, but scatter the non-ancillary data into a series of buffers instead of returning a new bytes object. The *buffers* argument must be an iterable of objects that export writable buffers (e.g. *bytearray* objects); these will be filled with successive chunks of the non-ancillary data until it has all been written or there are no more buffers. The operating system may set a limit (*sysconf()* value *SC_IOV_MAX*) on the number of buffers that can be used. The *ancbufsize* and *flags* arguments have the same meaning as for *recvmsg()*.

The return value is a 4-tuple: (*nbytes*, *ancdata*, *msg_flags*, *address*), where *nbytes* is the total number of bytes of non-ancillary data written into the buffers, and *ancdata*, *msg_flags* and *address* are the same as for *recvmsg()*.

示例:

```

>>> import socket
>>> s1, s2 = socket.socketpair()
>>> b1 = bytearray(b'----')
>>> b2 = bytearray(b'0123456789')
>>> b3 = bytearray(b'-----')
>>> s1.send(b'Mary had a little lamb')
22
>>> s2.recvmsg_into([b1, memoryview(b2)[2:9], b3])
(22, [], 0, None)
>>> [b1, b2, b3]
[bytearray(b'Mary'), bytearray(b'01 had a 9'), bytearray(b'little lamb---')]

```

Availability: most Unix platforms, possibly others.

3.3 新版功能.

`socket.recvfrom_into(buffer[, nbytes[, flags]])`

Receive data from the socket, writing it into *buffer* instead of creating a new bytestring. The return value is a pair (*nbytes*, *address*) where *nbytes* is the number of bytes received and *address* is the address of the socket sending the data. See the Unix manual page *recv(2)* for the meaning of the optional argument *flags*; it defaults to zero. (The format of *address* depends on the address family — see above.)

`socket.recv_into(buffer[, nbytes[, flags]])`

Receive up to *nbytes* bytes from the socket, storing the data into a buffer rather than creating a new bytestring. If *nbytes* is not specified (or 0), receive up to the size available in the given buffer. Returns the number of bytes received. See the Unix manual page *recv(2)* for the meaning of the optional argument *flags*; it defaults to zero.

`socket.send(bytes[, flags])`

Send data to the socket. The socket must be connected to a remote socket. The optional *flags* argument has the same meaning as for *recv()* above. Returns the number of bytes sent. Applications are responsible for checking that all data has been sent; if only some of the data was transmitted, the application needs to attempt delivery of the remaining data. For further information on this topic, consult the socket-howto.

在 3.5 版更改: If the system call is interrupted and the signal handler does not raise an exception, the method now retries the system call instead of raising an *InterruptedError* exception (see [PEP 475](#) for the rationale).

`socket.sendall(bytes[, flags])`

Send data to the socket. The socket must be connected to a remote socket. The optional *flags* argument has the same meaning as for *recv()* above. Unlike *send()*, this method continues to send data from *bytes* until either all data has been sent or an error occurs. *None* is returned on success. On error, an exception is raised, and there is no way to determine how much data, if any, was successfully sent.

在 3.5 版更改: The socket timeout is no more reset each time data is sent successfully. The socket timeout is now the maximum total duration to send all data.

在 3.5 版更改: If the system call is interrupted and the signal handler does not raise an exception, the method now retries the system call instead of raising an *InterruptedError* exception (see [PEP 475](#) for the rationale).

`socket.sendto(bytes, address)`

`socket.sendto(bytes, flags, address)`

Send data to the socket. The socket should not be connected to a remote socket, since the destination socket is specified by *address*. The optional *flags* argument has the same meaning as for *recv()* above. Return the number of bytes sent. (The format of *address* depends on the address family —see above.)

在 3.5 版更改: If the system call is interrupted and the signal handler does not raise an exception, the method now retries the system call instead of raising an *InterruptedError* exception (see [PEP 475](#) for the rationale).

`socket.sendmsg(buffers[, ancdata[, flags[, address]]])`

Send normal and ancillary data to the socket, gathering the non-ancillary data from a series of buffers and concatenating it into a single message. The *buffers* argument specifies the non-ancillary data as an iterable of *bytes-like objects* (e.g. *bytes* objects); the operating system may set a limit (*sysconf()* value *SC_IOV_MAX*) on the number of buffers that can be used. The *ancdata* argument specifies the ancillary data (control messages) as an iterable of zero or more tuples (*cmsg_level*, *cmsg_type*, *cmsg_data*), where *cmsg_level* and *cmsg_type* are integers specifying the protocol level and protocol-specific type respectively, and *cmsg_data* is a bytes-like object holding the associated data. Note that some systems (in particular, systems without *CMSG_SPACE()*) might support sending only one control message per call. The *flags* argument defaults to 0 and has the same meaning as for *send()*. If *address* is supplied and not *None*, it sets a destination address for the message. The return value is the number of bytes of non-ancillary data sent.

The following function sends the list of file descriptors *fds* over an *AF_UNIX* socket, on systems which support the *SCM_RIGHTS* mechanism. See also *recvmsg()*.

```
import socket, array

def send_fds(sock, msg, fds):
    return sock.sendmsg([msg], [(socket.SOL_SOCKET, socket.SCM_RIGHTS, array.
    ↪array("i", fds))])
```

Availability: most Unix platforms, possibly others.

3.3 新版功能.

在 3.5 版更改: If the system call is interrupted and the signal handler does not raise an exception, the method now retries the system call instead of raising an *InterruptedError* exception (see [PEP 475](#) for the rationale).

`socket.sendmsg_afalg([msg], *, op[, iv[, assoclen[, flags]]])`

Specialized version of *sendmsg()* for *AF_ALG* socket. Set mode, IV, AEAD associated data length and flags for *AF_ALG* socket.

Availability: Linux >= 2.6.38

3.6 新版功能.

`socket.sendfile(file, offset=0, count=None)`

Send a file until EOF is reached by using high-performance `os.sendfile` and return the total number of bytes which were sent. `file` must be a regular file object opened in binary mode. If `os.sendfile` is not available (e.g. Windows) or `file` is not a regular file `send()` will be used instead. `offset` tells from where to start reading the file. If specified, `count` is the total number of bytes to transmit as opposed to sending the file until EOF is reached. File position is updated on return or also in case of error in which case `file.tell()` can be used to figure out the number of bytes which were sent. The socket must be of `SOCK_STREAM` type. Non-blocking sockets are not supported.

3.5 新版功能.

`socket.set_inheritable(inheritable)`

Set the *inheritable flag* of the socket's file descriptor or socket's handle.

3.4 新版功能.

`socket.setblocking(flag)`

Set blocking or non-blocking mode of the socket: if `flag` is false, the socket is set to non-blocking, else to blocking mode.

This method is a shorthand for certain `settimeout()` calls:

- `sock.setblocking(True)` is equivalent to `sock.settimeout(None)`
- `sock.setblocking(False)` is equivalent to `sock.settimeout(0.0)`

`socket.settimeout(value)`

Set a timeout on blocking socket operations. The `value` argument can be a nonnegative floating point number expressing seconds, or `None`. If a non-zero value is given, subsequent socket operations will raise a `timeout` exception if the timeout period `value` has elapsed before the operation has completed. If zero is given, the socket is put in non-blocking mode. If `None` is given, the socket is put in blocking mode.

For further information, please consult the [notes on socket timeouts](#).

`socket.setsockopt(level, optname, value: int)`

`socket.setsockopt(level, optname, value: buffer)`

`socket.setsockopt(level, optname, None, optlen: int)`

Set the value of the given socket option (see the Unix manual page `setsockopt(2)`). The needed symbolic constants are defined in the `socket` module (`SO_*` etc.). The value can be an integer, `None` or a *bytes-like object* representing a buffer. In the later case it is up to the caller to ensure that the bytestring contains the proper bits (see the optional built-in module `struct` for a way to encode C structures as bytestrings). When value is set to `None`, `optlen` argument is required. It's equivalent to call `setsockopt` C function with `optval=NULL` and `optlen=optlen`.

在 3.5 版更改: Writable *bytes-like object* is now accepted.

在 3.6 版更改: `setsockopt(level, optname, None, optlen: int)` form added.

`socket.shutdown(how)`

Shut down one or both halves of the connection. If `how` is `SHUT_RD`, further receives are disallowed. If `how` is `SHUT_WR`, further sends are disallowed. If `how` is `SHUT_RDWR`, further sends and receives are disallowed.

`socket.share(process_id)`

Duplicate a socket and prepare it for sharing with a target process. The target process must be provided with `process_id`. The resulting bytes object can then be passed to the target process using some form of interprocess communication and the socket can be recreated there using `fromshare()`. Once this method has been called, it is safe to close the socket since the operating system has already duplicated it for the target process.

Availability: Windows.

3.3 新版功能.

Note that there are no methods `read()` or `write()`; use `recv()` and `send()` without *flags* argument instead.

Socket objects also have these (read-only) attributes that correspond to the values given to the `socket` constructor.

`socket.family`

The socket family.

`socket.type`

The socket type.

`socket.proto`

The socket protocol.

18.1.4 Notes on socket timeouts

A socket object can be in one of three modes: blocking, non-blocking, or timeout. Sockets are by default always created in blocking mode, but this can be changed by calling `setdefaulttimeout()`.

- In *blocking mode*, operations block until complete or the system returns an error (such as connection timed out).
- In *non-blocking mode*, operations fail (with an error that is unfortunately system-dependent) if they cannot be completed immediately: functions from the `select` can be used to know when and whether a socket is available for reading or writing.
- In *timeout mode*, operations fail if they cannot be completed within the timeout specified for the socket (they raise a `timeout` exception) or if the system returns an error.

注解: At the operating system level, sockets in *timeout mode* are internally set in non-blocking mode. Also, the blocking and timeout modes are shared between file descriptors and socket objects that refer to the same network endpoint. This implementation detail can have visible consequences if e.g. you decide to use the `fileno()` of a socket.

Timeouts and the `connect` method

The `connect()` operation is also subject to the timeout setting, and in general it is recommended to call `settimeout()` before calling `connect()` or pass a timeout parameter to `create_connection()`. However, the system network stack may also return a connection timeout error of its own regardless of any Python socket timeout setting.

Timeouts and the `accept` method

If `getdefaulttimeout()` is not `None`, sockets returned by the `accept()` method inherit that timeout. Otherwise, the behaviour depends on settings of the listening socket:

- if the listening socket is in *blocking mode* or in *timeout mode*, the socket returned by `accept()` is in *blocking mode*;
- if the listening socket is in *non-blocking mode*, whether the socket returned by `accept()` is in blocking or non-blocking mode is operating system-dependent. If you want to ensure cross-platform behaviour, it is recommended you manually override this setting.

18.1.5 示例

Here are four minimal example programs using the TCP/IP protocol: a server that echoes all data that it receives back (servicing only one client), and a client using it. Note that a server must perform the sequence `socket()`, `bind()`, `listen()`, `accept()` (possibly repeating the `accept()` to service more than one client), while a client only needs the sequence `socket()`, `connect()`. Also note that the server does not `sendall()/recv()` on the socket it is listening on but on the new socket returned by `accept()`.

The first two examples support IPv4 only.

```
# Echo server program
import socket

HOST = ''          # Symbolic name meaning all available interfaces
PORT = 50007       # Arbitrary non-privileged port
with socket.socket(socket.AF_INET, socket.SOCK_STREAM) as s:
    s.bind((HOST, PORT))
    s.listen(1)
    conn, addr = s.accept()
    with conn:
        print('Connected by', addr)
        while True:
            data = conn.recv(1024)
            if not data: break
            conn.sendall(data)
```

```
# Echo client program
import socket

HOST = 'daring.cwi.nl' # The remote host
PORT = 50007           # The same port as used by the server
with socket.socket(socket.AF_INET, socket.SOCK_STREAM) as s:
    s.connect((HOST, PORT))
    s.sendall(b'Hello, world')
    data = s.recv(1024)
print('Received', repr(data))
```

The next two examples are identical to the above two, but support both IPv4 and IPv6. The server side will listen to the first address family available (it should listen to both instead). On most of IPv6-ready systems, IPv6 will take precedence and the server may not accept IPv4 traffic. The client side will try to connect to the all addresses returned as a result of the name resolution, and sends traffic to the first one connected successfully.

```
# Echo server program
import socket
import sys

HOST = None        # Symbolic name meaning all available interfaces
PORT = 50007       # Arbitrary non-privileged port
s = None
for res in socket.getaddrinfo(HOST, PORT, socket.AF_UNSPEC,
                              socket.SOCK_STREAM, 0, socket.AI_PASSIVE):
    af, socktype, proto, canonname, sa = res
    try:
        s = socket.socket(af, socktype, proto)
    except OSError as msg:
        s = None
        continue
```

(下页继续)

(续上页)

```

    try:
        s.bind(sa)
        s.listen(1)
    except OSError as msg:
        s.close()
        s = None
        continue
    break
if s is None:
    print('could not open socket')
    sys.exit(1)
conn, addr = s.accept()
with conn:
    print('Connected by', addr)
    while True:
        data = conn.recv(1024)
        if not data: break
        conn.send(data)

```

```

# Echo client program
import socket
import sys

HOST = 'daring.cwi.nl'      # The remote host
PORT = 50007                # The same port as used by the server
s = None
for res in socket.getaddrinfo(HOST, PORT, socket.AF_UNSPEC, socket.SOCK_STREAM):
    af, socktype, proto, canonname, sa = res
    try:
        s = socket.socket(af, socktype, proto)
    except OSError as msg:
        s = None
        continue
    try:
        s.connect(sa)
    except OSError as msg:
        s.close()
        s = None
        continue
    break
if s is None:
    print('could not open socket')
    sys.exit(1)
with s:
    s.sendall(b'Hello, world')
    data = s.recv(1024)
print('Received', repr(data))

```

The next example shows how to write a very simple network sniffer with raw sockets on Windows. The example requires administrator privileges to modify the interface:

```

import socket

# the public network interface
HOST = socket.gethostbyname(socket.gethostname())

```

(下页继续)

(续上页)

```

# create a raw socket and bind it to the public interface
s = socket.socket(socket.AF_INET, socket.SOCK_RAW, socket.IPPROTO_IP)
s.bind((HOST, 0))

# Include IP headers
s.setsockopt(socket.IPPROTO_IP, socket.IP_HDRINCL, 1)

# receive all packages
s.ioctl(socket.SIO_RCVALL, socket.RCVALL_ON)

# receive a package
print(s.recvfrom(65565))

# disabled promiscuous mode
s.ioctl(socket.SIO_RCVALL, socket.RCVALL_OFF)

```

The last example shows how to use the socket interface to communicate to a CAN network using the raw socket protocol. To use CAN with the broadcast manager protocol instead, open a socket with:

```
socket.socket(socket.AF_CAN, socket.SOCK_DGRAM, socket.CAN_BCM)
```

After binding (`CAN_RAW`) or connecting (`CAN_BCM`) the socket, you can use the `socket.send()`, and the `socket.recv()` operations (and their counterparts) on the socket object as usual.

This example might require special privileges:

```

import socket
import struct

# CAN frame packing/unpacking (see 'struct can_frame' in <linux/can.h>)

can_frame_fmt = "=IB3x8s"
can_frame_size = struct.calcsize(can_frame_fmt)

def build_can_frame(can_id, data):
    can_dlc = len(data)
    data = data.ljust(8, b'\x00')
    return struct.pack(can_frame_fmt, can_id, can_dlc, data)

def dissect_can_frame(frame):
    can_id, can_dlc, data = struct.unpack(can_frame_fmt, frame)
    return (can_id, can_dlc, data[:can_dlc])

# create a raw socket and bind it to the 'vcan0' interface
s = socket.socket(socket.AF_CAN, socket.SOCK_RAW, socket.CAN_RAW)
s.bind(('vcan0',))

while True:
    cf, addr = s.recvfrom(can_frame_size)

    print('Received: can_id=%x, can_dlc=%x, data=%s' % dissect_can_frame(cf))

    try:
        s.send(cf)
    except OSError:

```

(下页继续)

(续上页)

```
print('Error sending CAN frame')

try:
    s.send(build_can_frame(0x01, b'\x01\x02\x03'))
except OSError:
    print('Error sending CAN frame')
```

Running an example several times with too small delay between executions, could lead to this error:

```
OSError: [Errno 98] Address already in use
```

This is because the previous execution has left the socket in a `TIME_WAIT` state, and can't be immediately reused.

There is a `socket` flag to set, in order to prevent this, `socket.SO_REUSEADDR`:

```
s = socket.socket(socket.AF_INET, socket.SOCK_STREAM)
s.setsockopt(socket.SOL_SOCKET, socket.SO_REUSEADDR, 1)
s.bind((HOST, PORT))
```

the `SO_REUSEADDR` flag tells the kernel to reuse a local socket in `TIME_WAIT` state, without waiting for its natural timeout to expire.

参见:

For an introduction to socket programming (in C), see the following papers:

- *An Introductory 4.3BSD Interprocess Communication Tutorial*, by Stuart Sechrest
- *An Advanced 4.3BSD Interprocess Communication Tutorial*, by Samuel J. Leffler et al,

both in the UNIX Programmer's Manual, Supplementary Documents 1 (sections PS1:7 and PS1:8). The platform-specific reference material for the various socket-related system calls are also a valuable source of information on the details of socket semantics. For Unix, refer to the manual pages; for Windows, see the WinSock (or Winsock 2) specification. For IPv6-ready APIs, readers may want to refer to [RFC 3493](#) titled Basic Socket Interface Extensions for IPv6.

18.2 `ssl` — TLS/SSL wrapper for socket objects

Source code: [Lib/ssl.py](#)

This module provides access to Transport Layer Security (often known as “Secure Sockets Layer”) encryption and peer authentication facilities for network sockets, both client-side and server-side. This module uses the OpenSSL library. It is available on all modern Unix systems, Windows, Mac OS X, and probably additional platforms, as long as OpenSSL is installed on that platform.

注解: Some behavior may be platform dependent, since calls are made to the operating system socket APIs. The installed version of OpenSSL may also cause variations in behavior. For example, TLSv1.1 and TLSv1.2 come with openssl version 1.0.1.

警告: Don't use this module without reading the [Security considerations](#). Doing so may lead to a false sense of security, as the default settings of the `ssl` module are not necessarily appropriate for your application.

This section documents the objects and functions in the `ssl` module; for more general information about TLS, SSL, and certificates, the reader is referred to the documents in the “See Also” section at the bottom.

This module provides a class, `ssl.SSLSocket`, which is derived from the `socket.socket` type, and provides a socket-like wrapper that also encrypts and decrypts the data going over the socket with SSL. It supports additional methods such as `getpeercert()`, which retrieves the certificate of the other side of the connection, and `cipher()`, which retrieves the cipher being used for the secure connection.

For more sophisticated applications, the `ssl.SSLContext` class helps manage settings and certificates, which can then be inherited by SSL sockets created through the `SSLContext.wrap_socket()` method.

在 3.5.3 版更改: Updated to support linking with OpenSSL 1.1.0

在 3.6 版更改: OpenSSL 0.9.8, 1.0.0 and 1.0.1 are deprecated and no longer supported. In the future the `ssl` module will require at least OpenSSL 1.0.2 or 1.1.0.

18.2.1 Functions, Constants, and Exceptions

exception `ssl.SSLError`

Raised to signal an error from the underlying SSL implementation (currently provided by the OpenSSL library). This signifies some problem in the higher-level encryption and authentication layer that’s superimposed on the underlying network connection. This error is a subtype of `OSError`. The error code and message of `SSLError` instances are provided by the OpenSSL library.

在 3.3 版更改: `SSLError` used to be a subtype of `socket.error`.

library

A string mnemonic designating the OpenSSL submodule in which the error occurred, such as `SSL`, `PEM` or `X509`. The range of possible values depends on the OpenSSL version.

3.3 新版功能.

reason

A string mnemonic designating the reason this error occurred, for example `CERTIFICATE_VERIFY_FAILED`. The range of possible values depends on the OpenSSL version.

3.3 新版功能.

exception `ssl.SSLZeroReturnError`

A subclass of `SSLError` raised when trying to read or write and the SSL connection has been closed cleanly. Note that this doesn’t mean that the underlying transport (read TCP) has been closed.

3.3 新版功能.

exception `ssl.SSLWantReadError`

A subclass of `SSLError` raised by a *non-blocking SSL socket* when trying to read or write data, but more data needs to be received on the underlying TCP transport before the request can be fulfilled.

3.3 新版功能.

exception `ssl.SSLWantWriteError`

A subclass of `SSLError` raised by a *non-blocking SSL socket* when trying to read or write data, but more data needs to be sent on the underlying TCP transport before the request can be fulfilled.

3.3 新版功能.

exception `ssl.SSLSyscallError`

A subclass of `SSLError` raised when a system error was encountered while trying to fulfill an operation on a SSL socket. Unfortunately, there is no easy way to inspect the original `errno` number.

3.3 新版功能.

exception `ssl.SSLEOFError`

A subclass of `SSL`[Error](#) raised when the SSL connection has been terminated abruptly. Generally, you shouldn't try to reuse the underlying transport when this error is encountered.

3.3 新版功能.

exception `ssl.CertificateError`

Raised to signal an error with a certificate (such as mismatching hostname). Certificate errors detected by OpenSSL, though, raise an `SSL`[Error](#).

Socket creation

The following function allows for standalone socket creation. Starting from Python 3.2, it can be more flexible to use `SSLContext.wrap_socket()` instead.

`ssl.wrap_socket(sock, keyfile=None, certfile=None, server_side=False, cert_reqs=CERT_NONE, ssl_version={see docs}, ca_certs=None, do_handshake_on_connect=True, suppress_ragged_eofs=True, ciphers=None)`

Takes an instance `sock` of `socket.socket`, and returns an instance of `ssl.SSLSocket`, a subtype of `socket.socket`, which wraps the underlying socket in an SSL context. `sock` must be a `SOCK_STREAM` socket; other socket types are unsupported.

For client-side sockets, the context construction is lazy; if the underlying socket isn't connected yet, the context construction will be performed after `connect()` is called on the socket. For server-side sockets, if the socket has no remote peer, it is assumed to be a listening socket, and the server-side SSL wrapping is automatically performed on client connections accepted via the `accept()` method. `wrap_socket()` may raise `SSL`[Error](#).

The `keyfile` and `certfile` parameters specify optional files which contain a certificate to be used to identify the local side of the connection. See the discussion of [Certificates](#) for more information on how the certificate is stored in the `certfile`.

The parameter `server_side` is a boolean which identifies whether server-side or client-side behavior is desired from this socket.

The parameter `cert_reqs` specifies whether a certificate is required from the other side of the connection, and whether it will be validated if provided. It must be one of the three values `CERT_NONE` (certificates ignored), `CERT_OPTIONAL` (not required, but validated if provided), or `CERT_REQUIRED` (required and validated). If the value of this parameter is not `CERT_NONE`, then the `ca_certs` parameter must point to a file of CA certificates.

The `ca_certs` file contains a set of concatenated “certification authority” certificates, which are used to validate certificates passed from the other end of the connection. See the discussion of [Certificates](#) for more information about how to arrange the certificates in this file.

The parameter `ssl_version` specifies which version of the SSL protocol to use. Typically, the server chooses a particular protocol version, and the client must adapt to the server's choice. Most of the versions are not interoperable with the other versions. If not specified, the default is `PROTOCOL_TLS`; it provides the most compatibility with other versions.

Here's a table showing which versions in a client (down the side) can connect to which versions in a server (along the top):

<i>client / server</i>	SSLv2	SSLv3	TLS ³	TLSv1	TLSv1.1	TLSv1.2
SSLv2	yes	no	no ¹	no	no	no
SSLv3	no	yes	no ²	no	no	no
TLS (SSLv23) ³	no ¹	no ²	yes	yes	yes	yes
TLSv1	no	no	yes	yes	no	no
TLSv1.1	no	no	yes	no	yes	no
TLSv1.2	no	no	yes	no	no	yes

注解： Which connections succeed will vary depending on the version of OpenSSL. For example, before OpenSSL 1.0.0, an SSLv23 client would always attempt SSLv2 connections.

The *ciphers* parameter sets the available ciphers for this SSL object. It should be a string in the [OpenSSL cipher list format](#).

The parameter *do_handshake_on_connect* specifies whether to do the SSL handshake automatically after doing a `socket.connect()`, or whether the application program will call it explicitly, by invoking the `SSLSocket.do_handshake()` method. Calling `SSLSocket.do_handshake()` explicitly gives the program control over the blocking behavior of the socket I/O involved in the handshake.

The parameter *suppress_ragged_eofs* specifies how the `SSLSocket.recv()` method should signal unexpected EOF from the other end of the connection. If specified as *True* (the default), it returns a normal EOF (an empty bytes object) in response to unexpected EOF errors raised from the underlying socket; if *False*, it will raise the exceptions back to the caller.

在 3.2 版更改: New optional argument *ciphers*.

Context creation

A convenience function helps create `SSLContext` objects for common purposes.

`ssl.create_default_context(purpose=Purpose.SERVER_AUTH, cafile=None, capath=None, cadata=None)`

Return a new `SSLContext` object with default settings for the given *purpose*. The settings are chosen by the `ssl` module, and usually represent a higher security level than when calling the `SSLContext` constructor directly.

cafile, *capath*, *cadata* represent optional CA certificates to trust for certificate verification, as in `SSLContext.load_verify_locations()`. If all three are *None*, this function can choose to trust the system's default CA certificates instead.

The settings are: `PROTOCOL_TLS`, `OP_NO_SSLv2`, and `OP_NO_SSLv3` with high encryption cipher suites without RC4 and without unauthenticated cipher suites. Passing `SERVER_AUTH` as *purpose* sets *verify_mode* to `CERT_REQUIRED` and either loads CA certificates (when at least one of *cafile*, *capath* or *cadata* is given) or uses `SSLContext.load_default_certs()` to load default CA certificates.

注解： The protocol, options, cipher and other settings may change to more restrictive values anytime without prior deprecation. The values represent a fair balance between compatibility and security.

If your application needs specific settings, you should create a `SSLContext` and apply the settings yourself.

注解： If you find that when certain older clients or servers attempt to connect with a `SSLContext` created by this function that they get an error stating “Protocol or cipher suite mismatch”, it may be that they only support SSL3.0 which this function excludes using the `OP_NO_SSLv3`. SSL3.0 is widely considered to be [completely broken](#). If you still wish to continue to use this function but still allow SSL 3.0 connections you can re-enable them using:

```
ctx = ssl.create_default_context(Purpose.CLIENT_AUTH)
ctx.options |= ~ssl.OP_NO_SSLv3
```

³ TLS 1.3 protocol will be available with `PROTOCOL_TLS` in OpenSSL \geq 1.1.1. There is no dedicated PROTOCOL constant for just TLS 1.3.

¹ `SSLContext` disables SSLv2 with `OP_NO_SSLv2` by default.

² `SSLContext` disables SSLv3 with `OP_NO_SSLv3` by default.

3.4 新版功能.

在 3.4.4 版更改: RC4 was dropped from the default cipher string.

在 3.6 版更改: ChaCha20/Poly1305 was added to the default cipher string.

3DES was dropped from the default cipher string.

Random generation

`ssl.RAND_bytes(num)`

Return *num* cryptographically strong pseudo-random bytes. Raises an *SSL**Error* if the PRNG has not been seeded with enough data or if the operation is not supported by the current RAND method. *RAND_status()* can be used to check the status of the PRNG and *RAND_add()* can be used to seed the PRNG.

For almost all applications *os.urandom()* is preferable.

Read the Wikipedia article, [Cryptographically secure pseudorandom number generator \(CSPRNG\)](#), to get the requirements of a cryptographically generator.

3.3 新版功能.

`ssl.RAND_pseudo_bytes(num)`

Return (bytes, is_cryptographic): bytes are *num* pseudo-random bytes, is_cryptographic is *True* if the bytes generated are cryptographically strong. Raises an *SSL**Error* if the operation is not supported by the current RAND method.

Generated pseudo-random byte sequences will be unique if they are of sufficient length, but are not necessarily unpredictable. They can be used for non-cryptographic purposes and for certain purposes in cryptographic protocols, but usually not for key generation etc.

For almost all applications *os.urandom()* is preferable.

3.3 新版功能.

3.6 版后已移除: OpenSSL has deprecated *ssl.RAND_pseudo_bytes()*, use *ssl.RAND_bytes()* instead.

`ssl.RAND_status()`

Return *True* if the SSL pseudo-random number generator has been seeded with ‘enough’ randomness, and *False* otherwise. You can use *ssl.RAND_egd()* and *ssl.RAND_add()* to increase the randomness of the pseudo-random number generator.

`ssl.RAND_egd(path)`

If you are running an entropy-gathering daemon (EGD) somewhere, and *path* is the pathname of a socket connection open to it, this will read 256 bytes of randomness from the socket, and add it to the SSL pseudo-random number generator to increase the security of generated secret keys. This is typically only necessary on systems without better sources of randomness.

See <http://egd.sourceforge.net/> or <http://prngd.sourceforge.net/> for sources of entropy-gathering daemons.

Availability: not available with LibreSSL and OpenSSL > 1.1.0

`ssl.RAND_add(bytes, entropy)`

Mix the given *bytes* into the SSL pseudo-random number generator. The parameter *entropy* (a float) is a lower bound on the entropy contained in string (so you can always use 0.0). See [RFC 1750](#) for more information on sources of entropy.

在 3.5 版更改: Writable *bytes-like object* is now accepted.

Certificate handling

`ssl.match_hostname(cert, hostname)`

Verify that *cert* (in decoded format as returned by `SSLSocket.getpeercert()`) matches the given *hostname*. The rules applied are those for checking the identity of HTTPS servers as outlined in [RFC 2818](#), [RFC 5280](#) and [RFC 6125](#). In addition to HTTPS, this function should be suitable for checking the identity of servers in various SSL-based protocols such as FTPS, IMAPS, POPS and others.

`CertificateError` is raised on failure. On success, the function returns nothing:

```
>>> cert = {'subject': (('commonName', 'example.com'),),)}
>>> ssl.match_hostname(cert, "example.com")
>>> ssl.match_hostname(cert, "example.org")
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
  File "/home/py3k/Lib/ssl.py", line 130, in match_hostname
ssl.CertificateError: hostname 'example.org' doesn't match 'example.com'
```

3.2 新版功能.

在 3.3.3 版更改: The function now follows [RFC 6125](#), section 6.4.3 and does neither match multiple wildcards (e.g. `*.*.com` or `*a*.example.org`) nor a wildcard inside an internationalized domain names (IDN) fragment. IDN A-labels such as `www*.xn--python-kva.org` are still supported, but `x*.python.org` no longer matches `xn--tda.python.org`.

在 3.5 版更改: Matching of IP addresses, when present in the `subjectAltName` field of the certificate, is now supported.

`ssl.cert_time_to_seconds(cert_time)`

Return the time in seconds since the Epoch, given the *cert_time* string representing the “notBefore” or “notAfter” date from a certificate in “%b %d %H:%M:%S %Y %Z” `strptime` format (C locale).

Here’s an example:

```
>>> import ssl
>>> timestamp = ssl.cert_time_to_seconds("Jan  5 09:34:43 2018 GMT")
>>> timestamp
1515144883
>>> from datetime import datetime
>>> print(datetime.utcfromtimestamp(timestamp))
2018-01-05 09:34:43
```

“notBefore” or “notAfter” dates must use GMT ([RFC 5280](#)).

在 3.5 版更改: Interpret the input time as a time in UTC as specified by ‘GMT’ timezone in the input string. Local timezone was used previously. Return an integer (no fractions of a second in the input format)

`ssl.get_server_certificate(addr, ssl_version=PROTOCOL_TLS, ca_certs=None)`

Given the address *addr* of an SSL-protected server, as a (*hostname*, *port-number*) pair, fetches the server’s certificate, and returns it as a PEM-encoded string. If *ssl_version* is specified, uses that version of the SSL protocol to attempt to connect to the server. If *ca_certs* is specified, it should be a file containing a list of root certificates, the same format as used for the same parameter in `wrap_socket()`. The call will attempt to validate the server certificate against that set of root certificates, and will fail if the validation attempt fails.

在 3.3 版更改: This function is now IPv6-compatible.

在 3.5 版更改: The default *ssl_version* is changed from `PROTOCOL_SSLv3` to `PROTOCOL_TLS` for maximum compatibility with modern servers.

`ssl.DER_cert_to_PEM_cert(DER_cert_bytes)`

Given a certificate as a DER-encoded blob of bytes, returns a PEM-encoded string version of the same certificate.

`ssl.PEM_cert_to_DER_cert (PEM_cert_string)`

Given a certificate as an ASCII PEM string, returns a DER-encoded sequence of bytes for that same certificate.

`ssl.get_default_verify_paths()`

Returns a named tuple with paths to OpenSSL's default cafile and capath. The paths are the same as used by `SSLContext.set_default_verify_paths()`. The return value is a *named tuple* `DefaultVerifyPaths`:

- `cafile` - resolved path to cafile or `None` if the file doesn't exist,
- `capath` - resolved path to capath or `None` if the directory doesn't exist,
- `openssl_cafile_env` - OpenSSL's environment key that points to a cafile,
- `openssl_cafile` - hard coded path to a cafile,
- `openssl_capath_env` - OpenSSL's environment key that points to a capath,
- `openssl_capath` - hard coded path to a capath directory

Availability: LibreSSL ignores the environment vars `openssl_cafile_env` and `openssl_capath_env`

3.4 新版功能.

`ssl.enum_certificates (store_name)`

Retrieve certificates from Windows' system cert store. `store_name` may be one of `CA`, `ROOT` or `MY`. Windows may provide additional cert stores, too.

The function returns a list of (cert_bytes, encoding_type, trust) tuples. The `encoding_type` specifies the encoding of cert_bytes. It is either `x509_asn` for X.509 ASN.1 data or `pkcs_7_asn` for PKCS#7 ASN.1 data. Trust specifies the purpose of the certificate as a set of OIDS or exactly `True` if the certificate is trustworthy for all purposes.

示例:

```
>>> ssl.enum_certificates("CA")
[(b'data...', 'x509_asn', {'1.3.6.1.5.5.7.3.1', '1.3.6.1.5.5.7.3.2'}),
 (b'data...', 'x509_asn', True)]
```

Availability: Windows.

3.4 新版功能.

`ssl.enum_crls (store_name)`

Retrieve CRLs from Windows' system cert store. `store_name` may be one of `CA`, `ROOT` or `MY`. Windows may provide additional cert stores, too.

The function returns a list of (cert_bytes, encoding_type, trust) tuples. The `encoding_type` specifies the encoding of cert_bytes. It is either `x509_asn` for X.509 ASN.1 data or `pkcs_7_asn` for PKCS#7 ASN.1 data.

Availability: Windows.

3.4 新版功能.

常数

All constants are now `enum.IntEnum` or `enum.IntFlag` collections.

3.6 新版功能.

ssl.CERT_NONE

Possible value for `SSLContext.verify_mode`, or the `cert_reqs` parameter to `wrap_socket()`. Except for `PROTOCOL_TLS_CLIENT`, it is the default mode. With client-side sockets, just about any cert is accepted. Validation errors, such as untrusted or expired cert, are ignored and do not abort the TLS/SSL handshake.

In server mode, no certificate is requested from the client, so the client does not send any for client cert authentication.

See the discussion of *Security considerations* below.

ssl.CERT_OPTIONAL

Possible value for `SSLContext.verify_mode`, or the `cert_reqs` parameter to `wrap_socket()`. In client mode, `CERT_OPTIONAL` has the same meaning as `CERT_REQUIRED`. It is recommended to use `CERT_REQUIRED` for client-side sockets instead.

In server mode, a client certificate request is sent to the client. The client may either ignore the request or send a certificate in order to perform TLS client cert authentication. If the client chooses to send a certificate, it is verified. Any verification error immediately aborts the TLS handshake.

Use of this setting requires a valid set of CA certificates to be passed, either to `SSLContext.load_verify_locations()` or as a value of the `ca_certs` parameter to `wrap_socket()`.

ssl.CERT_REQUIRED

Possible value for `SSLContext.verify_mode`, or the `cert_reqs` parameter to `wrap_socket()`. In this mode, certificates are required from the other side of the socket connection; an `SSLError` will be raised if no certificate is provided, or if its validation fails. This mode is **not** sufficient to verify a certificate in client mode as it does not match hostnames. `check_hostname` must be enabled as well to verify the authenticity of a cert. `PROTOCOL_TLS_CLIENT` uses `CERT_REQUIRED` and enables `check_hostname` by default.

With server socket, this mode provides mandatory TLS client cert authentication. A client certificate request is sent to the client and the client must provide a valid and trusted certificate.

Use of this setting requires a valid set of CA certificates to be passed, either to `SSLContext.load_verify_locations()` or as a value of the `ca_certs` parameter to `wrap_socket()`.

class ssl.VerifyMode

enum.IntEnum collection of `CERT_*` constants.

3.6 新版功能.

ssl.VERIFY_DEFAULT

Possible value for `SSLContext.verify_flags`. In this mode, certificate revocation lists (CRLs) are not checked. By default OpenSSL does neither require nor verify CRLs.

3.4 新版功能.

ssl.VERIFY_CRL_CHECK_LEAF

Possible value for `SSLContext.verify_flags`. In this mode, only the peer cert is checked but none of the intermediate CA certificates. The mode requires a valid CRL that is signed by the peer cert's issuer (its direct ancestor CA). If no proper has been loaded `SSLContext.load_verify_locations`, validation will fail.

3.4 新版功能.

ssl.VERIFY_CRL_CHECK_CHAIN

Possible value for `SSLContext.verify_flags`. In this mode, CRLs of all certificates in the peer cert chain are checked.

3.4 新版功能.

ssl.VERIFY_X509_STRICT

Possible value for `SSLContext.verify_flags` to disable workarounds for broken X.509 certificates.

3.4 新版功能.

`ssl.VERIFY_X509_TRUSTED_FIRST`

Possible value for `SSLContext.verify_flags`. It instructs OpenSSL to prefer trusted certificates when building the trust chain to validate a certificate. This flag is enabled by default.

3.4.4 新版功能.

class `ssl.VerifyFlags`

`enum.IntFlag` collection of `VERIFY_*` constants.

3.6 新版功能.

`ssl.PROTOCOL_TLS`

Selects the highest protocol version that both the client and server support. Despite the name, this option can select both “SSL” and “TLS” protocols.

3.6 新版功能.

`ssl.PROTOCOL_TLS_CLIENT`

Auto-negotiate the highest protocol version like `PROTOCOL_TLS`, but only support client-side `SSLSocket` connections. The protocol enables `CERT_REQUIRED` and `check_hostname` by default.

3.6 新版功能.

`ssl.PROTOCOL_TLS_SERVER`

Auto-negotiate the highest protocol version like `PROTOCOL_TLS`, but only support server-side `SSLSocket` connections.

3.6 新版功能.

`ssl.PROTOCOL_SSLv23`

Alias for `data:PROTOCOL_TLS`.

3.6 版后已移除: Use `PROTOCOL_TLS` instead.

`ssl.PROTOCOL_SSLv2`

Selects SSL version 2 as the channel encryption protocol.

This protocol is not available if OpenSSL is compiled with the `OPENSSL_NO_SSL2` flag.

警告: SSL version 2 is insecure. Its use is highly discouraged.

3.6 版后已移除: OpenSSL has removed support for SSLv2.

`ssl.PROTOCOL_SSLv3`

Selects SSL version 3 as the channel encryption protocol.

This protocol is not be available if OpenSSL is compiled with the `OPENSSL_NO_SSLv3` flag.

警告: SSL version 3 is insecure. Its use is highly discouraged.

3.6 版后已移除: OpenSSL has deprecated all version specific protocols. Use the default protocol `PROTOCOL_TLS` with flags like `OP_NO_SSLv3` instead.

`ssl.PROTOCOL_TLSv1`

Selects TLS version 1.0 as the channel encryption protocol.

3.6 版后已移除: OpenSSL has deprecated all version specific protocols. Use the default protocol `PROTOCOL_TLS` with flags like `OP_NO_SSLv3` instead.

ssl.PROTOCOL_TLSv1_1

Selects TLS version 1.1 as the channel encryption protocol. Available only with openssl version 1.0.1+.

3.4 新版功能.

3.6 版后已移除: OpenSSL has deprecated all version specific protocols. Use the default protocol *PROTOCOL_TLS* with flags like *OP_NO_SSLv3* instead.

ssl.PROTOCOL_TLSv1_2

Selects TLS version 1.2 as the channel encryption protocol. This is the most modern version, and probably the best choice for maximum protection, if both sides can speak it. Available only with openssl version 1.0.1+.

3.4 新版功能.

3.6 版后已移除: OpenSSL has deprecated all version specific protocols. Use the default protocol *PROTOCOL_TLS* with flags like *OP_NO_SSLv3* instead.

ssl.OP_ALL

Enables workarounds for various bugs present in other SSL implementations. This option is set by default. It does not necessarily set the same flags as OpenSSL's *SSL_OP_ALL* constant.

3.2 新版功能.

ssl.OP_NO_SSLv2

Prevents an SSLv2 connection. This option is only applicable in conjunction with *PROTOCOL_TLS*. It prevents the peers from choosing SSLv2 as the protocol version.

3.2 新版功能.

3.6 版后已移除: SSLv2 is deprecated

ssl.OP_NO_SSLv3

Prevents an SSLv3 connection. This option is only applicable in conjunction with *PROTOCOL_TLS*. It prevents the peers from choosing SSLv3 as the protocol version.

3.2 新版功能.

3.6 版后已移除: SSLv3 is deprecated

ssl.OP_NO_TLSv1

Prevents a TLSv1 connection. This option is only applicable in conjunction with *PROTOCOL_TLS*. It prevents the peers from choosing TLSv1 as the protocol version.

3.2 新版功能.

ssl.OP_NO_TLSv1_1

Prevents a TLSv1.1 connection. This option is only applicable in conjunction with *PROTOCOL_TLS*. It prevents the peers from choosing TLSv1.1 as the protocol version. Available only with openssl version 1.0.1+.

3.4 新版功能.

ssl.OP_NO_TLSv1_2

Prevents a TLSv1.2 connection. This option is only applicable in conjunction with *PROTOCOL_TLS*. It prevents the peers from choosing TLSv1.2 as the protocol version. Available only with openssl version 1.0.1+.

3.4 新版功能.

ssl.OP_NO_TLSv1_3

Prevents a TLSv1.3 connection. This option is only applicable in conjunction with *PROTOCOL_TLS*. It prevents the peers from choosing TLSv1.3 as the protocol version. TLS 1.3 is available with OpenSSL 1.1.1 or later. When Python has been compiled against an older version of OpenSSL, the flag defaults to 0.

3.6.3 新版功能.

`ssl.OP_CIPHER_SERVER_PREFERENCE`

Use the server's cipher ordering preference, rather than the client's. This option has no effect on client sockets and SSLv2 server sockets.

3.3 新版功能.

`ssl.OP_SINGLE_DH_USE`

Prevents re-use of the same DH key for distinct SSL sessions. This improves forward secrecy but requires more computational resources. This option only applies to server sockets.

3.3 新版功能.

`ssl.OP_SINGLE_ECDH_USE`

Prevents re-use of the same ECDH key for distinct SSL sessions. This improves forward secrecy but requires more computational resources. This option only applies to server sockets.

3.3 新版功能.

`ssl.OP_ENABLE_MIDDLEBOX_COMPAT`

Send dummy Change Cipher Spec (CCS) messages in TLS 1.3 handshake to make a TLS 1.3 connection look more like a TLS 1.2 connection.

This option is only available with OpenSSL 1.1.1 and later.

3.6.7 新版功能.

`ssl.OP_NO_COMPRESSION`

Disable compression on the SSL channel. This is useful if the application protocol supports its own compression scheme.

This option is only available with OpenSSL 1.0.0 and later.

3.3 新版功能.

`class ssl.Options`

`enum.IntFlag` collection of `OP_*` constants.

`ssl.OP_NO_TICKET`

Prevent client side from requesting a session ticket.

3.6 新版功能.

`ssl.HAS_ALPN`

Whether the OpenSSL library has built-in support for the *Application-Layer Protocol Negotiation* TLS extension as described in [RFC 7301](#).

3.5 新版功能.

`ssl.HAS_ECDH`

Whether the OpenSSL library has built-in support for Elliptic Curve-based Diffie-Hellman key exchange. This should be true unless the feature was explicitly disabled by the distributor.

3.3 新版功能.

`ssl.HAS_SNI`

Whether the OpenSSL library has built-in support for the *Server Name Indication* extension (as defined in [RFC 6066](#)).

3.2 新版功能.

`ssl.HAS_NPN`

Whether the OpenSSL library has built-in support for *Next Protocol Negotiation* as described in the [NPN draft specification](#). When true, you can use the `SSLContext.set_npn_protocols()` method to advertise which protocols you want to support.

3.3 新版功能.

`ssl.HAS_TLSv1_3`

Whether the OpenSSL library has built-in support for the TLS 1.3 protocol.

3.6.3 新版功能.

`ssl.CHANNEL_BINDING_TYPES`

List of supported TLS channel binding types. Strings in this list can be used as arguments to `SSLSocket.get_channel_binding()`.

3.3 新版功能.

`ssl.OPENSSSL_VERSION`

The version string of the OpenSSL library loaded by the interpreter:

```
>>> ssl.OPENSSSL_VERSION
'OpenSSL 1.0.2k  26 Jan 2017'
```

3.2 新版功能.

`ssl.OPENSSSL_VERSION_INFO`

A tuple of five integers representing version information about the OpenSSL library:

```
>>> ssl.OPENSSSL_VERSION_INFO
(1, 0, 2, 11, 15)
```

3.2 新版功能.

`ssl.OPENSSSL_VERSION_NUMBER`

The raw version number of the OpenSSL library, as a single integer:

```
>>> ssl.OPENSSSL_VERSION_NUMBER
268443839
>>> hex(ssl.OPENSSSL_VERSION_NUMBER)
'0x100020bf'
```

3.2 新版功能.

`ssl.ALERT_DESCRIPTION_HANDSHAKE_FAILURE`

`ssl.ALERT_DESCRIPTION_INTERNAL_ERROR`

`ALERT_DESCRIPTION_*`

Alert Descriptions from [RFC 5246](#) and others. The [IANA TLS Alert Registry](#) contains this list and references to the RFCs where their meaning is defined.

Used as the return value of the callback function in `SSLContext.set_servername_callback()`.

3.4 新版功能.

class `ssl.AlertDescription`

`enum.IntEnum` collection of `ALERT_DESCRIPTION_*` constants.

3.6 新版功能.

`Purpose.SERVER_AUTH`

Option for `create_default_context()` and `SSLContext.load_default_certs()`. This value indicates that the context may be used to authenticate Web servers (therefore, it will be used to create client-side sockets).

3.4 新版功能.

Purpose. **CLIENT_AUTH**

Option for `create_default_context()` and `SSLContext.load_default_certs()`. This value indicates that the context may be used to authenticate Web clients (therefore, it will be used to create server-side sockets).

3.4 新版功能.

class `ssl.SSLErrorNumber`

`enum.IntEnum` collection of `SSL_ERROR_*` constants.

3.6 新版功能.

18.2.2 SSL Sockets

class `ssl.SSLSocket` (`socket.socket`)

SSL sockets provide the following methods of *Socket Objects*:

- `accept()`
- `bind()`
- `close()`
- `connect()`
- `detach()`
- `fileno()`
- `getpeername()`, `getsockname()`
- `getsockopt()`, `setsockopt()`
- `gettimeout()`, `settimeout()`, `setblocking()`
- `listen()`
- `makefile()`
- `recv()`, `recv_into()` (but passing a non-zero `flags` argument is not allowed)
- `send()`, `sendall()` (with the same limitation)
- `sendfile()` (but `os.sendfile` will be used for plain-text sockets only, else `send()` will be used)
- `shutdown()`

However, since the SSL (and TLS) protocol has its own framing atop of TCP, the SSL sockets abstraction can, in certain respects, diverge from the specification of normal, OS-level sockets. See especially the *notes on non-blocking sockets*.

Usually, `SSLSocket` are not created directly, but using the `SSLContext.wrap_socket()` method.

在 3.5 版更改: The `sendfile()` method was added.

在 3.5 版更改: The `shutdown()` does not reset the socket timeout each time bytes are received or sent. The socket timeout is now to maximum total duration of the shutdown.

3.6 版后已移除: It is deprecated to create a `SSLSocket` instance directly, use `SSLContext.wrap_socket()` to wrap a socket.

SSL sockets also have the following additional methods and attributes:

`SSLSocket.read(len=1024, buffer=None)`

Read up to *len* bytes of data from the SSL socket and return the result as a `bytes` instance. If *buffer* is specified, then read into the buffer instead, and return the number of bytes read.

Raise `SSLWantReadError` or `SSLWantWriteError` if the socket is *non-blocking* and the read would block.

As at any time a re-negotiation is possible, a call to `read()` can also cause write operations.

在 3.5 版更改: The socket timeout is no more reset each time bytes are received or sent. The socket timeout is now to maximum total duration to read up to *len* bytes.

3.6 版后已移除: Use `recv()` instead of `read()`.

`SSLSocket.write(buf)`

Write *buf* to the SSL socket and return the number of bytes written. The *buf* argument must be an object supporting the buffer interface.

Raise `SSLWantReadError` or `SSLWantWriteError` if the socket is *non-blocking* and the write would block.

As at any time a re-negotiation is possible, a call to `write()` can also cause read operations.

在 3.5 版更改: The socket timeout is no more reset each time bytes are received or sent. The socket timeout is now to maximum total duration to write *buf*.

3.6 版后已移除: Use `send()` instead of `write()`.

注解: The `read()` and `write()` methods are the low-level methods that read and write unencrypted, application-level data and decrypt/encrypt it to encrypted, wire-level data. These methods require an active SSL connection, i.e. the handshake was completed and `SSLSocket.unwrap()` was not called.

Normally you should use the socket API methods like `recv()` and `send()` instead of these methods.

`SSLSocket.do_handshake()`

Perform the SSL setup handshake.

在 3.4 版更改: The handshake method also performs `match_hostname()` when the `check_hostname` attribute of the socket's `context` is true.

在 3.5 版更改: The socket timeout is no more reset each time bytes are received or sent. The socket timeout is now to maximum total duration of the handshake.

`SSLSocket.getpeercert(binary_form=False)`

If there is no certificate for the peer on the other end of the connection, return `None`. If the SSL handshake hasn't been done yet, raise `ValueError`.

If the `binary_form` parameter is `False`, and a certificate was received from the peer, this method returns a `dict` instance. If the certificate was not validated, the dict is empty. If the certificate was validated, it returns a dict with several keys, amongst them `subject` (the principal for which the certificate was issued) and `issuer` (the principal issuing the certificate). If a certificate contains an instance of the *Subject Alternative Name* extension (see [RFC 3280](#)), there will also be a `subjectAltName` key in the dictionary.

The `subject` and `issuer` fields are tuples containing the sequence of relative distinguished names (RDNs) given in the certificate's data structure for the respective fields, and each RDN is a sequence of name-value pairs. Here is a real-world example:

```
{'issuer': ((('countryName', 'IL'),),
            (('organizationName', 'StartCom Ltd.'),),
            (('organizationalUnitName',
              'Secure Digital Certificate Signing'),),
            (('commonName',
```

(下页继续)

(续上页)

```

        'StartCom Class 2 Primary Intermediate Server CA'),)),
'notAfter': 'Nov 22 08:15:19 2013 GMT',
'notBefore': 'Nov 21 03:09:52 2011 GMT',
'serialNumber': '95F0',
'subject': (((('description', '571208-SLe257oHY9fVQ07Z'),),
              (('countryName', 'US'),),
              (('stateOrProvinceName', 'California'),),
              (('localityName', 'San Francisco'),),
              (('organizationName', 'Electronic Frontier Foundation, Inc.'),),
              (('commonName', '*.eff.org'),),
              (('emailAddress', 'hostmaster@eff.org'),)),
'subjectAltName': (('DNS', '*.eff.org'), ('DNS', 'eff.org')),
'version': 3}

```

注解: To validate a certificate for a particular service, you can use the `match_hostname()` function.

If the `binary_form` parameter is `True`, and a certificate was provided, this method returns the DER-encoded form of the entire certificate as a sequence of bytes, or `None` if the peer did not provide a certificate. Whether the peer provides a certificate depends on the SSL socket's role:

- for a client SSL socket, the server will always provide a certificate, regardless of whether validation was required;
- for a server SSL socket, the client will only provide a certificate when requested by the server; therefore `getpeercert()` will return `None` if you used `CERT_NONE` (rather than `CERT_OPTIONAL` or `CERT_REQUIRED`).

在 3.2 版更改: The returned dictionary includes additional items such as `issuer` and `notBefore`.

在 3.4 版更改: `ValueError` is raised when the handshake isn't done. The returned dictionary includes additional X509v3 extension items such as `crlDistributionPoints`, `caIssuers` and `OCSP URIs`.

`SSLSocket.cipher()`

Returns a three-value tuple containing the name of the cipher being used, the version of the SSL protocol that defines its use, and the number of secret bits being used. If no connection has been established, returns `None`.

`SSLSocket.shared_ciphers()`

Return the list of ciphers shared by the client during the handshake. Each entry of the returned list is a three-value tuple containing the name of the cipher, the version of the SSL protocol that defines its use, and the number of secret bits the cipher uses. `shared_ciphers()` returns `None` if no connection has been established or the socket is a client socket.

3.5 新版功能.

`SSLSocket.compression()`

Return the compression algorithm being used as a string, or `None` if the connection isn't compressed.

If the higher-level protocol supports its own compression mechanism, you can use `OP_NO_COMPRESSION` to disable SSL-level compression.

3.3 新版功能.

`SSLSocket.get_channel_binding(cb_type="tls-unique")`

Get channel binding data for current connection, as a bytes object. Returns `None` if not connected or the handshake has not been completed.

The `cb_type` parameter allow selection of the desired channel binding type. Valid channel binding types are listed in the `CHANNEL_BINDING_TYPES` list. Currently only the 'tls-unique' channel binding, defined by [RFC 5929](#),

is supported. *ValueError* will be raised if an unsupported channel binding type is requested.

3.3 新版功能.

`SSLSocket.selected_alpn_protocol()`

Return the protocol that was selected during the TLS handshake. If `SSLContext.set_alpn_protocols()` was not called, if the other party does not support ALPN, if this socket does not support any of the client's proposed protocols, or if the handshake has not happened yet, `None` is returned.

3.5 新版功能.

`SSLSocket.selected_npn_protocol()`

Return the higher-level protocol that was selected during the TLS/SSL handshake. If `SSLContext.set_npn_protocols()` was not called, or if the other party does not support NPN, or if the handshake has not yet happened, this will return `None`.

3.3 新版功能.

`SSLSocket.unwrap()`

Performs the SSL shutdown handshake, which removes the TLS layer from the underlying socket, and returns the underlying socket object. This can be used to go from encrypted operation over a connection to unencrypted. The returned socket should always be used for further communication with the other side of the connection, rather than the original socket.

`SSLSocket.verify_client_post_handshake()`

Requests post-handshake authentication (PHA) from a TLS 1.3 client. PHA can only be initiated for a TLS 1.3 connection from a server-side socket, after the initial TLS handshake and with PHA enabled on both sides, see `SSLContext.post_handshake_auth`.

The method does not perform a cert exchange immediately. The server-side sends a `CertificateRequest` during the next write event and expects the client to respond with a certificate on the next read event.

If any precondition isn't met (e.g. not TLS 1.3, PHA not enabled), an *SSLError* is raised.

3.6.7 新版功能.

注解: Only available with OpenSSL 1.1.1 and TLS 1.3 enabled. Without TLS 1.3 support, the method raises *NotImplementedError*.

`SSLSocket.version()`

Return the actual SSL protocol version negotiated by the connection as a string, or `None` if no secure connection is established. As of this writing, possible return values include `"SSLv2"`, `"SSLv3"`, `"TLSv1"`, `"TLSv1.1"` and `"TLSv1.2"`. Recent OpenSSL versions may define more return values.

3.5 新版功能.

`SSLSocket.pending()`

Returns the number of already decrypted bytes available for read, pending on the connection.

`SSLSocket.context`

The *SSLContext* object this SSL socket is tied to. If the SSL socket was created using the top-level `wrap_socket()` function (rather than `SSLContext.wrap_socket()`), this is a custom context object created for this SSL socket.

3.2 新版功能.

`SSLSocket.server_side`

A boolean which is `True` for server-side sockets and `False` for client-side sockets.

3.2 新版功能.

SSLSocket.server_hostname

Hostname of the server: *str* type, or None for server-side socket or if the hostname was not specified in the constructor.

3.2 新版功能.

SSLSocket.session

The *SSLSession* for this SSL connection. The session is available for client and server side sockets after the TLS handshake has been performed. For client sockets the session can be set before *do_handshake()* has been called to reuse a session.

3.6 新版功能.

SSLSocket.session_reused

3.6 新版功能.

18.2.3 SSL Contexts

3.2 新版功能.

An SSL context holds various data longer-lived than single SSL connections, such as SSL configuration options, certificate(s) and private key(s). It also manages a cache of SSL sessions for server-side sockets, in order to speed up repeated connections from the same clients.

class ssl.SSLContext (protocol=PROTOCOL_TLS)

Create a new SSL context. You may pass *protocol* which must be one of the `PROTOCOL_*` constants defined in this module. *PROTOCOL_TLS* is currently recommended for maximum interoperability and default value.

参见:

create_default_context() lets the *ssl* module choose security settings for a given purpose.

在 3.6 版更改: The context is created with secure default values. The options *OP_NO_COMPRESSION*, *OP_CIPHER_SERVER_PREFERENCE*, *OP_SINGLE_DH_USE*, *OP_SINGLE_ECDH_USE*, *OP_NO_SSLv2* (except for *PROTOCOL_SSLv2*), and *OP_NO_SSLv3* (except for *PROTOCOL_SSLv3*) are set by default. The initial cipher suite list contains only HIGH ciphers, no NULL ciphers and no MD5 ciphers (except for *PROTOCOL_SSLv2*).

SSLContext objects have the following methods and attributes:

SSLContext.cert_store_stats()

Get statistics about quantities of loaded X.509 certificates, count of X.509 certificates flagged as CA certificates and certificate revocation lists as dictionary.

Example for a context with one CA cert and one other cert:

```
>>> context.cert_store_stats()
{'crl': 0, 'x509_ca': 1, 'x509': 2}
```

3.4 新版功能.

SSLContext.load_cert_chain (certfile, keyfile=None, password=None)

Load a private key and the corresponding certificate. The *certfile* string must be the path to a single file in PEM format containing the certificate as well as any number of CA certificates needed to establish the certificate's authenticity. The *keyfile* string, if present, must point to a file containing the private key in. Otherwise the private key will be taken from *certfile* as well. See the discussion of *Certificates* for more information on how the certificate is stored in the *certfile*.

The *password* argument may be a function to call to get the password for decrypting the private key. It will only be called if the private key is encrypted and a password is necessary. It will be called with no arguments, and it

should return a string, bytes, or bytearray. If the return value is a string it will be encoded as UTF-8 before using it to decrypt the key. Alternatively a string, bytes, or bytearray value may be supplied directly as the *password* argument. It will be ignored if the private key is not encrypted and no password is needed.

If the *password* argument is not specified and a password is required, OpenSSL's built-in password prompting mechanism will be used to interactively prompt the user for a password.

An *SSLError* is raised if the private key doesn't match with the certificate.

在 3.3 版更改: New optional argument *password*.

`SSLContext.load_default_certs(purpose=Purpose.SERVER_AUTH)`

Load a set of default “certification authority” (CA) certificates from default locations. On Windows it loads CA certs from the CA and ROOT system stores. On other systems it calls `SSLContext.set_default_verify_paths()`. In the future the method may load CA certificates from other locations, too.

The *purpose* flag specifies what kind of CA certificates are loaded. The default settings `Purpose.SERVER_AUTH` loads certificates, that are flagged and trusted for TLS web server authentication (client side sockets). `Purpose.CLIENT_AUTH` loads CA certificates for client certificate verification on the server side.

3.4 新版功能.

`SSLContext.load_verify_locations(cafile=None, capath=None, cadata=None)`

Load a set of “certification authority” (CA) certificates used to validate other peers' certificates when *verify_mode* is other than `CERT_NONE`. At least one of *cafile* or *capath* must be specified.

This method can also load certification revocation lists (CRLs) in PEM or DER format. In order to make use of CRLs, `SSLContext.verify_flags` must be configured properly.

The *cafile* string, if present, is the path to a file of concatenated CA certificates in PEM format. See the discussion of *Certificates* for more information about how to arrange the certificates in this file.

The *capath* string, if present, is the path to a directory containing several CA certificates in PEM format, following an *OpenSSL specific layout*.

The *cadata* object, if present, is either an ASCII string of one or more PEM-encoded certificates or a *bytes-like object* of DER-encoded certificates. Like with *capath* extra lines around PEM-encoded certificates are ignored but at least one certificate must be present.

在 3.4 版更改: New optional argument *cadata*

`SSLContext.get_ca_certs(binary_form=False)`

Get a list of loaded “certification authority” (CA) certificates. If the *binary_form* parameter is *False* each list entry is a dict like the output of `SSLSocket.getpeercert()`. Otherwise the method returns a list of DER-encoded certificates. The returned list does not contain certificates from *capath* unless a certificate was requested and loaded by a SSL connection.

注解: Certificates in a *capath* directory aren't loaded unless they have been used at least once.

3.4 新版功能.

`SSLContext.get_ciphers()`

Get a list of enabled ciphers. The list is in order of cipher priority. See `SSLContext.set_ciphers()`.

示例:

```
>>> ctx = ssl.SSLContext(ssl.PROTOCOL_SSLv23)
>>> ctx.set_ciphers('ECDHE+AESGCM:!ECDSA')
>>> ctx.get_ciphers() # OpenSSL 1.0.x
```

(下页继续)

(续上页)

```
[{'alg_bits': 256,
  'description': 'ECDHE-RSA-AES256-GCM-SHA384 TLSv1.2 Kx=ECDH      Au=RSA  '
                  'Enc=AESGCM(256) Mac=AEAD',
  'id': 50380848,
  'name': 'ECDHE-RSA-AES256-GCM-SHA384',
  'protocol': 'TLSv1/SSLv3',
  'strength_bits': 256},
 {'alg_bits': 128,
  'description': 'ECDHE-RSA-AES128-GCM-SHA256 TLSv1.2 Kx=ECDH      Au=RSA  '
                  'Enc=AESGCM(128) Mac=AEAD',
  'id': 50380847,
  'name': 'ECDHE-RSA-AES128-GCM-SHA256',
  'protocol': 'TLSv1/SSLv3',
  'strength_bits': 128}]
```

On OpenSSL 1.1 and newer the cipher dict contains additional fields::

```
>>> ctx.get_ciphers() # OpenSSL 1.1+
[{'aead': True,
  'alg_bits': 256,
  'auth': 'auth-rsa',
  'description': 'ECDHE-RSA-AES256-GCM-SHA384 TLSv1.2 Kx=ECDH      Au=RSA  '
                  'Enc=AESGCM(256) Mac=AEAD',
  'digest': None,
  'id': 50380848,
  'kea': 'kx-ecdh',
  'name': 'ECDHE-RSA-AES256-GCM-SHA384',
  'protocol': 'TLSv1.2',
  'strength_bits': 256,
  'symmetric': 'aes-256-gcm'},
 {'aead': True,
  'alg_bits': 128,
  'auth': 'auth-rsa',
  'description': 'ECDHE-RSA-AES128-GCM-SHA256 TLSv1.2 Kx=ECDH      Au=RSA  '
                  'Enc=AESGCM(128) Mac=AEAD',
  'digest': None,
  'id': 50380847,
  'kea': 'kx-ecdh',
  'name': 'ECDHE-RSA-AES128-GCM-SHA256',
  'protocol': 'TLSv1.2',
  'strength_bits': 128,
  'symmetric': 'aes-128-gcm'}]
```

Availability: OpenSSL 1.0.2+

3.6 新版功能.

`SSLContext.set_default_verify_paths()`

Load a set of default “certification authority” (CA) certificates from a filesystem path defined when building the OpenSSL library. Unfortunately, there’s no easy way to know whether this method succeeds: no error is returned if no certificates are to be found. When the OpenSSL library is provided as part of the operating system, though, it is likely to be configured properly.

`SSLContext.set_ciphers(ciphers)`

Set the available ciphers for sockets created with this context. It should be a string in the [OpenSSL cipher list format](#). If no cipher can be selected (because compile-time options or other configuration forbids use of all the specified ciphers), an `SSL_ERROR` will be raised.

注解: when connected, the `SSLSocket.cipher()` method of SSL sockets will give the currently selected cipher.

OpenSSL 1.1.1 has TLS 1.3 cipher suites enabled by default. The suites cannot be disabled with `set_ciphers()`.

`SSLContext.set_alpn_protocols(protocols)`

Specify which protocols the socket should advertise during the SSL/TLS handshake. It should be a list of ASCII strings, like `['http/1.1', 'spdy/2']`, ordered by preference. The selection of a protocol will happen during the handshake, and will play out according to [RFC 7301](#). After a successful handshake, the `SSLSocket.selected_alpn_protocol()` method will return the agreed-upon protocol.

This method will raise `NotImplementedError` if `HAS_ALPN` is False.

OpenSSL 1.1.0 to 1.1.0e will abort the handshake and raise `SSL_ERROR` when both sides support ALPN but cannot agree on a protocol. 1.1.0f+ behaves like 1.0.2, `SSLSocket.selected_alpn_protocol()` returns None.

3.5 新版功能.

`SSLContext.set_npn_protocols(protocols)`

Specify which protocols the socket should advertise during the SSL/TLS handshake. It should be a list of strings, like `['http/1.1', 'spdy/2']`, ordered by preference. The selection of a protocol will happen during the handshake, and will play out according to the [NPN draft specification](#). After a successful handshake, the `SSLSocket.selected_npn_protocol()` method will return the agreed-upon protocol.

This method will raise `NotImplementedError` if `HAS_NPN` is False.

3.3 新版功能.

`SSLContext.set_servername_callback(server_name_callback)`

Register a callback function that will be called after the TLS Client Hello handshake message has been received by the SSL/TLS server when the TLS client specifies a server name indication. The server name indication mechanism is specified in [RFC 6066](#) section 3 - Server Name Indication.

Only one callback can be set per `SSLContext`. If `server_name_callback` is None then the callback is disabled. Calling this function a subsequent time will disable the previously registered callback.

The callback function, `server_name_callback`, will be called with three arguments; the first being the `ssl.SSLSocket`, the second is a string that represents the server name that the client is intending to communicate (or None if the TLS Client Hello does not contain a server name) and the third argument is the original `SSLContext`. The server name argument is the IDNA decoded server name.

A typical use of this callback is to change the `ssl.SSLSocket`'s `SSLSocket.context` attribute to a new object of type `SSLContext` representing a certificate chain that matches the server name.

Due to the early negotiation phase of the TLS connection, only limited methods and attributes are usable like `SSLSocket.selected_alpn_protocol()` and `SSLSocket.context`. `SSLSocket.getpeercert()`, `SSLSocket.getpeername()`, `SSLSocket.cipher()` and `SSLSocket.compress()` methods require that the TLS connection has progressed beyond the TLS Client Hello and therefore will not contain return meaningful values nor can they be called safely.

The `server_name_callback` function must return None to allow the TLS negotiation to continue. If a TLS failure is required, a constant `ALERT_DESCRIPTION_*` can be returned. Other return values will result in a TLS fatal error with `ALERT_DESCRIPTION_INTERNAL_ERROR`.

If there is an IDNA decoding error on the server name, the TLS connection will terminate with an `ALERT_DESCRIPTION_INTERNAL_ERROR` fatal TLS alert message to the client.

If an exception is raised from the `server_name_callback` function the TLS connection will terminate with a fatal TLS alert message `ALERT_DESCRIPTION_HANDSHAKE_FAILURE`.

This method will raise `NotImplementedError` if the OpenSSL library had `OPENSSL_NO_TLSEXT` defined when it was built.

3.4 新版功能.

`SSLContext.load_dh_params(dhfile)`

Load the key generation parameters for Diffie-Hellman (DH) key exchange. Using DH key exchange improves forward secrecy at the expense of computational resources (both on the server and on the client). The *dhfile* parameter should be the path to a file containing DH parameters in PEM format.

This setting doesn't apply to client sockets. You can also use the `OP_SINGLE_DH_USE` option to further improve security.

3.3 新版功能.

`SSLContext.set_ecdh_curve(curve_name)`

Set the curve name for Elliptic Curve-based Diffie-Hellman (ECDH) key exchange. ECDH is significantly faster than regular DH while arguably as secure. The *curve_name* parameter should be a string describing a well-known elliptic curve, for example `prime256v1` for a widely supported curve.

This setting doesn't apply to client sockets. You can also use the `OP_SINGLE_ECDH_USE` option to further improve security.

This method is not available if `HAS_ECDH` is `False`.

3.3 新版功能.

参见:

SSL/TLS & Perfect Forward Secrecy Vincent Bernat.

`SSLContext.wrap_socket(sock, server_side=False, do_handshake_on_connect=True, suppress_ragged_eofs=True, server_hostname=None, session=None)`

Wrap an existing Python socket *sock* and return an `SSLSocket` object. *sock* must be a `SOCK_STREAM` socket; other socket types are unsupported.

The returned SSL socket is tied to the context, its settings and certificates. The parameters *server_side*, *do_handshake_on_connect* and *suppress_ragged_eofs* have the same meaning as in the top-level `wrap_socket()` function.

On client connections, the optional parameter *server_hostname* specifies the hostname of the service which we are connecting to. This allows a single server to host multiple SSL-based services with distinct certificates, quite similarly to HTTP virtual hosts. Specifying *server_hostname* will raise a `ValueError` if *server_side* is true.

session, see `session`.

在 3.5 版更改: Always allow a *server_hostname* to be passed, even if OpenSSL does not have SNI.

在 3.6 版更改: *session* argument was added.

`SSLContext.wrap_bio(incoming, outgoing, server_side=False, server_hostname=None, session=None)`

Create a new `SSLObject` instance by wrapping the BIO objects *incoming* and *outgoing*. The SSL routines will read input data from the incoming BIO and write data to the outgoing BIO.

The *server_side*, *server_hostname* and *session* parameters have the same meaning as in `SSLContext.wrap_socket()`.

在 3.6 版更改: *session* argument was added.

`SSLContext.session_stats()`

Get statistics about the SSL sessions created or managed by this context. A dictionary is returned which maps the names of each `piece of information` to their numeric values. For example, here is the total number of hits and misses in the session cache since the context was created:

```
>>> stats = context.session_stats()
>>> stats['hits'], stats['misses']
(0, 0)
```

SSLContext.check_hostname

Whether to match the peer cert's hostname with `match_hostname()` in `SSLSocket.do_handshake()`. The context's `verify_mode` must be set to `CERT_OPTIONAL` or `CERT_REQUIRED`, and you must pass `server_hostname` to `wrap_socket()` in order to match the hostname.

示例:

```
import socket, ssl

context = ssl.SSLContext()
context.verify_mode = ssl.CERT_REQUIRED
context.check_hostname = True
context.load_default_certs()

s = socket.socket(socket.AF_INET, socket.SOCK_STREAM)
ssl_sock = context.wrap_socket(s, server_hostname='www.verisign.com')
ssl_sock.connect(('www.verisign.com', 443))
```

3.4 新版功能.

注解: This features requires OpenSSL 0.9.8f or newer.

SSLContext.options

An integer representing the set of SSL options enabled on this context. The default value is `OP_ALL`, but you can specify other options such as `OP_NO_SSLv2` by ORing them together.

注解: With versions of OpenSSL older than 0.9.8m, it is only possible to set options, not to clear them. Attempting to clear an option (by resetting the corresponding bits) will raise a `ValueError`.

在 3.6 版更改: `SSLContext.options` returns `Options` flags:

```
>>> ssl.create_default_context().options
<Options.OP_ALL|OP_NO_SSLv3|OP_NO_SSLv2|OP_NO_COMPRESSION: 2197947391>
```

SSLContext.post_handshake_auth

Enable TLS 1.3 post-handshake client authentication. Post-handshake auth is disabled by default and a server can only request a TLS client certificate during the initial handshake. When enabled, a server may request a TLS client certificate at any time after the handshake.

When enabled on client-side sockets, the client signals the server that it supports post-handshake authentication.

When enabled on server-side sockets, `SSLContext.verify_mode` must be set to `CERT_OPTIONAL` or `CERT_REQUIRED`, too. The actual client cert exchange is delayed until `SSLSocket.verify_client_post_handshake()` is called and some I/O is performed.

3.6.7 新版功能.

注解: Only available with OpenSSL 1.1.1 and TLS 1.3 enabled. Without TLS 1.3 support, the property value is `None` and can't be modified

SSLContext.protocol

The protocol version chosen when constructing the context. This attribute is read-only.

SSLContext.verify_flags

The flags for certificate verification operations. You can set flags like `VERIFY_CRL_CHECK_LEAF` by ORing them together. By default OpenSSL does neither require nor verify certificate revocation lists (CRLs). Available only with openssl version 0.9.8+.

3.4 新版功能.

在 3.6 版更改: `SSLContext.verify_flags` returns `VerifyFlags` flags:

```
>>> ssl.create_default_context().verify_flags
<VerifyFlags.VERIFY_X509_TRUSTED_FIRST: 32768>
```

SSLContext.verify_mode

Whether to try to verify other peers' certificates and how to behave if verification fails. This attribute must be one of `CERT_NONE`, `CERT_OPTIONAL` or `CERT_REQUIRED`.

在 3.6 版更改: `SSLContext.verify_mode` returns `VerifyMode` enum:

```
>>> ssl.create_default_context().verify_mode
<VerifyMode.CERT_REQUIRED: 2>
```

18.2.4 Certificates

Certificates in general are part of a public-key / private-key system. In this system, each *principal*, (which may be a machine, or a person, or an organization) is assigned a unique two-part encryption key. One part of the key is public, and is called the *public key*; the other part is kept secret, and is called the *private key*. The two parts are related, in that if you encrypt a message with one of the parts, you can decrypt it with the other part, and **only** with the other part.

A certificate contains information about two principals. It contains the name of a *subject*, and the subject's public key. It also contains a statement by a second principal, the *issuer*, that the subject is who they claim to be, and that this is indeed the subject's public key. The issuer's statement is signed with the issuer's private key, which only the issuer knows. However, anyone can verify the issuer's statement by finding the issuer's public key, decrypting the statement with it, and comparing it to the other information in the certificate. The certificate also contains information about the time period over which it is valid. This is expressed as two fields, called "notBefore" and "notAfter".

In the Python use of certificates, a client or server can use a certificate to prove who they are. The other side of a network connection can also be required to produce a certificate, and that certificate can be validated to the satisfaction of the client or server that requires such validation. The connection attempt can be set to raise an exception if the validation fails. Validation is done automatically, by the underlying OpenSSL framework; the application need not concern itself with its mechanics. But the application does usually need to provide sets of certificates to allow this process to take place.

Python uses files to contain certificates. They should be formatted as "PEM" (see [RFC 1422](#)), which is a base-64 encoded form wrapped with a header line and a footer line:

```
-----BEGIN CERTIFICATE-----
... (certificate in base64 PEM encoding) ...
-----END CERTIFICATE-----
```

Certificate chains

The Python files which contain certificates can contain a sequence of certificates, sometimes called a *certificate chain*. This chain should start with the specific certificate for the principal who "is" the client or server, and then the certificate for the issuer of that certificate, and then the certificate for the issuer of *that* certificate, and so on up the chain till you

get to a certificate which is *self-signed*, that is, a certificate which has the same subject and issuer, sometimes called a *root certificate*. The certificates should just be concatenated together in the certificate file. For example, suppose we had a three certificate chain, from our server certificate to the certificate of the certification authority that signed our server certificate, to the root certificate of the agency which issued the certification authority's certificate:

```
-----BEGIN CERTIFICATE-----
... (certificate for your server)...
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
... (the certificate for the CA)...
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
... (the root certificate for the CA's issuer)...
-----END CERTIFICATE-----
```

CA certificates

If you are going to require validation of the other side of the connection's certificate, you need to provide a “CA certs” file, filled with the certificate chains for each issuer you are willing to trust. Again, this file just contains these chains concatenated together. For validation, Python will use the first chain it finds in the file which matches. The platform's certificates file can be used by calling `SSLContext.load_default_certs()`, this is done automatically with `create_default_context()`.

Combined key and certificate

Often the private key is stored in the same file as the certificate; in this case, only the `certfile` parameter to `SSLContext.load_cert_chain()` and `wrap_socket()` needs to be passed. If the private key is stored with the certificate, it should come before the first certificate in the certificate chain:

```
-----BEGIN RSA PRIVATE KEY-----
... (private key in base64 encoding) ...
-----END RSA PRIVATE KEY-----
-----BEGIN CERTIFICATE-----
... (certificate in base64 PEM encoding) ...
-----END CERTIFICATE-----
```

Self-signed certificates

If you are going to create a server that provides SSL-encrypted connection services, you will need to acquire a certificate for that service. There are many ways of acquiring appropriate certificates, such as buying one from a certification authority. Another common practice is to generate a self-signed certificate. The simplest way to do this is with the OpenSSL package, using something like the following:

```
% openssl req -new -x509 -days 365 -nodes -out cert.pem -keyout cert.pem
Generating a 1024 bit RSA private key
.....++++++
.....++++++
writing new private key to 'cert.pem'
-----
You are about to be asked to enter information that will be incorporated
into your certificate request.
What you are about to enter is what is called a Distinguished Name or a DN.
```

(下页继续)

(续上页)

```

There are quite a few fields but you can leave some blank
For some fields there will be a default value,
If you enter '.', the field will be left blank.
-----
Country Name (2 letter code) [AU]:US
State or Province Name (full name) [Some-State]:MyState
Locality Name (eg, city) []:Some City
Organization Name (eg, company) [Internet Widgits Pty Ltd]:My Organization, Inc.
Organizational Unit Name (eg, section) []:My Group
Common Name (eg, YOUR name) []:myserver.mygroup.myorganization.com
Email Address []:ops@myserver.mygroup.myorganization.com
%
```

The disadvantage of a self-signed certificate is that it is its own root certificate, and no one else will have it in their cache of known (and trusted) root certificates.

18.2.5 示例

Testing for SSL support

To test for the presence of SSL support in a Python installation, user code should use the following idiom:

```

try:
    import ssl
except ImportError:
    pass
else:
    ... # do something that requires SSL support
```

Client-side operation

This example creates a SSL context with the recommended security settings for client sockets, including automatic certificate verification:

```
>>> context = ssl.create_default_context()
```

If you prefer to tune security settings yourself, you might create a context from scratch (but beware that you might not get the settings right):

```

>>> context = ssl.SSLContext()
>>> context.verify_mode = ssl.CERT_REQUIRED
>>> context.check_hostname = True
>>> context.load_verify_locations("/etc/ssl/certs/ca-bundle.crt")
```

(this snippet assumes your operating system places a bundle of all CA certificates in /etc/ssl/certs/ca-bundle.crt; if not, you'll get an error and have to adjust the location)

When you use the context to connect to a server, `CERT_REQUIRED` validates the server certificate: it ensures that the server certificate was signed with one of the CA certificates, and checks the signature for correctness:

```

>>> conn = context.wrap_socket(socket.socket(socket.AF_INET),
...                             server_hostname="www.python.org")
>>> conn.connect(("www.python.org", 443))
```

You may then fetch the certificate:

```
>>> cert = conn.getpeercert()
```

Visual inspection shows that the certificate does identify the desired service (that is, the HTTPS host `www.python.org`):

```
>>> pprint.pprint(cert)
{'OCSP': ('http://ocsp.digicert.com',),
 'caIssuers': ('http://cacerts.digicert.com/DigiCertSHA2ExtendedValidationServerCA.crt',),
 'crlDistributionPoints': ('http://crl3.digicert.com/sha2-ev-server-g1.crl',
                           'http://crl4.digicert.com/sha2-ev-server-g1.crl'),
 'issuer': (((('countryName', 'US'),),
               (('organizationName', 'DigiCert Inc'),),
               (('organizationalUnitName', 'www.digicert.com'),),
               (('commonName', 'DigiCert SHA2 Extended Validation Server CA'),)),),
 'notAfter': 'Sep  9 12:00:00 2016 GMT',
 'notBefore': 'Sep  5 00:00:00 2014 GMT',
 'serialNumber': '01BB6F00122B177F36CAB49CEA8B6B26',
 'subject': (((('businessCategory', 'Private Organization'),),
                (('1.3.6.1.4.1.311.60.2.1.3', 'US'),),
                (('1.3.6.1.4.1.311.60.2.1.2', 'Delaware'),),
                (('serialNumber', '3359300'),),
                (('streetAddress', '16 Allen Rd'),),
                (('postalCode', '03894-4801'),),
                (('countryName', 'US'),),
                (('stateOrProvinceName', 'NH'),),
                (('localityName', 'Wolfeboro,'),),
                (('organizationName', 'Python Software Foundation'),),
                (('commonName', 'www.python.org'),)),),
 'subjectAltName': (('DNS', 'www.python.org'),
                    ('DNS', 'python.org'),
                    ('DNS', 'pypi.org'),
                    ('DNS', 'docs.python.org'),
                    ('DNS', 'testpypi.org'),
                    ('DNS', 'bugs.python.org'),
                    ('DNS', 'wiki.python.org'),
                    ('DNS', 'hg.python.org'),
                    ('DNS', 'mail.python.org'),
                    ('DNS', 'packaging.python.org'),
                    ('DNS', 'pythonhosted.org'),
                    ('DNS', 'www.pythonhosted.org'),
                    ('DNS', 'test.pythonhosted.org'),
                    ('DNS', 'us.pycon.org'),
                    ('DNS', 'id.python.org')),
 'version': 3}
```

Now the SSL channel is established and the certificate verified, you can proceed to talk with the server:

```
>>> conn.sendall(b"HEAD / HTTP/1.0\r\nHost: linuxfr.org\r\n\r\n")
>>> pprint.pprint(conn.recv(1024).split(b"\r\n"))
[b'HTTP/1.1 200 OK',
 b'Date: Sat, 18 Oct 2014 18:27:20 GMT',
 b'Server: nginx',
 b'Content-Type: text/html; charset=utf-8',
 b'X-Frame-Options: SAMEORIGIN',
 b'Content-Length: 45679',
```

(下页继续)

(续上页)

```

b'Accept-Ranges: bytes',
b'Via: 1.1 varnish',
b'Age: 2188',
b'X-Served-By: cache-lcy1134-LCY',
b'X-Cache: HIT',
b'X-Cache-Hits: 11',
b'Vary: Cookie',
b'Strict-Transport-Security: max-age=63072000; includeSubDomains',
b'Connection: close',
b'',
b'']

```

See the discussion of *Security considerations* below.

Server-side operation

For server operation, typically you'll need to have a server certificate, and private key, each in a file. You'll first create a context holding the key and the certificate, so that clients can check your authenticity. Then you'll open a socket, bind it to a port, call `listen()` on it, and start waiting for clients to connect:

```

import socket, ssl

context = ssl.create_default_context(ssl.Purpose.CLIENT_AUTH)
context.load_cert_chain(certfile="mycertfile", keyfile="mykeyfile")

bindsocket = socket.socket()
bindsocket.bind(('myaddr.mydomain.com', 10023))
bindsocket.listen(5)

```

When a client connects, you'll call `accept()` on the socket to get the new socket from the other end, and use the context's `SSLContext.wrap_socket()` method to create a server-side SSL socket for the connection:

```

while True:
    newsocket, fromaddr = bindsocket.accept()
    connstream = context.wrap_socket(newsocket, server_side=True)
    try:
        deal_with_client(connstream)
    finally:
        connstream.shutdown(socket.SHUT_RDWR)
        connstream.close()

```

Then you'll read data from the `connstream` and do something with it till you are finished with the client (or the client is finished with you):

```

def deal_with_client(connstream):
    data = connstream.recv(1024)
    # empty data means the client is finished with us
    while data:
        if not do_something(connstream, data):
            # we'll assume do_something returns False
            # when we're finished with client
            break
        data = connstream.recv(1024)
    # finished with client

```

And go back to listening for new client connections (of course, a real server would probably handle each client connection in a separate thread, or put the sockets in *non-blocking mode* and use an event loop).

18.2.6 Notes on non-blocking sockets

SSL sockets behave slightly different than regular sockets in non-blocking mode. When working with non-blocking sockets, there are thus several things you need to be aware of:

- Most `SSLSocket` methods will raise either `SSLWantWriteError` or `SSLWantReadError` instead of `BlockingIOError` if an I/O operation would block. `SSLWantReadError` will be raised if a read operation on the underlying socket is necessary, and `SSLWantWriteError` for a write operation on the underlying socket. Note that attempts to *write* to an SSL socket may require *reading* from the underlying socket first, and attempts to *read* from the SSL socket may require a prior *write* to the underlying socket.

在 3.5 版更改: In earlier Python versions, the `SSLSocket.send()` method returned zero instead of raising `SSLWantWriteError` or `SSLWantReadError`.

- Calling `select()` tells you that the OS-level socket can be read from (or written to), but it does not imply that there is sufficient data at the upper SSL layer. For example, only part of an SSL frame might have arrived. Therefore, you must be ready to handle `SSLSocket.recv()` and `SSLSocket.send()` failures, and retry after another call to `select()`.
- Conversely, since the SSL layer has its own framing, a SSL socket may still have data available for reading without `select()` being aware of it. Therefore, you should first call `SSLSocket.recv()` to drain any potentially available data, and then only block on a `select()` call if still necessary.

(of course, similar provisions apply when using other primitives such as `poll()`, or those in the `selectors` module)

- The SSL handshake itself will be non-blocking: the `SSLSocket.do_handshake()` method has to be retried until it returns successfully. Here is a synopsis using `select()` to wait for the socket's readiness:

```
while True:
    try:
        sock.do_handshake()
        break
    except ssl.SSLWantReadError:
        select.select([sock], [], [])
    except ssl.SSLWantWriteError:
        select.select([], [sock], [])
```

参见:

The `asyncio` module supports *non-blocking SSL sockets* and provides a higher level API. It polls for events using the `selectors` module and handles `SSLWantWriteError`, `SSLWantReadError` and `BlockingIOError` exceptions. It runs the SSL handshake asynchronously as well.

18.2.7 Memory BIO Support

3.5 新版功能.

Ever since the SSL module was introduced in Python 2.6, the `SSLSocket` class has provided two related but distinct areas of functionality:

- SSL protocol handling
- Network IO

The network IO API is identical to that provided by `socket.socket`, from which `SSLSocket` also inherits. This allows an SSL socket to be used as a drop-in replacement for a regular socket, making it very easy to add SSL support to an existing application.

Combining SSL protocol handling and network IO usually works well, but there are some cases where it doesn't. An example is async IO frameworks that want to use a different IO multiplexing model than the “select/poll on a file descriptor” (readiness based) model that is assumed by `socket.socket` and by the internal OpenSSL socket IO routines. This is mostly relevant for platforms like Windows where this model is not efficient. For this purpose, a reduced scope variant of `SSLSocket` called `SSLObject` is provided.

class `ssl.SSLObject`

A reduced-scope variant of `SSLSocket` representing an SSL protocol instance that does not contain any network IO methods. This class is typically used by framework authors that want to implement asynchronous IO for SSL through memory buffers.

This class implements an interface on top of a low-level SSL object as implemented by OpenSSL. This object captures the state of an SSL connection but does not provide any network IO itself. IO needs to be performed through separate “BIO” objects which are OpenSSL's IO abstraction layer.

An `SSLObject` instance can be created using the `wrap_bio()` method. This method will create the `SSLObject` instance and bind it to a pair of BIOs. The *incoming* BIO is used to pass data from Python to the SSL protocol instance, while the *outgoing* BIO is used to pass data the other way around.

The following methods are available:

- `context`
- `server_side`
- `server_hostname`
- `session`
- `session_reused`
- `read()`
- `write()`
- `getpeercert()`
- `selected_npn_protocol()`
- `cipher()`
- `shared_ciphers()`
- `compression()`
- `pending()`
- `do_handshake()`
- `unwrap()`
- `get_channel_binding()`

When compared to `SSLSocket`, this object lacks the following features:

- Any form of network IO; `recv()` and `send()` read and write only to the underlying `MemoryBIO` buffers.
- There is no `do_handshake_on_connect` machinery. You must always manually call `do_handshake()` to start the handshake.
- There is no handling of `suppress_ragged_eofs`. All end-of-file conditions that are in violation of the protocol are reported via the `SSLEOFError` exception.

- The method `unwrap()` call does not return anything, unlike for an SSL socket where it returns the underlying socket.
- The `server_name_callback` callback passed to `SSLContext.set_servername_callback()` will get an `SSLObject` instance instead of a `SSLSocket` instance as its first parameter.

Some notes related to the use of `SSLObject`:

- All IO on an `SSLObject` is *non-blocking*. This means that for example `read()` will raise an `SSLWantReadError` if it needs more data than the incoming BIO has available.
- There is no module-level `wrap_bio()` call like there is for `wrap_socket()`. An `SSLObject` is always created via an `SSLContext`.

An `SSLObject` communicates with the outside world using memory buffers. The class `MemoryBIO` provides a memory buffer that can be used for this purpose. It wraps an OpenSSL memory BIO (Basic IO) object:

class `ssl.MemoryBIO`

A memory buffer that can be used to pass data between Python and an SSL protocol instance.

pending

Return the number of bytes currently in the memory buffer.

eof

A boolean indicating whether the memory BIO is current at the end-of-file position.

read (*n=-1*)

Read up to *n* bytes from the memory buffer. If *n* is not specified or negative, all bytes are returned.

write (*buf*)

Write the bytes from *buf* to the memory BIO. The *buf* argument must be an object supporting the buffer protocol.

The return value is the number of bytes written, which is always equal to the length of *buf*.

write_eof ()

Write an EOF marker to the memory BIO. After this method has been called, it is illegal to call `write()`. The attribute `eof` will become true after all data currently in the buffer has been read.

18.2.8 SSL session

3.6 新版功能.

class `ssl.SSLSession`

Session object used by `session`.

id

time

timeout

ticket_lifetime_hint

has_ticket

18.2.9 Security considerations

Best defaults

For **client use**, if you don't have any special requirements for your security policy, it is highly recommended that you use the `create_default_context()` function to create your SSL context. It will load the system's trusted CA certificates, enable certificate validation and hostname checking, and try to choose reasonably secure protocol and cipher settings.

For example, here is how you would use the `smtplib.SMTP` class to create a trusted, secure connection to a SMTP server:

```
>>> import ssl, smtplib
>>> smtp = smtplib.SMTP("mail.python.org", port=587)
>>> context = ssl.create_default_context()
>>> smtp.starttls(context=context)
(220, b'2.0.0 Ready to start TLS')
```

If a client certificate is needed for the connection, it can be added with `SSLContext.load_cert_chain()`.

By contrast, if you create the SSL context by calling the `SSLContext` constructor yourself, it will not have certificate validation nor hostname checking enabled by default. If you do so, please read the paragraphs below to achieve a good security level.

Manual settings

Verifying certificates

When calling the `SSLContext` constructor directly, `CERT_NONE` is the default. Since it does not authenticate the other peer, it can be insecure, especially in client mode where most of time you would like to ensure the authenticity of the server you're talking to. Therefore, when in client mode, it is highly recommended to use `CERT_REQUIRED`. However, it is in itself not sufficient; you also have to check that the server certificate, which can be obtained by calling `SSLSocket.getpeercert()`, matches the desired service. For many protocols and applications, the service can be identified by the hostname; in this case, the `match_hostname()` function can be used. This common check is automatically performed when `SSLContext.check_hostname` is enabled.

In server mode, if you want to authenticate your clients using the SSL layer (rather than using a higher-level authentication mechanism), you'll also have to specify `CERT_REQUIRED` and similarly check the client certificate.

Protocol versions

SSL versions 2 and 3 are considered insecure and are therefore dangerous to use. If you want maximum compatibility between clients and servers, it is recommended to use `PROTOCOL_TLS_CLIENT` or `PROTOCOL_TLS_SERVER` as the protocol version. SSLv2 and SSLv3 are disabled by default.

```
>>> client_context = ssl.SSLContext(ssl.PROTOCOL_TLS_CLIENT)
>>> client_context.options |= ssl.OP_NO_TLSv1
>>> client_context.options |= ssl.OP_NO_TLSv1_1
```

The SSL context created above will only allow TLSv1.2 and later (if supported by your system) connections to a server. `PROTOCOL_TLS_CLIENT` implies certificate validation and hostname checks by default. You have to load certificates into the context.

Cipher selection

If you have advanced security requirements, fine-tuning of the ciphers enabled when negotiating a SSL session is possible through the `SSLContext.set_ciphers()` method. Starting from Python 3.2.3, the `ssl` module disables certain weak ciphers by default, but you may want to further restrict the cipher choice. Be sure to read OpenSSL's documentation about the [cipher list format](#). If you want to check which ciphers are enabled by a given cipher list, use `SSLContext.get_ciphers()` or the `openssl ciphers` command on your system.

Multi-processing

If using this module as part of a multi-processed application (using, for example the `multiprocessing` or `concurrent.futures` modules), be aware that OpenSSL's internal random number generator does not properly handle forked processes. Applications must change the PRNG state of the parent process if they use any SSL feature with `os.fork()`. Any successful call of `RAND_add()`, `RAND_bytes()` or `RAND_pseudo_bytes()` is sufficient.

18.2.10 LibreSSL support

LibreSSL is a fork of OpenSSL 1.0.1. The `ssl` module has limited support for LibreSSL. Some features are not available when the `ssl` module is compiled with LibreSSL.

- LibreSSL >= 2.6.1 no longer supports NPN. The methods `SSLContext.set_npn_protocols()` and `SSLSocket.selected_npn_protocol()` are not available.
- `SSLContext.set_default_verify_paths()` ignores the env vars `SSL_CERT_FILE` and `SSL_CERT_PATH` although `get_default_verify_paths()` still reports them.

参见:

Class `socket.socket` Documentation of underlying `socket` class

SSL/TLS Strong Encryption: An Introduction Intro from the Apache HTTP Server documentation

RFC 1422: Privacy Enhancement for Internet Electronic Mail: Part II: Certificate-Based Key Management

Steve Kent

RFC 4086: Randomness Requirements for Security Donald E., Jeffrey I. Schiller

RFC 5280: Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile

D. Cooper

RFC 5246: The Transport Layer Security (TLS) Protocol Version 1.2 T. Dierks et. al.

RFC 6066: Transport Layer Security (TLS) Extensions D. Eastlake

IANA TLS: Transport Layer Security (TLS) Parameters IANA

RFC 7525: Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)
IETF

Mozilla's Server Side TLS recommendations Mozilla

18.3 select —Waiting for I/O completion

This module provides access to the `select()` and `poll()` functions available in most operating systems, `devpoll()` available on Solaris and derivatives, `epoll()` available on Linux 2.5+ and `kqueue()` available on most BSD. Note

that on Windows, it only works for sockets; on other operating systems, it also works for other file types (in particular, on Unix, it works on pipes). It cannot be used on regular files to determine whether a file has grown since it was last read.

注解: The `selectors` module allows high-level and efficient I/O multiplexing, built upon the `select` module primitives. Users are encouraged to use the `selectors` module instead, unless they want precise control over the OS-level primitives used.

The module defines the following:

exception `select.error`

A deprecated alias of `OSError`.

在 3.3 版更改: Following [PEP 3151](#), this class was made an alias of `OSError`.

`select.devpoll()`

(Only supported on Solaris and derivatives.) Returns a `/dev/poll` polling object; see section [/dev/poll Polling Objects](#) below for the methods supported by devpoll objects.

`devpoll()` objects are linked to the number of file descriptors allowed at the time of instantiation. If your program reduces this value, `devpoll()` will fail. If your program increases this value, `devpoll()` may return an incomplete list of active file descriptors.

The new file descriptor is *non-inheritable*.

3.3 新版功能.

在 3.4 版更改: 新的文件描述符现在是不可继承的。

`select.epoll(sizehint=-1, flags=0)`

(Only supported on Linux 2.5.44 and newer.) Return an edge polling object, which can be used as Edge or Level Triggered interface for I/O events.

sizehint informs `epoll` about the expected number of events to be registered. It must be positive, or `-1` to use the default. It is only used on older systems where `epoll_create1()` is not available; otherwise it has no effect (though its value is still checked).

flags is deprecated and completely ignored. However, when supplied, its value must be `0` or `select.EPOLL_CLOEXEC`, otherwise `OSError` is raised.

See the [Edge and Level Trigger Polling \(epoll\) Objects](#) section below for the methods supported by epolling objects.

`epoll` objects support the context management protocol: when used in a `with` statement, the new file descriptor is automatically closed at the end of the block.

The new file descriptor is *non-inheritable*.

在 3.3 版更改: Added the *flags* parameter.

在 3.4 版更改: Support for the `with` statement was added. The new file descriptor is now non-inheritable.

3.4 版后已移除: The *flags* parameter. `select.EPOLL_CLOEXEC` is used by default now. Use `os.set_inheritable()` to make the file descriptor inheritable.

`select.poll()`

(Not supported by all operating systems.) Returns a polling object, which supports registering and unregistering file descriptors, and then polling them for I/O events; see section [Polling Objects](#) below for the methods supported by polling objects.

`select.kqueue()`

(Only supported on BSD.) Returns a kernel queue object; see section [Kqueue Objects](#) below for the methods supported by kqueue objects.

The new file descriptor is *non-inheritable*.

在 3.4 版更改: 新的文件描述符现在是不可继承的。

`select.kevent` (*ident*, *filter*=`KQ_FILTER_READ`, *flags*=`KQ_EV_ADD`, *f*flags=0, *data*=0, *udata*=0)
(Only supported on BSD.) Returns a kernel event object; see section [Kevent Objects](#) below for the methods supported by kevent objects.

`select.select` (*rlist*, *wlist*, *xlist*[, *timeout*])

This is a straightforward interface to the Unix `select()` system call. The first three arguments are sequences of ‘waitable objects’: either integers representing file descriptors or objects with a parameterless method named `fileno()` returning such an integer:

- *rlist*: wait until ready for reading
- *wlist*: wait until ready for writing
- *xlist*: wait for an “exceptional condition” (see the manual page for what your system considers such a condition)

Empty sequences are allowed, but acceptance of three empty sequences is platform-dependent. (It is known to work on Unix but not on Windows.) The optional *timeout* argument specifies a time-out as a floating point number in seconds. When the *timeout* argument is omitted the function blocks until at least one file descriptor is ready. A time-out value of zero specifies a poll and never blocks.

The return value is a triple of lists of objects that are ready: subsets of the first three arguments. When the time-out is reached without a file descriptor becoming ready, three empty lists are returned.

Among the acceptable object types in the sequences are Python *file objects* (e.g. `sys.stdin`, or objects returned by `open()` or `os.popen()`), socket objects returned by `socket.socket()`. You may also define a *wrapper* class yourself, as long as it has an appropriate `fileno()` method (that really returns a file descriptor, not just a random integer).

注解: File objects on Windows are not acceptable, but sockets are. On Windows, the underlying `select()` function is provided by the WinSock library, and does not handle file descriptors that don’t originate from WinSock.

在 3.5 版更改: The function is now retried with a recomputed timeout when interrupted by a signal, except if the signal handler raises an exception (see [PEP 475](#) for the rationale), instead of raising `InterruptedError`.

`select.PIPE_BUF`

The minimum number of bytes which can be written without blocking to a pipe when the pipe has been reported as ready for writing by `select()`, `poll()` or another interface in this module. This doesn’t apply to other kind of file-like objects such as sockets.

This value is guaranteed by POSIX to be at least 512. Availability: Unix.

3.2 新版功能.

18.3.1 /dev/poll Polling Objects

Solaris and derivatives have `/dev/poll`. While `select()` is $O(\text{highest file descriptor})$ and `poll()` is $O(\text{number of file descriptors})$, `/dev/poll` is $O(\text{active file descriptors})$.

`/dev/poll` behaviour is very close to the standard `poll()` object.

`devpoll.close()`

Close the file descriptor of the polling object.

3.4 新版功能.

`devpoll.closed`

True if the polling object is closed.

3.4 新版功能.

`devpoll.fileno()`

Return the file descriptor number of the polling object.

3.4 新版功能.

`devpoll.register(fd[, eventmask])`

Register a file descriptor with the polling object. Future calls to the `poll()` method will then check whether the file descriptor has any pending I/O events. `fd` can be either an integer, or an object with a `fileno()` method that returns an integer. File objects implement `fileno()`, so they can also be used as the argument.

`eventmask` is an optional bitmask describing the type of events you want to check for. The constants are the same that with `poll()` object. The default value is a combination of the constants `POLLIN`, `POLLPRIO`, and `POLLOUT`.

警告: Registering a file descriptor that's already registered is not an error, but the result is undefined. The appropriate action is to unregister or modify it first. This is an important difference compared with `poll()`.

`devpoll.modify(fd[, eventmask])`

This method does an `unregister()` followed by a `register()`. It is (a bit) more efficient than doing the same explicitly.

`devpoll.unregister(fd)`

Remove a file descriptor being tracked by a polling object. Just like the `register()` method, `fd` can be an integer or an object with a `fileno()` method that returns an integer.

Attempting to remove a file descriptor that was never registered is safely ignored.

`devpoll.poll([timeout])`

Polls the set of registered file descriptors, and returns a possibly-empty list containing `(fd, event)` 2-tuples for the descriptors that have events or errors to report. `fd` is the file descriptor, and `event` is a bitmask with bits set for the reported events for that descriptor — `POLLIN` for waiting input, `POLLOUT` to indicate that the descriptor can be written to, and so forth. An empty list indicates that the call timed out and no file descriptors had any events to report. If `timeout` is given, it specifies the length of time in milliseconds which the system will wait for events before returning. If `timeout` is omitted, `-1`, or `None`, the call will block until there is an event for this poll object.

在 3.5 版更改: The function is now retried with a recomputed timeout when interrupted by a signal, except if the signal handler raises an exception (see [PEP 475](#) for the rationale), instead of raising `InterruptedError`.

18.3.2 Edge and Level Trigger Polling (epoll) Objects

<http://linux.die.net/man/4/epoll>

`eventmask`

常数	意义
<code>EPOLLIN</code>	Available for read
<code>EPOLLOUT</code>	Available for write
<code>EPOLLPRI</code>	Urgent data for read
<code>EPOLLERR</code>	Error condition happened on the assoc. fd
<code>EPOLLHUP</code>	Hang up happened on the assoc. fd
<code>EPOLLET</code>	Set Edge Trigger behavior, the default is Level Trigger behavior
<code>EPOLLONESHOT</code>	Set one-shot behavior. After one event is pulled out, the fd is internally disabled
<code>EPOLLEXCLUSIVE</code>	Wake only one epoll object when the associated fd has an event. The default (if this flag is not set) is to wake all epoll objects polling on a fd.
<code>EPOLLRDHUP</code>	Stream socket peer closed connection or shut down writing half of connection.
<code>EPOLLRDNONE</code>	Equivalent to <code>EPOLLIN</code>
<code>EPOLLRDBAND</code>	Priority data band can be read.
<code>EPOLLWRNONE</code>	Equivalent to <code>EPOLLOUT</code>
<code>EPOLLWRBAND</code>	Priority data may be written.
<code>EPOLLMSG</code>	Ignored.

`epoll.close()`

Close the control file descriptor of the epoll object.

`epoll.closed`

True if the epoll object is closed.

`epoll.fileno()`

Return the file descriptor number of the control fd.

`epoll.fromfd(fd)`

Create an epoll object from a given file descriptor.

`epoll.register(fd[, eventmask])`

Register a fd descriptor with the epoll object.

`epoll.modify(fd, eventmask)`

Modify a registered file descriptor.

`epoll.unregister(fd)`

Remove a registered file descriptor from the epoll object.

`epoll.poll(timeout=-1, maxevents=-1)`

Wait for events. timeout in seconds (float)

在 3.5 版更改: The function is now retried with a recomputed timeout when interrupted by a signal, except if the signal handler raises an exception (see [PEP 475](#) for the rationale), instead of raising `InterruptedError`.

18.3.3 Polling Objects

The `poll()` system call, supported on most Unix systems, provides better scalability for network servers that service many, many clients at the same time. `poll()` scales better because the system call only requires listing the file descriptors of interest, while `select()` builds a bitmap, turns on bits for the fds of interest, and then afterward the whole bitmap has to be linearly scanned again. `select()` is $O(\text{highest file descriptor})$, while `poll()` is $O(\text{number of file descriptors})$.

`poll.register(fd[, eventmask])`

Register a file descriptor with the polling object. Future calls to the `poll()` method will then check whether the file descriptor has any pending I/O events. `fd` can be either an integer, or an object with a `fileno()` method that returns an integer. File objects implement `fileno()`, so they can also be used as the argument.

eventmask is an optional bitmask describing the type of events you want to check for, and can be a combination of the constants `POLLIN`, `POLLPRI`, and `POLLOUT`, described in the table below. If not specified, the default value used will check for all 3 types of events.

常数	意义
<code>POLLIN</code>	There is data to read
<code>POLLPRI</code>	There is urgent data to read
<code>POLLOUT</code>	Ready for output: writing will not block
<code>POLLERR</code>	Error condition of some sort
<code>POLLHUP</code>	Hung up
<code>POLLRDHUP</code>	Stream socket peer closed connection, or shut down writing half of connection
<code>POLLNVAL</code>	Invalid request: descriptor not open

Registering a file descriptor that's already registered is not an error, and has the same effect as registering the descriptor exactly once.

`poll.modify(fd, eventmask)`

Modifies an already registered *fd*. This has the same effect as `register(fd, eventmask)`. Attempting to modify a file descriptor that was never registered causes an `OSError` exception with `errno` `ENOENT` to be raised.

`poll.unregister(fd)`

Remove a file descriptor being tracked by a polling object. Just like the `register()` method, *fd* can be an integer or an object with a `fileno()` method that returns an integer.

Attempting to remove a file descriptor that was never registered causes a `KeyError` exception to be raised.

`poll.poll([timeout])`

Polls the set of registered file descriptors, and returns a possibly-empty list containing (*fd*, *event*) 2-tuples for the descriptors that have events or errors to report. *fd* is the file descriptor, and *event* is a bitmask with bits set for the reported events for that descriptor — `POLLIN` for waiting input, `POLLOUT` to indicate that the descriptor can be written to, and so forth. An empty list indicates that the call timed out and no file descriptors had any events to report. If *timeout* is given, it specifies the length of time in milliseconds which the system will wait for events before returning. If *timeout* is omitted, negative, or `None`, the call will block until there is an event for this poll object.

在 3.5 版更改: The function is now retried with a recomputed timeout when interrupted by a signal, except if the signal handler raises an exception (see [PEP 475](#) for the rationale), instead of raising `InterruptedError`.

18.3.4 Kqueue Objects

`kqueue.close()`

Close the control file descriptor of the kqueue object.

`kqueue.closed`

True if the kqueue object is closed.

`kqueue.fileno()`

Return the file descriptor number of the control fd.

`kqueue.fromfd(fd)`

Create a kqueue object from a given file descriptor.

`kqueue.control(changelist, max_events[, timeout=None])` → eventlist

Low level interface to kevent

- *changelist* must be an iterable of kevent object or `None`
- *max_events* must be 0 or a positive integer

- timeout in seconds (floats possible)

在 3.5 版更改: The function is now retried with a recomputed timeout when interrupted by a signal, except if the signal handler raises an exception (see [PEP 475](#) for the rationale), instead of raising `InterruptedError`.

18.3.5 Kevent Objects

<https://www.freebsd.org/cgi/man.cgi?query=kqueue&sektion=2>

`kevent.ident`

Value used to identify the event. The interpretation depends on the filter but it's usually the file descriptor. In the constructor `ident` can either be an int or an object with a `fileno()` method. `kevent` stores the integer internally.

`kevent.filter`

Name of the kernel filter.

常数	意义
<code>KQ_FILTER_READ</code>	Takes a descriptor and returns whenever there is data available to read
<code>KQ_FILTER_WRITE</code>	Takes a descriptor and returns whenever there is data available to write
<code>KQ_FILTER_AIO</code>	AIO requests
<code>KQ_FILTER_VNODE</code>	Returns when one or more of the requested events watched in <i>fflag</i> occurs
<code>KQ_FILTER_PROC</code>	Watch for events on a process id
<code>KQ_FILTER_NETDEV</code>	Watch for events on a network device [not available on Mac OS X]
<code>KQ_FILTER_SIGNAL</code>	Returns whenever the watched signal is delivered to the process
<code>KQ_FILTER_TIMER</code>	Establishes an arbitrary timer

`kevent.flags`

Filter action.

常数	意义
<code>KQ_EV_ADD</code>	Adds or modifies an event
<code>KQ_EV_DELETE</code>	Removes an event from the queue
<code>KQ_EV_ENABLE</code>	Permits <code>control()</code> to return the event
<code>KQ_EV_DISABLE</code>	Disable event
<code>KQ_EV_ONESHOT</code>	Removes event after first occurrence
<code>KQ_EV_CLEAR</code>	Reset the state after an event is retrieved
<code>KQ_EV_SYSFLAGS</code>	internal event
<code>KQ_EV_FLAG1</code>	internal event
<code>KQ_EV_EOF</code>	Filter specific EOF condition
<code>KQ_EV_ERROR</code>	See return values

`kevent.fflags`

Filter specific flags.

`KQ_FILTER_READ` and `KQ_FILTER_WRITE` filter flags:

常数	意义
<code>KQ_NOTE_LOWAT</code>	low water mark of a socket buffer

`KQ_FILTER_VNODE` filter flags:

常数	意义
KQ_NOTE_DELETE	<i>unlink()</i> was called
KQ_NOTE_WRITE	a write occurred
KQ_NOTE_EXTEND	the file was extended
KQ_NOTE_ATTRIB	an attribute was changed
KQ_NOTE_LINK	the link count has changed
KQ_NOTE_RENAME	the file was renamed
KQ_NOTE_REVOKE	access to the file was revoked

KQ_FILTER_PROC filter flags:

常数	意义
KQ_NOTE_EXIT	the process has exited
KQ_NOTE_FORK	the process has called <i>fork()</i>
KQ_NOTE_EXEC	the process has executed a new process
KQ_NOTE_PCTRLMASK	internal filter flag
KQ_NOTE_PDATAMASK	internal filter flag
KQ_NOTE_TRACK	follow a process across <i>fork()</i>
KQ_NOTE_CHILD	returned on the child process for <i>NOTE_TRACK</i>
KQ_NOTE_TRACKERR	unable to attach to a child

KQ_FILTER_NETDEV filter flags (not available on Mac OS X):

常数	意义
KQ_NOTE_LINKUP	link is up
KQ_NOTE_LINKDOWN	link is down
KQ_NOTE_LINKINV	link state is invalid

`kevent.data`
Filter specific data.

`kevent.udata`
User defined value.

18.4 selectors —高级 I/O 复用库

3.4 新版功能.

源码: [Lib/selectors.py](#)

18.4.1 概述

This module allows high-level and efficient I/O multiplexing, built upon the *select* module primitives. Users are encouraged to use this module instead, unless they want precise control over the OS-level primitives used.

It defines a *BaseSelector* abstract base class, along with several concrete implementations (*KqueueSelector*, *EpollSelector*…), that can be used to wait for I/O readiness notification on multiple file objects. In the following, “file object” refers to any object with a `fileno()` method, or a raw file descriptor. See *file object*.

`DefaultSelector` is an alias to the most efficient implementation available on the current platform: this should be the default choice for most users.

注解： The type of file objects supported depends on the platform: on Windows, sockets are supported, but not pipes, whereas on Unix, both are supported (some other types may be supported as well, such as fifos or special file devices).

参见：

`select` Low-level I/O multiplexing module.

18.4.2 类

Classes hierarchy:

```
BaseSelector
+-- SelectSelector
+-- PollSelector
+-- EpollSelector
+-- DevpollSelector
+-- KqueueSelector
```

In the following, *events* is a bitwise mask indicating which I/O events should be waited for on a given file object. It can be a combination of the modules constants below:

常数	意义
EVENT_READ	Available for read
EVENT_WRITE	Available for write

class selectors.SelectorKey

A *SelectorKey* is a *namedtuple* used to associate a file object to its underlying file descriptor, selected event mask and attached data. It is returned by several *BaseSelector* methods.

fileobj

File object registered.

fd

Underlying file descriptor.

events

Events that must be waited for on this file object.

data

Optional opaque data associated to this file object: for example, this could be used to store a per-client session ID.

class selectors.BaseSelector

A *BaseSelector* is used to wait for I/O event readiness on multiple file objects. It supports file stream registration, unregistration, and a method to wait for I/O events on those streams, with an optional timeout. It's an abstract base class, so cannot be instantiated. Use *DefaultSelector* instead, or one of *SelectSelector*, *KqueueSelector* etc. if you want to specifically use an implementation, and your platform supports it. *BaseSelector* and its concrete implementations support the *context manager* protocol.

abstractmethod register (*fileobj*, *events*, *data=None*)

Register a file object for selection, monitoring it for I/O events.

fileobj is the file object to monitor. It may either be an integer file descriptor or an object with a `fileno()` method. *events* is a bitwise mask of events to monitor. *data* is an opaque object.

This returns a new `SelectorKey` instance, or raises a `ValueError` in case of invalid event mask or file descriptor, or `KeyError` if the file object is already registered.

abstractmethod unregister (*fileobj*)

Unregister a file object from selection, removing it from monitoring. A file object shall be unregistered prior to being closed.

fileobj must be a file object previously registered.

This returns the associated `SelectorKey` instance, or raises a `KeyError` if *fileobj* is not registered. It will raise `ValueError` if *fileobj* is invalid (e.g. it has no `fileno()` method or its `fileno()` method has an invalid return value).

modify (*fileobj*, *events*, *data=None*)

Change a registered file object's monitored events or attached data.

This is equivalent to `BaseSelector.unregister(fileobj)()` followed by `BaseSelector.register(fileobj, events, data)()`, except that it can be implemented more efficiently.

This returns a new `SelectorKey` instance, or raises a `ValueError` in case of invalid event mask or file descriptor, or `KeyError` if the file object is not registered.

abstractmethod select (*timeout=None*)

Wait until some registered file objects become ready, or the timeout expires.

If *timeout* > 0, this specifies the maximum wait time, in seconds. If *timeout* ≤ 0, the call won't block, and will report the currently ready file objects. If *timeout* is `None`, the call will block until a monitored file object becomes ready.

This returns a list of (*key*, *events*) tuples, one for each ready file object.

key is the `SelectorKey` instance corresponding to a ready file object. *events* is a bitmask of events ready on this file object.

注解: This method can return before any file object becomes ready or the timeout has elapsed if the current process receives a signal: in this case, an empty list will be returned.

在 3.5 版更改: The selector is now retried with a recomputed timeout when interrupted by a signal if the signal handler did not raise an exception (see [PEP 475](#) for the rationale), instead of returning an empty list of events before the timeout.

close ()

Close the selector.

This must be called to make sure that any underlying resource is freed. The selector shall not be used once it has been closed.

get_key (*fileobj*)

Return the key associated with a registered file object.

This returns the `SelectorKey` instance associated to this file object, or raises `KeyError` if the file object is not registered.

abstractmethod get_map ()

Return a mapping of file objects to selector keys.

This returns a `Mapping` instance mapping registered file objects to their associated `SelectorKey` instance.

class `selectors.DefaultSelector`

The default selector class, using the most efficient implementation available on the current platform. This should be the default choice for most users.

class `selectors.SelectSelector`

`select.select()`-based selector.

class `selectors.PollSelector`

`select.poll()`-based selector.

class `selectors.EpollSelector`

`select.epoll()`-based selector.

fileno()

This returns the file descriptor used by the underlying `select.epoll()` object.

class `selectors.DevpollSelector`

`select.devpoll()`-based selector.

fileno()

This returns the file descriptor used by the underlying `select.devpoll()` object.

3.5 新版功能.

class `selectors.KqueueSelector`

`select.kqueue()`-based selector.

fileno()

This returns the file descriptor used by the underlying `select.kqueue()` object.

18.4.3 示例

Here is a simple echo server implementation:

```
import selectors
import socket

sel = selectors.DefaultSelector()

def accept(sock, mask):
    conn, addr = sock.accept() # Should be ready
    print('accepted', conn, 'from', addr)
    conn.setblocking(False)
    sel.register(conn, selectors.EVENT_READ, read)

def read(conn, mask):
    data = conn.recv(1000) # Should be ready
    if data:
        print('echoing', repr(data), 'to', conn)
        conn.send(data) # Hope it won't block
    else:
        print('closing', conn)
        sel.unregister(conn)
        conn.close()

sock = socket.socket()
sock.bind(('localhost', 1234))
sock.listen(100)
sock.setblocking(False)
```

(下页继续)

(续上页)

```
sel.register(sock, selectors.EVENT_READ, accept)

while True:
    events = sel.select()
    for key, mask in events:
        callback = key.data
        callback(key.fileobj, mask)
```

18.5 `asyncio` —Asynchronous I/O, event loop, coroutines and tasks

3.4 新版功能.

Source code: [Lib/asyncio/](#)

This module provides infrastructure for writing single-threaded concurrent code using coroutines, multiplexing I/O access over sockets and other resources, running network clients and servers, and other related primitives. Here is a more detailed list of the package contents:

- a pluggable *event loop* with various system-specific implementations;
- *transport* and *protocol* abstractions (similar to those in *Twisted*);
- concrete support for TCP, UDP, SSL, subprocess pipes, delayed calls, and others (some may be system-dependent);
- a *Future* class that mimics the one in the *concurrent.futures* module, but adapted for use with the event loop;
- coroutines and tasks based on `yield from` ([PEP 380](#)), to help write concurrent code in a sequential fashion;
- cancellation support for *Futures* and coroutines;
- *synchronization primitives* for use between coroutines in a single thread, mimicking those in the *threading* module;
- an interface for passing work off to a threadpool, for times when you absolutely, positively have to use a library that makes blocking I/O calls.

Asynchronous programming is more complex than classical “sequential” programming: see the *Develop with asyncio* page which lists common traps and explains how to avoid them. *Enable the debug mode* during development to detect common issues.

Table of contents:

18.5.1 Base Event Loop

Source code: [Lib/asyncio/events.py](#)

The event loop is the central execution device provided by *asyncio*. It provides multiple facilities, including:

- Registering, executing and cancelling delayed calls (timeouts).
- Creating client and server *transports* for various kinds of communication.
- Launching subprocesses and the associated *transports* for communication with an external program.
- Delegating costly function calls to a pool of threads.

class `asyncio.BaseEventLoop`

This class is an implementation detail. It is a subclass of `AbstractEventLoop` and may be a base class of concrete event loop implementations found in `asyncio`. It should not be used directly; use `AbstractEventLoop` instead. `BaseEventLoop` should not be subclassed by third-party code; the internal interface is not stable.

class `asyncio.AbstractEventLoop`

Abstract base class of event loops.

This class is *not thread safe*.

Run an event loop`AbstractEventLoop.run_forever()`

Run until `stop()` is called. If `stop()` is called before `run_forever()` is called, this polls the I/O selector once with a timeout of zero, runs all callbacks scheduled in response to I/O events (and those that were already scheduled), and then exits. If `stop()` is called while `run_forever()` is running, this will run the current batch of callbacks and then exit. Note that callbacks scheduled by callbacks will not run in that case; they will run the next time `run_forever()` is called.

在 3.5.1 版更改.

`AbstractEventLoop.run_until_complete(future)`

Run until the `Future` is done.

If the argument is a *coroutine object*, it is wrapped by `ensure_future()`.

Return the `Future`'s result, or raise its exception.

`AbstractEventLoop.is_running()`

Returns running status of event loop.

`AbstractEventLoop.stop()`

Stop running the event loop.

This causes `run_forever()` to exit at the next suitable opportunity (see there for more details).

在 3.5.1 版更改.

`AbstractEventLoop.is_closed()`

Returns True if the event loop was closed.

3.4.2 新版功能.

`AbstractEventLoop.close()`

Close the event loop. The loop must not be running. Pending callbacks will be lost.

This clears the queues and shuts down the executor, but does not wait for the executor to finish.

This is idempotent and irreversible. No other methods should be called after this one.

coroutine `AbstractEventLoop.shutdown_asyncgens()`

Schedule all currently open *asynchronous generator* objects to close with an `aclose()` call. After calling this method, the event loop will issue a warning whenever a new asynchronous generator is iterated. Should be used to finalize all scheduled asynchronous generators reliably. Example:

```
try:
    loop.run_forever()
finally:
    loop.run_until_complete(loop.shutdown_asyncgens())
    loop.close()
```

3.6 新版功能.

Calls

Most *asyncio* functions don't accept keywords. If you want to pass keywords to your callback, use *functools.partial()*. For example, `loop.call_soon(functools.partial(print, "Hello", flush=True))` will call `print("Hello", flush=True)`.

注解: *functools.partial()* is better than lambda functions, because *asyncio* can inspect *functools.partial()* object to display parameters in debug mode, whereas lambda functions have a poor representation.

`AbstractEventLoop.call_soon(callback, *args)`

Arrange for a callback to be called as soon as possible. The callback is called after *call_soon()* returns, when control returns to the event loop.

This operates as a FIFO queue, callbacks are called in the order in which they are registered. Each callback will be called exactly once.

Any positional arguments after the callback will be passed to the callback when it is called.

An instance of *asyncio.Handle* is returned, which can be used to cancel the callback.

Use functools.partial to pass keywords to the callback.

`AbstractEventLoop.call_soon_threadsafe(callback, *args)`

Like *call_soon()*, but thread safe.

查看并发和多线程 章节的文档。

Delayed calls

The event loop has its own internal clock for computing timeouts. Which clock is used depends on the (platform-specific) event loop implementation; ideally it is a monotonic clock. This will generally be a different clock than *time.time()*.

注解: Timeouts (relative *delay* or absolute *when*) should not exceed one day.

`AbstractEventLoop.call_later(delay, callback, *args)`

Arrange for the *callback* to be called after the given *delay* seconds (either an int or float).

An instance of *asyncio.Handle* is returned, which can be used to cancel the callback.

callback will be called exactly once per call to *call_later()*. If two callbacks are scheduled for exactly the same time, it is undefined which will be called first.

The optional positional *args* will be passed to the callback when it is called. If you want the callback to be called with some named arguments, use a closure or *functools.partial()*.

Use functools.partial to pass keywords to the callback.

`AbstractEventLoop.call_at(when, callback, *args)`

Arrange for the *callback* to be called at the given absolute timestamp *when* (an int or float), using the same time reference as *AbstractEventLoop.time()*.

这个函数的行为与 *call_later()* 相同。

An instance of *asyncio.Handle* is returned, which can be used to cancel the callback.

Use functools.partial to pass keywords to the callback.

`AbstractEventLoop.time()`

Return the current time, as a *float* value, according to the event loop's internal clock.

参见:

`asyncio.sleep()` 函数

Futures

`AbstractEventLoop.create_future()`

Create an `asyncio.Future` object attached to the loop.

This is a preferred way to create futures in asyncio, as event loop implementations can provide alternative implementations of the Future class (with better performance or instrumentation).

3.5.2 新版功能.

Tasks

`AbstractEventLoop.create_task(coro)`

Schedule the execution of a *coroutine object*: wrap it in a future. Return a `Task` object.

三方的事件循环可以使用它们自己定义的`Task`类的子类来实现互操作性。这个例子里，返回值的类型是`Task`的子类。

This method was added in Python 3.4.2. Use the `async()` function to support also older Python versions.

3.4.2 新版功能.

`AbstractEventLoop.set_task_factory(factory)`

Set a task factory that will be used by `AbstractEventLoop.create_task()`.

If `factory` is `None` the default task factory will be set.

If `factory` is a *callable*, it should have a signature matching `(loop, coro)`, where `loop` will be a reference to the active event loop, `coro` will be a coroutine object. The callable must return an `asyncio.Future` compatible object.

3.4.4 新版功能.

`AbstractEventLoop.get_task_factory()`

Return a task factory, or `None` if the default one is in use.

3.4.4 新版功能.

Creating connections

coroutine `AbstractEventLoop.create_connection(protocol_factory, host=None, port=None, *, ssl=None, family=0, proto=0, flags=0, sock=None, local_addr=None, server_hostname=None)`

Create a streaming transport connection to a given Internet *host* and *port*: socket family `AF_INET` or `AF_INET6` depending on *host* (or *family* if specified), socket type `SOCK_STREAM`. *protocol_factory* must be a callable returning a *protocol* instance.

This method is a *coroutine* which will try to establish the connection in the background. When successful, the coroutine returns a `(transport, protocol)` pair.

基本操作的时间顺序如下:

1. The connection is established, and a *transport* is created to represent it.
2. *protocol_factory* is called without arguments and must return a *protocol* instance.

3. The protocol instance is tied to the transport, and its `connection_made()` method is called.
4. The coroutine returns successfully with the `(transport, protocol)` pair.

创建的 `transport` 是一个实现相关的双向流。

注解: `protocol_factory` can be any kind of callable, not necessarily a class. For example, if you want to use a pre-created protocol instance, you can pass `lambda: my_protocol`.

Options that change how the connection is created:

- `ssl`: if given and not false, a SSL/TLS transport is created (by default a plain TCP transport is created). If `ssl` is a `ssl.SSLContext` object, this context is used to create the transport; if `ssl` is `True`, a context with some unspecified default settings is used.

参见:

[SSL/TLS 安全事项](#)

- `server_hostname`, is only for use together with `ssl`, and sets or overrides the hostname that the target server's certificate will be matched against. By default the value of the `host` argument is used. If `host` is empty, there is no default and you must pass a value for `server_hostname`. If `server_hostname` is an empty string, hostname matching is disabled (which is a serious security risk, allowing for man-in-the-middle-attacks).
- `family`, `proto`, `flags` 是可选的地址族, 协议和标志, 通过传递给 `getaddrinfo()` 来解析 `host`. 如果给出, 这些应该都是来自 `socket` 模块相应的常量的整数。
- `sock`, 如果给出的话, 应该是一个现有的, 已经连接的 `socket.socket` 对象, 这个对象将被 `transport` 使用。如果给出 `sock`, 则 `host`, `port`, `family`, `proto`, `flags` 和 `local_addr` 中的任何一个都不应该指定。
- `local_addr`, if given, is a `(local_host, local_port)` tuple used to bind the socket to locally. The `local_host` and `local_port` are looked up using `getaddrinfo()`, similarly to `host` and `port`.

在 3.5 版更改: On Windows with `ProactorEventLoop`, SSL/TLS is now supported.

参见:

The `open_connection()` function can be used to get a pair of `(StreamReader, StreamWriter)` instead of a protocol.

```
coroutine AbstractEventLoop.create_datagram_endpoint(protocol_factory, local_addr=None, remote_addr=None, *, family=0, proto=0, flags=0, reuse_address=None, reuse_port=None, allow_broadcast=None, sock=None)
```

Create datagram connection: socket family `AF_INET` or `AF_INET6` depending on `host` (or `family` if specified), socket type `SOCK_DGRAM`. `protocol_factory` must be a callable returning a `protocol` instance.

This method is a *coroutine* which will try to establish the connection in the background. When successful, the coroutine returns a `(transport, protocol)` pair.

Options changing how the connection is created:

- `local_addr`, 如果指定的话, 就是一个 `(local_host, local_port)` 元组, 用于在本地绑定套接字。 `local_host` 和 `local_port` 是使用 `getaddrinfo()` 来查找的。
- `remote_addr`, 如果指定的话, 就是一个 `(remote_host, remote_port)` 元组, 用于同一个远程地址连接。 `remote_host` 和 `remote_port` 是使用 `getaddrinfo()` 来查找的。

- *family, proto, flags* 是可选的地址族, 协议和标志, 其会被传递给 `getaddrinfo()` 来完成 *host* 的解析。如果要指定的话, 这些都应该是来自于 `socket` 模块的对应常量。
- *reuse_address* tells the kernel to reuse a local socket in `TIME_WAIT` state, without waiting for its natural timeout to expire. If not specified will automatically be set to `True` on UNIX.
- *reuse_port* tells the kernel to allow this endpoint to be bound to the same port as other existing endpoints are bound to, so long as they all set this flag when being created. This option is not supported on Windows and some UNIX' s. If the `SO_REUSEPORT` constant is not defined then this capability is unsupported.
- *allow_broadcast* 告知内核允许此端点向广播地址发送消息。
- *sock* 可选择通过指定此值用于使用一个预先存在的, 已经处于连接状态的 `socket.socket` 对象, 并将其提供给此传输对象使用。如果指定了这个值, *local_addr* 和 *remote_addr* 就应该被忽略 (必须为 `None`)。

On Windows with `ProactorEventLoop`, this method is not supported.

参见 `UDP echo` 客户端协议 和 `UDP echo` 服务端协议 的例子。

在 3.4.4 版更改: 添加了 *family, proto, flags, reuse_address, reuse_port, allow_broadcast* 和 *sock* 等参数。

```
coroutine AbstractEventLoop.create_unix_connection(protocol_factory, path,
                                                    *, ssl=None, sock=None,
                                                    server_hostname=None)
```

Create UNIX connection: socket family `AF_UNIX`, socket type `SOCK_STREAM`. The `AF_UNIX` socket family is used to communicate between processes on the same machine efficiently.

This method is a *coroutine* which will try to establish the connection in the background. When successful, the coroutine returns a `(transport, protocol)` pair.

path is the name of a UNIX domain socket, and is required unless a *sock* parameter is specified. Abstract UNIX sockets, *str*, and *bytes* paths are supported.

See the `AbstractEventLoop.create_connection()` method for parameters.

Availability: UNIX.

Creating listening connections

```
coroutine AbstractEventLoop.create_server(protocol_factory, host=None, port=None,
                                           *, family=socket.AF_UNSPEC,
                                           flags=socket.AI_PASSIVE, sock=None, backlog=100,
                                           ssl=None, reuse_address=None, reuse_port=None)
```

Create a TCP server (socket type `SOCK_STREAM`) bound to *host* and *port*.

Return a `Server` object, its *sockets* attribute contains created sockets. Use the `Server.close()` method to stop the server: close listening sockets.

参数:

- The *host* parameter can be a string, in that case the TCP server is bound to *host* and *port*. The *host* parameter can also be a sequence of strings and in that case the TCP server is bound to all hosts of the sequence. If *host* is an empty string or `None`, all interfaces are assumed and a list of multiple sockets will be returned (most likely one for IPv4 and another one for IPv6).
- *family* can be set to either `socket.AF_INET` or `AF_INET6` to force the socket to use IPv4 or IPv6. If not set it will be determined from host (defaults to `socket.AF_UNSPEC`).
- *flags* 是用于 `getaddrinfo()` 的位掩码。

- *sock* can optionally be specified in order to use a preexisting socket object. If specified, *host* and *port* should be omitted (must be *None*).
- *backlog* 是传递给 *listen()* 的最大排队连接的数量 (默认为 100)。
- *ssl* 可被设置为一个 *SSLContext* 以在接受的连接上启用 SSL。
- *reuse_address* tells the kernel to reuse a local socket in TIME_WAIT state, without waiting for its natural timeout to expire. If not specified will automatically be set to *True* on UNIX.
- *reuse_port* 告知内核, 只要在创建的时候都设置了这个标志, 就允许此端点绑定到其它端点列表所绑定的同样的端口上。这个选项在 Windows 上是不支持的。

This method is a *coroutine*.

在 3.5 版更改: On Windows with *ProactorEventLoop*, SSL/TLS is now supported.

参见:

The function *start_server()* creates a (*StreamReader*, *StreamWriter*) pair and calls back a function with this pair.

在 3.5.1 版更改: The *host* parameter can now be a sequence of strings.

coroutine *AbstractEventLoop.create_unix_server* (*protocol_factory*, *path=None*, *,
sock=None, *backlog=100*, *ssl=None*)
Similar to *AbstractEventLoop.create_server()*, but specific to the socket family *AF_UNIX*.

This method is a *coroutine*.

Availability: UNIX.

coroutine *BaseEventLoop.connect_accepted_socket* (*protocol_factory*, *sock*, *, *ssl=None*)
Handle an accepted connection.

This is used by servers that accept connections outside of *asyncio* but that use *asyncio* to handle them.

参数:

- *sock* is a preexisting socket object returned from an *accept* call.
- *ssl* 可被设置为一个 *SSLContext* 以在接受的连接上启用 SSL。

This method is a *coroutine*. When completed, the coroutine returns a (*transport*, *protocol*) pair.

3.5.3 新版功能.

Watch file descriptors

On Windows with *SelectorEventLoop*, only socket handles are supported (ex: pipe file descriptors are not supported).

On Windows with *ProactorEventLoop*, these methods are not supported.

AbstractEventLoop.add_reader (*fd*, *callback*, **args*)
Start watching the file descriptor for read availability and then call the *callback* with specified arguments.
Use functools.partial to pass keywords to the callback.

AbstractEventLoop.remove_reader (*fd*)
Stop watching the file descriptor for read availability.

AbstractEventLoop.add_writer (*fd*, *callback*, **args*)
Start watching the file descriptor for write availability and then call the *callback* with specified arguments.
Use functools.partial to pass keywords to the callback.

`AbstractEventLoop.remove_writer(fd)`

Stop watching the file descriptor for write availability.

The *watch a file descriptor for read events* example uses the low-level `AbstractEventLoop.add_reader()` method to register the file descriptor of a socket.

Low-level socket operations

coroutine `AbstractEventLoop.sock_recv(sock, nbytes)`

Receive data from the socket. Modeled after blocking `socket.socket.recv()` method.

The return value is a bytes object representing the data received. The maximum amount of data to be received at once is specified by `nbytes`.

With `SelectorEventLoop` event loop, the socket `sock` must be non-blocking.

This method is a *coroutine*.

coroutine `AbstractEventLoop.sock_sendall(sock, data)`

Send data to the socket. Modeled after blocking `socket.socket.sendall()` method.

The socket must be connected to a remote socket. This method continues to send data from `data` until either all data has been sent or an error occurs. `None` is returned on success. On error, an exception is raised, and there is no way to determine how much data, if any, was successfully processed by the receiving end of the connection.

With `SelectorEventLoop` event loop, the socket `sock` must be non-blocking.

This method is a *coroutine*.

coroutine `AbstractEventLoop.sock_connect(sock, address)`

Connect to a remote socket at `address`. Modeled after blocking `socket.socket.connect()` method.

With `SelectorEventLoop` event loop, the socket `sock` must be non-blocking.

This method is a *coroutine*.

在 3.5.2 版更改: `address` no longer needs to be resolved. `sock_connect` will try to check if the `address` is already resolved by calling `socket.inet_pton()`. If not, `AbstractEventLoop.getaddrinfo()` will be used to resolve the `address`.

参见:

`AbstractEventLoop.create_connection()` and `asyncio.open_connection()`.

coroutine `AbstractEventLoop.sock_accept(sock)`

Accept a connection. Modeled after blocking `socket.socket.accept()`.

此 `socket` 必须绑定到一个地址上并且监听连接。返回值是一个 `(conn, address)` 对, 其中 `conn` 是一个新 * 的套接字对象, 用于在此连接上收发数据, *`address` 是连接的另一端的套接字所绑定的地址。

The socket `sock` must be non-blocking.

This method is a *coroutine*.

参见:

`AbstractEventLoop.create_server()` and `start_server()`.

Resolve host name

coroutine `AbstractEventLoop.getaddrinfo(host, port, *, family=0, type=0, proto=0, flags=0)`

This method is a *coroutine*, similar to `socket.getaddrinfo()` function but non-blocking.

coroutine `AbstractEventLoop.getnameinfo(sockaddr, flags=0)`

This method is a *coroutine*, similar to `socket.getnameinfo()` function but non-blocking.

Connect pipes

On Windows with *SelectorEventLoop*, these methods are not supported. Use *ProactorEventLoop* to support pipes on Windows.

coroutine `AbstractEventLoop.connect_read_pipe(protocol_factory, pipe)`

Register read pipe in eventloop.

protocol_factory should instantiate object with *Protocol* interface. *pipe* is a *file-like object*. Return pair `(transport, protocol)`, where *transport* supports the *ReadTransport* interface.

使用 *SelectorEventLoop* 事件循环, *pipe* 被设置为非阻塞模式。

This method is a *coroutine*.

coroutine `AbstractEventLoop.connect_write_pipe(protocol_factory, pipe)`

Register write pipe in eventloop.

protocol_factory should instantiate object with *BaseProtocol* interface. *pipe* is *file-like object*. Return pair `(transport, protocol)`, where *transport* supports *WriteTransport* interface.

使用 *SelectorEventLoop* 事件循环, *pipe* 被设置为非阻塞模式。

This method is a *coroutine*.

参见:

The `AbstractEventLoop.subprocess_exec()` and `AbstractEventLoop.subprocess_shell()` methods.

UNIX signals

Availability: UNIX only.

`AbstractEventLoop.add_signal_handler(signum, callback, *args)`

Add a handler for a signal.

如果信号数字非法或者不可捕获, 就抛出一个 *ValueError*。如果建立处理器的过程中出现问题, 会抛出一个 *RuntimeError*。

Use *functools.partial* to pass keywords to the callback.

`AbstractEventLoop.remove_signal_handler(sig)`

Remove a handler for a signal.

Return True if a signal handler was removed, False if not.

参见:

signal 模块。

Executor

Call a function in an *Executor* (pool of threads or pool of processes). By default, an event loop uses a thread pool executor (*ThreadPoolExecutor*).

coroutine `AbstractEventLoop.run_in_executor(executor, func, *args)`

Arrange for a *func* to be called in the specified executor.

The *executor* argument should be an *Executor* instance. The default executor is used if *executor* is `None`.

Use *functools.partial* to pass keywords to the **func**.

This method is a *coroutine*.

在 3.5.3 版更改: `BaseEventLoop.run_in_executor()` no longer configures the `max_workers` of the thread pool executor it creates, instead leaving it up to the thread pool executor (*ThreadPoolExecutor*) to set the default.

`AbstractEventLoop.set_default_executor(executor)`

Set the default executor used by *run_in_executor()*.

错误处理 API

允许自定义事件循环中如何去处理异常。

`AbstractEventLoop.set_exception_handler(handler)`

将 *handler* 设置为新的事件循环异常处理器。

If *handler* is `None`, the default exception handler will be set.

If *handler* is a callable object, it should have a matching signature to `(loop, context)`, where *loop* will be a reference to the active event loop, *context* will be a dict object (see *call_exception_handler()* documentation for details about context).

`AbstractEventLoop.get_exception_handler()`

Return the exception handler, or `None` if the default one is in use.

3.5.2 新版功能.

`AbstractEventLoop.default_exception_handler(context)`

默认的异常处理器。

This is called when an exception occurs and no exception handler is set, and can be called by a custom exception handler that wants to defer to the default behavior.

context 参数和 *call_exception_handler()* 中的同名参数完全相同。

`AbstractEventLoop.call_exception_handler(context)`

调用当前事件循环异常处理器。

context is a dict object containing the following keys (new keys may be introduced later):

- ‘message’: 错误消息;
- ‘exception’ (可选): 异常对象;
- ‘future’ (可选): *asyncio.Future* 实例。
- ‘handle’ (可选): *asyncio.Handle* 实例;
- ‘protocol’ (可选): *Protocol* 实例;
- ‘transport’ (可选): *Transport* 实例;
- ‘socket’ (可选): *socket.socket* 实例。

注解: Note: this method should not be overloaded in subclassed event loops. For any custom exception handling, use *set_exception_handler()* method.

Debug mode

`AbstractEventLoop.get_debug()`
获取事件循环调试模式状态 (*bool*)。

如果环境变量 `PYTHONASYNCIODEBUG` 是一个非空字符串，就返回 `True`，否则就返回 `False`。

3.4.2 新版功能。

`AbstractEventLoop.set_debug(enabled: bool)`
设置事件循环的调试模式。

3.4.2 新版功能。

参见：

debug mode of asyncio.

Server

class `asyncio.Server`
Server listening on sockets.

Object created by the `AbstractEventLoop.create_server()` method and the `start_server()` function. Don't instantiate the class directly.

close()
停止服务：关闭监听的套接字并且设置 `sockets` 属性为 `None`。

用于表示已经连进来的客户端连接会保持打开的状态。

服务器是被异步关闭的，使用 `wait_closed()` 协程来等待服务器关闭。

coroutine wait_closed()
等待 `close()` 方法执行完毕。

This method is a *coroutine*.

sockets
List of `socket.socket` objects the server is listening to, or `None` if the server is closed.

Handle

class `asyncio.Handle`
A callback wrapper object returned by `AbstractEventLoop.call_soon()`, `AbstractEventLoop.call_soon_threadsafe()`, `AbstractEventLoop.call_later()`, and `AbstractEventLoop.call_at()`.

cancel()
Cancel the call. If the callback is already canceled or executed, this method has no effect.

Event loop examples

call_soon() 的 Hello World 示例。

Example using the `AbstractEventLoop.call_soon()` method to schedule a callback. The callback displays "Hello World" and then stops the event loop:

```
import asyncio

def hello_world(loop):
    print('Hello World')
    loop.stop()

loop = asyncio.get_event_loop()

# Schedule a call to hello_world()
loop.call_soon(hello_world, loop)

# Blocking call interrupted by loop.stop()
loop.run_forever()
loop.close()
```

参见:

The *Hello World coroutine* example uses a *coroutine*.

使用 `call_later()` 来展示当前的日期

Example of callback displaying the current date every second. The callback uses the *AbstractEventLoop.call_later()* method to reschedule itself during 5 seconds, and then stops the event loop:

```
import asyncio
import datetime

def display_date(end_time, loop):
    print(datetime.datetime.now())
    if (loop.time() + 1.0) < end_time:
        loop.call_later(1, display_date, end_time, loop)
    else:
        loop.stop()

loop = asyncio.get_event_loop()

# Schedule the first call to display_date()
end_time = loop.time() + 5.0
loop.call_soon(display_date, end_time, loop)

# Blocking call interrupted by loop.stop()
loop.run_forever()
loop.close()
```

参见:

The *coroutine displaying the current date* example uses a *coroutine*.

监控一个文件描述符的读事件

Wait until a file descriptor received some data using the *AbstractEventLoop.add_reader()* method and then close the event loop:

```
import asyncio
try:
```

(下页继续)

(续上页)

```

    from socket import socketpair
except ImportError:
    from asyncio.windows_utils import socketpair

# Create a pair of connected file descriptors
rsock, wsock = socketpair()
loop = asyncio.get_event_loop()

def reader():
    data = rsock.recv(100)
    print("Received:", data.decode())
    # We are done: unregister the file descriptor
    loop.remove_reader(rsock)
    # Stop the event loop
    loop.stop()

# Register the file descriptor for read event
loop.add_reader(rsock, reader)

# Simulate the reception of data from the network
loop.call_soon(wsock.send, 'abc'.encode())

# Run the event loop
loop.run_forever()

# We are done, close sockets and the event loop
rsock.close()
wsock.close()
loop.close()

```

参见:

The *register an open socket to wait for data using a protocol* example uses a low-level protocol created by the `AbstractEventLoop.create_connection()` method.

The *register an open socket to wait for data using streams* example uses high-level streams created by the `open_connection()` function in a coroutine.

为 SIGINT 和 SIGTERM 设置信号处理器

Register handlers for signals SIGINT and SIGTERM using the `AbstractEventLoop.add_signal_handler()` method:

```

import asyncio
import functools
import os
import signal

def ask_exit(signame):
    print("got signal %s: exit" % signame)
    loop.stop()

loop = asyncio.get_event_loop()
for signame in ('SIGINT', 'SIGTERM'):
    loop.add_signal_handler(getattr(signal, signame),

```

(下页继续)

(续上页)

```

                                functools.partial(ask_exit, signame))

print("Event loop running forever, press Ctrl+C to interrupt.")
print("pid %s: send SIGINT or SIGTERM to exit." % os.getpid())
try:
    loop.run_forever()
finally:
    loop.close()

```

This example only works on UNIX.

18.5.2 事件循环

源码: `Lib/asyncio/events.py`

事件循环函数

下面的函数是访问全局策略方法的便利快捷方式。注意这只是提供了对默认策略的访问，除非要在进程开始之前设置替代的策略。

```

asyncio.get_event_loop()
    与调用 get_event_loop_policy().get_event_loop() 等价。

asyncio.set_event_loop(loop)
    等价于调用 get_event_loop_policy().set_event_loop(loop)。

asyncio.new_event_loop()
    等价于调用 get_event_loop_policy().new_event_loop()。

```

可用的事件循环

`asyncio` 当前提供两种事件循环实现: *SelectorEventLoop* 和 *ProactorEventLoop*。

class `asyncio.SelectorEventLoop`

基于 *selectors* 模块的事件循环。*AbstractEventLoop* 的子类。

使用对应平台上最高效的选择器。

在 Windows 上，只支持套接字（举例：不支持管线）：参见 *MSDN* 关于 *select* 的文档 <https://msdn.microsoft.com/en-us/library/windows/desktop/ms740141%28v=vs.85%29.aspx>。

class `asyncio.ProactorEventLoop`

Windows 上的 Proactor 事件循环，使用的是“I/O 完成端口”，也叫做 IOCP。为 *AbstractEventLoop* 的子类。

可用性: Windows。

参见:

[MSDN 上关于 I/O 完成端口的文档](#)。

在 Windows 上使用 *ProactorEventLoop* 的例子:

```
import asyncio, sys

if sys.platform == 'win32':
    loop = asyncio.ProactorEventLoop()
    asyncio.set_event_loop(loop)
```

平台支持

`asyncio` 模块被设计为可移植的，但是每个平台仍旧有细微的不同，而且不一定支持 `asyncio` 的所有特性。

Windows

Windows 事件循环的一般限制：

- 不支持 `create_unix_connection()` 和 `create_unix_server()`： `socket.AF_UNIX` 套接字族只在 UNIX 上可用
- 不支持 `add_signal_handler()` 和 `remove_signal_handler()`
- 不支持 `EventLoopPolicy.set_child_watcher()`。`ProactorEventLoop` 支持子进程。监控子进程的实现只有一种，不需要对其进行设置。

特定于 `SelectorEventLoop` 的限制：

- `SelectSelector` 只被用于支持套接字，并且最多支持 512 个套接字。
- `add_reader()` 和 `add_writer()` 只接受套接字的文件描述符
- 不支持管线（示例： `connect_read_pipe()`, `connect_write_pipe()`）
- 不支持 `Subprocesses`（示例： `subprocess_exec()`, `subprocess_shell()`）

特定于 `ProactorEventLoop` 的限制：

- 不支持 `create_datagram_endpoint()` (UDP)
- 不支持 `add_reader()` 和 `add_writer()`

Windows 上单调时钟的分辨率大约为 15.6 毫秒。最佳的分辨率是 0.5 毫秒。分辨率依赖于具体的硬件（HPET 的可用性）和 Windows 的设置。参见：[ref:asyncio delayed calls <asyncio-delayed-calls>](#)。

在 3.5 版更改：`ProactorEventLoop` 现在支持 SSL 了。

Mac OS X

类似于 PTY 的字符设备在 Mavericks（Mac OS 10.9）之前并未有很好的支持。在 Mac OS 10.5 以及更老的版本上则根本不支持。

在 Mac OS 10.6、10.7 和 10.8 上，默认的事件循环是 `SelectorEventLoop`，其使用了 `selectors.KqueueSelector`。`selectors.KqueueSelector` 在这些版本上并不支持字符设备。`SelectorEventLoop` 可用同 `SelectSelector` 或者 `PollSelector` 一同使用来在这些版本的 Mac OS X 上支持块设备。例如：

```
import asyncio
import selectors

selector = selectors.SelectSelector()
```

(下页继续)

(续上页)

```
loop = asyncio.SelectorEventLoop(selector)
asyncio.set_event_loop(loop)
```

事件循环策略和默认策略

使用一个策略模式对事件循环进行管理，用以为特定平台和框架提供最大限度的灵活性。在一个进程的执行中，唯一的全局策略对象，基于调用的上下文，管理着此进程所有可用的事件循环。一个策略就是一个实现了 *AbstractEventLoopPolicy* 的对象。

对于 *asyncio* 的多数使用者来说，从来都不需要显式地处理策略，因为默认的全局策略就已经足够了（看下面）。

模块级别的函数 *get_event_loop()* 和 *set_event_loop()* 对于默认策略管理的事件循环的便利访问方式。

事件循环策略接口

一个事件循环必须实现如下的接口：

```
class asyncio.AbstractEventLoopPolicy
```

事件循环策略。

```
get_event_loop()
```

为当前上下文获取事件循环。

返回一个实现了 *AbstractEventLoop* 接口的事件循环对象。如果是从一个协程中调用，就会返回当前正在运行的事件循环。

如果没有为当前上下文设置一个事件循环，并且当前策略不能创建一个事件循环，就会抛出一个错误。其务必不能返回 *None*。

在 3.6 版更改。

```
set_event_loop(loop)
```

将当前上下文的事件循环设置为 *loop*。

```
new_event_loop()
```

根据策略的规则，创建并返回一个新的事件循环对象。

如果有需要设置这个循环作为当前上下文的事件循环，就必须显式调用 *set_event_loop()*。

The default policy defines context as the current thread, and manages an event loop per thread that interacts with *asyncio*. If the current thread doesn't already have an event loop associated with it, the default policy's *get_event_loop()* method creates one when called from the main thread, but raises *RuntimeError* otherwise.

访问全局循环策略

```
asyncio.get_event_loop_policy()
```

获取当前事件循环策略。

```
asyncio.set_event_loop_policy(policy)
```

设置当前事件循环策略。如果 *policy* 为 *None*，就会恢复使用默认的策略。

自定义事件循环策略

要实现一个新的事件循环策略，推荐你继承具体的默认事件循环策略 `DefaultEventLoopPolicy` 并且重写你想要改变行为的方法，例如：

```
class MyEventLoopPolicy(asyncio.DefaultEventLoopPolicy):

    def get_event_loop(self):
        """Get the event loop.

        This may be None or an instance of EventLoop.
        """
        loop = super().get_event_loop()
        # Do something with loop ...
        return loop

asyncio.set_event_loop_policy(MyEventLoopPolicy())
```

18.5.3 Tasks and coroutines

Source code: `Lib/asyncio/tasks.py`

Source code: `Lib/asyncio/coroutines.py`

协程

Coroutines used with `asyncio` may be implemented using the `async def` statement, or by using *generators*. The `async def` type of coroutine was added in Python 3.5, and is recommended if there is no need to support older Python versions.

Generator-based coroutines should be decorated with `@asyncio.coroutine`, although this is not strictly enforced. The decorator enables compatibility with `async def` coroutines, and also serves as documentation. Generator-based coroutines use the `yield from` syntax introduced in **PEP 380**, instead of the original `yield` syntax.

The word “coroutine”, like the word “generator”, is used for two different (though related) concepts:

- The function that defines a coroutine (a function definition using `async def` or decorated with `@asyncio.coroutine`). If disambiguation is needed we will call this a *coroutine function* (`iscoroutinefunction()` returns True).
- The object obtained by calling a coroutine function. This object represents a computation or an I/O operation (usually a combination) that will complete eventually. If disambiguation is needed we will call it a *coroutine object* (`iscoroutine()` returns True).

Things a coroutine can do:

- `result = await future` or `result = yield from future` – suspends the coroutine until the future is done, then returns the future’s result, or raises an exception, which will be propagated. (If the future is cancelled, it will raise a `CancelledError` exception.) Note that tasks are futures, and everything said about futures also applies to tasks.
- `result = await coroutine` or `result = yield from coroutine` – wait for another coroutine to produce a result (or raise an exception, which will be propagated). The `coroutine` expression must be a *call* to another coroutine.
- `return expression` – produce a result to the coroutine that is waiting for this one using `await` or `yield from`.

- `raise exception` – raise an exception in the coroutine that is waiting for this one using `await` or `yield from`.

Calling a coroutine does not start its code running – the coroutine object returned by the call doesn't do anything until you schedule its execution. There are two basic ways to start it running: call `await coroutine` or `yield from coroutine` from another coroutine (assuming the other coroutine is already running!), or schedule its execution using the `ensure_future()` function or the `AbstractEventLoop.create_task()` method.

Coroutines (and tasks) can only run when the event loop is running.

`@asyncio.coroutine`

Decorator to mark generator-based coroutines. This enables the generator use `yield from` to call `async def` coroutines, and also enables the generator to be called by `async def` coroutines, for instance using an `await` expression.

There is no need to decorate `async def` coroutines themselves.

If the generator is not yielded from before it is destroyed, an error message is logged. See *Detect coroutines never scheduled*.

注解: In this documentation, some methods are documented as coroutines, even if they are plain Python functions returning a `Future`. This is intentional to have a freedom of tweaking the implementation of these functions in the future. If such a function is needed to be used in a callback-style code, wrap its result with `ensure_future()`.

Example: Hello World coroutine

Example of coroutine displaying "Hello World":

```
import asyncio

async def hello_world():
    print("Hello World!")

loop = asyncio.get_event_loop()
# Blocking call which returns when the hello_world() coroutine is done
loop.run_until_complete(hello_world())
loop.close()
```

参见:

The *Hello World with call_soon()* example uses the `AbstractEventLoop.call_soon()` method to schedule a callback.

Example: Coroutine displaying the current date

Example of coroutine displaying the current date every second during 5 seconds using the `sleep()` function:

```
import asyncio
import datetime

async def display_date(loop):
    end_time = loop.time() + 5.0
    while True:
        print(datetime.datetime.now())
```

(下页继续)

(续上页)

```
        if (loop.time() + 1.0) >= end_time:
            break
        await asyncio.sleep(1)

loop = asyncio.get_event_loop()
# Blocking call which returns when the display_date() coroutine is done
loop.run_until_complete(display_date(loop))
loop.close()
```

参见:

The *display the current date with call_later()* example uses a callback with the *AbstractEventLoop.call_later()* method.

Example: Chain coroutines

Example chaining coroutines:

```
import asyncio

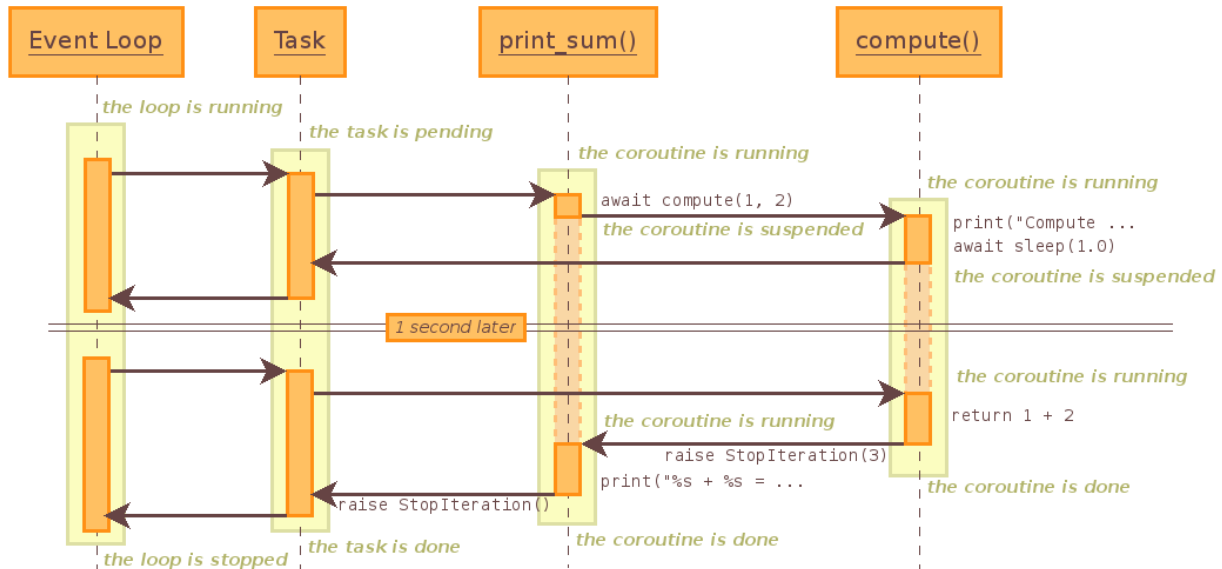
async def compute(x, y):
    print("Compute %s + %s ..." % (x, y))
    await asyncio.sleep(1.0)
    return x + y

async def print_sum(x, y):
    result = await compute(x, y)
    print("%s + %s = %s" % (x, y, result))

loop = asyncio.get_event_loop()
loop.run_until_complete(print_sum(1, 2))
loop.close()
```

`compute()` is chained to `print_sum()`: `print_sum()` coroutine waits until `compute()` is completed before returning its result.

Sequence diagram of the example:



The “Task” is created by the `AbstractEventLoop.run_until_complete()` method when it gets a coroutine object instead of a task.

The diagram shows the control flow, it does not describe exactly how things work internally. For example, the sleep coroutine creates an internal future which uses `AbstractEventLoop.call_later()` to wake up the task in 1 second.

InvalidStateError

exception `asyncio.InvalidStateError`
The operation is not allowed in this state.

TimeoutError

exception `asyncio.TimeoutError`
The operation exceeded the given deadline.

注解: This exception is different from the builtin `TimeoutError` exception!

Future

class `asyncio.Future` (*, loop=None)
This class is *almost* compatible with `concurrent.futures.Future`.

Differences:

- `result()` and `exception()` do not take a timeout argument and raise an exception when the future isn't done yet.
- Callbacks registered with `add_done_callback()` are always called via the event loop's `call_soon()`.

- This class is not compatible with the `wait()` and `as_completed()` functions in the `concurrent.futures` package.

This class is *not thread safe*.

cancel()

Cancel the future and schedule callbacks.

If the future is already done or cancelled, return `False`. Otherwise, change the future's state to cancelled, schedule the callbacks and return `True`.

cancelled()

Return `True` if the future was cancelled.

done()

Return `True` if the future is done.

Done means either that a result / exception are available, or that the future was cancelled.

result()

Return the result this future represents.

If the future has been cancelled, raises `CancelledError`. If the future's result isn't yet available, raises `InvalidStateError`. If the future is done and has an exception set, this exception is raised.

exception()

Return the exception that was set on this future.

The exception (or `None` if no exception was set) is returned only if the future is done. If the future has been cancelled, raises `CancelledError`. If the future isn't done yet, raises `InvalidStateError`.

add_done_callback(fn)

Add a callback to be run when the future becomes done.

The callback is called with a single argument - the future object. If the future is already done when this is called, the callback is scheduled with `call_soon()`.

Use `functools.partial` to pass parameters to the callback. For example, `fut.add_done_callback(functools.partial(print, "Future:", flush=True))` will call `print("Future:", fut, flush=True)`.

remove_done_callback(fn)

Remove all instances of a callback from the "call when done" list.

Returns the number of callbacks removed.

set_result(result)

Mark the future done and set its result.

If the future is already done when this method is called, raises `InvalidStateError`.

set_exception(exception)

Mark the future done and set an exception.

If the future is already done when this method is called, raises `InvalidStateError`.

Example: Future with `run_until_complete()`

Example combining a *Future* and a *coroutine function*:

```
import asyncio

async def slow_operation(future):
    await asyncio.sleep(1)
    future.set_result('Future is done!')

loop = asyncio.get_event_loop()
future = asyncio.Future()
asyncio.ensure_future(slow_operation(future))
loop.run_until_complete(future)
print(future.result())
loop.close()
```

The coroutine function is responsible for the computation (which takes 1 second) and it stores the result into the future. The `run_until_complete()` method waits for the completion of the future.

注解: The `run_until_complete()` method uses internally the `add_done_callback()` method to be notified when the future is done.

Example: Future with `run_forever()`

The previous example can be written differently using the `Future.add_done_callback()` method to describe explicitly the control flow:

```
import asyncio

async def slow_operation(future):
    await asyncio.sleep(1)
    future.set_result('Future is done!')

def got_result(future):
    print(future.result())
    loop.stop()

loop = asyncio.get_event_loop()
future = asyncio.Future()
asyncio.ensure_future(slow_operation(future))
future.add_done_callback(got_result)
try:
    loop.run_forever()
finally:
    loop.close()
```

In this example, the future is used to link `slow_operation()` to `got_result()`: when `slow_operation()` is done, `got_result()` is called with the result.

Task

class `asyncio.Task` (*coro*, *, *loop=None*)

Schedule the execution of a *coroutine*: wrap it in a future. A task is a subclass of *Future*.

A task is responsible for executing a coroutine object in an event loop. If the wrapped coroutine yields from a future, the task suspends the execution of the wrapped coroutine and waits for the completion of the future. When

the future is done, the execution of the wrapped coroutine restarts with the result or the exception of the future.

Event loops use cooperative scheduling: an event loop only runs one task at a time. Other tasks may run in parallel if other event loops are running in different threads. While a task waits for the completion of a future, the event loop executes a new task.

The cancellation of a task is different from the cancellation of a future. Calling `cancel()` will throw a `CancelledError` to the wrapped coroutine. `cancelled()` only returns `True` if the wrapped coroutine did not catch the `CancelledError` exception, or raised a `CancelledError` exception.

If a pending task is destroyed, the execution of its wrapped *coroutine* did not complete. It is probably a bug and a warning is logged: see *Pending task destroyed*.

Don't directly create `Task` instances: use the `ensure_future()` function or the `AbstractEventLoop.create_task()` method.

This class is *not thread safe*.

classmethod `all_tasks(loop=None)`

返回一个事件循环中所有任务的集合。

By default all tasks for the current event loop are returned.

classmethod `current_task(loop=None)`

Return the currently running task in an event loop or `None`.

By default the current task for the current event loop is returned.

`None` is returned when called not in the context of a `Task`.

cancel()

Request that this task cancel itself.

This arranges for a `CancelledError` to be thrown into the wrapped coroutine on the next cycle through the event loop. The coroutine then has a chance to clean up or even deny the request using `try/except/finally`.

Unlike `Future.cancel()`, this does not guarantee that the task will be cancelled: the exception might be caught and acted upon, delaying cancellation of the task or preventing cancellation completely. The task may also return a value or raise a different exception.

Immediately after this method is called, `cancelled()` will not return `True` (unless the task was already cancelled). A task will be marked as cancelled when the wrapped coroutine terminates with a `CancelledError` exception (even if `cancel()` was not called).

get_stack(*, limit=None)

Return the list of stack frames for this task's coroutine.

If the coroutine is not done, this returns the stack where it is suspended. If the coroutine has completed successfully or was cancelled, this returns an empty list. If the coroutine was terminated by an exception, this returns the list of traceback frames.

框架总是从按从旧到新排序。

The optional `limit` gives the maximum number of frames to return; by default all available frames are returned. Its meaning differs depending on whether a stack or a traceback is returned: the newest frames of a stack are returned, but the oldest frames of a traceback are returned. (This matches the behavior of the traceback module.)

For reasons beyond our control, only one stack frame is returned for a suspended coroutine.

print_stack(*, limit=None, file=None)

Print the stack or traceback for this task's coroutine.

This produces output similar to that of the `traceback` module, for the frames retrieved by `get_stack()`. The `limit` argument is passed to `get_stack()`. The `file` argument is an I/O stream to which the output is written; by default output is written to `sys.stderr`.

Example: Parallel execution of tasks

Example executing 3 tasks (A, B, C) in parallel:

```
import asyncio

async def factorial(name, number):
    f = 1
    for i in range(2, number+1):
        print("Task %s: Compute factorial(%s)..." % (name, i))
        await asyncio.sleep(1)
        f *= i
    print("Task %s: factorial(%s) = %s" % (name, number, f))

loop = asyncio.get_event_loop()
loop.run_until_complete(asyncio.gather(
    factorial("A", 2),
    factorial("B", 3),
    factorial("C", 4),
))
loop.close()
```

Output:

```
Task A: Compute factorial(2)...
Task B: Compute factorial(2)...
Task C: Compute factorial(2)...
Task A: factorial(2) = 2
Task B: Compute factorial(3)...
Task C: Compute factorial(3)...
Task B: factorial(3) = 6
Task C: Compute factorial(4)...
Task C: factorial(4) = 24
```

A task is automatically scheduled for execution when it is created. The event loop stops when all tasks are done.

Task functions

注解: In the functions below, the optional `loop` argument allows explicitly setting the event loop object used by the underlying task or coroutine. If it's not provided, the default event loop is used.

`asyncio.as_completed(fs, *, loop=None, timeout=None)`

Return an iterator whose values, when waited for, are `Future` instances.

如果在所有 `Future` 对象完成前发生超时则将引发 `asyncio.TimeoutError`。

示例:

```
for f in as_completed(fs):
    result = yield from f # The 'yield from' may raise
    # Use result
```

注解: The futures `f` are not necessarily members of `fs`.

`asyncio.ensure_future(coro_or_future, *, loop=None)`

Schedule the execution of a *coroutine object*: wrap it in a future. Return a *Task* object.

If the argument is a *Future*, it is returned directly.

3.4.4 新版功能.

在 3.5.1 版更改: The function accepts any *awaitable* object.

参见:

The `AbstractEventLoop.create_task()` method.

`asyncio.async(coro_or_future, *, loop=None)`

A deprecated alias to `ensure_future()`.

3.4.4 版后已移除.

`asyncio.wrap_future(future, *, loop=None)`

Wrap a `concurrent.futures.Future` object in a *Future* object.

`asyncio.gather(*coros_or_futures, loop=None, return_exceptions=False)`

Return a future aggregating results from the given coroutine objects or futures.

All futures must share the same event loop. If all the tasks are done successfully, the returned future's result is the list of results (in the order of the original sequence, not necessarily the order of results arrival). If `return_exceptions` is true, exceptions in the tasks are treated the same as successful results, and gathered in the result list; otherwise, the first raised exception will be immediately propagated to the returned future.

Cancellation: if the outer Future is cancelled, all children (that have not completed yet) are also cancelled. If any child is cancelled, this is treated as if it raised `CancelledError` – the outer Future is *not* cancelled in this case. (This is to prevent the cancellation of one child to cause other children to be cancelled.)

在 3.6.6 版更改: 如果 `gather` 本身被取消, 则无论 `return_exceptions` 取值为何, 消息都会被传播。

`asyncio.iscoroutine(obj)`

Return True if `obj` is a *coroutine object*, which may be based on a generator or an `async def` coroutine.

`asyncio.iscoroutinefunction(func)`

Return True if `func` is determined to be a *coroutine function*, which may be a decorated generator function or an `async def` function.

`asyncio.run_coroutine_threadsafe(coro, loop)`

Submit a *coroutine object* to a given event loop.

Return a `concurrent.futures.Future` to access the result.

This function is meant to be called from a different thread than the one where the event loop is running. Usage:

```
# Create a coroutine
coro = asyncio.sleep(1, result=3)
# Submit the coroutine to a given loop
future = asyncio.run_coroutine_threadsafe(coro, loop)
# Wait for the result with an optional timeout argument
assert future.result(timeout) == 3
```

If an exception is raised in the coroutine, the returned future will be notified. It can also be used to cancel the task in the event loop:

```
try:
    result = future.result(timeout)
except asyncio.TimeoutError:
    print('The coroutine took too long, cancelling the task...')
    future.cancel()
except Exception as exc:
    print('The coroutine raised an exception: {!r}'.format(exc))
else:
    print('The coroutine returned: {!r}'.format(result))
```

查看并发和多线程 章节的文档。

注解： Unlike other functions from the module, `run_coroutine_threadsafe()` requires the `loop` argument to be passed explicitly.

3.5.1 新版功能.

coroutine `asyncio.sleep(delay, result=None, *, loop=None)`

Create a *coroutine* that completes after a given time (in seconds). If *result* is provided, it is produced to the caller when the coroutine completes.

The resolution of the sleep depends on the *granularity of the event loop*.

This function is a *coroutine*.

coroutine `asyncio.shield(arg, *, loop=None)`

Wait for a future, shielding it from cancellation.

以下语句:

```
res = yield from shield(something())
```

is exactly equivalent to the statement:

```
res = yield from something()
```

except that if the coroutine containing it is cancelled, the task running in `something()` is not cancelled. From the point of view of `something()`, the cancellation did not happen. But its caller is still cancelled, so the `yield-from` expression still raises `CancelledError`. Note: If `something()` is cancelled by other means this will still cancel `shield()`.

If you want to completely ignore cancellation (not recommended) you can combine `shield()` with a `try/except` clause, as follows:

```
try:
    res = yield from shield(something())
except CancelledError:
    res = None
```

coroutine `asyncio.wait(futures, *, loop=None, timeout=None, return_when=ALL_COMPLETED)`

Wait for the Futures and coroutine objects given by the sequence *futures* to complete. Coroutines will be wrapped in Tasks. Returns two sets of *Future*: (done, pending).

The sequence *futures* must not be empty.

timeout can be used to control the maximum number of seconds to wait before returning. *timeout* can be an int or float. If *timeout* is not specified or `None`, there is no limit to the wait time.

return_when indicates when this function should return. It must be one of the following constants of the `concurrent.futures` module:

常数	描述
<code>FIRST_COMPLETED</code>	函数将在任意可等待对象结束或取消时返回。
<code>FIRST_EXCEPTION</code>	函数将在任意可等待对象因引发异常而结束时返回。当没有引发任何异常时它就相当于 <code>ALL_COMPLETED</code> 。
<code>ALL_COMPLETED</code>	函数将在所有可等待对象结束或取消时返回。

This function is a *coroutine*.

用法:

```
done, pending = yield from asyncio.wait(fs)
```

注解: This does not raise `asyncio.TimeoutError`! Futures that aren't done when the timeout occurs are returned in the second set.

coroutine `asyncio.wait_for(fut, timeout, *, loop=None)`

Wait for the single *Future* or *coroutine object* to complete with timeout. If *timeout* is `None`, block until the future completes.

Coroutine will be wrapped in *Task*.

Returns result of the Future or coroutine. When a timeout occurs, it cancels the task and raises `asyncio.TimeoutError`. To avoid the task cancellation, wrap it in `shield()`.

If the wait is cancelled, the future *fut* is also cancelled.

This function is a *coroutine*, usage:

```
result = yield from asyncio.wait_for(fut, 60.0)
```

在 3.4.3 版更改: If the wait is cancelled, the future *fut* is now also cancelled.

18.5.4 Transports and protocols (callback based API)

Source code: `Lib/asyncio/transports.py`

Source code: `Lib/asyncio/protocols.py`

传输

Transports are classes provided by `asyncio` in order to abstract various kinds of communication channels. You generally won't instantiate a transport yourself; instead, you will call an `AbstractEventLoop` method which will create the transport and try to initiate the underlying communication channel, calling you back when it succeeds.

Once the communication channel is established, a transport is always paired with a *protocol* instance. The protocol can then call the transport's methods for various purposes.

`asyncio` currently implements transports for TCP, UDP, SSL, and subprocess pipes. The methods available on a transport depend on the transport's kind.

传输类属于线程不安全。

在 3.6 版更改: The socket option `TCP_NODELAY` is now set by default.

BaseTransport

class `asyncio.BaseTransport`

Base class for transports.

close()

Close the transport. If the transport has a buffer for outgoing data, buffered data will be flushed asynchronously. No more data will be received. After all buffered data is flushed, the protocol's `connection_lost()` method will be called with `None` as its argument.

is_closing()

返回 `True`，如果传输正在关闭或已经关闭。

3.5.1 新版功能.

get_extra_info(*name*, *default=None*)

Return optional transport information. *name* is a string representing the piece of transport-specific information to get, *default* is the value to return if the information doesn't exist.

This method allows transport implementations to easily expose channel-specific information.

- socket:
 - 'peername': the remote address to which the socket is connected, result of `socket.socket.getpeername()` (None on error)
 - 'socket': `socket.socket` instance
 - 'sockname': the socket's own address, result of `socket.socket.getsockname()`
- SSL socket:
 - 'compression': the compression algorithm being used as a string, or None if the connection isn't compressed; result of `ssl.SSLSocket.compression()`
 - 'cipher': a three-value tuple containing the name of the cipher being used, the version of the SSL protocol that defines its use, and the number of secret bits being used; result of `ssl.SSLSocket.cipher()`
 - 'peercert': peer certificate; result of `ssl.SSLSocket.getpeercert()`
 - 'sslcontext': `ssl.SSLContext` instance
 - 'ssl_object': `ssl.SSLObject` or `ssl.SSLSocket` instance
- pipe:
 - 'pipe': pipe object
- subprocess:
 - 'subprocess': `subprocess.Popen` instance

set_protocol(*protocol*)

Set a new protocol. Switching protocol should only be done when both protocols are documented to support the switch.

3.5.3 新版功能.

get_protocol()
返回当前协议。

3.5.3 新版功能。

在 3.5.1 版更改: 'ssl_object' info was added to SSL sockets.

ReadTransport

class `asyncio.ReadTransport`

Interface for read-only transports.

pause_reading()
Pause the receiving end of the transport. No data will be passed to the protocol's `data_received()` method until `resume_reading()` is called.

在 3.6.7 版更改: The method is idempotent, i.e. it can be called when the transport is already paused or closed.

resume_reading()
Resume the receiving end. The protocol's `data_received()` method will be called once again if some data is available for reading.

在 3.6.7 版更改: The method is idempotent, i.e. it can be called when the transport is already reading.

WriteTransport

class `asyncio.WriteTransport`

Interface for write-only transports.

abort()
Close the transport immediately, without waiting for pending operations to complete. Buffered data will be lost. No more data will be received. The protocol's `connection_lost()` method will eventually be called with `None` as its argument.

can_write_eof()
Return `True` if the transport supports `write_eof()`, `False` if not.

get_write_buffer_size()
Return the current size of the output buffer used by the transport.

get_write_buffer_limits()
Get the *high*- and *low*-water limits for write flow control. Return a tuple (*low*, *high*) where *low* and *high* are positive number of bytes.

Use `set_write_buffer_limits()` to set the limits.

3.4.2 新版功能。

set_write_buffer_limits (*high=None*, *low=None*)
Set the *high*- and *low*-water limits for write flow control.

These two values (measured in number of bytes) control when the protocol's `pause_writing()` and `resume_writing()` methods are called. If specified, the low-water limit must be less than or equal to the high-water limit. Neither *high* nor *low* can be negative.

`pause_writing()` is called when the buffer size becomes greater than or equal to the *high* value. If writing has been paused, `resume_writing()` is called when the buffer size becomes less than or equal to the *low* value.

The defaults are implementation-specific. If only the high-water limit is given, the low-water limit defaults to an implementation-specific value less than or equal to the high-water limit. Setting *high* to zero forces *low* to zero as well, and causes `pause_writing()` to be called whenever the buffer becomes non-empty. Setting *low* to zero causes `resume_writing()` to be called only once the buffer is empty. Use of zero for either limit is generally sub-optimal as it reduces opportunities for doing I/O and computation concurrently.

Use `get_write_buffer_limits()` to get the limits.

write (*data*)

Write some *data* bytes to the transport.

This method does not block; it buffers the data and arranges for it to be sent out asynchronously.

writelines (*list_of_data*)

Write a list (or any iterable) of data bytes to the transport. This is functionally equivalent to calling `write()` on each element yielded by the iterable, but may be implemented more efficiently.

write_eof ()

Close the write end of the transport after flushing buffered data. Data may still be received.

This method can raise `NotImplementedError` if the transport (e.g. SSL) doesn't support half-closes.

DatagramTransport

`DatagramTransport.sendto` (*data*, *addr=None*)

Send the *data* bytes to the remote peer given by *addr* (a transport-dependent target address). If *addr* is *None*, the data is sent to the target address given on transport creation.

This method does not block; it buffers the data and arranges for it to be sent out asynchronously.

`DatagramTransport.abort` ()

Close the transport immediately, without waiting for pending operations to complete. Buffered data will be lost. No more data will be received. The protocol's `connection_lost()` method will eventually be called with *None* as its argument.

BaseSubprocessTransport

class `asyncio.BaseSubprocessTransport`

get_pid ()

Return the subprocess process id as an integer.

get_pipe_transport (*fd*)

Return the transport for the communication pipe corresponding to the integer file descriptor *fd*:

- 0: readable streaming transport of the standard input (*stdin*), or *None* if the subprocess was not created with `stdin=PIPE`
- 1: writable streaming transport of the standard output (*stdout*), or *None* if the subprocess was not created with `stdout=PIPE`
- 2: writable streaming transport of the standard error (*stderr*), or *None* if the subprocess was not created with `stderr=PIPE`
- other *fd*: *None*

get_returncode()

Return the subprocess returncode as an integer or *None* if it hasn't returned, similarly to the *subprocess.Popen.returncode* attribute.

kill()

Kill the subprocess, as in *subprocess.Popen.kill()*.

On POSIX systems, the function sends SIGKILL to the subprocess. On Windows, this method is an alias for *terminate()*.

send_signal(signal)

Send the *signal* number to the subprocess, as in *subprocess.Popen.send_signal()*.

terminate()

Ask the subprocess to stop, as in *subprocess.Popen.terminate()*. This method is an alias for the *close()* method.

On POSIX systems, this method sends SIGTERM to the subprocess. On Windows, the Windows API function *TerminateProcess()* is called to stop the subprocess.

close()

Ask the subprocess to stop by calling the *terminate()* method if the subprocess hasn't returned yet, and close transports of all pipes (*stdin*, *stdout* and *stderr*).

协议

asyncio provides base classes that you can subclass to implement your network protocols. Those classes are used in conjunction with *transports* (see below): the protocol parses incoming data and asks for the writing of outgoing data, while the transport is responsible for the actual I/O and buffering.

When subclassing a protocol class, it is recommended you override certain methods. Those methods are callbacks: they will be called by the transport on certain events (for example when some data is received); you shouldn't call them yourself, unless you are implementing a transport.

注解: All callbacks have default implementations, which are empty. Therefore, you only need to implement the callbacks for the events in which you are interested.

Protocol classes

class asyncio.Protocol

The base class for implementing streaming protocols (for use with e.g. TCP and SSL transports).

class asyncio.DatagramProtocol

The base class for implementing datagram protocols (for use with e.g. UDP transports).

class asyncio.SubprocessProtocol

The base class for implementing protocols communicating with child processes (through a set of unidirectional pipes).

Connection callbacks

These callbacks may be called on *Protocol*, *DatagramProtocol* and *SubprocessProtocol* instances:

`BaseProtocol.connection_made(transport)`

链接建立时被调用。

The *transport* argument is the transport representing the connection. You are responsible for storing it somewhere (e.g. as an attribute) if you need to.

`BaseProtocol.connection_lost(exc)`

链接丢失或关闭时被调用。

The argument is either an exception object or *None*. The latter means a regular EOF is received, or the connection was aborted or closed by this side of the connection.

`connection_made()` and `connection_lost()` are called exactly once per successful connection. All other callbacks will be called between those two methods, which allows for easier resource management in your protocol implementation.

The following callbacks may be called only on *SubprocessProtocol* instances:

`SubprocessProtocol.pipe_data_received(fd, data)`

Called when the child process writes data into its stdout or stderr pipe. *fd* is the integer file descriptor of the pipe. *data* is a non-empty bytes object containing the data.

`SubprocessProtocol.pipe_connection_lost(fd, exc)`

Called when one of the pipes communicating with the child process is closed. *fd* is the integer file descriptor that was closed.

`SubprocessProtocol.process_exited()`

子进程退出时被调用。

Streaming protocols

The following callbacks are called on *Protocol* instances:

`Protocol.data_received(data)`

Called when some data is received. *data* is a non-empty bytes object containing the incoming data.

注解: Whether the data is buffered, chunked or reassembled depends on the transport. In general, you shouldn't rely on specific semantics and instead make your parsing generic and flexible enough. However, data is always received in the correct order.

`Protocol.eof_received()`

Called when the other end signals it won't send any more data (for example by calling `write_eof()`, if the other end also uses `asyncio`).

This method may return a false value (including *None*), in which case the transport will close itself. Conversely, if this method returns a true value, closing the transport is up to the protocol. Since the default implementation returns *None*, it implicitly closes the connection.

注解: Some transports such as SSL don't support half-closed connections, in which case returning true from this method will not prevent closing the connection.

`data_received()` can be called an arbitrary number of times during a connection. However, `eof_received()` is called at most once and, if called, `data_received()` won't be called after it.

State machine:

```
start -> connection_made() [-> data_received() *] [-> eof_received() ?] ->
connection_lost() -> end
```

Datagram protocols

The following callbacks are called on *DatagramProtocol* instances.

DatagramProtocol.**datagram_received**(*data*, *addr*)

Called when a datagram is received. *data* is a bytes object containing the incoming data. *addr* is the address of the peer sending the data; the exact format depends on the transport.

DatagramProtocol.**error_received**(*exc*)

Called when a previous send or receive operation raises an *OSError*. *exc* is the *OSError* instance.

This method is called in rare conditions, when the transport (e.g. UDP) detects that a datagram couldn't be delivered to its recipient. In many conditions though, undeliverable datagrams will be silently dropped.

Flow control callbacks

These callbacks may be called on *Protocol*, *DatagramProtocol* and *SubprocessProtocol* instances:

BaseProtocol.**pause_writing**()

Called when the transport's buffer goes over the high-water mark.

BaseProtocol.**resume_writing**()

Called when the transport's buffer drains below the low-water mark.

pause_writing() and *resume_writing()* calls are paired – *pause_writing()* is called once when the buffer goes strictly over the high-water mark (even if subsequent writes increases the buffer size even more), and eventually *resume_writing()* is called once when the buffer size reaches the low-water mark.

注解: If the buffer size equals the high-water mark, *pause_writing()* is not called – it must go strictly over. Conversely, *resume_writing()* is called when the buffer size is equal or lower than the low-water mark. These end conditions are important to ensure that things go as expected when either mark is zero.

注解: On BSD systems (OS X, FreeBSD, etc.) flow control is not supported for *DatagramProtocol*, because send failures caused by writing too many packets cannot be detected easily. The socket always appears 'ready' and excess packets are dropped; an *OSError* with *errno* set to *errno.ENOBUFS* may or may not be raised; if it is raised, it will be reported to *DatagramProtocol.error_received()* but otherwise ignored.

Coroutines and protocols

Coroutines can be scheduled in a protocol method using *ensure_future()*, but there is no guarantee made about the execution order. Protocols are not aware of coroutines created in protocol methods and so will not wait for them.

To have a reliable execution order, use *stream objects* in a coroutine with *yield from*. For example, the *StreamWriter.drain()* coroutine can be used to wait until the write buffer is flushed.

Protocol examples

TCP echo client protocol

TCP echo client using the `AbstractEventLoop.create_connection()` method, send data and wait until the connection is closed:

```
import asyncio

class EchoClientProtocol(asyncio.Protocol):
    def __init__(self, message, loop):
        self.message = message
        self.loop = loop

    def connection_made(self, transport):
        transport.write(self.message.encode())
        print('Data sent: {!r}'.format(self.message))

    def data_received(self, data):
        print('Data received: {!r}'.format(data.decode()))

    def connection_lost(self, exc):
        print('The server closed the connection')
        print('Stop the event loop')
        self.loop.stop()

loop = asyncio.get_event_loop()
message = 'Hello World!'
coro = loop.create_connection(lambda: EchoClientProtocol(message, loop),
                              '127.0.0.1', 8888)
loop.run_until_complete(coro)
loop.run_forever()
loop.close()
```

The event loop is running twice. The `run_until_complete()` method is preferred in this short example to raise an exception if the server is not listening, instead of having to write a short coroutine to handle the exception and stop the running loop. At `run_until_complete()` exit, the loop is no longer running, so there is no need to stop the loop in case of an error.

参见:

The *TCP echo client using streams* example uses the `asyncio.open_connection()` function.

TCP echo server protocol

TCP echo server using the `AbstractEventLoop.create_server()` method, send back received data and close the connection:

```
import asyncio

class EchoServerClientProtocol(asyncio.Protocol):
    def connection_made(self, transport):
        peername = transport.get_extra_info('peername')
        print('Connection from {}'.format(peername))
        self.transport = transport
```

(下页继续)

(续上页)

```

def data_received(self, data):
    message = data.decode()
    print('Data received: {!r}'.format(message))

    print('Send: {!r}'.format(message))
    self.transport.write(data)

    print('Close the client socket')
    self.transport.close()

loop = asyncio.get_event_loop()
# Each client connection will create a new protocol instance
coro = loop.create_server(EchoServerClientProtocol, '127.0.0.1', 8888)
server = loop.run_until_complete(coro)

# Serve requests until Ctrl+C is pressed
print('Serving on {}'.format(server.sockets[0].getsockname()))
try:
    loop.run_forever()
except KeyboardInterrupt:
    pass

# Close the server
server.close()
loop.run_until_complete(server.wait_closed())
loop.close()

```

`Transport.close()` can be called immediately after `WriteTransport.write()` even if data are not sent yet on the socket: both methods are asynchronous. `yield from` is not needed because these transport methods are not coroutines.

参见:

The *TCP echo server using streams* example uses the `asyncio.start_server()` function.

UDP echo client protocol

UDP echo client using the `AbstractEventLoop.create_datagram_endpoint()` method, send data and close the transport when we received the answer:

```

import asyncio

class EchoClientProtocol:
    def __init__(self, message, loop):
        self.message = message
        self.loop = loop
        self.transport = None

    def connection_made(self, transport):
        self.transport = transport
        print('Send:', self.message)
        self.transport.sendto(self.message.encode())

    def datagram_received(self, data, addr):
        print("Received:", data.decode())

```

(下页继续)

(续上页)

```

        print("Close the socket")
        self.transport.close()

    def error_received(self, exc):
        print('Error received:', exc)

    def connection_lost(self, exc):
        print("Socket closed, stop the event loop")
        loop = asyncio.get_event_loop()
        loop.stop()

loop = asyncio.get_event_loop()
message = "Hello World!"
connect = loop.create_datagram_endpoint(
    lambda: EchoClientProtocol(message, loop),
    remote_addr=('127.0.0.1', 9999))
transport, protocol = loop.run_until_complete(connect)
loop.run_forever()
transport.close()
loop.close()

```

UDP echo server protocol

UDP echo server using the `AbstractEventLoop.create_datagram_endpoint()` method, send back received data:

```

import asyncio

class EchoServerProtocol:
    def connection_made(self, transport):
        self.transport = transport

    def datagram_received(self, data, addr):
        message = data.decode()
        print('Received %r from %s' % (message, addr))
        print('Send %r to %s' % (message, addr))
        self.transport.sendto(data, addr)

loop = asyncio.get_event_loop()
print("Starting UDP server")
# One protocol instance will be created to serve all client requests
listen = loop.create_datagram_endpoint(
    EchoServerProtocol, local_addr=('127.0.0.1', 9999))
transport, protocol = loop.run_until_complete(listen)

try:
    loop.run_forever()
except KeyboardInterrupt:
    pass

transport.close()
loop.close()

```

Register an open socket to wait for data using a protocol

Wait until a socket receives data using the `AbstractEventLoop.create_connection()` method with a protocol, and then close the event loop

```
import asyncio
try:
    from socket import socketpair
except ImportError:
    from asyncio.windows_utils import socketpair

# Create a pair of connected sockets
rsock, wsock = socketpair()
loop = asyncio.get_event_loop()

class MyProtocol(asyncio.Protocol):
    transport = None

    def connection_made(self, transport):
        self.transport = transport

    def data_received(self, data):
        print("Received:", data.decode())

        # We are done: close the transport (it will call connection_lost())
        self.transport.close()

    def connection_lost(self, exc):
        # The socket has been closed, stop the event loop
        loop.stop()

# Register the socket to wait for data
connect_coro = loop.create_connection(MyProtocol, sock=rsock)
transport, protocol = loop.run_until_complete(connect_coro)

# Simulate the reception of data from the network
loop.call_soon(wsock.send, 'abc'.encode())

# Run the event loop
loop.run_forever()

# We are done, close sockets and the event loop
rsock.close()
wsock.close()
loop.close()
```

参见:

The *watch a file descriptor for read events* example uses the low-level `AbstractEventLoop.add_reader()` method to register the file descriptor of a socket.

The *register an open socket to wait for data using streams* example uses high-level streams created by the `open_connection()` function in a coroutine.

18.5.5 Streams (coroutine based API)

Source code: `Lib/asyncio/streams.py`

Stream functions

注解: The top-level functions in this module are meant as convenience wrappers only; there's really nothing special there, and if they don't do exactly what you want, feel free to copy their code.

coroutine `asyncio.open_connection` (*host=None, port=None, *, loop=None, limit=None, **kws*)

A wrapper for `create_connection()` returning a (reader, writer) pair.

The reader returned is a `StreamReader` instance; the writer is a `StreamWriter` instance.

The arguments are all the usual arguments to `AbstractEventLoop.create_connection()` except `protocol_factory`; most common are positional `host` and `port`, with various optional keyword arguments following.

Additional optional keyword arguments are `loop` (to set the event loop instance to use) and `limit` (to set the buffer limit passed to the `StreamReader`).

This function is a *coroutine*.

coroutine `asyncio.start_server` (*client_connected_cb, host=None, port=None, *, loop=None, limit=None, **kws*)

Start a socket server, with a callback for each client connected. The return value is the same as `create_server()`.

The `client_connected_cb` parameter is called with two parameters: `client_reader`, `client_writer`. `client_reader` is a `StreamReader` object, while `client_writer` is a `StreamWriter` object. The `client_connected_cb` parameter can either be a plain callback function or a *coroutine function*; if it is a coroutine function, it will be automatically converted into a *Task*.

The rest of the arguments are all the usual arguments to `create_server()` except `protocol_factory`; most common are positional `host` and `port`, with various optional keyword arguments following.

Additional optional keyword arguments are `loop` (to set the event loop instance to use) and `limit` (to set the buffer limit passed to the `StreamReader`).

This function is a *coroutine*.

coroutine `asyncio.open_unix_connection` (*path=None, *, loop=None, limit=None, **kws*)

A wrapper for `create_unix_connection()` returning a (reader, writer) pair.

See `open_connection()` for information about return value and other details.

This function is a *coroutine*.

Availability: UNIX.

coroutine `asyncio.start_unix_server` (*client_connected_cb, path=None, *, loop=None, limit=None, **kws*)

Start a UNIX Domain Socket server, with a callback for each client connected.

See `start_server()` for information about return value and other details.

This function is a *coroutine*.

Availability: UNIX.

StreamReader

class `asyncio.StreamReader` (*limit=_DEFAULT_LIMIT, loop=None*)

This class is *not thread safe*.

The `limit` argument's default value is set to `_DEFAULT_LIMIT` which is `2**16` (64 KiB)

exception()

Get the exception.

feed_eof()

Acknowledge the EOF.

feed_data(data)

Feed *data* bytes in the internal buffer. Any operations waiting for the data will be resumed.

set_exception(exc)

Set the exception.

set_transport(transport)

Set the transport.

coroutine read(n=-1)

Read up to *n* bytes. If *n* is not provided, or set to `-1`, read until EOF and return all read bytes.

If the EOF was received and the internal buffer is empty, return an empty `bytes` object.

This method is a *coroutine*.

coroutine readline()

Read one line, where “line” is a sequence of bytes ending with `\n`.

If EOF is received, and `\n` was not found, the method will return the partial read bytes.

If the EOF was received and the internal buffer is empty, return an empty `bytes` object.

This method is a *coroutine*.

coroutine readexactly(n)

Read exactly *n* bytes. Raise an *IncompleteReadError* if the end of the stream is reached before *n* can be read, the *IncompleteReadError.partial* attribute of the exception contains the partial read bytes.

This method is a *coroutine*.

coroutine readuntil(separator=b'\n')

Read data from the stream until *separator* is found.

On success, the data and separator will be removed from the internal buffer (consumed). Returned data will include the separator at the end.

Configured stream limit is used to check result. Limit sets the maximal length of data that can be returned, not counting the separator.

If an EOF occurs and the complete separator is still not found, an *IncompleteReadError* exception will be raised, and the internal buffer will be reset. The *IncompleteReadError.partial* attribute may contain the separator partially.

If the data cannot be read because of over limit, a *LimitOverrunError* exception will be raised, and the data will be left in the internal buffer, so it can be read again.

3.5.2 新版功能.

at_eof()

Return `True` if the buffer is empty and *feed_eof()* was called.

StreamWriter

class `asyncio.StreamWriter(transport, protocol, reader, loop)`

Wraps a Transport.

This exposes `write()`, `writelines()`, `can_write_eof()`, `write_eof()`, `get_extra_info()` and `close()`. It adds `drain()` which returns an optional *Future* on which you can wait for flow control. It also adds a transport attribute which references the Transport directly.

This class is *not thread safe*.

transport

Transport.

can_write_eof()

Return *True* if the transport supports `write_eof()`, *False* if not. See *WriteTransport.can_write_eof()*.

close()

Close the transport: see *BaseTransport.close()*.

coroutine drain()

Let the write buffer of the underlying transport a chance to be flushed.

The intended use is to write:

```
w.write(data)
yield from w.drain()
```

When the size of the transport buffer reaches the high-water limit (the protocol is paused), block until the size of the buffer is drained down to the low-water limit and the protocol is resumed. When there is nothing to wait for, the yield-from continues immediately.

Yielding from `drain()` gives the opportunity for the loop to schedule the write operation and flush the buffer. It should especially be used when a possibly large amount of data is written to the transport, and the coroutine does not yield-from between calls to `write()`.

This method is a *coroutine*.

get_extra_info (name, default=None)

Return optional transport information: see *BaseTransport.get_extra_info()*.

write (data)

Write some *data* bytes to the transport: see *WriteTransport.write()*.

writelines (data)

Write a list (or any iterable) of data bytes to the transport: see *WriteTransport.writelines()*.

write_eof()

Close the write end of the transport after flushing buffered data: see *WriteTransport.write_eof()*.

StreamReaderProtocol

class `asyncio.StreamReaderProtocol (stream_reader, client_connected_cb=None, loop=None)`

Trivial helper class to adapt between *Protocol* and *StreamReader*. Subclass of *Protocol*.

stream_reader is a *StreamReader* instance, *client_connected_cb* is an optional function called with (stream_reader, stream_writer) when a connection is made, *loop* is the event loop instance to use.

(This is a helper class instead of making *StreamReader* itself a *Protocol* subclass, because the *StreamReader* has other potential uses, and to prevent the user of the *StreamReader* from accidentally calling inappropriate methods of the protocol.)

IncompleteReadError

exception `asyncio.IncompleteReadError`

Incomplete read error, subclass of `EOFError`.

expected

Total number of expected bytes (*int*).

partial

Read bytes string before the end of stream was reached (*bytes*).

LimitOverrunError

exception `asyncio.LimitOverrunError`

Reached the buffer limit while looking for a separator.

consumed

Total number of to be consumed bytes.

Stream examples

TCP echo client using streams

TCP echo client using the `asyncio.open_connection()` function:

```
import asyncio

@asyncio.coroutine
def tcp_echo_client(message, loop):
    reader, writer = yield from asyncio.open_connection('127.0.0.1', 8888,
                                                         loop=loop)

    print('Send: %r' % message)
    writer.write(message.encode())

    data = yield from reader.read(100)
    print('Received: %r' % data.decode())

    print('Close the socket')
    writer.close()

message = 'Hello World!'
loop = asyncio.get_event_loop()
loop.run_until_complete(tcp_echo_client(message, loop))
loop.close()
```

参见:

The *TCP echo client protocol* example uses the `AbstractEventLoop.create_connection()` method.

TCP echo server using streams

TCP echo server using the `asyncio.start_server()` function:

```

import asyncio

@asyncio.coroutine
def handle_echo(reader, writer):
    data = yield from reader.read(100)
    message = data.decode()
    addr = writer.get_extra_info('peername')
    print("Received %r from %r" % (message, addr))

    print("Send: %r" % message)
    writer.write(data)
    yield from writer.drain()

    print("Close the client socket")
    writer.close()

loop = asyncio.get_event_loop()
coro = asyncio.start_server(handle_echo, '127.0.0.1', 8888, loop=loop)
server = loop.run_until_complete(coro)

# Serve requests until Ctrl+C is pressed
print('Serving on {}'.format(server.sockets[0].getsockname()))
try:
    loop.run_forever()
except KeyboardInterrupt:
    pass

# Close the server
server.close()
loop.run_until_complete(server.wait_closed())
loop.close()

```

参见:

The *TCP echo server protocol* example uses the `AbstractEventLoop.create_server()` method.

Get HTTP headers

Simple example querying HTTP headers of the URL passed on the command line:

```

import asyncio
import urllib.parse
import sys

@asyncio.coroutine
def print_http_headers(url):
    url = urllib.parse.urlsplit(url)
    if url.scheme == 'https':
        connect = asyncio.open_connection(url.hostname, 443, ssl=True)
    else:
        connect = asyncio.open_connection(url.hostname, 80)
    reader, writer = yield from connect
    query = ('HEAD {path} HTTP/1.0\r\n'
            'Host: {hostname}\r\n'
            '\r\n').format(path=url.path or '/', hostname=url.hostname)
    writer.write(query.encode('latin-1'))

```

(下页继续)

(续上页)

```

while True:
    line = yield from reader.readline()
    if not line:
        break
    line = line.decode('latin1').rstrip()
    if line:
        print('HTTP header> %s' % line)

# Ignore the body, close the socket
writer.close()

url = sys.argv[1]
loop = asyncio.get_event_loop()
task = asyncio.ensure_future(print_http_headers(url))
loop.run_until_complete(task)
loop.close()

```

用法:

```
python example.py http://example.com/path/page.html
```

or with HTTPS:

```
python example.py https://example.com/path/page.html
```

Register an open socket to wait for data using streams

Coroutine waiting until a socket receives data using the `open_connection()` function:

```

import asyncio
try:
    from socket import socketpair
except ImportError:
    from asyncio.windows_utils import socketpair

@asyncio.coroutine
def wait_for_data(loop):
    # Create a pair of connected sockets
    rsock, wsock = socketpair()

    # Register the open socket to wait for data
    reader, writer = yield from asyncio.open_connection(sock=rsock, loop=loop)

    # Simulate the reception of data from the network
    loop.call_soon(wsock.send, 'abc'.encode())

    # Wait for data
    data = yield from reader.read(100)

    # Got data, we are done: close the socket
    print("Received:", data.decode())
    writer.close()

    # Close the second socket
    wsock.close()

```

(下页继续)

(续上页)

```

loop = asyncio.get_event_loop()
loop.run_until_complete(wait_for_data(loop))
loop.close()

```

参见:

The *register an open socket to wait for data using a protocol* example uses a low-level protocol created by the `AbstractEventLoop.create_connection()` method.

The *watch a file descriptor for read events* example uses the low-level `AbstractEventLoop.add_reader()` method to register the file descriptor of a socket.

18.5.6 Subprocess

Source code: `Lib/asyncio/subprocess.py`

Windows event loop

On Windows, the default event loop is `SelectorEventLoop` which does not support subprocesses. `ProactorEventLoop` should be used instead. Example to use it on Windows:

```

import asyncio, sys

if sys.platform == 'win32':
    loop = asyncio.ProactorEventLoop()
    asyncio.set_event_loop(loop)

```

参见:

Available event loops and *Platform support*.

Create a subprocess: high-level API using Process

coroutine `asyncio.create_subprocess_exec(*args, stdin=None, stdout=None, stderr=None, loop=None, limit=None, **kwds)`

创建一个子进程。

The *limit* parameter sets the buffer limit passed to the `StreamReader`. See `AbstractEventLoop.subprocess_exec()` for other parameters.

Return a `Process` instance.

This function is a *coroutine*.

coroutine `asyncio.create_subprocess_shell(cmd, stdin=None, stdout=None, stderr=None, loop=None, limit=None, **kwds)`

Run the shell command *cmd*.

The *limit* parameter sets the buffer limit passed to the `StreamReader`. See `AbstractEventLoop.subprocess_shell()` for other parameters.

Return a `Process` instance.

It is the application's responsibility to ensure that all whitespace and metacharacters are quoted appropriately to avoid *shell injection* vulnerabilities. The `shlex.quote()` function can be used to properly escape whitespace and shell metacharacters in strings that are going to be used to construct shell commands.

This function is a *coroutine*.

Use the `AbstractEventLoop.connect_read_pipe()` and `AbstractEventLoop.connect_write_pipe()` methods to connect pipes.

Create a subprocess: low-level API using `subprocess.Popen`

Run subprocesses asynchronously using the `subprocess` module.

```
coroutine AbstractEventLoop.subprocess_exec(protocol_factory,          *args,
                                             stdin=subprocess.PIPE,      std-
                                             out=subprocess.PIPE, stderr=subprocess.PIPE,
                                             **kwargs)
```

Create a subprocess from one or more string arguments (character strings or bytes strings encoded to the *filesystem encoding*), where the first string specifies the program to execute, and the remaining strings specify the program's arguments. (Thus, together the string arguments form the `sys.argv` value of the program, assuming it is a Python script.) This is similar to the standard library `subprocess.Popen` class called with `shell=False` and the list of strings passed as the first argument; however, where `Popen` takes a single argument which is list of strings, `subprocess_exec()` takes multiple string arguments.

The `protocol_factory` must instantiate a subclass of the `asyncio.SubprocessProtocol` class.

Other parameters:

- `stdin`: Either a file-like object representing the pipe to be connected to the subprocess's standard input stream using `connect_write_pipe()`, or the constant `subprocess.PIPE` (the default). By default a new pipe will be created and connected.
- `stdout`: Either a file-like object representing the pipe to be connected to the subprocess's standard output stream using `connect_read_pipe()`, or the constant `subprocess.PIPE` (the default). By default a new pipe will be created and connected.
- `stderr`: Either a file-like object representing the pipe to be connected to the subprocess's standard error stream using `connect_read_pipe()`, or one of the constants `subprocess.PIPE` (the default) or `subprocess.STDOUT`. By default a new pipe will be created and connected. When `subprocess.STDOUT` is specified, the subprocess's standard error stream will be connected to the same pipe as the standard output stream.
- All other keyword arguments are passed to `subprocess.Popen` without interpretation, except for `bufsize`, `universal_newlines` and `shell`, which should not be specified at all.

Returns a pair of (transport, protocol), where `transport` is an instance of `BaseSubprocessTransport`.

This method is a *coroutine*.

See the constructor of the `subprocess.Popen` class for parameters.

```
coroutine AbstractEventLoop.subprocess_shell(protocol_factory,      cmd,
                                             *,                      stdin=subprocess.PIPE,
                                             stdout=subprocess.PIPE,
                                             stderr=subprocess.PIPE, **kwargs)
```

Create a subprocess from `cmd`, which is a character string or a bytes string encoded to the *filesystem encoding*, using the platform's "shell" syntax. This is similar to the standard library `subprocess.Popen` class called with `shell=True`.

The `protocol_factory` must instantiate a subclass of the `asyncio.SubprocessProtocol` class.

See `subprocess_exec()` for more details about the remaining arguments.

Returns a pair of `(transport, protocol)`, where *transport* is an instance of `BaseSubprocessTransport`.

It is the application's responsibility to ensure that all whitespace and metacharacters are quoted appropriately to avoid [shell injection](#) vulnerabilities. The `shlex.quote()` function can be used to properly escape whitespace and shell metacharacters in strings that are going to be used to construct shell commands.

This method is a *coroutine*.

参见:

The `AbstractEventLoop.connect_read_pipe()` and `AbstractEventLoop.connect_write_pipe()` methods.

常数

`asyncio.subprocess.PIPE`

Special value that can be used as the *stdin*, *stdout* or *stderr* argument to `create_subprocess_shell()` and `create_subprocess_exec()` and indicates that a pipe to the standard stream should be opened.

`asyncio.subprocess.STDOUT`

Special value that can be used as the *stderr* argument to `create_subprocess_shell()` and `create_subprocess_exec()` and indicates that standard error should go into the same handle as standard output.

`asyncio.subprocess.DEVNULL`

Special value that can be used as the *stdin*, *stdout* or *stderr* argument to `create_subprocess_shell()` and `create_subprocess_exec()` and indicates that the special file `os.devnull` will be used.

Process

class `asyncio.subprocess.Process`

A subprocess created by the `create_subprocess_exec()` or the `create_subprocess_shell()` function.

The API of the `Process` class was designed to be close to the API of the `subprocess.Popen` class, but there are some differences:

- There is no explicit `poll()` method
- The `communicate()` and `wait()` methods don't take a *timeout* parameter: use the `wait_for()` function
- The *universal_newlines* parameter is not supported (only bytes strings are supported)
- The `wait()` method of the `Process` class is asynchronous whereas the `wait()` method of the `Popen` class is implemented as a busy loop.

This class is *not thread safe*. See also the *Subprocess and threads* section.

coroutine `wait()`

Wait for child process to terminate. Set and return *returncode* attribute.

This method is a *coroutine*.

注解: This will deadlock when using `stdout=PIPE` or `stderr=PIPE` and the child process generates enough output to a pipe such that it blocks waiting for the OS pipe buffer to accept more data. Use the `communicate()` method when using pipes to avoid that.

coroutine communicate (*input=None*)

Interact with process: Send data to stdin. Read data from stdout and stderr, until end-of-file is reached. Wait for process to terminate. The optional *input* argument should be data to be sent to the child process, or `None`, if no data should be sent to the child. The type of *input* must be bytes.

communicate() returns a tuple (`stdout_data`, `stderr_data`).

If a `BrokenPipeError` or `ConnectionResetError` exception is raised when writing *input* into stdin, the exception is ignored. It occurs when the process exits before all data are written into stdin.

Note that if you want to send data to the process's stdin, you need to create the `Process` object with `stdin=PIPE`. Similarly, to get anything other than `None` in the result tuple, you need to give `stdout=PIPE` and/or `stderr=PIPE` too.

This method is a *coroutine*.

注解: The data read is buffered in memory, so do not use this method if the data size is large or unlimited.

在 3.4.2 版更改: The method now ignores `BrokenPipeError` and `ConnectionResetError`.

send_signal (*signal*)

Sends the signal *signal* to the child process.

注解: On Windows, `SIGTERM` is an alias for `terminate()`. `CTRL_C_EVENT` and `CTRL_BREAK_EVENT` can be sent to processes started with a *creationflags* parameter which includes `CREATE_NEW_PROCESS_GROUP`.

terminate ()

Stop the child. On Posix OSs the method sends `signal.SIGTERM` to the child. On Windows the `Win32 API` function `TerminateProcess()` is called to stop the child.

kill ()

Kills the child. On Posix OSs the function sends `SIGKILL` to the child. On Windows *kill()* is an alias for *terminate()*.

stdin

Standard input stream (*StreamWriter*), `None` if the process was created with `stdin=None`.

stdout

Standard output stream (*StreamReader*), `None` if the process was created with `stdout=None`.

stderr

Standard error stream (*StreamReader*), `None` if the process was created with `stderr=None`.

警告: Use the *communicate()* method rather than *.stdin.write*, *.stdout.read* or *.stderr.read* to avoid deadlocks due to streams pausing reading or writing and blocking the child process.

pid

The identifier of the process.

Note that for processes created by the *create_subprocess_shell()* function, this attribute is the process identifier of the spawned shell.

returncode

Return code of the process when it exited. A `None` value indicates that the process has not terminated yet.

A negative value `-N` indicates that the child was terminated by signal `N` (Unix only).

Subprocess and threads

`asyncio` supports running subprocesses from different threads, but there are limits:

- An event loop must run in the main thread
- The child watcher must be instantiated in the main thread, before executing subprocesses from other threads. Call the `get_child_watcher()` function in the main thread to instantiate the child watcher.

The `asyncio.subprocess.Process` class is not thread safe.

参见:

The *Concurrency and multithreading in asyncio* section.

Subprocess examples

Subprocess using transport and protocol

Example of a subprocess protocol using to get the output of a subprocess and to wait for the subprocess exit. The subprocess is created by the `AbstractEventLoop.subprocess_exec()` method:

```
import asyncio
import sys

class DateProtocol(asyncio.SubprocessProtocol):
    def __init__(self, exit_future):
        self.exit_future = exit_future
        self.output = bytearray()

    def pipe_data_received(self, fd, data):
        self.output.extend(data)

    def process_exited(self):
        self.exit_future.set_result(True)

@asyncio.coroutine
def get_date(loop):
    code = 'import datetime; print(datetime.datetime.now())'
    exit_future = asyncio.Future(loop=loop)

    # Create the subprocess controlled by the protocol DateProtocol,
    # redirect the standard output into a pipe
    create = loop.subprocess_exec(lambda: DateProtocol(exit_future),
                                  sys.executable, '-c', code,
                                  stdin=None, stderr=None)

    transport, protocol = yield from create

    # Wait for the subprocess exit using the process_exited() method
    # of the protocol
    yield from exit_future

    # Close the stdout pipe
    transport.close()
```

(下页继续)

(续上页)

```

    # Read the output which was collected by the pipe_data_received()
    # method of the protocol
    data = bytes(protocol.output)
    return data.decode('ascii').rstrip()

if sys.platform == "win32":
    loop = asyncio.ProactorEventLoop()
    asyncio.set_event_loop(loop)
else:
    loop = asyncio.get_event_loop()

date = loop.run_until_complete(get_date(loop))
print("Current date: %s" % date)
loop.close()

```

Subprocess using streams

Example using the *Process* class to control the subprocess and the *StreamReader* class to read from the standard output. The subprocess is created by the *create_subprocess_exec()* function:

```

import asyncio.subprocess
import sys

@asyncio.coroutine
def get_date():
    code = 'import datetime; print(datetime.datetime.now())'

    # Create the subprocess, redirect the standard output into a pipe
    create = asyncio.create_subprocess_exec(sys.executable, '-c', code,
                                           stdout=asyncio.subprocess.PIPE)

    proc = yield from create

    # Read one line of output
    data = yield from proc.stdout.readline()
    line = data.decode('ascii').rstrip()

    # Wait for the subprocess exit
    yield from proc.wait()
    return line

if sys.platform == "win32":
    loop = asyncio.ProactorEventLoop()
    asyncio.set_event_loop(loop)
else:
    loop = asyncio.get_event_loop()

date = loop.run_until_complete(get_date())
print("Current date: %s" % date)
loop.close()

```

18.5.7 Synchronization primitives

Source code: [Lib/asyncio/locks.py](#)

Locks:

- `Lock`
- `Event`
- `Condition`

Semaphores:

- `Semaphore`
- `BoundedSemaphore`

asyncio lock API was designed to be close to classes of the `threading` module (`Lock`, `Event`, `Condition`, `Semaphore`, `BoundedSemaphore`), but it has no `timeout` parameter. The `asyncio.wait_for()` function can be used to cancel a task after a timeout.

Locks

Lock

class `asyncio.Lock` (*, `loop=None`)

Primitive lock objects.

A primitive lock is a synchronization primitive that is not owned by a particular coroutine when locked. A primitive lock is in one of two states, ‘locked’ or ‘unlocked’.

It is created in the unlocked state. It has two basic methods, `acquire()` and `release()`. When the state is unlocked, `acquire()` changes the state to locked and returns immediately. When the state is locked, `acquire()` blocks until a call to `release()` in another coroutine changes it to unlocked, then the `acquire()` call resets it to locked and returns. The `release()` method should only be called in the locked state; it changes the state to unlocked and returns immediately. If an attempt is made to release an unlocked lock, a `RuntimeError` will be raised.

When more than one coroutine is blocked in `acquire()` waiting for the state to turn to unlocked, only one coroutine proceeds when a `release()` call resets the state to unlocked; first coroutine which is blocked in `acquire()` is being processed.

`acquire()` is a coroutine and should be called with `yield from`.

Locks also support the context management protocol. `(yield from lock)` should be used as the context manager expression.

This class is *not thread safe*.

Usage:

```
lock = Lock()
...
yield from lock
try:
    ...
finally:
    lock.release()
```

Context manager usage:

```
lock = Lock()
...
with (yield from lock):
    ...
```

Lock objects can be tested for locking state:

```
if not lock.locked():
    yield from lock
else:
    # lock is acquired
    ...
```

locked()

Return True if the lock is acquired.

coroutine acquire()

Acquire a lock.

This method blocks until the lock is unlocked, then sets it to locked and returns True.

This method is a *coroutine*.

release()

Release a lock.

When the lock is locked, reset it to unlocked, and return. If any other coroutines are blocked waiting for the lock to become unlocked, allow exactly one of them to proceed.

在未锁定的锁调用时，会引发`RuntimeError`异常。

There is no return value.

Event

class `asyncio.Event` (*, loop=None)

An Event implementation, asynchronous equivalent to `threading.Event`.

Class implementing event objects. An event manages a flag that can be set to true with the `set()` method and reset to false with the `clear()` method. The `wait()` method blocks until the flag is true. The flag is initially false.

This class is *not thread safe*.

clear()

Reset the internal flag to false. Subsequently, coroutines calling `wait()` will block until `set()` is called to set the internal flag to true again.

is_set()

Return True if and only if the internal flag is true.

set()

Set the internal flag to true. All coroutines waiting for it to become true are awakened. Coroutine that call `wait()` once the flag is true will not block at all.

coroutine wait()

Block until the internal flag is true.

If the internal flag is true on entry, return True immediately. Otherwise, block until another coroutine calls `set()` to set the flag to true, then return True.

This method is a *coroutine*.

Condition

class `asyncio.Condition` (*lock=None*, *, *loop=None*)

A Condition implementation, asynchronous equivalent to `threading.Condition`.

This class implements condition variable objects. A condition variable allows one or more coroutines to wait until they are notified by another coroutine.

If the *lock* argument is given and not `None`, it must be a `Lock` object, and it is used as the underlying lock. Otherwise, a new `Lock` object is created and used as the underlying lock.

This class is *not thread safe*.

coroutine `acquire()`

Acquire the underlying lock.

This method blocks until the lock is unlocked, then sets it to locked and returns `True`.

This method is a *coroutine*.

notify (*n=1*)

By default, wake up one coroutine waiting on this condition, if any. If the calling coroutine has not acquired the lock when this method is called, a `RuntimeError` is raised.

This method wakes up at most *n* of the coroutines waiting for the condition variable; it is a no-op if no coroutines are waiting.

注解: An awakened coroutine does not actually return from its `wait()` call until it can reacquire the lock. Since `notify()` does not release the lock, its caller should.

locked()

Return `True` if the underlying lock is acquired.

notify_all()

Wake up all coroutines waiting on this condition. This method acts like `notify()`, but wakes up all waiting coroutines instead of one. If the calling coroutine has not acquired the lock when this method is called, a `RuntimeError` is raised.

release()

Release the underlying lock.

When the lock is locked, reset it to unlocked, and return. If any other coroutines are blocked waiting for the lock to become unlocked, allow exactly one of them to proceed.

在未锁定的锁调用时，会引发`RuntimeError`异常。

There is no return value.

coroutine `wait()`

Wait until notified.

If the calling coroutine has not acquired the lock when this method is called, a `RuntimeError` is raised.

This method releases the underlying lock, and then blocks until it is awakened by a `notify()` or `notify_all()` call for the same condition variable in another coroutine. Once awakened, it re-acquires the lock and returns `True`.

This method is a *coroutine*.

coroutine `wait_for(predicate)`

Wait until a predicate becomes true.

The predicate should be a callable which result will be interpreted as a boolean value. The final predicate value is the return value.

This method is a *coroutine*.

Semaphores

Semaphore

class `asyncio.Semaphore` (*value=1, *, loop=None*)

A Semaphore implementation.

A semaphore manages an internal counter which is decremented by each `acquire()` call and incremented by each `release()` call. The counter can never go below zero; when `acquire()` finds that it is zero, it blocks, waiting until some other coroutine calls `release()`.

Semaphores also support the context management protocol.

The optional argument gives the initial value for the internal counter; it defaults to 1. If the value given is less than 0, `ValueError` is raised.

This class is *not thread safe*.

coroutine `acquire()`

获取一个信号量。

If the internal counter is larger than zero on entry, decrement it by one and return `True` immediately. If it is zero on entry, block, waiting until some other coroutine has called `release()` to make it larger than 0, and then return `True`.

This method is a *coroutine*.

locked()

Returns `True` if semaphore can not be acquired immediately.

release()

Release a semaphore, incrementing the internal counter by one. When it was zero on entry and another coroutine is waiting for it to become larger than zero again, wake up that coroutine.

BoundedSemaphore

class `asyncio.BoundedSemaphore` (*value=1, *, loop=None*)

A bounded semaphore implementation. Inherit from `Semaphore`.

This raises `ValueError` in `release()` if it would increase the value above the initial value.

18.5.8 队列集

Source code: `Lib/asyncio/queues.py`

Queues:

- `Queue`
- `PriorityQueue`
- `LifoQueue`

asyncio queue API was designed to be close to classes of the `queue` module (`Queue`, `PriorityQueue`, `LifoQueue`), but it has no `timeout` parameter. The `asyncio.wait_for()` function can be used to cancel a task after a timeout.

队列

class `asyncio.Queue` (*maxsize=0*, *, *loop=None*)

A queue, useful for coordinating producer and consumer coroutines.

If *maxsize* is less than or equal to zero, the queue size is infinite. If it is an integer greater than 0, then `yield` from `put()` will block when the queue reaches *maxsize*, until an item is removed by `get()`.

Unlike the standard library `queue`, you can reliably know this `Queue`'s size with `qsize()`, since your single-threaded asyncio application won't be interrupted between calling `qsize()` and doing an operation on the `Queue`.

这个类不是线程安全的。

在 3.4.4 版更改: New `join()` and `task_done()` methods.

empty()

如果队列为空返回 `True`，否则返回 `False`。

full()

如果有 *maxsize* 个条目在队列中，则返回 `True`。

注解: If the `Queue` was initialized with *maxsize*=0 (the default), then `full()` is never `True`.

coroutine `get()`

从队列中删除并返回一个元素。如果队列为空，则等待，直到队列中有元素。

This method is a *coroutine*.

参见:

The `empty()` method.

get_nowait()

Remove and return an item from the queue.

立即返回一个队列中的元素，如果队列内有值，否则引发异常 `QueueEmpty`。

coroutine `join()`

Block until all items in the queue have been gotten and processed.

The count of unfinished tasks goes up whenever an item is added to the queue. The count goes down whenever a consumer thread calls `task_done()` to indicate that the item was retrieved and all work on it is complete. When the count of unfinished tasks drops to zero, `join()` unblocks.

This method is a *coroutine*.

3.4.4 新版功能.

coroutine `put(item)`

Put an item into the queue. If the queue is full, wait until a free slot is available before adding item.

This method is a *coroutine*.

参见:

The `full()` method.

put_nowait (item)

不阻塞的放一个元素入队列。

如果没有立即可用的空闲槽，引发 `QueueFull` 异常。

qsize ()

Number of items in the queue.

task_done ()

表明前面排队的任务已经完成，即 `get` 出来的元素相关操作已经完成。

由队列使用者控制。每个 `get ()` 用于获取一个任务，任务最后调用 `task_done ()` 告诉队列，这个任务已经完成。

如果 `join ()` 当前正在阻塞，在所有条目都被处理后，将解除阻塞 (意味着每个 `put ()` 进队列的条目的 `task_done ()` 都被收到)。

如果被调用的次数多于放入队列中的项目数量，将引发 `ValueError`。

3.4.4 新版功能.

maxsize

队列中可存放的元素数量。

PriorityQueue

class `asyncio.PriorityQueue`

A subclass of `Queue`; retrieves entries in priority order (lowest first).

Entries are typically tuples of the form: (priority number, data).

LifoQueue

class `asyncio.LifoQueue`

A subclass of `Queue` that retrieves most recently added entries first.

异常

exception `asyncio.QueueEmpty`

Exception raised when the `get_nowait ()` method is called on a `Queue` object which is empty.

exception `asyncio.QueueFull`

Exception raised when the `put_nowait ()` method is called on a `Queue` object which is full.

18.5.9 Develop with asyncio

Asynchronous programming is different than classical “sequential” programming. This page lists common traps and explains how to avoid them.

Debug mode of asyncio

The implementation of `asyncio` has been written for performance. In order to ease the development of asynchronous code, you may wish to enable *debug mode*.

To enable all debug checks for an application:

- Enable the asyncio debug mode globally by setting the environment variable `PYTHONASYNCIODEBUG` to 1, or by calling `AbstractEventLoop.set_debug()`.
- Set the log level of the *asyncio logger* to `logging.DEBUG`. For example, call `logging.basicConfig(level=logging.DEBUG)` at startup.
- Configure the *warnings* module to display *ResourceWarning* warnings. For example, use the `-Wdefault` command line option of Python to display them.

Examples debug checks:

- Log *coroutines defined but never “yielded from”*
- `call_soon()` and `call_at()` methods raise an exception if they are called from the wrong thread.
- Log the execution time of the selector
- Log callbacks taking more than 100 ms to be executed. The `AbstractEventLoop.slow_callback_duration` attribute is the minimum duration in seconds of “slow” callbacks.
- *ResourceWarning* warnings are emitted when transports and event loops are *not closed explicitly*.

参见:

The `AbstractEventLoop.set_debug()` method and the *asyncio logger*.

Cancellation

Cancellation of tasks is not common in classic programming. In asynchronous programming, not only is it something common, but you have to prepare your code to handle it.

Futures and tasks can be cancelled explicitly with their `Future.cancel()` method. The `wait_for()` function cancels the waited task when the timeout occurs. There are many other cases where a task can be cancelled indirectly.

Don't call `set_result()` or `set_exception()` method of *Future* if the future is cancelled: it would fail with an exception. For example, write:

```
if not fut.cancelled():
    fut.set_result('done')
```

Don't schedule directly a call to the `set_result()` or the `set_exception()` method of a future with `AbstractEventLoop.call_soon()`: the future can be cancelled before its method is called.

If you wait for a future, you should check early if the future was cancelled to avoid useless operations. Example:

```
@coroutine
def slow_operation(fut):
    if fut.cancelled():
        return
    # ... slow computation ...
    yield from fut
    # ...
```

The `shield()` function can also be used to ignore cancellation.

Concurrency and multithreading

An event loop runs in a thread and executes all callbacks and tasks in the same thread. While a task is running in the event loop, no other task is running in the same thread. But when the task uses `yield from`, the task is suspended and the event loop executes the next task.

To schedule a callback from a different thread, the `AbstractEventLoop.call_soon_threadsafe()` method should be used. Example:

```
loop.call_soon_threadsafe(callback, *args)
```

Most `asyncio` objects are not thread safe. You should only worry if you access objects outside the event loop. For example, to cancel a future, don't call directly its `Future.cancel()` method, but:

```
loop.call_soon_threadsafe(fut.cancel)
```

为了能够处理信号和执行子进程，事件循环必须运行于主线程中。

To schedule a coroutine object from a different thread, the `run_coroutine_threadsafe()` function should be used. It returns a `concurrent.futures.Future` to access the result:

```
future = asyncio.run_coroutine_threadsafe(coro_func(), loop)
result = future.result(timeout) # Wait for the result with a timeout
```

The `AbstractEventLoop.run_in_executor()` method can be used with a thread pool executor to execute a callback in different thread to not block the thread of the event loop.

参见:

The *Synchronization primitives* section describes ways to synchronize tasks.

The *Subprocess and threads* section lists `asyncio` limitations to run subprocesses from different threads.

Handle blocking functions correctly

Blocking functions should not be called directly. For example, if a function blocks for 1 second, other tasks are delayed by 1 second which can have an important impact on reactivity.

For networking and subprocesses, the `asyncio` module provides high-level APIs like *protocols*.

An executor can be used to run a task in a different thread or even in a different process, to not block the thread of the event loop. See the `AbstractEventLoop.run_in_executor()` method.

参见:

The *Delayed calls* section details how the event loop handles time.

日志

The `asyncio` module logs information with the `logging` module in the logger `'asyncio'`.

The default log level for the `asyncio` module is `logging.INFO`. For those not wanting such verbosity from `asyncio` the log level can be changed. For example, to change the level to `logging.WARNING`:

```
logging.getLogger('asyncio').setLevel(logging.WARNING)
```

Detect coroutine objects never scheduled

When a coroutine function is called and its result is not passed to `ensure_future()` or to the `AbstractEventLoop.create_task()` method, the execution of the coroutine object will never be scheduled which is probably a bug. *Enable the debug mode of asyncio* to *log a warning* to detect it.

Example with the bug:

```
import asyncio

@asyncio.coroutine
def test():
    print("never scheduled")

test()
```

调试模式的输出:

```
Coroutine test() at test.py:3 was never yielded from
Coroutine object created at (most recent call last):
  File "test.py", line 7, in <module>
    test()
```

The fix is to call the `ensure_future()` function or the `AbstractEventLoop.create_task()` method with the coroutine object.

参见:

Pending task destroyed.

Detect exceptions never consumed

Python usually calls `sys.excepthook()` on unhandled exceptions. If `Future.set_exception()` is called, but the exception is never consumed, `sys.excepthook()` is not called. Instead, *a log is emitted* when the future is deleted by the garbage collector, with the traceback where the exception was raised.

Example of unhandled exception:

```
import asyncio

@asyncio.coroutine
def bug():
    raise Exception("not consumed")

loop = asyncio.get_event_loop()
asyncio.ensure_future(bug())
loop.run_forever()
loop.close()
```

输出:

```
Task exception was never retrieved
future: <Task finished coro=<coro() done, defined at asyncio/coroutines.py:139>_
↳exception=Exception('not consumed',)>
Traceback (most recent call last):
  File "asyncio/tasks.py", line 237, in _step
    result = next(coro)
  File "asyncio/coroutines.py", line 141, in coro
    res = func(*args, **kw)
  File "test.py", line 5, in bug
    raise Exception("not consumed")
Exception: not consumed
```

Enable the debug mode of asyncio to get the traceback where the task was created. Output in debug mode:

```

Task exception was never retrieved
future: <Task finished coro=<bug() done, defined at test.py:3> exception=Exception(
↪ 'not consumed',) created at test.py:8>
source_traceback: Object created at (most recent call last):
  File "test.py", line 8, in <module>
    asyncio.ensure_future(bug())
Traceback (most recent call last):
  File "asyncio/tasks.py", line 237, in _step
    result = next(coro)
  File "asyncio/coroutines.py", line 79, in __next__
    return next(self.gen)
  File "asyncio/coroutines.py", line 141, in coro
    res = func(*args, **kw)
  File "test.py", line 5, in bug
    raise Exception("not consumed")
Exception: not consumed

```

There are different options to fix this issue. The first option is to chain the coroutine in another coroutine and use classic try/except:

```

@asyncio.coroutine
def handle_exception():
    try:
        yield from bug()
    except Exception:
        print("exception consumed")

loop = asyncio.get_event_loop()
asyncio.ensure_future(handle_exception())
loop.run_forever()
loop.close()

```

Another option is to use the `AbstractEventLoop.run_until_complete()` function:

```

task = asyncio.ensure_future(bug())
try:
    loop.run_until_complete(task)
except Exception:
    print("exception consumed")

```

参见:

The `Future.exception()` method.

Chain coroutines correctly

When a coroutine function calls other coroutine functions and tasks, they should be chained explicitly with `yield from`. Otherwise, the execution is not guaranteed to be sequential.

Example with different bugs using `asyncio.sleep()` to simulate slow operations:

```

import asyncio

@asyncio.coroutine
def create():
    yield from asyncio.sleep(3.0)

```

(下页继续)

(续上页)

```

    print("(1) create file")

@asyncio.coroutine
def write():
    yield from asyncio.sleep(1.0)
    print("(2) write into file")

@asyncio.coroutine
def close():
    print("(3) close file")

@asyncio.coroutine
def test():
    asyncio.ensure_future(create())
    asyncio.ensure_future(write())
    asyncio.ensure_future(close())
    yield from asyncio.sleep(2.0)
    loop.stop()

loop = asyncio.get_event_loop()
asyncio.ensure_future(test())
loop.run_forever()
print("Pending tasks at exit: %s" % asyncio.Task.all_tasks(loop))
loop.close()

```

Expected output:

```

(1) create file
(2) write into file
(3) close file
Pending tasks at exit: set()

```

Actual output:

```

(3) close file
(2) write into file
Pending tasks at exit: {<Task pending create() at test.py:7 wait_for=<Future pending_
↳cb=[Task._wakeup()]>>>}
Task was destroyed but it is pending!
task: <Task pending create() done at test.py:5 wait_for=<Future pending cb=[Task._
↳wakeup()]>>>

```

The loop stopped before the `create()` finished, `close()` has been called before `write()`, whereas coroutine functions were called in this order: `create()`, `write()`, `close()`.

To fix the example, tasks must be marked with `yield from`:

```

@asyncio.coroutine
def test():
    yield from asyncio.ensure_future(create())
    yield from asyncio.ensure_future(write())
    yield from asyncio.ensure_future(close())
    yield from asyncio.sleep(2.0)
    loop.stop()

```

Or without `asyncio.ensure_future()`:

```
@asyncio.coroutine
def test():
    yield from create()
    yield from write()
    yield from close()
    yield from asyncio.sleep(2.0)
    loop.stop()
```

Pending task destroyed

If a pending task is destroyed, the execution of its wrapped *coroutine* did not complete. It is probably a bug and so a warning is logged.

Example of log:

```
Task was destroyed but it is pending!
task: <Task pending coro=<kill_me() done, defined at test.py:5> wait_for=<Future_
↳pending cb=[Task._wakeup()]>>
```

Enable the debug mode of asyncio to get the traceback where the task was created. Example of log in debug mode:

```
Task was destroyed but it is pending!
source_traceback: Object created at (most recent call last):
  File "test.py", line 15, in <module>
    task = asyncio.ensure_future(coro, loop=loop)
task: <Task pending coro=<kill_me() done, defined at test.py:5> wait_for=<Future_
↳pending cb=[Task._wakeup()] created at test.py:7> created at test.py:15>
```

参见:

Detect coroutine objects never scheduled.

Close transports and event loops

When a transport is no more needed, call its `close()` method to release resources. Event loops must also be closed explicitly.

If a transport or an event loop is not closed explicitly, a *ResourceWarning* warning will be emitted in its destructor. By default, *ResourceWarning* warnings are ignored. The *Debug mode of asyncio* section explains how to display them.

参见:

The *asyncio* module was designed in **PEP 3156**. For a motivational primer on transports and protocols, see **PEP 3153**.

18.6 asyncore — 异步 socket 处理器

源码: `Lib/asyncore.py`

3.6 版后已移除: 请使用 *asyncio* 替代。

注解：该模块仅为提供向后兼容。我们推荐在新代码中使用 `asyncio`。

该模块提供用于编写异步套接字服务客户端与服务端的基础构件。

只有两种方法让单个处理器上的程序“同一时间完成不止一件事”。多线程编程是最简单和最流行的方法，但是还有另一种非常不同的技术，它可以让你拥有多线程的几乎所有优点，而无需实际使用多线程。它仅仅在你的程序主要受 I/O 限制时有用，那么。如果你的程序受处理器限制，那么先发制人的预定线程可能就是你需要。但是，网络服务器很少受处理器限制。

如果你的操作系统在其 I/O 库中支持 `select()` 系统调用（几乎所有操作系统），那么你可以使用它来同时处理多个通信通道；在 I/O 正在“后台”时进行其他工作。虽然这种策略看起来很奇怪和复杂，特别是起初，它在很多方面比多线程编程更容易理解和控制。`asyncore` 模块为您解决了许多难题，使得构建复杂的高性能网络服务器和客户端的任务变得轻而易举。对于“会话”应用程序和协议，伴侣 `asynchat` 模块是非常宝贵的。

这两个模块背后的基本思想是创建一个或多个网络通道，类的实例 `asyncore.dispatcher` 和 `asynchat.async_chat`。创建通道会将它们添加到全局映射中，如果你不为它提供自己的映射，则由 `loop()` 函数使用。

一旦创建了初始通道，调用 `loop()` 函数将激活通道服务，该服务将一直持续到最后一个通道（包括在异步服务期间已添加到映射中的任何通道）关闭。

```
asyncore.loop([timeout[, use_poll[, map[, count]]]])
```

进入一个轮询循环，其在循环计数超出或所有打开的通道关闭后终止。所有参数都是可选的。`count` 形参默认为 `None`，导致循环仅在所有通道关闭时终止。`timeout` 形参为适当的 `select()` 或 `poll()` 调用设置超时参数，以秒为单位；默认值为 30 秒。`use_poll` 形参，如果为 `True`，则表示 `poll()` 应优先使用 `select`（默认为 `False`）。

The `map` parameter is a dictionary whose items are the channels to watch. As channels are closed they are deleted from their map. If `map` is omitted, a global map is used. Channels (instances of `asyncore.dispatcher`, `asynchat.async_chat` and subclasses thereof) can freely be mixed in the map.

class `asyncore.dispatcher`

The `dispatcher` class is a thin wrapper around a low-level socket object. To make it more useful, it has a few methods for event-handling which are called from the asynchronous loop. Otherwise, it can be treated as a normal non-blocking socket object.

The firing of low-level events at certain times or in certain connection states tells the asynchronous loop that certain higher-level events have taken place. For example, if we have asked for a socket to connect to another host, we know that the connection has been made when the socket becomes writable for the first time (at this point you know that you may write to it with the expectation of success). The implied higher-level events are:

Event	描述
<code>handle_connect()</code>	Implied by the first read or write event
<code>handle_close()</code>	Implied by a read event with no data available
<code>handle_accepted()</code>	Implied by a read event on a listening socket

During asynchronous processing, each mapped channel's `readable()` and `writable()` methods are used to determine whether the channel's socket should be added to the list of channels `select()`ed or `poll()`ed for read and write events.

Thus, the set of channel events is larger than the basic socket events. The full set of methods that can be overridden in your subclass follows:

handle_read()

Called when the asynchronous loop detects that a `read()` call on the channel's socket will succeed.

handle_write()

Called when the asynchronous loop detects that a writable socket can be written. Often this method will implement the necessary buffering for performance. For example:

```
def handle_write(self):
    sent = self.send(self.buffer)
    self.buffer = self.buffer[sent:]
```

handle_expt()

Called when there is out of band (OOB) data for a socket connection. This will almost never happen, as OOB is tenuously supported and rarely used.

handle_connect()

Called when the active opener's socket actually makes a connection. Might send a "welcome" banner, or initiate a protocol negotiation with the remote endpoint, for example.

handle_close()

Called when the socket is closed.

handle_error()

Called when an exception is raised and not otherwise handled. The default version prints a condensed trace-back.

handle_accept()

Called on listening channels (passive openers) when a connection can be established with a new remote endpoint that has issued a `connect()` call for the local endpoint. Deprecated in version 3.2; use `handle_accepted()` instead.

3.2 版后已移除.

handle_accepted(sock, addr)

Called on listening channels (passive openers) when a connection has been established with a new remote endpoint that has issued a `connect()` call for the local endpoint. `sock` is a new socket object usable to send and receive data on the connection, and `addr` is the address bound to the socket on the other end of the connection.

3.2 新版功能.

readable()

Called each time around the asynchronous loop to determine whether a channel's socket should be added to the list on which read events can occur. The default method simply returns `True`, indicating that by default, all channels will be interested in read events.

writable()

Called each time around the asynchronous loop to determine whether a channel's socket should be added to the list on which write events can occur. The default method simply returns `True`, indicating that by default, all channels will be interested in write events.

In addition, each channel delegates or extends many of the socket methods. Most of these are nearly identical to their socket partners.

create_socket(family=socket.AF_INET, type=socket.SOCK_STREAM)

This is identical to the creation of a normal socket, and will use the same options for creation. Refer to the `socket` documentation for information on creating sockets.

在 3.3 版更改: `family` and `type` arguments can be omitted.

connect(address)

As with the normal socket object, `address` is a tuple with the first element the host to connect to, and the second the port number.

send (*data*)

Send *data* to the remote end-point of the socket.

recv (*buffer_size*)

Read at most *buffer_size* bytes from the socket's remote end-point. An empty bytes object implies that the channel has been closed from the other end.

Note that *recv()* may raise *BlockingIOError*, even though *select.select()* or *select.poll()* has reported the socket ready for reading.

listen (*backlog*)

Listen for connections made to the socket. The *backlog* argument specifies the maximum number of queued connections and should be at least 1; the maximum value is system-dependent (usually 5).

bind (*address*)

Bind the socket to *address*. The socket must not already be bound. (The format of *address* depends on the address family—refer to the *socket* documentation for more information.) To mark the socket as re-usable (setting the *SO_REUSEADDR* option), call the *dispatcher* object's *set_reuse_addr()* method.

accept ()

Accept a connection. The socket must be bound to an address and listening for connections. The return value can be either *None* or a pair (*conn*, *address*) where *conn* is a *new* socket object usable to send and receive data on the connection, and *address* is the address bound to the socket on the other end of the connection. When *None* is returned it means the connection didn't take place, in which case the server should just ignore this event and keep listening for further incoming connections.

close ()

Close the socket. All future operations on the socket object will fail. The remote end-point will receive no more data (after queued data is flushed). Sockets are automatically closed when they are garbage-collected.

class *asyncore.dispatcher_with_send*

A *dispatcher* subclass which adds simple buffered output capability, useful for simple clients. For more sophisticated usage use *asynchat.async_chat*.

class *asyncore.file_dispatcher*

A *file_dispatcher* takes a file descriptor or *file object* along with an optional *map* argument and wraps it for use with the *poll()* or *loop()* functions. If provided a file object or anything with a *fileno()* method, that method will be called and passed to the *file_wrapper* constructor. Availability: UNIX.

class *asyncore.file_wrapper*

A *file_wrapper* takes an integer file descriptor and calls *os.dup()* to duplicate the handle so that the original handle may be closed independently of the *file_wrapper*. This class implements sufficient methods to emulate a socket for use by the *file_dispatcher* class. Availability: UNIX.

18.6.1 *asyncore* Example basic HTTP client

Here is a very basic HTTP client that uses the *dispatcher* class to implement its socket handling:

```
import asyncore

class HTTPClient(asyncore.dispatcher):

    def __init__(self, host, path):
        asyncore.dispatcher.__init__(self)
        self.create_socket()
        self.connect((host, 80))
        self.buffer = bytes('GET %s HTTP/1.0\r\nHost: %s\r\n\r\n' %
                             (path, host), 'ascii')
```

(下页继续)

(续上页)

```

def handle_connect(self):
    pass

def handle_close(self):
    self.close()

def handle_read(self):
    print(self.recv(8192))

def writable(self):
    return (len(self.buffer) > 0)

def handle_write(self):
    sent = self.send(self.buffer)
    self.buffer = self.buffer[sent:]

client = HTTPClient('www.python.org', '/')
asyncore.loop()

```

18.6.2 asyncore Example basic echo server

Here is a basic echo server that uses the *dispatcher* class to accept connections and dispatches the incoming connections to a handler:

```

import asyncore

class EchoHandler(asyncore.dispatcher_with_send):

    def handle_read(self):
        data = self.recv(8192)
        if data:
            self.send(data)

class EchoServer(asyncore.dispatcher):

    def __init__(self, host, port):
        asyncore.dispatcher.__init__(self)
        self.create_socket()
        self.set_reuse_addr()
        self.bind((host, port))
        self.listen(5)

    def handle_accepted(self, sock, addr):
        print('Incoming connection from %s' % repr(addr))
        handler = EchoHandler(sock)

server = EchoServer('localhost', 8080)
asyncore.loop()

```

18.7 `asyncchat` — 异步 socket 指令/响应处理器

Source code: `Lib/asyncchat.py`

3.6 版后已移除: 请使用 `asyncio` 替代。

注解: 该模块仅为提供向后兼容。我们推荐在新代码中使用 `asyncio`。

This module builds on the `asyncore` infrastructure, simplifying asynchronous clients and servers and making it easier to handle protocols whose elements are terminated by arbitrary strings, or are of variable length. `asyncchat` defines the abstract class `async_chat` that you subclass, providing implementations of the `collect_incoming_data()` and `found_terminator()` methods. It uses the same asynchronous loop as `asyncore`, and the two types of channel, `asyncore.dispatcher` and `asyncchat.async_chat`, can freely be mixed in the channel map. Typically an `asyncore.dispatcher` server channel generates new `asyncchat.async_chat` channel objects as it receives incoming connection requests.

`class asyncchat.async_chat`

This class is an abstract subclass of `asyncore.dispatcher`. To make practical use of the code you must subclass `async_chat`, providing meaningful `collect_incoming_data()` and `found_terminator()` methods. The `asyncore.dispatcher` methods can be used, although not all make sense in a message/response context.

Like `asyncore.dispatcher`, `async_chat` defines a set of events that are generated by an analysis of socket conditions after a `select()` call. Once the polling loop has been started the `async_chat` object's methods are called by the event-processing framework with no action on the part of the programmer.

Two class attributes can be modified, to improve performance, or possibly even to conserve memory.

`ac_in_buffer_size`

The asynchronous input buffer size (default 4096).

`ac_out_buffer_size`

The asynchronous output buffer size (default 4096).

Unlike `asyncore.dispatcher`, `async_chat` allows you to define a FIFO queue of *producers*. A producer need have only one method, `more()`, which should return data to be transmitted on the channel. The producer indicates exhaustion (*i.e.* that it contains no more data) by having its `more()` method return the empty bytes object. At this point the `async_chat` object removes the producer from the queue and starts using the next producer, if any. When the producer queue is empty the `handle_write()` method does nothing. You use the channel object's `set_terminator()` method to describe how to recognize the end of, or an important breakpoint in, an incoming transmission from the remote endpoint.

To build a functioning `async_chat` subclass your input methods `collect_incoming_data()` and `found_terminator()` must handle the data that the channel receives asynchronously. The methods are described below.

`async_chat.close_when_done()`

Pushes a `None` on to the producer queue. When this producer is popped off the queue it causes the channel to be closed.

`async_chat.collect_incoming_data(data)`

Called with `data` holding an arbitrary amount of received data. The default method, which must be overridden, raises a `NotImplementedError` exception.

`async_chat.discard_buffers()`

In emergencies this method will discard any data held in the input and/or output buffers and the producer queue.

`async_chat.found_terminator()`

Called when the incoming data stream matches the termination condition set by `set_terminator()`. The default method, which must be overridden, raises a `NotImplementedError` exception. The buffered input data should be available via an instance attribute.

`async_chat.get_terminator()`

Returns the current terminator for the channel.

`async_chat.push(data)`

Pushes data on to the channel's queue to ensure its transmission. This is all you need to do to have the channel write the data out to the network, although it is possible to use your own producers in more complex schemes to implement encryption and chunking, for example.

`async_chat.push_with_producer(producer)`

Takes a producer object and adds it to the producer queue associated with the channel. When all currently-pushed producers have been exhausted the channel will consume this producer's data by calling its `more()` method and send the data to the remote endpoint.

`async_chat.set_terminator(term)`

Sets the terminating condition to be recognized on the channel. `term` may be any of three types of value, corresponding to three different ways to handle incoming protocol data.

term	描述
<i>string</i>	Will call <code>found_terminator()</code> when the string is found in the input stream
<i>integer</i>	Will call <code>found_terminator()</code> when the indicated number of characters have been received
<code>None</code>	The channel continues to collect data forever

Note that any data following the terminator will be available for reading by the channel after `found_terminator()` is called.

18.7.1 asynchat Example

The following partial example shows how HTTP requests can be read with `asynchat`. A web server might create an `http_request_handler` object for each incoming client connection. Notice that initially the channel terminator is set to match the blank line at the end of the HTTP headers, and a flag indicates that the headers are being read.

Once the headers have been read, if the request is of type POST (indicating that further data are present in the input stream) then the `Content-Length:` header is used to set a numeric terminator to read the right amount of data from the channel.

The `handle_request()` method is called once all relevant input has been marshalled, after setting the channel terminator to `None` to ensure that any extraneous data sent by the web client are ignored.

```
import asynchat

class http_request_handler(asynchat.async_chat):

    def __init__(self, sock, addr, sessions, log):
        asynchat.async_chat.__init__(self, sock=sock)
        self.addr = addr
        self.sessions = sessions
        self.ibuffer = []
        self.obuffer = b""
        self.set_terminator(b"\r\n\r\n")
        self.reading_headers = True
        self.handling = False
```

(下页继续)

(续上页)

```

self.cgi_data = None
self.log = log

def collect_incoming_data(self, data):
    """Buffer the data"""
    self.ibuffer.append(data)

def found_terminator(self):
    if self.reading_headers:
        self.reading_headers = False
        self.parse_headers(b"".join(self.ibuffer))
        self.ibuffer = []
        if self.op.upper() == b"POST":
            clen = self.headers.getheader("content-length")
            self.set_terminator(int(clen))
        else:
            self.handling = True
            self.set_terminator(None)
            self.handle_request()
    elif not self.handling:
        self.set_terminator(None) # browsers sometimes over-send
        self.cgi_data = parse(self.headers, b"".join(self.ibuffer))
        self.handling = True
        self.ibuffer = []
        self.handle_request()

```

18.8 signal — Set handlers for asynchronous events

This module provides mechanisms to use signal handlers in Python.

18.8.1 General rules

The `signal.signal()` function allows defining custom handlers to be executed when a signal is received. A small number of default handlers are installed: `SIGPIPE` is ignored (so write errors on pipes and sockets can be reported as ordinary Python exceptions) and `SIGINT` is translated into a `KeyboardInterrupt` exception.

A handler for a particular signal, once set, remains installed until it is explicitly reset (Python emulates the BSD style interface regardless of the underlying implementation), with the exception of the handler for `SIGCHLD`, which follows the underlying implementation.

Execution of Python signal handlers

A Python signal handler does not get executed inside the low-level (C) signal handler. Instead, the low-level signal handler sets a flag which tells the *virtual machine* to execute the corresponding Python signal handler at a later point (for example at the next *bytecode* instruction). This has consequences:

- It makes little sense to catch synchronous errors like `SIGFPE` or `SIGSEGV` that are caused by an invalid operation in C code. Python will return from the signal handler to the C code, which is likely to raise the same signal again, causing Python to apparently hang. From Python 3.3 onwards, you can use the `faulthandler` module to report on synchronous errors.

- A long-running calculation implemented purely in C (such as regular expression matching on a large body of text) may run uninterrupted for an arbitrary amount of time, regardless of any signals received. The Python signal handlers will be called when the calculation finishes.

Signals and threads

Python signal handlers are always executed in the main Python thread, even if the signal was received in another thread. This means that signals can't be used as a means of inter-thread communication. You can use the synchronization primitives from the *threading* module instead.

Besides, only the main thread is allowed to set a new signal handler.

18.8.2 Module contents

在 3.5 版更改: `signal` (`SIG*`), `handler` (`SIG_DFL`, `SIG_IGN`) and `sigmask` (`SIG_BLOCK`, `SIG_UNBLOCK`, `SIG_SETMASK`) related constants listed below were turned into *enums*. `getsignal()`, `pthread_sigmask()`, `sigpending()` and `sigwait()` functions return human-readable *enums*.

The variables defined in the *signal* module are:

`signal.SIG_DFL`

This is one of two standard signal handling options; it will simply perform the default function for the signal. For example, on most systems the default action for `SIGQUIT` is to dump core and exit, while the default action for `SIGCHLD` is to simply ignore it.

`signal.SIG_IGN`

This is another standard signal handler, which will simply ignore the given signal.

SIG*

All the signal numbers are defined symbolically. For example, the hangup signal is defined as `signal.SIGHUP`; the variable names are identical to the names used in C programs, as found in `<signal.h>`. The Unix man page for '`signal()`' lists the existing signals (on some systems this is `signal(2)`, on others the list is in `signal(7)`). Note that not all systems define the same set of signal names; only those names defined by the system are defined by this module.

`signal.CTRL_C_EVENT`

The signal corresponding to the `Ctrl+C` keystroke event. This signal can only be used with `os.kill()`.

Availability: Windows.

3.2 新版功能.

`signal.CTRL_BREAK_EVENT`

The signal corresponding to the `Ctrl+Break` keystroke event. This signal can only be used with `os.kill()`.

Availability: Windows.

3.2 新版功能.

`signal.NSIG`

One more than the number of the highest signal number.

`signal.ITIMER_REAL`

Decrements interval timer in real time, and delivers `SIGALRM` upon expiration.

`signal.ITIMER_VIRTUAL`

Decrements interval timer only when the process is executing, and delivers `SIGVTALRM` upon expiration.

signal.ITIMER_PROF

Decrements interval timer both when the process executes and when the system is executing on behalf of the process. Coupled with `ITIMER_VIRTUAL`, this timer is usually used to profile the time spent by the application in user and kernel space. `SIGPROF` is delivered upon expiration.

signal.SIG_BLOCK

A possible value for the *how* parameter to `pthread_sigmask()` indicating that signals are to be blocked.

3.3 新版功能.

signal.SIG_UNBLOCK

A possible value for the *how* parameter to `pthread_sigmask()` indicating that signals are to be unblocked.

3.3 新版功能.

signal.SIG_SETMASK

A possible value for the *how* parameter to `pthread_sigmask()` indicating that the signal mask is to be replaced.

3.3 新版功能.

The `signal` module defines one exception:

exception signal.ItimerError

Raised to signal an error from the underlying `setitimer()` or `getitimer()` implementation. Expect this error if an invalid interval timer or a negative time is passed to `setitimer()`. This error is a subtype of `OSError`.

3.3 新版功能: This error used to be a subtype of `IOError`, which is now an alias of `OSError`.

The `signal` module defines the following functions:

signal.alarm(*time*)

If *time* is non-zero, this function requests that a `SIGALRM` signal be sent to the process in *time* seconds. Any previously scheduled alarm is canceled (only one alarm can be scheduled at any time). The returned value is then the number of seconds before any previously set alarm was to have been delivered. If *time* is zero, no alarm is scheduled, and any scheduled alarm is canceled. If the return value is zero, no alarm is currently scheduled. (See the Unix man page `alarm(2)`.) Availability: Unix.

signal.getsignal(*signalnum*)

Return the current signal handler for the signal *signalnum*. The returned value may be a callable Python object, or one of the special values `signal.SIG_IGN`, `signal.SIG_DFL` or `None`. Here, `signal.SIG_IGN` means that the signal was previously ignored, `signal.SIG_DFL` means that the default way of handling the signal was previously in use, and `None` means that the previous signal handler was not installed from Python.

signal.pause()

Cause the process to sleep until a signal is received; the appropriate handler will then be called. Returns nothing. Not on Windows. (See the Unix man page `signal(2)`.)

See also `sigwait()`, `sigwaitinfo()`, `sigtimedwait()` and `sigpending()`.

signal.pthread_kill(*thread_id*, *signalnum*)

Send the signal *signalnum* to the thread *thread_id*, another thread in the same process as the caller. The target thread can be executing any code (Python or not). However, if the target thread is executing the Python interpreter, the Python signal handlers will be *executed by the main thread*. Therefore, the only point of sending a signal to a particular Python thread would be to force a running system call to fail with `InterruptedError`.

Use `threading.get_ident()` or the `ident` attribute of `threading.Thread` objects to get a suitable value for *thread_id*.

If *signalnum* is 0, then no signal is sent, but error checking is still performed; this can be used to check if the target thread is still running.

Availability: Unix (see the man page `pthread_kill(3)` for further information).

See also `os.kill()`.

3.3 新版功能.

`signal.thread_sigmask(how, mask)`

Fetch and/or change the signal mask of the calling thread. The signal mask is the set of signals whose delivery is currently blocked for the caller. Return the old signal mask as a set of signals.

The behavior of the call is dependent on the value of *how*, as follows.

- `SIG_BLOCK`: The set of blocked signals is the union of the current set and the *mask* argument.
- `SIG_UNBLOCK`: The signals in *mask* are removed from the current set of blocked signals. It is permissible to attempt to unblock a signal which is not blocked.
- `SIG_SETMASK`: The set of blocked signals is set to the *mask* argument.

mask is a set of signal numbers (e.g. `{signal.SIGINT, signal.SIGTERM}`). Use `range(1, signal.NSIG)` for a full mask including all signals.

For example, `signal.thread_sigmask(signal.SIG_BLOCK, [])` reads the signal mask of the calling thread.

Availability: Unix. See the man page `sigprocmask(3)` and `pthread_sigmask(3)` for further information.

See also `pause()`, `sigpending()` and `sigwait()`.

3.3 新版功能.

`signal.setitimer(which, seconds[, interval])`

Sets given interval timer (one of `signal.ITIMER_REAL`, `signal.ITIMER_VIRTUAL` or `signal.ITIMER_PROF`) specified by *which* to fire after *seconds* (float is accepted, different from `alarm()`) and after that every *interval* seconds. The interval timer specified by *which* can be cleared by setting *seconds* to zero.

When an interval timer fires, a signal is sent to the process. The signal sent is dependent on the timer being used; `signal.ITIMER_REAL` will deliver `SIGALRM`, `signal.ITIMER_VIRTUAL` sends `SIGVTALRM`, and `signal.ITIMER_PROF` will deliver `SIGPROF`.

The old values are returned as a tuple: (delay, interval).

Attempting to pass an invalid interval timer will cause an `ItimerError`. Availability: Unix.

`signal.getitimer(which)`

Returns current value of a given interval timer specified by *which*. Availability: Unix.

`signal.set_wakeup_fd(fd)`

Set the wakeup file descriptor to *fd*. When a signal is received, the signal number is written as a single byte into the *fd*. This can be used by a library to wakeup a poll or select call, allowing the signal to be fully processed.

The old wakeup fd is returned (or -1 if file descriptor wakeup was not enabled). If *fd* is -1, file descriptor wakeup is disabled. If not -1, *fd* must be non-blocking. It is up to the library to remove any bytes from *fd* before calling poll or select again.

Use for example `struct.unpack('%uB' % len(data), data)` to decode the signal numbers list.

When threads are enabled, this function can only be called from the main thread; attempting to call it from other threads will cause a `ValueError` exception to be raised.

在 3.5 版更改: On Windows, the function now also supports socket handles.

`signal.siginterrupt(signalnum, flag)`

Change system call restart behaviour: if *flag* is `False`, system calls will be restarted when interrupted by signal *signalnum*, otherwise system calls will be interrupted. Returns nothing. Availability: Unix (see the man page `siginterrupt(3)` for further information).

Note that installing a signal handler with `signal()` will reset the restart behaviour to interruptible by implicitly calling `siginterrupt()` with a true *flag* value for the given signal.

`signal.signal(signalnum, handler)`

Set the handler for signal *signalnum* to the function *handler*. *handler* can be a callable Python object taking two arguments (see below), or one of the special values `signal.SIG_IGN` or `signal.SIG_DFL`. The previous signal handler will be returned (see the description of `getsignal()` above). (See the Unix man page `signal(2)`.)

When threads are enabled, this function can only be called from the main thread; attempting to call it from other threads will cause a `ValueError` exception to be raised.

The *handler* is called with two arguments: the signal number and the current stack frame (`None` or a frame object; for a description of frame objects, see the description in the type hierarchy or see the attribute descriptions in the `inspect` module).

On Windows, `signal()` can only be called with `SIGABRT`, `SIGFPE`, `SIGILL`, `SIGINT`, `SIGSEGV`, `SIGTERM`, or `SIGBREAK`. A `ValueError` will be raised in any other case. Note that not all systems define the same set of signal names; an `AttributeError` will be raised if a signal name is not defined as `SIG*` module level constant.

`signal.sigpending()`

Examine the set of signals that are pending for delivery to the calling thread (i.e., the signals which have been raised while blocked). Return the set of the pending signals.

Availability: Unix (see the man page `sigpending(2)` for further information).

See also `pause()`, `pthread_sigmask()` and `sigwait()`.

3.3 新版功能.

`signal.sigwait(sigset)`

Suspend execution of the calling thread until the delivery of one of the signals specified in the signal set *sigset*. The function accepts the signal (removes it from the pending list of signals), and returns the signal number.

Availability: Unix (see the man page `sigwait(3)` for further information).

See also `pause()`, `pthread_sigmask()`, `sigpending()`, `sigwaitinfo()` and `sigtimedwait()`.

3.3 新版功能.

`signal.sigwaitinfo(sigset)`

Suspend execution of the calling thread until the delivery of one of the signals specified in the signal set *sigset*. The function accepts the signal and removes it from the pending list of signals. If one of the signals in *sigset* is already pending for the calling thread, the function will return immediately with information about that signal. The signal handler is not called for the delivered signal. The function raises an `InterruptedError` if it is interrupted by a signal that is not in *sigset*.

The return value is an object representing the data contained in the `siginfo_t` structure, namely: `si_signo`, `si_code`, `si_errno`, `si_pid`, `si_uid`, `si_status`, `si_band`.

Availability: Unix (see the man page `sigwaitinfo(2)` for further information).

See also `pause()`, `sigwait()` and `sigtimedwait()`.

3.3 新版功能.

在 3.5 版更改: The function is now retried if interrupted by a signal not in *sigset* and the signal handler does not raise an exception (see [PEP 475](#) for the rationale).

`signal.sigtimedwait(sigset, timeout)`

Like `sigwaitinfo()`, but takes an additional *timeout* argument specifying a timeout. If *timeout* is specified as 0, a poll is performed. Returns `None` if a timeout occurs.

Availability: Unix (see the man page *sigtimedwait(2)* for further information).

See also *pause()*, *sigwait()* and *sigwaitinfo()*.

3.3 新版功能.

在 3.5 版更改: The function is now retried with the recomputed *timeout* if interrupted by a signal not in *sigset* and the signal handler does not raise an exception (see [PEP 475](#) for the rationale).

18.8.3 示例

Here is a minimal example program. It uses the *alarm()* function to limit the time spent waiting to open a file; this is useful if the file is for a serial device that may not be turned on, which would normally cause the *os.open()* to hang indefinitely. The solution is to set a 5-second alarm before opening the file; if the operation takes too long, the alarm signal will be sent, and the handler raises an exception.

```
import signal, os

def handler(signum, frame):
    print('Signal handler called with signal', signum)
    raise OSError("Couldn't open device!")

# Set the signal handler and a 5-second alarm
signal.signal(signal.SIGALRM, handler)
signal.alarm(5)

# This open() may hang indefinitely
fd = os.open('/dev/ttyS0', os.O_RDWR)

signal.alarm(0)          # Disable the alarm
```

18.9 mmap — 内存映射文件支持

内存映射 (mmap) 文件对象的行为既像 *bytearray* 又像 *文件对象*。你可以在大部分接受 *bytearray* 的地方使用 *mmap* 对象；例如，你可以使用 *re* 模块来搜索一个内存映射文件。你也可以通过执行 `obj[index] = 97` 来修改单个字节，或者通过对切片赋值来修改一个子序列: `obj[i1:i2] = b'...'`。你还可以在文件的当前位置开始读取和写入数据，并使用 *seek()* 前往另一个位置。

A memory-mapped file is created by the *mmap* constructor, which is different on Unix and on Windows. In either case you must provide a file descriptor for a file opened for update. If you wish to map an existing Python file object, use its *fileno()* method to obtain the correct value for the *fileno* parameter. Otherwise, you can open the file using the *os.open()* function, which returns a file descriptor directly (the file still needs to be closed when done).

注解: If you want to create a memory-mapping for a writable, buffered file, you should *flush()* the file first. This is necessary to ensure that local modifications to the buffers are actually available to the mapping.

For both the Unix and Windows versions of the constructor, *access* may be specified as an optional keyword parameter. *access* accepts one of three values: *ACCESS_READ*, *ACCESS_WRITE*, or *ACCESS_COPY* to specify read-only, write-through or copy-on-write memory respectively. *access* can be used on both Unix and Windows. If *access* is not specified, Windows *mmap* returns a write-through mapping. The initial memory values for all three access types are taken from the specified file. Assignment to an *ACCESS_READ* memory map raises a *TypeError* exception. Assignment to an

ACCESS_WRITE memory map affects both memory and the underlying file. Assignment to an ACCESS_COPY memory map affects memory but does not update the underlying file.

To map anonymous memory, -1 should be passed as the *fileno* along with the length.

class `mmap.mmap` (*fileno*, *length*, *tagname*=None, *access*=ACCESS_DEFAULT[, *offset*])

(Windows version) Maps *length* bytes from the file specified by the file handle *fileno*, and creates a `mmap` object. If *length* is larger than the current size of the file, the file is extended to contain *length* bytes. If *length* is 0, the maximum length of the map is the current size of the file, except that if the file is empty Windows raises an exception (you cannot create an empty mapping on Windows).

tagname, if specified and not None, is a string giving a tag name for the mapping. Windows allows you to have many different mappings against the same file. If you specify the name of an existing tag, that tag is opened, otherwise a new tag of this name is created. If this parameter is omitted or None, the mapping is created without a name. Avoiding the use of the tag parameter will assist in keeping your code portable between Unix and Windows.

offset may be specified as a non-negative integer offset. `mmap` references will be relative to the offset from the beginning of the file. *offset* defaults to 0. *offset* must be a multiple of the ALLOCATIONGRANULARITY.

class `mmap.mmap` (*fileno*, *length*, *flags*=MAP_SHARED, *prot*=PROT_WRITE|PROT_READ, *access*=ACCESS_DEFAULT[, *offset*])

(Unix version) Maps *length* bytes from the file specified by the file descriptor *fileno*, and returns a `mmap` object. If *length* is 0, the maximum length of the map will be the current size of the file when `mmap` is called.

flags specifies the nature of the mapping. MAP_PRIVATE creates a private copy-on-write mapping, so changes to the contents of the `mmap` object will be private to this process, and MAP_SHARED creates a mapping that's shared with all other processes mapping the same areas of the file. The default value is MAP_SHARED.

prot, if specified, gives the desired memory protection; the two most useful values are PROT_READ and PROT_WRITE, to specify that the pages may be read or written. *prot* defaults to PROT_READ | PROT_WRITE.

access may be specified in lieu of *flags* and *prot* as an optional keyword parameter. It is an error to specify both *flags*, *prot* and *access*. See the description of *access* above for information on how to use this parameter.

offset may be specified as a non-negative integer offset. `mmap` references will be relative to the offset from the beginning of the file. *offset* defaults to 0. *offset* must be a multiple of ALLOCATIONGRANULARITY which is equal to PAGE_SIZE on Unix systems.

To ensure validity of the created memory mapping the file specified by the descriptor *fileno* is internally automatically synchronized with physical backing store on Mac OS X and OpenVMS.

This example shows a simple way of using `mmap`:

```
import mmap

# write a simple example file
with open("hello.txt", "wb") as f:
    f.write(b"Hello Python!\n")

with open("hello.txt", "r+b") as f:
    # memory-map the file, size 0 means whole file
    mm = mmap.mmap(f.fileno(), 0)
    # read content via standard file methods
    print(mm.readline()) # prints b"Hello Python!\n"
    # read content via slice notation
    print(mm[:5]) # prints b"Hello"
    # update content using slice notation;
    # note that new content must have same size
    mm[6:] = b" world!\n"
    # ... and read again using standard file methods
```

(下页继续)

(续上页)

```
mm.seek(0)
print(mm.readline()) # prints b"Hello world!\n"
# close the map
mm.close()
```

`mmap` can also be used as a context manager in a `with` statement:

```
import mmap

with mmap.mmap(-1, 13) as mm:
    mm.write(b"Hello world!")
```

3.2 新版功能: Context manager support.

The next example demonstrates how to create an anonymous map and exchange data between the parent and child processes:

```
import mmap
import os

mm = mmap.mmap(-1, 13)
mm.write(b"Hello world!")

pid = os.fork()

if pid == 0: # In a child process
    mm.seek(0)
    print(mm.readline())

    mm.close()
```

Memory-mapped file objects support the following methods:

close()

Closes the `mmap`. Subsequent calls to other methods of the object will result in a `ValueError` exception being raised. This will not close the open file.

closed

True if the file is closed.

3.2 新版功能.

find(sub[, start[, end]])

Returns the lowest index in the object where the subsequence *sub* is found, such that *sub* is contained in the range *[start, end]*. Optional arguments *start* and *end* are interpreted as in slice notation. Returns `-1` on failure.

在 3.5 版更改: Writable *bytes-like object* is now accepted.

flush([offset[, size]])

Flushes changes made to the in-memory copy of a file back to disk. Without use of this call there is no guarantee that changes are written back before the object is destroyed. If *offset* and *size* are specified, only changes to the given range of bytes will be flushed to disk; otherwise, the whole extent of the mapping is flushed. *offset* must be a multiple of the `PAGESIZE` or `ALLOCATIONGRANULARITY`.

(Windows version) A nonzero value returned indicates success; zero indicates failure.

(Unix version) A zero value is returned to indicate success. An exception is raised when the call failed.

move (*dest, src, count*)

Copy the *count* bytes starting at offset *src* to the destination index *dest*. If the mmap was created with `ACCESS_READ`, then calls to move will raise a `TypeError` exception.

read (*[n]*)

Return a *bytes* containing up to *n* bytes starting from the current file position. If the argument is omitted, `None` or negative, return all bytes from the current file position to the end of the mapping. The file position is updated to point after the bytes that were returned.

在 3.3 版更改: Argument can be omitted or `None`.

read_byte ()

Returns a byte at the current file position as an integer, and advances the file position by 1.

readline ()

Returns a single line, starting at the current file position and up to the next newline.

resize (*newsize*)

Resizes the map and the underlying file, if any. If the mmap was created with `ACCESS_READ` or `ACCESS_COPY`, resizing the map will raise a `TypeError` exception.

rfind (*sub* [, *start* [, *end*]])

Returns the highest index in the object where the subsequence *sub* is found, such that *sub* is contained in the range [*start*, *end*]. Optional arguments *start* and *end* are interpreted as in slice notation. Returns `-1` on failure.

在 3.5 版更改: Writable *bytes-like object* is now accepted.

seek (*pos* [, *whence*])

Set the file's current position. *whence* argument is optional and defaults to `os.SEEK_SET` or 0 (absolute file positioning); other values are `os.SEEK_CUR` or 1 (seek relative to the current position) and `os.SEEK_END` or 2 (seek relative to the file's end).

size ()

Return the length of the file, which can be larger than the size of the memory-mapped area.

tell ()

Returns the current position of the file pointer.

write (*bytes*)

Write the bytes in *bytes* into memory at the current position of the file pointer and return the number of bytes written (never less than `len(bytes)`, since if the write fails, a `ValueError` will be raised). The file position is updated to point after the bytes that were written. If the mmap was created with `ACCESS_READ`, then writing to it will raise a `TypeError` exception.

在 3.5 版更改: Writable *bytes-like object* is now accepted.

在 3.6 版更改: The number of bytes written is now returned.

write_byte (*byte*)

Write the integer *byte* into memory at the current position of the file pointer; the file position is advanced by 1. If the mmap was created with `ACCESS_READ`, then writing to it will raise a `TypeError` exception.

本章介绍了支持处理互联网上常用数据格式的模块。

19.1 `email` — 电子邮件与 MIME 处理包

源码: `Lib/email/__init__.py`

The `email` package is a library for managing email messages. It is specifically *not* designed to do any sending of email messages to SMTP ([RFC 2821](#)), NNTP, or other servers; those are functions of modules such as `smtplib` and `nntplib`. The `email` package attempts to be as RFC-compliant as possible, supporting [RFC 5233](#) and [RFC 6532](#), as well as such MIME-related RFCs as [RFC 2045](#), [RFC 2046](#), [RFC 2047](#), [RFC 2183](#), and [RFC 2231](#).

The overall structure of the email package can be divided into three major components, plus a fourth component that controls the behavior of the other components.

The central component of the package is an “object model” that represents email messages. An application interacts with the package primarily through the object model interface defined in the `message` sub-module. The application can use this API to ask questions about an existing email, to construct a new email, or to add or remove email subcomponents that themselves use the same object model interface. That is, following the nature of email messages and their MIME subcomponents, the email object model is a tree structure of objects that all provide the `EmailMessage` API.

The other two major components of the package are the `parser` and the `generator`. The parser takes the serialized version of an email message (a stream of bytes) and converts it into a tree of `EmailMessage` objects. The generator takes an `EmailMessage` and turns it back into a serialized byte stream. (The parser and generator also handle streams of text characters, but this usage is discouraged as it is too easy to end up with messages that are not valid in one way or another.)

The control component is the `policy` module. Every `EmailMessage`, every `generator`, and every `parser` has an associated `policy` object that controls its behavior. Usually an application only needs to specify the policy when an `EmailMessage` is created, either by directly instantiating an `EmailMessage` to create a new email, or by parsing an input stream using a `parser`. But the policy can be changed when the message is serialized using a `generator`.

This allows, for example, a generic email message to be parsed from disk, but to serialize it using standard SMTP settings when sending it to an email server.

The email package does its best to hide the details of the various governing RFCs from the application. Conceptually the application should be able to treat the email message as a structured tree of unicode text and binary attachments, without having to worry about how these are represented when serialized. In practice, however, it is often necessary to be aware of at least some of the rules governing MIME messages and their structure, specifically the names and nature of the MIME “content types” and how they identify multipart documents. For the most part this knowledge should only be required for more complex applications, and even then it should only be the high level structure in question, and not the details of how those structures are represented. Since MIME content types are used widely in modern internet software (not just email), this will be a familiar concept to many programmers.

The following sections describe the functionality of the `email` package. We start with the `message` object model, which is the primary interface an application will use, and follow that with the `parser` and `generator` components. Then we cover the `policy` controls, which completes the treatment of the main components of the library.

The next three sections cover the exceptions the package may raise and the defects (non-compliance with the RFCs) that the `parser` may detect. Then we cover the `headerregistry` and the `contentmanager` sub-components, which provide tools for doing more detailed manipulation of headers and payloads, respectively. Both of these components contain features relevant to consuming and producing non-trivial messages, but also document their extensibility APIs, which will be of interest to advanced applications.

Following those is a set of examples of using the fundamental parts of the APIs covered in the preceding sections.

The foregoing represent the modern (unicode friendly) API of the email package. The remaining sections, starting with the `Message` class, cover the legacy `compat32` API that deals much more directly with the details of how email messages are represented. The `compat32` API does *not* hide the details of the RFCs from the application, but for applications that need to operate at that level, they can be useful tools. This documentation is also relevant for applications that are still using the `compat32` API for backward compatibility reasons.

在 3.6 版更改: Docs reorganized and rewritten to promote the new `EmailMessage/EmailPolicy` API.

Contents of the `email` package documentation:

19.1.1 `email.message`: Representing an email message

Source code: `Lib/email/message.py`

3.6 新版功能:¹

The central class in the `email` package is the `EmailMessage` class, imported from the `email.message` module. It is the base class for the `email` object model. `EmailMessage` provides the core functionality for setting and querying header fields, for accessing message bodies, and for creating or modifying structured messages.

An email message consists of *headers* and a *payload* (which is also referred to as the *content*). Headers are **RFC 5322** or **RFC 6532** style field names and values, where the field name and value are separated by a colon. The colon is not part of either the field name or the field value. The payload may be a simple text message, or a binary object, or a structured sequence of sub-messages each with their own set of headers and their own payload. The latter type of payload is indicated by the message having a MIME type such as `multipart/*` or `message/rfc822`.

The conceptual model provided by an `EmailMessage` object is that of an ordered dictionary of headers coupled with a *payload* that represents the **RFC 5322** body of the message, which might be a list of sub-`EmailMessage` objects. In addition to the normal dictionary methods for accessing the header names and values, there are methods for accessing specialized information from the headers (for example the MIME content type), for operating on the payload, for generating a serialized version of the message, and for recursively walking over the object tree.

¹ Originally added in 3.4 as a *provisional module*. Docs for legacy message class moved to `email.message.Message: Representing an email message using the compat32 API`.

The `EmailMessage` dictionary-like interface is indexed by the header names, which must be ASCII values. The values of the dictionary are strings with some extra methods. Headers are stored and returned in case-preserving form, but field names are matched case-insensitively. Unlike a real dict, there is an ordering to the keys, and there can be duplicate keys. Additional methods are provided for working with headers that have duplicate keys.

The *payload* is either a string or bytes object, in the case of simple message objects, or a list of `EmailMessage` objects, for MIME container documents such as `multipart/*` and `message/rfc822` message objects.

class `email.message.EmailMessage` (*policy=default*)

If *policy* is specified use the rules it specifies to update and serialize the representation of the message. If *policy* is not set, use the `default` policy, which follows the rules of the email RFCs except for line endings (instead of the RFC mandated `\r\n`, it uses the Python standard `\n` line endings). For more information see the *policy* documentation.

as_string (*unixfrom=False, maxheaderlen=None, policy=None*)

Return the entire message flattened as a string. When optional *unixfrom* is true, the envelope header is included in the returned string. *unixfrom* defaults to `False`. For backward compatibility with the base `Message` class *maxheaderlen* is accepted, but defaults to `None`, which means that by default the line length is controlled by the *max_line_length* of the policy. The *policy* argument may be used to override the default policy obtained from the message instance. This can be used to control some of the formatting produced by the method, since the specified *policy* will be passed to the `Generator`.

Flattening the message may trigger changes to the `EmailMessage` if defaults need to be filled in to complete the transformation to a string (for example, MIME boundaries may be generated or modified).

Note that this method is provided as a convenience and may not be the most useful way to serialize messages in your application, especially if you are dealing with multiple messages. See `email.generator.Generator` for a more flexible API for serializing messages. Note also that this method is restricted to producing messages serialized as “7 bit clean” when *utf8* is `False`, which is the default.

在 3.6 版更改: the default behavior when *maxheaderlen* is not specified was changed from defaulting to 0 to defaulting to the value of *max_line_length* from the policy.

__str__ ()

Equivalent to `as_string(policy=self.policy.clone(utf8=True))`. Allows `str(msg)` to produce a string containing the serialized message in a readable format.

在 3.4 版更改: the method was changed to use *utf8=True*, thus producing an **RFC 6531**-like message representation, instead of being a direct alias for `as_string()`.

as_bytes (*unixfrom=False, policy=None*)

Return the entire message flattened as a bytes object. When optional *unixfrom* is true, the envelope header is included in the returned string. *unixfrom* defaults to `False`. The *policy* argument may be used to override the default policy obtained from the message instance. This can be used to control some of the formatting produced by the method, since the specified *policy* will be passed to the `BytesGenerator`.

Flattening the message may trigger changes to the `EmailMessage` if defaults need to be filled in to complete the transformation to a string (for example, MIME boundaries may be generated or modified).

Note that this method is provided as a convenience and may not be the most useful way to serialize messages in your application, especially if you are dealing with multiple messages. See `email.generator.BytesGenerator` for a more flexible API for serializing messages.

__bytes__ ()

Equivalent to `as_bytes()`. Allows `bytes(msg)` to produce a bytes object containing the serialized message.

is_multipart ()

Return `True` if the message's payload is a list of sub-`EmailMessage` objects, otherwise return `False`. When `is_multipart()` returns `False`, the payload should be a string object (which might be a CTE

encoded binary payload). Note that `is_multipart()` returning `True` does not necessarily mean that `msg.get_content_maintype() == 'multipart'` will return the `True`. For example, `is_multipart` will return `True` when the `EmailMessage` is of type `message/rfc822`.

set_unixfrom (*unixfrom*)

Set the message's envelope header to *unixfrom*, which should be a string. (See `mbboxMessage` for a brief description of this header.)

get_unixfrom ()

Return the message's envelope header. Defaults to `None` if the envelope header was never set.

The following methods implement the mapping-like interface for accessing the message's headers. Note that there are some semantic differences between these methods and a normal mapping (i.e. dictionary) interface. For example, in a dictionary there are no duplicate keys, but here there may be duplicate message headers. Also, in dictionaries there is no guaranteed order to the keys returned by `keys()`, but in an `EmailMessage` object, headers are always returned in the order they appeared in the original message, or in which they were added to the message later. Any header deleted and then re-added is always appended to the end of the header list.

These semantic differences are intentional and are biased toward convenience in the most common use cases.

Note that in all cases, any envelope header present in the message is not included in the mapping interface.

__len__ ()

Return the total number of headers, including duplicates.

__contains__ (*name*)

Return `true` if the message object has a field named *name*. Matching is done without regard to case and *name* does not include the trailing colon. Used for the `in` operator. For example:

```
if 'message-id' in myMessage:
    print('Message-ID:', myMessage['message-id'])
```

__getitem__ (*name*)

Return the value of the named header field. *name* does not include the colon field separator. If the header is missing, `None` is returned; a `KeyError` is never raised.

Note that if the named field appears more than once in the message's headers, exactly which of those field values will be returned is undefined. Use the `get_all()` method to get the values of all the extant headers named *name*.

Using the standard (non-`compat32`) policies, the returned value is an instance of a subclass of `email.headerregistry.BaseHeader`.

__setitem__ (*name*, *val*)

Add a header to the message with field name *name* and value *val*. The field is appended to the end of the message's existing headers.

Note that this does *not* overwrite or delete any existing header with the same name. If you want to ensure that the new header is the only one present in the message with field name *name*, delete the field first, e.g.:

```
del msg['subject']
msg['subject'] = 'Python roolz!'
```

If the policy defines certain headers to be unique (as the standard policies do), this method may raise a `ValueError` when an attempt is made to assign a value to such a header when one already exists. This behavior is intentional for consistency's sake, but do not depend on it as we may choose to make such assignments do an automatic deletion of the existing header in the future.

__delitem__ (*name*)

Delete all occurrences of the field with name *name* from the message's headers. No exception is raised if the named field isn't present in the headers.

keys()

Return a list of all the message's header field names.

values()

Return a list of all the message's field values.

items()

Return a list of 2-tuples containing all the message's field headers and values.

get(name, failobj=None)

Return the value of the named header field. This is identical to `__getitem__()` except that optional *failobj* is returned if the named header is missing (*failobj* defaults to `None`).

Here are some additional useful header related methods:

get_all(name, failobj=None)

Return a list of all the values for the field named *name*. If there are no such named headers in the message, *failobj* is returned (defaults to `None`).

add_header(_name, _value, **_params)

Extended header setting. This method is similar to `__setitem__()` except that additional header parameters can be provided as keyword arguments. *_name* is the header field to add and *_value* is the *primary* value for the header.

For each item in the keyword argument dictionary *_params*, the key is taken as the parameter name, with underscores converted to dashes (since dashes are illegal in Python identifiers). Normally, the parameter will be added as `key="value"` unless the value is `None`, in which case only the key will be added.

If the value contains non-ASCII characters, the charset and language may be explicitly controlled by specifying the value as a three tuple in the format `(CHARSET, LANGUAGE, VALUE)`, where *CHARSET* is a string naming the charset to be used to encode the value, *LANGUAGE* can usually be set to `None` or the empty string (see [RFC 2231](#) for other possibilities), and *VALUE* is the string value containing non-ASCII code points. If a three tuple is not passed and the value contains non-ASCII characters, it is automatically encoded in [RFC 2231](#) format using a *CHARSET* of `utf-8` and a *LANGUAGE* of `None`.

Here is an example:

```
msg.add_header('Content-Disposition', 'attachment', filename='bud.gif')
```

This will add a header that looks like

```
Content-Disposition: attachment; filename="bud.gif"
```

An example of the extended interface with non-ASCII characters:

```
msg.add_header('Content-Disposition', 'attachment',
               filename=('iso-8859-1', '', 'Fußballer.ppt'))
```

replace_header(_name, _value)

Replace a header. Replace the first header found in the message that matches *_name*, retaining header order and field name case of the original header. If no matching header is found, raise a *KeyError*.

get_content_type()

Return the message's content type, coerced to lower case of the form *maintype/subtype*. If there is no *Content-Type* header in the message return the value returned by `get_default_type()`. If the *Content-Type* header is invalid, return `text/plain`.

(According to [RFC 2045](#), messages always have a default type, `get_content_type()` will always return a value. [RFC 2045](#) defines a message's default type to be `text/plain` unless it appears inside a

multipart/digest container, in which case it would be *message/rfc822*. If the *Content-Type* header has an invalid type specification, **RFC 2045** mandates that the default type be *text/plain*.)

get_content_maintype()

Return the message's main content type. This is the *maintype* part of the string returned by *get_content_type()*.

get_content_subtype()

Return the message's sub-content type. This is the *subtype* part of the string returned by *get_content_type()*.

get_default_type()

Return the default content type. Most messages have a default content type of *text/plain*, except for messages that are subparts of *multipart/digest* containers. Such subparts have a default content type of *message/rfc822*.

set_default_type(ctype)

Set the default content type. *ctype* should either be *text/plain* or *message/rfc822*, although this is not enforced. The default content type is not stored in the *Content-Type* header, so it only affects the return value of the *get_content_type* methods when no *Content-Type* header is present in the message.

set_param(param, value, header='Content-Type', requote=True, charset=None, language="", replace=False)

Set a parameter in the *Content-Type* header. If the parameter already exists in the header, replace its value with *value*. When *header* is *Content-Type* (the default) and the header does not yet exist in the message, add it, set its value to *text/plain*, and append the new parameter value. Optional *header* specifies an alternative header to *Content-Type*.

If the value contains non-ASCII characters, the charset and language may be explicitly specified using the optional *charset* and *language* parameters. Optional *language* specifies the **RFC 2231** language, defaulting to the empty string. Both *charset* and *language* should be strings. The default is to use the *utf8* charset and *None* for the *language*.

If *replace* is *False* (the default) the header is moved to the end of the list of headers. If *replace* is *True*, the header will be updated in place.

Use of the *requote* parameter with *EmailMessage* objects is deprecated.

Note that existing parameter values of headers may be accessed through the *params* attribute of the header value (for example, `msg['Content-Type'].params['charset']`).

在 3.4 版更改: *replace* keyword was added.

del_param(param, header='content-type', requote=True)

Remove the given parameter completely from the *Content-Type* header. The header will be re-written in place without the parameter or its value. Optional *header* specifies an alternative to *Content-Type*.

Use of the *requote* parameter with *EmailMessage* objects is deprecated.

get_filename(failobj=None)

Return the value of the *filename* parameter of the *Content-Disposition* header of the message. If the header does not have a *filename* parameter, this method falls back to looking for the *name* parameter on the *Content-Type* header. If neither is found, or the header is missing, then *failobj* is returned. The returned string will always be unquoted as per *email.utils.unquote()*.

get_boundary(failobj=None)

Return the value of the *boundary* parameter of the *Content-Type* header of the message, or *failobj* if either the header is missing, or has no *boundary* parameter. The returned string will always be unquoted as per *email.utils.unquote()*.

set_boundary (*boundary*)

Set the boundary parameter of the *Content-Type* header to *boundary*. *set_boundary()* will always quote *boundary* if necessary. A *HeaderParseError* is raised if the message object has no *Content-Type* header.

Note that using this method is subtly different from deleting the old *Content-Type* header and adding a new one with the new boundary via *add_header()*, because *set_boundary()* preserves the order of the *Content-Type* header in the list of headers.

get_content_charset (*failobj=None*)

Return the charset parameter of the *Content-Type* header, coerced to lower case. If there is no *Content-Type* header, or if that header has no charset parameter, *failobj* is returned.

get_charsets (*failobj=None*)

Return a list containing the character set names in the message. If the message is a *multipart*, then the list will contain one element for each subpart in the payload, otherwise, it will be a list of length 1.

Each item in the list will be a string which is the value of the charset parameter in the *Content-Type* header for the represented subpart. If the subpart has no *Content-Type* header, no charset parameter, or is not of the *text* main MIME type, then that item in the returned list will be *failobj*.

is_attachment ()

Return True if there is a *Content-Disposition* header and its (case insensitive) value is attachment, False otherwise.

在 3.4.2 版更改: *is_attachment* is now a method instead of a property, for consistency with *is_multipart()*.

get_content_disposition ()

Return the lowercased value (without parameters) of the message's *Content-Disposition* header if it has one, or None. The possible values for this method are *inline*, *attachment* or None if the message follows [RFC 2183](#).

3.5 新版功能.

The following methods relate to interrogating and manipulating the content (payload) of the message.

walk ()

The *walk()* method is an all-purpose generator which can be used to iterate over all the parts and subparts of a message object tree, in depth-first traversal order. You will typically use *walk()* as the iterator in a for loop; each iteration returns the next subpart.

Here's an example that prints the MIME type of every part of a multipart message structure:

```
>>> for part in msg.walk():
...     print(part.get_content_type())
multipart/report
text/plain
message/delivery-status
text/plain
text/plain
message/rfc822
text/plain
```

walk iterates over the subparts of any part where *is_multipart()* returns True, even though *msg.get_content_maintype() == 'multipart'* may return False. We can see this in our example by making use of the *_structure* debug helper function:

```
>>> for part in msg.walk():
...     print(part.get_content_maintype() == 'multipart',
```

(下页继续)

(续上页)

```

...         part.is_multipart()
True True
False False
False True
False False
False False
False True
False False
>>> _structure(msg)
multipart/report
    text/plain
    message/delivery-status
        text/plain
        text/plain
    message/rfc822
        text/plain

```

Here the message parts are not multipart, but they do contain subparts. `is_multipart()` returns True and walk descends into the subparts.

get_body (*preferencelist*=('related', 'html', 'plain'))

Return the MIME part that is the best candidate to be the “body” of the message.

preferencelist must be a sequence of strings from the set `related`, `html`, and `plain`, and indicates the order of preference for the content type of the part returned.

Start looking for candidate matches with the object on which the `get_body` method is called.

If `related` is not included in *preferencelist*, consider the root part (or subpart of the root part) of any related encountered as a candidate if the (sub-)part matches a preference.

When encountering a `multipart/related`, check the `start` parameter and if a part with a matching *Content-ID* is found, consider only it when looking for candidate matches. Otherwise consider only the first (default root) part of the `multipart/related`.

If a part has a *Content-Disposition* header, only consider the part a candidate match if the value of the header is `inline`.

If none of the candidates matches any of the preferences in *preferencelist*, return `None`.

Notes: (1) For most applications the only *preferencelist* combinations that really make sense are `('plain',)`, `('html', 'plain')`, and the default `('related', 'html', 'plain')`. (2) Because matching starts with the object on which `get_body` is called, calling `get_body` on a `multipart/related` will return the object itself unless *preferencelist* has a non-default value. (3) Messages (or message parts) that do not specify a *Content-Type* or whose *Content-Type* header is invalid will be treated as if they are of type `text/plain`, which may occasionally cause `get_body` to return unexpected results.

iter_attachments ()

Return an iterator over all of the immediate sub-parts of the message that are not candidate “body” parts. That is, skip the first occurrence of each of `text/plain`, `text/html`, `multipart/related`, or `multipart/alternative` (unless they are explicitly marked as attachments via *Content-Disposition: attachment*), and return all remaining parts. When applied directly to a `multipart/related`, return an iterator over the all the related parts except the root part (ie: the part pointed to by the `start` parameter, or the first part if there is no `start` parameter or the `start` parameter doesn't match the *Content-ID* of any of the parts). When applied directly to a `multipart/alternative` or a non-multipart, return an empty iterator.

iter_parts ()

Return an iterator over all of the immediate sub-parts of the message, which will be empty for a non-multipart. (See also `walk()`.)

get_content (*args, content_manager=None, **kw)

Call the `get_content()` method of the `content_manager`, passing self as the message object, and passing along any other arguments or keywords as additional arguments. If `content_manager` is not specified, use the `content_manager` specified by the current `policy`.

set_content (*args, content_manager=None, **kw)

Call the `set_content()` method of the `content_manager`, passing self as the message object, and passing along any other arguments or keywords as additional arguments. If `content_manager` is not specified, use the `content_manager` specified by the current `policy`.

make_related (boundary=None)

Convert a non-multipart message into a multipart/related message, moving any existing `Content-` headers and payload into a (new) first part of the multipart. If `boundary` is specified, use it as the boundary string in the multipart, otherwise leave the boundary to be automatically created when it is needed (for example, when the message is serialized).

make_alternative (boundary=None)

Convert a non-multipart or a multipart/related into a multipart/alternative, moving any existing `Content-` headers and payload into a (new) first part of the multipart. If `boundary` is specified, use it as the boundary string in the multipart, otherwise leave the boundary to be automatically created when it is needed (for example, when the message is serialized).

make_mixed (boundary=None)

Convert a non-multipart, a multipart/related, or a multipart-alternative into a multipart/mixed, moving any existing `Content-` headers and payload into a (new) first part of the multipart. If `boundary` is specified, use it as the boundary string in the multipart, otherwise leave the boundary to be automatically created when it is needed (for example, when the message is serialized).

add_related (*args, content_manager=None, **kw)

If the message is a multipart/related, create a new message object, pass all of the arguments to its `set_content()` method, and `attach()` it to the multipart. If the message is a non-multipart, call `make_related()` and then proceed as above. If the message is any other type of multipart, raise a `TypeError`. If `content_manager` is not specified, use the `content_manager` specified by the current `policy`. If the added part has no `Content-Disposition` header, add one with the value `inline`.

add_alternative (*args, content_manager=None, **kw)

If the message is a multipart/alternative, create a new message object, pass all of the arguments to its `set_content()` method, and `attach()` it to the multipart. If the message is a non-multipart or multipart/related, call `make_alternative()` and then proceed as above. If the message is any other type of multipart, raise a `TypeError`. If `content_manager` is not specified, use the `content_manager` specified by the current `policy`.

add_attachment (*args, content_manager=None, **kw)

If the message is a multipart/mixed, create a new message object, pass all of the arguments to its `set_content()` method, and `attach()` it to the multipart. If the message is a non-multipart, multipart/related, or multipart/alternative, call `make_mixed()` and then proceed as above. If `content_manager` is not specified, use the `content_manager` specified by the current `policy`. If the added part has no `Content-Disposition` header, add one with the value `attachment`. This method can be used both for explicit attachments (`Content-Disposition: attachment`) and inline attachments (`Content-Disposition: inline`), by passing appropriate options to the `content_manager`.

clear ()

Remove the payload and all of the headers.

clear_content()

Remove the payload and all of the `Content-` headers, leaving all other headers intact and in their original order.

EmailMessage objects have the following instance attributes:

preamble

The format of a MIME document allows for some text between the blank line following the headers, and the first multipart boundary string. Normally, this text is never visible in a MIME-aware mail reader because it falls outside the standard MIME armor. However, when viewing the raw text of the message, or when viewing the message in a non-MIME aware reader, this text can become visible.

The *preamble* attribute contains this leading extra-armor text for MIME documents. When the *Parser* discovers some text after the headers but before the first boundary string, it assigns this text to the message's *preamble* attribute. When the *Generator* is writing out the plain text representation of a MIME message, and it finds the message has a *preamble* attribute, it will write this text in the area between the headers and the first boundary. See *email.parser* and *email.generator* for details.

Note that if the message object has no preamble, the *preamble* attribute will be `None`.

epilogue

The *epilogue* attribute acts the same way as the *preamble* attribute, except that it contains text that appears between the last boundary and the end of the message. As with the *preamble*, if there is no epilog text this attribute will be `None`.

defects

The *defects* attribute contains a list of all the problems found when parsing this message. See *email.errors* for a detailed description of the possible parsing defects.

class email.message.MIMEPart (policy=default)

This class represents a subpart of a MIME message. It is identical to *EmailMessage*, except that no *MIME-Version* headers are added when *set_content()* is called, since sub-parts do not need their own *MIME-Version* headers.

19.1.2 email.parser: Parsing email messages

Source code: [Lib/email/parser.py](#)

Message object structures can be created in one of two ways: they can be created from whole cloth by creating an *EmailMessage* object, adding headers using the dictionary interface, and adding payload(s) using *set_content()* and related methods, or they can be created by parsing a serialized representation of the email message.

The *email* package provides a standard parser that understands most email document structures, including MIME documents. You can pass the parser a bytes, string or file object, and the parser will return to you the root *EmailMessage* instance of the object structure. For simple, non-MIME messages the payload of this root object will likely be a string containing the text of the message. For MIME messages, the root object will return `True` from its *is_multipart()* method, and the subparts can be accessed via the payload manipulation methods, such as *get_body()*, *iter_parts()*, and *walk()*.

There are actually two parser interfaces available for use, the *Parser* API and the incremental *FeedParser* API. The *Parser* API is most useful if you have the entire text of the message in memory, or if the entire message lives in a file on the file system. *FeedParser* is more appropriate when you are reading the message from a stream which might block waiting for more input (such as reading an email message from a socket). The *FeedParser* can consume and parse the message incrementally, and only returns the root object when you close the parser.

Note that the parser can be extended in limited ways, and of course you can implement your own parser completely from scratch. All of the logic that connects the *email* package's bundled parser and the *EmailMessage* class

is embodied in the `policy` class, so a custom parser can create message object trees any way it finds necessary by implementing custom versions of the appropriate `policy` methods.

FeedParser API

The `BytesFeedParser`, imported from the `email.feedparser` module, provides an API that is conducive to incremental parsing of email messages, such as would be necessary when reading the text of an email message from a source that can block (such as a socket). The `BytesFeedParser` can of course be used to parse an email message fully contained in a *bytes-like object*, string, or file, but the `BytesParser` API may be more convenient for such use cases. The semantics and results of the two parser APIs are identical.

The `BytesFeedParser`'s API is simple; you create an instance, feed it a bunch of bytes until there's no more to feed it, then close the parser to retrieve the root message object. The `BytesFeedParser` is extremely accurate when parsing standards-compliant messages, and it does a very good job of parsing non-compliant messages, providing information about how a message was deemed broken. It will populate a message object's `defects` attribute with a list of any problems it found in a message. See the `email.errors` module for the list of defects that it can find.

Here is the API for the `BytesFeedParser`:

class `email.parser.BytesFeedParser` (`_factory=None`, *, `policy=policy.compat32`)

Create a `BytesFeedParser` instance. Optional `_factory` is a no-argument callable; if not specified use the `message_factory` from the `policy`. Call `_factory` whenever a new message object is needed.

If `policy` is specified use the rules it specifies to update the representation of the message. If `policy` is not set, use the `compat32` policy, which maintains backward compatibility with the Python 3.2 version of the email package and provides `Message` as the default factory. All other policies provide `EmailMessage` as the default `_factory`. For more information on what else `policy` controls, see the `policy` documentation.

Note: **The `policy` keyword should always be specified**; The default will change to `email.policy.default` in a future version of Python.

3.2 新版功能.

在 3.3 版更改: Added the `policy` keyword.

在 3.6 版更改: `_factory` defaults to the `policy.message_factory`.

feed (`data`)

Feed the parser some more data. `data` should be a *bytes-like object* containing one or more lines. The lines can be partial and the parser will stitch such partial lines together properly. The lines can have any of the three common line endings: carriage return, newline, or carriage return and newline (they can even be mixed).

close ()

Complete the parsing of all previously fed data and return the root message object. It is undefined what happens if `feed()` is called after this method has been called.

class `email.parser.FeedParser` (`_factory=None`, *, `policy=policy.compat32`)

Works like `BytesFeedParser` except that the input to the `feed()` method must be a string. This is of limited utility, since the only way for such a message to be valid is for it to contain only ASCII text or, if `utf8` is `True`, no binary attachments.

在 3.3 版更改: Added the `policy` keyword.

Parser API

The `BytesParser` class, imported from the `email.parser` module, provides an API that can be used to parse a message when the complete contents of the message are available in a *bytes-like object* or file. The `email.parser` module also provides `Parser` for parsing strings, and header-only parsers, `BytesHeaderParser` and

HeaderParser, which can be used if you're only interested in the headers of the message. *BytesHeaderParser* and *HeaderParser* can be much faster in these situations, since they do not attempt to parse the message body, instead setting the payload to the raw body.

class email.parser.BytesParser (*_class=None*, *, *policy=policy.compat32*)

Create a *BytesParser* instance. The *_class* and *policy* arguments have the same meaning and semantics as the *_factory* and *policy* arguments of *BytesFeedParser*.

Note: **The policy keyword should always be specified**; The default will change to *email.policy.default* in a future version of Python.

在 3.3 版更改: Removed the *strict* argument that was deprecated in 2.4. Added the *policy* keyword.

在 3.6 版更改: *_class* defaults to the *policy message_factory*.

parse (*fp*, *headersonly=False*)

Read all the data from the binary file-like object *fp*, parse the resulting bytes, and return the message object. *fp* must support both the *readline()* and the *read()* methods.

The bytes contained in *fp* must be formatted as a block of **RFC 5322** (or, if *utf8* is *True*, **RFC 6532**) style headers and header continuation lines, optionally preceded by an envelope header. The header block is terminated either by the end of the data or by a blank line. Following the header block is the body of the message (which may contain MIME-encoded subparts, including subparts with a *Content-Transfer-Encoding* of 8bit).

Optional *headersonly* is a flag specifying whether to stop parsing after reading the headers or not. The default is *False*, meaning it parses the entire contents of the file.

parsebytes (*bytes*, *headersonly=False*)

Similar to the *parse()* method, except it takes a *bytes-like object* instead of a file-like object. Calling this method on a *bytes-like object* is equivalent to wrapping *bytes* in a *BytesIO* instance first and calling *parse()*.

Optional *headersonly* is as with the *parse()* method.

3.2 新版功能.

class email.parser.BytesHeaderParser (*_class=None*, *, *policy=policy.compat32*)

Exactly like *BytesParser*, except that *headersonly* defaults to *True*.

3.3 新版功能.

class email.parser.Parser (*_class=None*, *, *policy=policy.compat32*)

This class is parallel to *BytesParser*, but handles string input.

在 3.3 版更改: Removed the *strict* argument. Added the *policy* keyword.

在 3.6 版更改: *_class* defaults to the *policy message_factory*.

parse (*fp*, *headersonly=False*)

Read all the data from the text-mode file-like object *fp*, parse the resulting text, and return the root message object. *fp* must support both the *readline()* and the *read()* methods on file-like objects.

Other than the text mode requirement, this method operates like *BytesParser.parse()*.

parsestr (*text*, *headersonly=False*)

Similar to the *parse()* method, except it takes a string object instead of a file-like object. Calling this method on a string is equivalent to wrapping *text* in a *StringIO* instance first and calling *parse()*.

Optional *headersonly* is as with the *parse()* method.

class email.parser.HeaderParser (*_class=None*, *, *policy=policy.compat32*)

Exactly like *Parser*, except that *headersonly* defaults to *True*.

Since creating a message object structure from a string or a file object is such a common task, four functions are provided as a convenience. They are available in the top-level `email` package namespace.

`email.message_from_bytes(s, _class=None, *, policy=policy.compat32)`

Return a message object structure from a *bytes-like object*. This is equivalent to `BytesParser().parsebytes(s)`. Optional `_class` and `policy` are interpreted as with the `BytesParser` class constructor.

3.2 新版功能.

在 3.3 版更改: Removed the `strict` argument. Added the `policy` keyword.

`email.message_from_binary_file(fp, _class=None, *, policy=policy.compat32)`

Return a message object structure tree from an open binary *file object*. This is equivalent to `BytesParser().parse(fp)`. `_class` and `policy` are interpreted as with the `BytesParser` class constructor.

3.2 新版功能.

在 3.3 版更改: Removed the `strict` argument. Added the `policy` keyword.

`email.message_from_string(s, _class=None, *, policy=policy.compat32)`

Return a message object structure from a string. This is equivalent to `Parser().parsestr(s)`. `_class` and `policy` are interpreted as with the `Parser` class constructor.

在 3.3 版更改: Removed the `strict` argument. Added the `policy` keyword.

`email.message_from_file(fp, _class=None, *, policy=policy.compat32)`

Return a message object structure tree from an open *file object*. This is equivalent to `Parser().parse(fp)`. `_class` and `policy` are interpreted as with the `Parser` class constructor.

在 3.3 版更改: Removed the `strict` argument. Added the `policy` keyword.

在 3.6 版更改: `_class` defaults to the `policy` message_factory.

Here's an example of how you might use `message_from_bytes()` at an interactive Python prompt:

```
>>> import email
>>> msg = email.message_from_bytes(myBytes)
```

Additional notes

Here are some notes on the parsing semantics:

- Most non-*multipart* type messages are parsed as a single message object with a string payload. These objects will return `False` for `is_multipart()`, and `iter_parts()` will yield an empty list.
- All *multipart* type messages will be parsed as a container message object with a list of sub-message objects for their payload. The outer container message will return `True` for `is_multipart()`, and `iter_parts()` will yield a list of subparts.
- Most messages with a content type of *message/** (such as *message/delivery-status* and *message/rfc822*) will also be parsed as container object containing a list payload of length 1. Their `is_multipart()` method will return `True`. The single element yielded by `iter_parts()` will be a sub-message object.
- Some non-standards-compliant messages may not be internally consistent about their *multipart*-edness. Such messages may have a *Content-Type* header of type *multipart*, but their `is_multipart()` method may return `False`. If such messages were parsed with the `FeedParser`, they will have an instance of the `MultipartInvariantViolationDefect` class in their `defects` attribute list. See `email.errors` for details.

19.1.3 email.generator: Generating MIME documents

Source code: [Lib/email/generator.py](#)

One of the most common tasks is to generate the flat (serialized) version of the email message represented by a message object structure. You will need to do this if you want to send your message via `smtplib.SMTP.sendmail()` or the `nntplib` module, or print the message on the console. Taking a message object structure and producing a serialized representation is the job of the generator classes.

As with the `email.parser` module, you aren't limited to the functionality of the bundled generator; you could write one from scratch yourself. However the bundled generator knows how to generate most email in a standards-compliant way, should handle MIME and non-MIME email messages just fine, and is designed so that the bytes-oriented parsing and generation operations are inverses, assuming the same non-transforming `policy` is used for both. That is, parsing the serialized byte stream via the `BytesParser` class and then regenerating the serialized byte stream using `BytesGenerator` should produce output identical to the input¹. (On the other hand, using the generator on an `EmailMessage` constructed by program may result in changes to the `EmailMessage` object as defaults are filled in.)

The `Generator` class can be used to flatten a message into a text (as opposed to binary) serialized representation, but since Unicode cannot represent binary data directly, the message is of necessity transformed into something that contains only ASCII characters, using the standard email RFC Content Transfer Encoding techniques for encoding email messages for transport over channels that are not "8 bit clean".

class email.generator.BytesGenerator(outfp, mangle_from_=None, maxheaderlen=None, *, policy=None)

Return a `BytesGenerator` object that will write any message provided to the `flatten()` method, or any surrogateescape encoded text provided to the `write()` method, to the file-like object `outfp`. `outfp` must support a `write` method that accepts binary data.

If optional `mangle_from_` is `True`, put a > character in front of any line in the body that starts with the exact string "From ", that is From followed by a space at the beginning of a line. `mangle_from_` defaults to the value of the `mangle_from_` setting of the `policy` (which is `True` for the `compat32` policy and `False` for all others). `mangle_from_` is intended for use when messages are stored in unix mbox format (see [mailbox](#) and [WHY THE CONTENT-LENGTH FORMAT IS BAD](#)).

If `maxheaderlen` is not `None`, refold any header lines that are longer than `maxheaderlen`, or if 0, do not rewrap any headers. If `manheaderlen` is `None` (the default), wrap headers and other message lines according to the `policy` settings.

If `policy` is specified, use that policy to control message generation. If `policy` is `None` (the default), use the policy associated with the `Message` or `EmailMessage` object passed to `flatten` to control the message generation. See [email.policy](#) for details on what `policy` controls.

3.2 新版功能.

在 3.3 版更改: Added the `policy` keyword.

在 3.6 版更改: The default behavior of the `mangle_from_` and `maxheaderlen` parameters is to follow the `policy`.

flatten (msg, unixfrom=False, linesep=None)

Print the textual representation of the message object structure rooted at `msg` to the output file specified when the `BytesGenerator` instance was created.

If the `policy` option `cte_type` is `8bit` (the default), copy any headers in the original parsed message that have not been modified to the output with any bytes with the high bit set reproduced

¹ This statement assumes that you use the appropriate setting for `unixfrom`, and that there are no `policy` settings calling for automatic adjustments (for example, `refold_source` must be `none`, which is *not* the default). It is also not 100% true, since if the message does not conform to the RFC standards occasionally information about the exact original text is lost during parsing error recovery. It is a goal to fix these latter edge cases when possible.

as in the original, and preserve the non-ASCII *Content-Transfer-Encoding* of any body parts that have them. If `cte_type` is `7bit`, convert the bytes with the high bit set as needed using an ASCII-compatible *Content-Transfer-Encoding*. That is, transform parts with non-ASCII *Content-Transfer-Encoding* (*Content-Transfer-Encoding: 8bit*) to an ASCII compatible *Content-Transfer-Encoding*, and encode RFC-invalid non-ASCII bytes in headers using the MIME unknown-8bit character set, thus rendering them RFC-compliant.

If `unixfrom` is `True`, print the envelope header delimiter used by the Unix mailbox format (see [mailbox](#)) before the first of the [RFC 5322](#) headers of the root message object. If the root object has no envelope header, craft a standard one. The default is `False`. Note that for subparts, no envelope header is ever printed.

If `linesep` is not `None`, use it as the separator character between all the lines of the flattened message. If `linesep` is `None` (the default), use the value specified in the *policy*.

clone (*fp*)

Return an independent clone of this *BytesGenerator* instance with the exact same option settings, and *fp* as the new *outfp*.

write (*s*)

Encode *s* using the ASCII codec and the surrogateescape error handler, and pass it to the *write* method of the *outfp* passed to the *BytesGenerator*'s constructor.

As a convenience, *EmailMessage* provides the methods *as_bytes()* and *bytes(aMessage)* (a.k.a. *__bytes__()*), which simplify the generation of a serialized binary representation of a message object. For more detail, see [email.message](#).

Because strings cannot represent binary data, the *Generator* class must convert any binary data in any message it flattens to an ASCII compatible format, by converting them to an ASCII compatible *Content-Transfer-Encoding*. Using the terminology of the email RFCs, you can think of this as *Generator* serializing to an I/O stream that is not “8 bit clean”. In other words, most applications will want to be using *BytesGenerator*, and not *Generator*.

class `email.generator.Generator` (*outfp*, *mangle_from_=None*, *maxheaderlen=None*, *, *policy=None*)

Return a *Generator* object that will write any message provided to the *flatten()* method, or any text provided to the *write()* method, to the *file-like object* *outfp*. *outfp* must support a *write* method that accepts string data.

If optional *mangle_from_* is `True`, put a > character in front of any line in the body that starts with the exact string "From ", that is From followed by a space at the beginning of a line. *mangle_from_* defaults to the value of the *mangle_from_* setting of the *policy* (which is `True` for the *compat32* policy and `False` for all others). *mangle_from_* is intended for use when messages are stored in unix mbox format (see [mailbox](#) and [WHY THE CONTENT-LENGTH FORMAT IS BAD](#)).

If *maxheaderlen* is not `None`, refold any header lines that are longer than *maxheaderlen*, or if 0, do not rewrap any headers. If *manheaderlen* is `None` (the default), wrap headers and other message lines according to the *policy* settings.

If *policy* is specified, use that policy to control message generation. If *policy* is `None` (the default), use the policy associated with the *Message* or *EmailMessage* object passed to *flatten* to control the message generation. See [email.policy](#) for details on what *policy* controls.

在 3.3 版更改: Added the *policy* keyword.

在 3.6 版更改: The default behavior of the *mangle_from_* and *maxheaderlen* parameters is to follow the *policy*.

flatten (*msg*, *unixfrom=False*, *linesep=None*)

Print the textual representation of the message object structure rooted at *msg* to the output file specified when the *Generator* instance was created.

If the *policy* option *cte_type* is `8bit`, generate the message as if the option were set to `7bit`. (This is required because strings cannot represent non-ASCII bytes.) Convert any bytes with the high bit set as needed

using an ASCII-compatible *Content-Transfer-Encoding*. That is, transform parts with non-ASCII *Content-Transfer-Encoding* (*Content-Transfer-Encoding: 8bit*) to an ASCII compatible *Content-Transfer-Encoding*, and encode RFC-invalid non-ASCII bytes in headers using the MIME unknown-8bit character set, thus rendering them RFC-compliant.

If *unixfrom* is `True`, print the envelope header delimiter used by the Unix mailbox format (see *mailbox*) before the first of the [RFC 5322](#) headers of the root message object. If the root object has no envelope header, craft a standard one. The default is `False`. Note that for subparts, no envelope header is ever printed.

If *linesep* is not `None`, use it as the separator character between all the lines of the flattened message. If *linesep* is `None` (the default), use the value specified in the *policy*.

在 3.2 版更改: Added support for re-encoding 8bit message bodies, and the *linesep* argument.

clone (*fp*)

Return an independent clone of this *Generator* instance with the exact same options, and *fp* as the new *outfp*.

write (*s*)

Write *s* to the *write* method of the *outfp* passed to the *Generator*'s constructor. This provides just enough file-like API for *Generator* instances to be used in the *print()* function.

As a convenience, *EmailMessage* provides the methods *as_string()* and *str(aMessage)* (a.k.a. *__str__()*), which simplify the generation of a formatted string representation of a message object. For more detail, see *email.message*.

The *email.generator* module also provides a derived class, *DecodedGenerator*, which is like the *Generator* base class, except that non-*text* parts are not serialized, but are instead represented in the output stream by a string derived from a template filled in with information about the part.

class *email.generator.DecodedGenerator* (*outfp*, *mangle_from_=None*, *maxheaderlen=None*,
fmt=None, *, *policy=None*)

Act like *Generator*, except that for any subpart of the message passed to *Generator.flatten()*, if the subpart is of main type *text*, print the decoded payload of the subpart, and if the main type is not *text*, instead of printing it fill in the string *fmt* using information from the part and print the resulting filled-in string.

To fill in *fmt*, execute *fmt % part_info*, where *part_info* is a dictionary composed of the following keys and values:

- *type* – Full MIME type of the non-*text* part
- *maintype* – Main MIME type of the non-*text* part
- *subtype* – Sub-MIME type of the non-*text* part
- *filename* – Filename of the non-*text* part
- *description* – Description associated with the non-*text* part
- *encoding* – Content transfer encoding of the non-*text* part

If *fmt* is `None`, use the following default *fmt*:

“[Non-text %(type)s) part of message omitted, filename %(filename)s]”

Optional *_mangle_from_* and *maxheaderlen* are as with the *Generator* base class.

19.1.4 email.policy: Policy Objects

3.3 新版功能.

Source code: [Lib/email/policy.py](#)

The *email* package's prime focus is the handling of email messages as described by the various email and MIME RFCs. However, the general format of email messages (a block of header fields each consisting of a name followed by a colon followed by a value, the whole block followed by a blank line and an arbitrary 'body'), is a format that has found utility outside of the realm of email. Some of these uses conform fairly closely to the main email RFCs, some do not. Even when working with email, there are times when it is desirable to break strict compliance with the RFCs, such as generating emails that interoperate with email servers that do not themselves follow the standards, or that implement extensions you want to use in ways that violate the standards.

Policy objects give the email package the flexibility to handle all these disparate use cases.

A *Policy* object encapsulates a set of attributes and methods that control the behavior of various components of the email package during use. *Policy* instances can be passed to various classes and methods in the email package to alter the default behavior. The settable values and their defaults are described below.

There is a default policy used by all classes in the email package. For all of the *parser* classes and the related convenience functions, and for the *Message* class, this is the *Compat32* policy, via its corresponding pre-defined instance *compat32*. This policy provides for complete backward compatibility (in some cases, including bug compatibility) with the pre-Python3.3 version of the email package.

This default value for the *policy* keyword to *EmailMessage* is the *EmailPolicy* policy, via its pre-defined instance *default*.

When a *Message* or *EmailMessage* object is created, it acquires a policy. If the message is created by a *parser*, a policy passed to the parser will be the policy used by the message it creates. If the message is created by the program, then the policy can be specified when it is created. When a message is passed to a *generator*, the generator uses the policy from the message by default, but you can also pass a specific policy to the generator that will override the one stored on the message object.

The default value for the *policy* keyword for the *email.parser* classes and the parser convenience functions **will be changing** in a future version of Python. Therefore you should **always specify explicitly which policy you want to use** when calling any of the classes and functions described in the *parser* module.

The first part of this documentation covers the features of *Policy*, an *abstract base class* that defines the features that are common to all policy objects, including *compat32*. This includes certain hook methods that are called internally by the email package, which a custom policy could override to obtain different behavior. The second part describes the concrete classes *EmailPolicy* and *Compat32*, which implement the hooks that provide the standard behavior and the backward compatible behavior and features, respectively.

Policy instances are immutable, but they can be cloned, accepting the same keyword arguments as the class constructor and returning a new *Policy* instance that is a copy of the original but with the specified attributes values changed.

As an example, the following code could be used to read an email message from a file on disk and pass it to the system *sendmail* program on a Unix system:

```
>>> from email import message_from_binary_file
>>> from email.generator import BytesGenerator
>>> from email import policy
>>> from subprocess import Popen, PIPE
>>> with open('mymsg.txt', 'rb') as f:
...     msg = message_from_binary_file(f, policy=policy.default)
>>> p = Popen(['sendmail', msg['To'].addresses[0]], stdin=PIPE)
>>> g = BytesGenerator(p.stdin, policy=msg.policy.clone(linesep='\r\n'))
>>> g.flatten(msg)
>>> p.stdin.close()
>>> rc = p.wait()
```

Here we are telling *BytesGenerator* to use the RFC correct line separator characters when creating the binary string to feed into *sendmail*'s *stdin*, where the default policy would use `\n` line separators.

Some email package methods accept a *policy* keyword argument, allowing the policy to be overridden for that method. For example, the following code uses the `as_bytes()` method of the *msg* object from the previous example and writes the message to a file using the native line separators for the platform on which it is running:

```
>>> import os
>>> with open('converted.txt', 'wb') as f:
...     f.write(msg.as_bytes(policy=msg.policy.clone(linesep=os.linesep)))
17
```

Policy objects can also be combined using the addition operator, producing a policy object whose settings are a combination of the non-default values of the summed objects:

```
>>> compat SMTP = policy.compat32.clone(linesep='\r\n')
>>> compat_strict = policy.compat32.clone(raise_on_defect=True)
>>> compat_strict SMTP = compat SMTP + compat_strict
```

This operation is not commutative; that is, the order in which the objects are added matters. To illustrate:

```
>>> policy100 = policy.compat32.clone(max_line_length=100)
>>> policy80 = policy.compat32.clone(max_line_length=80)
>>> apolicy = policy100 + policy80
>>> apolicy.max_line_length
80
>>> apolicy = policy80 + policy100
>>> apolicy.max_line_length
100
```

class email.policy.Policy(**kw)

This is the *abstract base class* for all policy classes. It provides default implementations for a couple of trivial methods, as well as the implementation of the immutability property, the `clone()` method, and the constructor semantics.

The constructor of a policy class can be passed various keyword arguments. The arguments that may be specified are any non-method properties on this class, plus any additional non-method properties on the concrete class. A value specified in the constructor will override the default value for the corresponding attribute.

This class defines the following properties, and thus values for the following may be passed in the constructor of any policy class:

max_line_length

The maximum length of any line in the serialized output, not counting the end of line character(s). Default is 78, per [RFC 5322](#). A value of 0 or *None* indicates that no line wrapping should be done at all.

linesep

The string to be used to terminate lines in serialized output. The default is `\n` because that's the internal end-of-line discipline used by Python, though `\r\n` is required by the RFCs.

cte_type

Controls the type of Content Transfer Encodings that may be or are required to be used. The possible values are:

7bit	all data must be “7 bit clean” (ASCII-only). This means that where necessary data will be encoded using either quoted-printable or base64 encoding.
8bit	data is not constrained to be 7 bit clean. Data in headers is still required to be ASCII-only and so will be encoded (see <code>fold_binary()</code> and <code>utf8</code> below for exceptions), but body parts may use the 8bit CTE.

A `cte_type` value of 8bit only works with `BytesGenerator`, not `Generator`, because strings

cannot contain binary data. If a `Generator` is operating under a policy that specifies `cte_type=8bit`, it will act as if `cte_type` is `7bit`.

raise_on_defect

If `True`, any defects encountered will be raised as errors. If `False` (the default), defects will be passed to the `register_defect()` method.

mangle_from_

If `True`, lines starting with “*From* “ in the body are escaped by putting a `>` in front of them. This parameter is used when the message is being serialized by a generator. Default: `False`.

3.5 新版功能: The `mangle_from_` parameter.

message_factory

A factory function for constructing a new empty message object. Used by the parser when building messages. Defaults to `None`, in which case `Message` is used.

3.6 新版功能.

The following `Policy` method is intended to be called by code using the email library to create policy instances with custom settings:

clone (***kw*)

Return a new `Policy` instance whose attributes have the same values as the current instance, except where those attributes are given new values by the keyword arguments.

The remaining `Policy` methods are called by the email package code, and are not intended to be called by an application using the email package. A custom policy must implement all of these methods.

handle_defect (*obj*, *defect*)

Handle a *defect* found on *obj*. When the email package calls this method, *defect* will always be a subclass of `Defect`.

The default implementation checks the `raise_on_defect` flag. If it is `True`, *defect* is raised as an exception. If it is `False` (the default), *obj* and *defect* are passed to `register_defect()`.

register_defect (*obj*, *defect*)

Register a *defect* on *obj*. In the email package, *defect* will always be a subclass of `Defect`.

The default implementation calls the `append` method of the `defects` attribute of *obj*. When the email package calls `handle_defect`, *obj* will normally have a `defects` attribute that has an `append` method. Custom object types used with the email package (for example, custom `Message` objects) should also provide such an attribute, otherwise defects in parsed messages will raise unexpected errors.

header_max_count (*name*)

Return the maximum allowed number of headers named *name*.

Called when a header is added to an `EmailMessage` or `Message` object. If the returned value is not 0 or `None`, and there are already a number of headers with the name *name* greater than or equal to the value returned, a `ValueError` is raised.

Because the default behavior of `Message.__setitem__` is to append the value to the list of headers, it is easy to create duplicate headers without realizing it. This method allows certain headers to be limited in the number of instances of that header that may be added to a `Message` programmatically. (The limit is not observed by the parser, which will faithfully produce as many headers as exist in the message being parsed.)

The default implementation returns `None` for all header names.

header_source_parse (*sourcelines*)

The email package calls this method with a list of strings, each string ending with the line separation characters found in the source being parsed. The first line includes the field header name and separator. All whitespace

in the source is preserved. The method should return the `(name, value)` tuple that is to be stored in the `Message` to represent the parsed header.

If an implementation wishes to retain compatibility with the existing email package policies, *name* should be the case preserved name (all characters up to the `:` separator), while *value* should be the unfolded value (all line separator characters removed, but whitespace kept intact), stripped of leading whitespace.

sourcelines may contain surrogateescaped binary data.

There is no default implementation

header_store_parse (*name, value*)

The email package calls this method with the *name* and *value* provided by the application program when the application program is modifying a `Message` programmatically (as opposed to a `Message` created by a parser). The method should return the `(name, value)` tuple that is to be stored in the `Message` to represent the header.

If an implementation wishes to retain compatibility with the existing email package policies, the *name* and *value* should be strings or string subclasses that do not change the content of the passed in arguments.

There is no default implementation

header_fetch_parse (*name, value*)

The email package calls this method with the *name* and *value* currently stored in the `Message` when that header is requested by the application program, and whatever the method returns is what is passed back to the application as the value of the header being retrieved. Note that there may be more than one header with the same name stored in the `Message`; the method is passed the specific name and value of the header destined to be returned to the application.

value may contain surrogateescaped binary data. There should be no surrogateescaped binary data in the value returned by the method.

There is no default implementation

fold (*name, value*)

The email package calls this method with the *name* and *value* currently stored in the `Message` for a given header. The method should return a string that represents that header “folded” correctly (according to the policy settings) by composing the *name* with the *value* and inserting *lineseq* characters at the appropriate places. See [RFC 5322](#) for a discussion of the rules for folding email headers.

value may contain surrogateescaped binary data. There should be no surrogateescaped binary data in the string returned by the method.

fold_binary (*name, value*)

The same as *fold()*, except that the returned value should be a bytes object rather than a string.

value may contain surrogateescaped binary data. These could be converted back into binary data in the returned bytes object.

class email.policy.**EmailPolicy** (**kw)

This concrete *Policy* provides behavior that is intended to be fully compliant with the current email RFCs. These include (but are not limited to) [RFC 5322](#), [RFC 2047](#), and the current MIME RFCs.

This policy adds new header parsing and folding algorithms. Instead of simple strings, headers are `str` subclasses with attributes that depend on the type of the field. The parsing and folding algorithm fully implement [RFC 2047](#) and [RFC 5322](#).

The default value for the *message_factory* attribute is *EmailMessage*.

In addition to the settable attributes listed above that apply to all policies, this policy adds the following additional attributes:

3.6 新版功能:¹

utf8

If `False`, follow [RFC 5322](#), supporting non-ASCII characters in headers by encoding them as “encoded words” . If `True`, follow [RFC 6532](#) and use `utf-8` encoding for headers. Messages formatted in this way may be passed to SMTP servers that support the SMTPUTF8 extension ([RFC 6531](#)).

refold_source

If the value for a header in the `Message` object originated from a *parser* (as opposed to being set by a program), this attribute indicates whether or not a generator should refold that value when transforming the message back into serialized form. The possible values are:

<code>none</code>	all source values use original folding
<code>long</code>	source values that have any line that is longer than <code>max_line_length</code> will be refolded
<code>all</code>	all values are refolded.

The default is `long`.

header_factory

A callable that takes two arguments, `name` and `value`, where `name` is a header field name and `value` is an unfolded header field value, and returns a string subclass that represents that header. A default `header_factory` (see *headerregistry*) is provided that supports custom parsing for the various address and date [RFC 5322](#) header field types, and the major MIME header field stypes. Support for additional custom parsing will be added in the future.

content_manager

An object with at least two methods: `get_content` and `set_content`. When the `get_content()` or `set_content()` method of an *EmailMessage* object is called, it calls the corresponding method of this object, passing it the message object as its first argument, and any arguments or keywords that were passed to it as additional arguments. By default `content_manager` is set to *raw_data_manager*.

3.4 新版功能.

The class provides the following concrete implementations of the abstract methods of *Policy*:

header_max_count (*name*)

Returns the value of the *max_count* attribute of the specialized class used to represent the header with the given name.

header_source_parse (*sourcelines*)

The name is parsed as everything up to the ‘:’ and returned unmodified. The value is determined by stripping leading whitespace off the remainder of the first line, joining all subsequent lines together, and stripping any trailing carriage return or linefeed characters.

header_store_parse (*name*, *value*)

The name is returned unchanged. If the input value has a *name* attribute and it matches *name* ignoring case, the value is returned unchanged. Otherwise the *name* and *value* are passed to `header_factory`, and the resulting header object is returned as the value. In this case a `ValueError` is raised if the input value contains CR or LF characters.

header_fetch_parse (*name*, *value*)

If the value has a *name* attribute, it is returned to unmodified. Otherwise the *name*, and the *value* with any CR or LF characters removed, are passed to the `header_factory`, and the resulting header object is returned. Any surrogateescaped bytes get turned into the unicode unknown-character glyph.

fold (*name*, *value*)

Header folding is controlled by the *refold_source* policy setting. A value is considered to be a ‘source

¹ Originally added in 3.3 as a *provisional feature*.

value' if and only if it does not have a `name` attribute (having a `name` attribute means it is a header object of some sort). If a source value needs to be refolded according to the policy, it is converted into a header object by passing the *name* and the *value* with any CR and LF characters removed to the `header_factory`. Folding of a header object is done by calling its `fold` method with the current policy.

Source values are split into lines using `splitlines()`. If the value is not to be refolded, the lines are rejoined using the `linesep` from the policy and returned. The exception is lines containing non-ascii binary data. In that case the value is refolded regardless of the `refold_source` setting, which causes the binary data to be CTE encoded using the `unknown-8bit` charset.

`fold_binary(name, value)`

The same as `fold()` if `cte_type` is 7bit, except that the returned value is bytes.

If `cte_type` is 8bit, non-ASCII binary data is converted back into bytes. Headers with binary data are not refolded, regardless of the `refold_header` setting, since there is no way to know whether the binary data consists of single byte characters or multibyte characters.

The following instances of *EmailPolicy* provide defaults suitable for specific application domains. Note that in the future the behavior of these instances (in particular the HTTP instance) may be adjusted to conform even more closely to the RFCs relevant to their domains.

`email.policy.default`

An instance of *EmailPolicy* with all defaults unchanged. This policy uses the standard Python `\n` line endings rather than the RFC-correct `\r\n`.

`email.policy.SMTP`

Suitable for serializing messages in conformance with the email RFCs. Like `default`, but with `linesep` set to `\r\n`, which is RFC compliant.

`email.policy.SMTPUTF8`

The same as SMTP except that `utf8` is `True`. Useful for serializing messages to a message store without using encoded words in the headers. Should only be used for SMTP transmission if the sender or recipient addresses have non-ASCII characters (the `smtplib.SMTP.send_message()` method handles this automatically).

`email.policy.HTTP`

Suitable for serializing headers with for use in HTTP traffic. Like SMTP except that `max_line_length` is set to `None` (unlimited).

`email.policy.strict`

Convenience instance. The same as `default` except that `raise_on_defect` is set to `True`. This allows any policy to be made strict by writing:

```
somepolicy + policy.strict
```

With all of these *EmailPolicies*, the effective API of the email package is changed from the Python 3.2 API in the following ways:

- Setting a header on a *Message* results in that header being parsed and a header object created.
- Fetching a header value from a *Message* results in that header being parsed and a header object created and returned.
- Any header object, or any header that is refolded due to the policy settings, is folded using an algorithm that fully implements the RFC folding algorithms, including knowing where encoded words are required and allowed.

From the application view, this means that any header obtained through the *EmailMessage* is a header object with extra attributes, whose string value is the fully decoded unicode value of the header. Likewise, a header may be assigned a new value, or a new header created, using a unicode string, and the policy will take care of converting the unicode string into the correct RFC encoded form.

The header objects and their attributes are described in *headerregistry*.

class email.policy.Compat32 (**kw)

This concrete *Policy* is the backward compatibility policy. It replicates the behavior of the email package in Python 3.2. The *policy* module also defines an instance of this class, *compat32*, that is used as the default policy. Thus the default behavior of the email package is to maintain compatibility with Python 3.2.

The following attributes have values that are different from the *Policy* default:

mangle_from_

The default is True.

The class provides the following concrete implementations of the abstract methods of *Policy*:

header_source_parse (sourcelines)

The name is parsed as everything up to the ‘:’ and returned unmodified. The value is determined by stripping leading whitespace off the remainder of the first line, joining all subsequent lines together, and stripping any trailing carriage return or linefeed characters.

header_store_parse (name, value)

The name and value are returned unmodified.

header_fetch_parse (name, value)

If the value contains binary data, it is converted into a *Header* object using the unknown-8bit charset. Otherwise it is returned unmodified.

fold (name, value)

Headers are folded using the *Header* folding algorithm, which preserves existing line breaks in the value, and wraps each resulting line to the `max_line_length`. Non-ASCII binary data are CTE encoded using the unknown-8bit charset.

fold_binary (name, value)

Headers are folded using the *Header* folding algorithm, which preserves existing line breaks in the value, and wraps each resulting line to the `max_line_length`. If `cte_type` is 7bit, non-ascii binary data is CTE encoded using the unknown-8bit charset. Otherwise the original source header is used, with its existing line breaks and any (RFC invalid) binary data it may contain.

email.policy.compat32

An instance of *Compat32*, providing backward compatibility with the behavior of the email package in Python 3.2.

19.1.5 email.errors: 异常和缺陷类

源代码: [Lib/email/errors.py](#)

以下异常类在 *email.errors* 模块中定义:

exception email.errors.MessageError

这是 *email* 包可以引发的所有异常的基类。它源自标准异常 *Exception* 类, 这个类没有定义其他方法。

exception email.errors.MessageParseError

这是由 *Parser* 类引发的异常的基类。它派生自 *MessageError*。 *headerregistry* 使用的解析器也在内部使用这个类。

exception email.errors.HeaderParseError

在解析消息的 **RFC 5322** 标头时, 某些错误条件下会触发, 此类派生自 *MessageParseError*。如果在调用方法时内容类型未知, 则 *set_boundary()* 方法将引发此错误。当尝试创建一个看起来包含嵌入式标头的标头时 *Header* 可能会针对某些 base64 解码错误引发此错误 (也就是说, 应该是一个没有前导空格并且看起来像标题的延续行)。

exception `email.errors.BoundaryError`

已弃用和不再使用的。

exception `email.errors.MultipartConversionError`

当使用 `add_payload()` 将有效负载添加到 `Message` 对象时，有效负载已经是一个标量，而消息的 `Content-Type` 主类型不是 `multipart` 或者缺少时触发该异常。`MultipartConversionError` 多重继承自 `MessageError` 和内置的 `TypeError`。

Since `Message.add_payload()` is deprecated, this exception is rarely raised in practice. However the exception may also be raised if the `attach()` method is called on an instance of a class derived from `MIMENonMultipart` (e.g. `MIMEImage`).

Here is the list of the defects that the `FeedParser` can find while parsing messages. Note that the defects are added to the message where the problem was found, so for example, if a message nested inside a `multipart/alternative` had a malformed header, that nested message object would have a defect, but the containing messages would not.

All defect classes are subclassed from `email.errors.MessageDefect`.

- `NoBoundaryInMultipartDefect` – A message claimed to be a multipart, but had no *boundary* parameter.
- `StartBoundaryNotFoundDefect` – The start boundary claimed in the *Content-Type* header was never found.
- `CloseBoundaryNotFoundDefect` – A start boundary was found, but no corresponding close boundary was ever found.

3.3 新版功能.

- `FirstHeaderLineIsContinuationDefect` – The message had a continuation line as its first header line.
- `MisplacedEnvelopeHeaderDefect` - A “Unix From” header was found in the middle of a header block.
- `MissingHeaderBodySeparatorDefect` - A line was found while parsing headers that had no leading white space but contained no `:`. Parsing continues assuming that the line represents the first line of the body.

3.3 新版功能.

- `MalformedHeaderDefect` – A header was found that was missing a colon, or was otherwise malformed.

3.3 版后已移除: This defect has not been used for several Python versions.

- `MultipartInvariantViolationDefect` – A message claimed to be a *multipart*, but no subparts were found. Note that when a message has this defect, its `is_multipart()` method may return false even though its content type claims to be *multipart*.
- `InvalidBase64PaddingDefect` – When decoding a block of base64 encoded bytes, the padding was not correct. Enough padding is added to perform the decode, but the resulting decoded bytes may be invalid.
- `InvalidBase64CharactersDefect` – When decoding a block of base64 encoded bytes, characters outside the base64 alphabet were encountered. The characters are ignored, but the resulting decoded bytes may be invalid.
- `InvalidBase64LengthDefect` – When decoding a block of base64 encoded bytes, the number of non-padding base64 characters was invalid (1 more than a multiple of 4). The encoded block was kept as-is.

19.1.6 `email.headerregistry`: Custom Header Objects

Source code: `Lib/email/headerregistry.py`

3.6 新版功能:¹

¹ Originally added in 3.3 as a *provisional module*

Headers are represented by customized subclasses of `str`. The particular class used to represent a given header is determined by the `header_factory` of the `policy` in effect when the headers are created. This section documents the particular `header_factory` implemented by the email package for handling [RFC 5322](#) compliant email messages, which not only provides customized header objects for various header types, but also provides an extension mechanism for applications to add their own custom header types.

When using any of the policy objects derived from `EmailPolicy`, all headers are produced by `HeaderRegistry` and have `BaseHeader` as their last base class. Each header class has an additional base class that is determined by the type of the header. For example, many headers have the class `UnstructuredHeader` as their other base class. The specialized second class for a header is determined by the name of the header, using a lookup table stored in the `HeaderRegistry`. All of this is managed transparently for the typical application program, but interfaces are provided for modifying the default behavior for use by more complex applications.

The sections below first document the header base classes and their attributes, followed by the API for modifying the behavior of `HeaderRegistry`, and finally the support classes used to represent the data parsed from structured headers.

class `email.headerregistry.BaseHeader` (*name*, *value*)

name and *value* are passed to `BaseHeader` from the `header_factory` call. The string value of any header object is the *value* fully decoded to unicode.

This base class defines the following read-only properties:

name

The name of the header (the portion of the field before the `:`). This is exactly the value passed in the `header_factory` call for *name*; that is, case is preserved.

defects

A tuple of `HeaderDefect` instances reporting any RFC compliance problems found during parsing. The email package tries to be complete about detecting compliance issues. See the `errors` module for a discussion of the types of defects that may be reported.

max_count

The maximum number of headers of this type that can have the same *name*. A value of `None` means unlimited. The `BaseHeader` value for this attribute is `None`; it is expected that specialized header classes will override this value as needed.

`BaseHeader` also provides the following method, which is called by the email library code and should not in general be called by application programs:

fold (***, *policy*)

Return a string containing `linesep` characters as required to correctly fold the header according to *policy*. A `cte_type` of `8bit` will be treated as if it were `7bit`, since headers may not contain arbitrary binary data. If `utf8` is `False`, non-ASCII data will be [RFC 2047](#) encoded.

`BaseHeader` by itself cannot be used to create a header object. It defines a protocol that each specialized header cooperates with in order to produce the header object. Specifically, `BaseHeader` requires that the specialized class provide a `classmethod()` named `parse`. This method is called as follows:

```
parse(string, kwds)
```

kwds is a dictionary containing one pre-initialized key, `defects`. `defects` is an empty list. The `parse` method should append any detected defects to this list. On return, the *kwds* dictionary *must* contain values for at least the keys `decoded` and `defects`. `decoded` should be the string value for the header (that is, the header value fully decoded to unicode). The `parse` method should assume that *string* may contain content-transfer-encoded parts, but should correctly handle all valid unicode characters as well so that it can parse un-encoded header values.

`BaseHeader`'s `__new__` then creates the header instance, and calls its `init` method. The specialized class only needs to provide an `init` method if it wishes to set additional attributes beyond those provided by `BaseHeader` itself. Such an `init` method should look like this:

```
def init(self, *args, **kw):
    self._myattr = kw.pop('myattr')
    super().init(*args, **kw)
```

That is, anything extra that the specialized class puts in to the `kwds` dictionary should be removed and handled, and the remaining contents of `kw` (and `args`) passed to the `BaseHeader` `init` method.

class email.headerregistry.UnstructuredHeader

An “unstructured” header is the default type of header in [RFC 5322](#). Any header that does not have a specified syntax is treated as unstructured. The classic example of an unstructured header is the *Subject* header.

In [RFC 5322](#), an unstructured header is a run of arbitrary text in the ASCII character set. [RFC 2047](#), however, has an [RFC 5322](#) compatible mechanism for encoding non-ASCII text as ASCII characters within a header value. When a *value* containing encoded words is passed to the constructor, the `UnstructuredHeader` parser converts such encoded words into unicode, following the [RFC 2047](#) rules for unstructured text. The parser uses heuristics to attempt to decode certain non-compliant encoded words. Defects are registered in such cases, as well as defects for issues such as invalid characters within the encoded words or the non-encoded text.

This header type provides no additional attributes.

class email.headerregistry.DateHeader

[RFC 5322](#) specifies a very specific format for dates within email headers. The `DateHeader` parser recognizes that date format, as well as recognizing a number of variant forms that are sometimes found “in the wild”.

This header type provides the following additional attributes:

datetime

If the header value can be recognized as a valid date of one form or another, this attribute will contain a *datetime* instance representing that date. If the timezone of the input date is specified as `-0000` (indicating it is in UTC but contains no information about the source timezone), then *datetime* will be a naive *datetime*. If a specific timezone offset is found (including `+0000`), then *datetime* will contain an aware *datetime* that uses *datetime.timezone* to record the timezone offset.

The decoded value of the header is determined by formatting the *datetime* according to the [RFC 5322](#) rules; that is, it is set to:

```
email.utils.format_datetime(self.datetime)
```

When creating a `DateHeader`, *value* may be *datetime* instance. This means, for example, that the following code is valid and does what one would expect:

```
msg['Date'] = datetime(2011, 7, 15, 21)
```

Because this is a naive *datetime* it will be interpreted as a UTC timestamp, and the resulting value will have a timezone of `-0000`. Much more useful is to use the *localtime()* function from the *utils* module:

```
msg['Date'] = utils.localtime()
```

This example sets the date header to the current time and date using the current timezone offset.

class email.headerregistry.AddressHeader

Address headers are one of the most complex structured header types. The `AddressHeader` class provides a generic interface to any address header.

This header type provides the following additional attributes:

groups

A tuple of *Group* objects encoding the addresses and groups found in the header value. Addresses that are not part of a group are represented in this list as single-address Groups whose *display_name* is `None`.

addresses

A tuple of *Address* objects encoding all of the individual addresses from the header value. If the header value contains any groups, the individual addresses from the group are included in the list at the point where the group occurs in the value (that is, the list of addresses is “flattened” into a one dimensional list).

The decoded value of the header will have all encoded words decoded to unicode. *idna* encoded domain names are also decoded to unicode. The decoded value is set by *joining* the *str* value of the elements of the *groups* attribute with ' , '.

A list of *Address* and *Group* objects in any combination may be used to set the value of an address header. Group objects whose *display_name* is *None* will be interpreted as single addresses, which allows an address list to be copied with groups intact by using the list obtained from the *groups* attribute of the source header.

class email.headerregistry.**SingleAddressHeader**

A subclass of *AddressHeader* that adds one additional attribute:

address

The single address encoded by the header value. If the header value actually contains more than one address (which would be a violation of the RFC under the default *policy*), accessing this attribute will result in a *ValueError*.

Many of the above classes also have a Unique variant (for example, *UniqueUnstructuredHeader*). The only difference is that in the Unique variant, *max_count* is set to 1.

class email.headerregistry.**MIMEVersionHeader**

There is really only one valid value for the *MIME-Version* header, and that is 1.0. For future proofing, this header class supports other valid version numbers. If a version number has a valid value per **RFC 2045**, then the header object will have non-*None* values for the following attributes:

version

The version number as a string, with any whitespace and/or comments removed.

major

The major version number as an integer

minor

The minor version number as an integer

class email.headerregistry.**ParameterizedMIMEHeader**

MIME headers all start with the prefix ‘Content-’. Each specific header has a certain value, described under the class for that header. Some can also take a list of supplemental parameters, which have a common format. This class serves as a base for all the MIME headers that take parameters.

params

A dictionary mapping parameter names to parameter values.

class email.headerregistry.**ContentTypeHeader**

A *ParameterizedMIMEHeader* class that handles the *Content-Type* header.

content_type

The content type string, in the form maintype/subtype.

maintype**subtype**

class email.headerregistry.**ContentDispositionHeader**

A *ParameterizedMIMEHeader* class that handles the *Content-Disposition* header.

content-disposition

inline and *attachment* are the only valid values in common use.

class email.headerregistry.**ContentTransferEncoding**

Handles the *Content-Transfer-Encoding* header.

cte

Valid values are 7bit, 8bit, base64, and quoted-printable. See [RFC 2045](#) for more information.

class email.headerregistry.**HeaderRegistry** (*base_class=BaseHeader*, *de-*
fault_class=UnstructuredHeader,
use_default_map=True)

This is the factory used by *EmailPolicy* by default. *HeaderRegistry* builds the class used to create a header instance dynamically, using *base_class* and a specialized class retrieved from a registry that it holds. When a given header name does not appear in the registry, the class specified by *default_class* is used as the specialized class. When *use_default_map* is *True* (the default), the standard mapping of header names to classes is copied in to the registry during initialization. *base_class* is always the last class in the generated class' s `__bases__` list.

The default mappings are:

subject UniqueUnstructuredHeader

date UniqueDateHeader

resent-date DateHeader

orig-date UniqueDateHeader

sender UniqueSingleAddressHeader

resent-sender SingleAddressHeader

to UniqueAddressHeader

resent-to AddressHeader

cc UniqueAddressHeader

resent-cc AddressHeader

from UniqueAddressHeader

resent-from AddressHeader

reply-to UniqueAddressHeader

HeaderRegistry has the following methods:

map_to_type (*self*, *name*, *cls*)

name is the name of the header to be mapped. It will be converted to lower case in the registry. *cls* is the specialized class to be used, along with *base_class*, to create the class used to instantiate headers that match *name*.

__getitem__ (*name*)

Construct and return a class to handle creating a *name* header.

__call__ (*name*, *value*)

Retrieves the specialized header associated with *name* from the registry (using *default_class* if *name* does not appear in the registry) and composes it with *base_class* to produce a class, calls the constructed class' s constructor, passing it the same argument list, and finally returns the class instance created thereby.

The following classes are the classes used to represent data parsed from structured headers and can, in general, be used by an application program to construct structured values to assign to specific headers.

class email.headerregistry.**Address** (*display_name=*, *username=*, *domain=*,
addr_spec=None)

The class used to represent an email address. The general form of an address is:

```
[display_name] <username@domain>
```

或者:

```
username@domain
```

where each part must conform to specific syntax rules spelled out in [RFC 5322](#).

As a convenience *addr_spec* can be specified instead of *username* and *domain*, in which case *username* and *domain* will be parsed from the *addr_spec*. An *addr_spec* must be a properly RFC quoted string; if it is not *Address* will raise an error. Unicode characters are allowed and will be properly encoded when serialized. However, per the RFCs, unicode is *not* allowed in the username portion of the address.

display_name

The display name portion of the address, if any, with all quoting removed. If the address does not have a display name, this attribute will be an empty string.

username

The username portion of the address, with all quoting removed.

domain

The domain portion of the address.

addr_spec

The username@domain portion of the address, correctly quoted for use as a bare address (the second form shown above). This attribute is not mutable.

__str__()

The *str* value of the object is the address quoted according to [RFC 5322](#) rules, but with no Content Transfer Encoding of any non-ASCII characters.

To support SMTP ([RFC 5321](#)), *Address* handles one special case: if *username* and *domain* are both the empty string (or *None*), then the string value of the *Address* is *<>*.

class `email.headerregistry.Group` (*display_name=None, addresses=None*)

The class used to represent an address group. The general form of an address group is:

```
display_name: [address-list];
```

As a convenience for processing lists of addresses that consist of a mixture of groups and single addresses, a *Group* may also be used to represent single addresses that are not part of a group by setting *display_name* to *None* and providing a list of the single address as *addresses*.

display_name

The *display_name* of the group. If it is *None* and there is exactly one *Address* in *addresses*, then the *Group* represents a single address that is not in a group.

addresses

A possibly empty tuple of *Address* objects representing the addresses in the group.

__str__()

The *str* value of a *Group* is formatted according to [RFC 5322](#), but with no Content Transfer Encoding of any non-ASCII characters. If *display_name* is *None* and there is a single *Address* in the *addresses* list, the *str* value will be the same as the *str* of that single *Address*.

19.1.7 email.contentmanager: Managing MIME Content

Source code: [Lib/email/contentmanager.py](#)

3.6 新版功能:¹

class email.contentmanager.ContentManager

Base class for content managers. Provides the standard registry mechanisms to register converters between MIME content and other representations, as well as the `get_content` and `set_content` dispatch methods.

get_content (*msg*, **args*, ***kw*)

Look up a handler function based on the `mimetype` of *msg* (see next paragraph), call it, passing through all arguments, and return the result of the call. The expectation is that the handler will extract the payload from *msg* and return an object that encodes information about the extracted data.

To find the handler, look for the following keys in the registry, stopping with the first one found:

- the string representing the full MIME type (`maintype/subtype`)
- the string representing the `maintype`
- the empty string

If none of these keys produce a handler, raise a `KeyError` for the full MIME type.

set_content (*msg*, *obj*, **args*, ***kw*)

If the `maintype` is `multipart`, raise a `TypeError`; otherwise look up a handler function based on the type of *obj* (see next paragraph), call `clear_content()` on the *msg*, and call the handler function, passing through all arguments. The expectation is that the handler will transform and store *obj* into *msg*, possibly making other changes to *msg* as well, such as adding various MIME headers to encode information needed to interpret the stored data.

To find the handler, obtain the type of *obj* (`typ = type(obj)`), and look for the following keys in the registry, stopping with the first one found:

- the type itself (`typ`)
- the type's fully qualified name (`typ.__module__ + '.' + typ.__qualname__`).
- the type's `qualname` (`typ.__qualname__`)
- the type's `name` (`typ.__name__`).

If none of the above match, repeat all of the checks above for each of the types in the `MRO` (`typ.__mro__`). Finally, if no other key yields a handler, check for a handler for the key `None`. If there is no handler for `None`, raise a `KeyError` for the fully qualified name of the type.

Also add a `MIME-Version` header if one is not present (see also `MIMEPart`).

add_get_handler (*key*, *handler*)

Record the function *handler* as the handler for *key*. For the possible values of *key*, see `get_content()`.

add_set_handler (*typekey*, *handler*)

Record *handler* as the function to call when an object of a type matching *typekey* is passed to `set_content()`. For the possible values of *typekey*, see `set_content()`.

Content Manager Instances

Currently the email package provides only one concrete content manager, `raw_data_manager`, although more may be added in the future. `raw_data_manager` is the `content_manager` provided by `EmailPolicy` and its derivatives.

¹ Originally added in 3.4 as a *provisional module*

`email.contentmanager.raw_data_manager`

This content manager provides only a minimum interface beyond that provided by `Message` itself: it deals only with text, raw byte strings, and `Message` objects. Nevertheless, it provides significant advantages compared to the base API: `get_content` on a text part will return a unicode string without the application needing to manually decode it, `set_content` provides a rich set of options for controlling the headers added to a part and controlling the content transfer encoding, and it enables the use of the various `add_` methods, thereby simplifying the creation of multipart messages.

`email.contentmanager.get_content(msg, errors='replace')`

Return the payload of the part as either a string (for text parts), an `EmailMessage` object (for message/rfc822 parts), or a bytes object (for all other non-multipart types). Raise a `KeyError` if called on a multipart. If the part is a text part and `errors` is specified, use it as the error handler when decoding the payload to unicode. The default error handler is `replace`.

`email.contentmanager.set_content(msg, <'str'>, subtype="plain", charset='utf-8' cte=None, disposition=None, filename=None, cid=None, params=None, headers=None)`

`email.contentmanager.set_content(msg, <'bytes'>, maintype, subtype, cte="base64", disposition=None, filename=None, cid=None, params=None, headers=None)`

`email.contentmanager.set_content(msg, <'EmailMessage'>, cte=None, disposition=None, filename=None, cid=None, params=None, headers=None)`

Add headers and payload to `msg`:

Add a `Content-Type` header with a maintype/subtype value.

- For `str`, set the MIME maintype to `text`, and set the subtype to `subtype` if it is specified, or `plain` if it is not.
- For `bytes`, use the specified `maintype` and `subtype`, or raise a `TypeError` if they are not specified.
- For `EmailMessage` objects, set the maintype to `message`, and set the subtype to `subtype` if it is specified or `rfc822` if it is not. If `subtype` is partial, raise an error (bytes objects must be used to construct message/partial parts).

If `charset` is provided (which is valid only for `str`), encode the string to bytes using the specified character set. The default is `utf-8`. If the specified `charset` is a known alias for a standard MIME charset name, use the standard charset instead.

If `cte` is set, encode the payload using the specified content transfer encoding, and set the `Content-Transfer-Encoding` header to that value. Possible values for `cte` are `quoted-printable`, `base64`, `7bit`, `8bit`, and `binary`. If the input cannot be encoded in the specified encoding (for example, specifying a `cte` of `7bit` for an input that contains non-ASCII values), raise a `ValueError`.

- For `str` objects, if `cte` is not set use heuristics to determine the most compact encoding.
- For `EmailMessage`, per [RFC 2046](#), raise an error if a `cte` of `quoted-printable` or `base64` is requested for `subtype rfc822`, and for any `cte` other than `7bit` for `subtype external-body`. For `message/rfc822`, use `8bit` if `cte` is not specified. For all other values of `subtype`, use `7bit`.

注解: A `cte` of `binary` does not actually work correctly yet. The `EmailMessage` object as modified by `set_content` is correct, but `BytesGenerator` does not serialize it correctly.

If `disposition` is set, use it as the value of the `Content-Disposition` header. If not specified, and `filename` is specified, add the header with the value `attachment`. If `disposition` is not specified and `filename` is also not specified, do not add the header. The only valid values for `disposition` are `attachment` and `inline`.

If *filename* is specified, use it as the value of the *filename* parameter of the *Content-Disposition* header.

If *cid* is specified, add a *Content-ID* header with *cid* as its value.

If *params* is specified, iterate its *items* method and use the resulting (*key*, *value*) pairs to set additional parameters on the *Content-Type* header.

If *headers* is specified and is a list of strings of the form *headername: headervalue* or a list of header objects (distinguished from strings by having a *name* attribute), add the headers to *msg*.

19.1.8 email: 示例

以下是一些如何使用 *email* 包来读取、写入和发送简单电子邮件以及更复杂的 MIME 邮件的示例。

首先, 让我们看看如何创建和发送简单的文本消息 (文本内容和地址都可能包含 *unicode* 字符):

```
# Import smtplib for the actual sending function
import smtplib

# Import the email modules we'll need
from email.message import EmailMessage

# Open the plain text file whose name is in textfile for reading.
with open(textfile) as fp:
    # Create a text/plain message
    msg = EmailMessage()
    msg.set_content(fp.read())

# me == the sender's email address
# you == the recipient's email address
msg['Subject'] = 'The contents of %s' % textfile
msg['From'] = me
msg['To'] = you

# Send the message via our own SMTP server.
s = smtplib.SMTP('localhost')
s.send_message(msg)
s.quit()
```

解析 **RFC 822** 标题可以通过使用 *parser* 模块中的类来轻松完成:

```
# Import the email modules we'll need
from email.parser import BytesParser, Parser
from email.policy import default

# If the e-mail headers are in a file, uncomment these two lines:
# with open(messagefile, 'rb') as fp:
#     headers = BytesParser(policy=default).parse(fp)

# Or for parsing headers in a string (this is an uncommon operation), use:
headers = Parser(policy=default).parsestr(
    'From: Foo Bar <user@example.com>\n'
    'To: <someone_else@example.com>\n'
    'Subject: Test message\n'
    '\n'
    'Body would go here\n')
```

(下页继续)

(续上页)

```
# Now the header items can be accessed as a dictionary:
print('To: {}'.format(headers['to']))
print('From: {}'.format(headers['from']))
print('Subject: {}'.format(headers['subject']))

# You can also access the parts of the addresses:
print('Recipient username: {}'.format(headers['to'].addresses[0].username))
print('Sender name: {}'.format(headers['from'].addresses[0].display_name))
```

以下是如何发送包含可能在目录中的一系列家庭照片的 MIME 消息示例：

```
# Import smtplib for the actual sending function
import smtplib

# And imghdr to find the types of our images
import imghdr

# Here are the email package modules we'll need
from email.message import EmailMessage

# Create the container email message.
msg = EmailMessage()
msg['Subject'] = 'Our family reunion'
# me == the sender's email address
# family = the list of all recipients' email addresses
msg['From'] = me
msg['To'] = ', '.join(family)
msg.preamble = 'Our family reunion'

# Open the files in binary mode. Use imghdr to figure out the
# MIME subtype for each specific image.
for file in pngfiles:
    with open(file, 'rb') as fp:
        img_data = fp.read()
        msg.add_attachment(img_data, maintype='image',
                           subtype=imghdr.what(None, img_data))

# Send the email via our own SMTP server.
with smtplib.SMTP('localhost') as s:
    s.send_message(msg)
```

以下是如何将目录的全部内容作为电子邮件消息发送的示例：¹

```
#!/usr/bin/env python3

"""Send the contents of a directory as a MIME message."""

import os
import smtplib
# For guessing MIME type based on file name extension
import mimetypes

from argparse import ArgumentParser

from email.message import EmailMessage
```

(下页继续)

¹ 感谢 Matthew Dixon Cowles 提供最初的灵感和示例。

(续上页)

```

from email.policy import SMTP

def main():
    parser = ArgumentParser(description="""\
Send the contents of a directory as a MIME message.
Unless the -o option is given, the email is sent by forwarding to your local
SMTP server, which then does the normal delivery process. Your local machine
must be running an SMTP server.
""")
    parser.add_argument('-d', '--directory',
                        help="""Mail the contents of the specified directory,
otherwise use the current directory. Only the regular
files in the directory are sent, and we don't recurse to
subdirectories.""")
    parser.add_argument('-o', '--output',
                        metavar='FILE',
                        help="""Print the composed message to FILE instead of
sending the message to the SMTP server.""")
    parser.add_argument('-s', '--sender', required=True,
                        help='The value of the From: header (required)')
    parser.add_argument('-r', '--recipient', required=True,
                        action='append', metavar='RECIPIENT',
                        default=[], dest='recipients',
                        help='A To: header value (at least one required)')

    args = parser.parse_args()
    directory = args.directory
    if not directory:
        directory = '.'
    # Create the message
    msg = EmailMessage()
    msg['Subject'] = 'Contents of directory %s' % os.path.abspath(directory)
    msg['To'] = ', '.join(args.recipients)
    msg['From'] = args.sender
    msg.preamble = 'You will not see this in a MIME-aware mail reader.\n'

    for filename in os.listdir(directory):
        path = os.path.join(directory, filename)
        if not os.path.isfile(path):
            continue
        # Guess the content type based on the file's extension. Encoding
        # will be ignored, although we should check for simple things like
        # gzip'd or compressed files.
        ctype, encoding = mimetypes.guess_type(path)
        if ctype is None or encoding is not None:
            # No guess could be made, or the file is encoded (compressed), so
            # use a generic bag-of-bits type.
            ctype = 'application/octet-stream'
        maintype, subtype = ctype.split('/', 1)
        with open(path, 'rb') as fp:
            msg.add_attachment(fp.read(),
                              maintype=maintype,
                              subtype=subtype,
                              filename=filename)

    # Now send or store the message
    if args.output:

```

(下页继续)

(续上页)

```

        with open(args.output, 'wb') as fp:
            fp.write(msg.as_bytes(policy=SMTP))
    else:
        with smtplib.SMTP('localhost') as s:
            s.send_message(msg)

if __name__ == '__main__':
    main()

```

以下是如何将上述 MIME 消息解压缩到文件目录中的示例：

```

#!/usr/bin/env python3

"""Unpack a MIME message into a directory of files."""

import os
import email
import mimetypes

from email.policy import default
from argparse import ArgumentParser

def main():
    parser = ArgumentParser(description="""\
Unpack a MIME message into a directory of files.
""")
    parser.add_argument('-d', '--directory', required=True,
                        help="""Unpack the MIME message into the named
                        directory, which will be created if it doesn't already
                        exist.""")
    parser.add_argument('msgfile')
    args = parser.parse_args()

    with open(args.msgfile, 'rb') as fp:
        msg = email.message_from_binary_file(fp, policy=default)

    try:
        os.mkdir(args.directory)
    except FileExistsError:
        pass

    counter = 1
    for part in msg.walk():
        # multipart/* are just containers
        if part.get_content_maintype() == 'multipart':
            continue
        # Applications should really sanitize the given filename so that an
        # email message can't be used to overwrite important files
        filename = part.get_filename()
        if not filename:
            ext = mimetypes.guess_extension(part.get_content_type())
            if not ext:
                # Use a generic bag-of-bits extension

```

(下页继续)

(续上页)

```

        ext = '.bin'
        filename = 'part-%03d%s' % (counter, ext)
        counter += 1
        with open(os.path.join(args.directory, filename), 'wb') as fp:
            fp.write(part.get_payload(decode=True))

if __name__ == '__main__':
    main()

```

以下是如何使用备用纯文本版本创建 HTML 消息的示例。为了让事情变得更有趣，我们在 html 部分中包含了一个相关的图像，我们保存了一份我们要发送的内容到硬盘中，然后发送它。

```

#!/usr/bin/env python3

import smtplib

from email.message import EmailMessage
from email.headerregistry import Address
from email.utils import make_msgid

# Create the base text message.
msg = EmailMessage()
msg['Subject'] = "Ayons asperges pour le déjeuner"
msg['From'] = Address("Pepé Le Pew", "pepe", "example.com")
msg['To'] = (Address("Penelope Pussycat", "penelope", "example.com"),
            Address("Fabrette Pussycat", "fabrette", "example.com"))
msg.set_content("""\
Salut!

Cela ressemble à un excellent recipie[1] déjeuner.

[1] http://www.yummly.com/recipe/Roasted-Asparagus-Epicurious-203718

--Pepé
""")

# Add the html version. This converts the message into a multipart/alternative
# container, with the original text message as the first part and the new html
# message as the second part.
asparagus_cid = make_msgid()
msg.add_alternative("""\
<html>
  <head></head>
  <body>
    <p>Salut!</p>
    <p>Cela ressemble à un excellent
      <a href="http://www.yummly.com/recipe/Roasted-Asparagus-Epicurious-203718">
        recipie
      </a> déjeuner.
    </p>
    
  </body>
</html>
""".format(asparagus_cid=asparagus_cid[1:-1]), subtype='html')
# note that we needed to peel the <> off the msgid for use in the html.

```

(下页继续)

(续上页)

```

# Now add the related image to the html part.
with open("roasted-asparagus.jpg", 'rb') as img:
    msg.get_payload()[1].add_related(img.read(), 'image', 'jpeg',
                                     cid=asparagus_cid)

# Make a local copy of what we are going to send.
with open('outgoing.msg', 'wb') as f:
    f.write(bytes(msg))

# Send the message via local SMTP server.
with smtplib.SMTP('localhost') as s:
    s.send_message(msg)

```

如果我们发送最后一个示例中的消息，这是我们可以处理它的一种方法：

```

import os
import sys
import tempfile
import mimetypes
import webbrowser

# Import the email modules we'll need
from email import policy
from email.parser import BytesParser

# An imaginary module that would make this work and be safe.
from imaginary import magic_html_parser

# In a real program you'd get the filename from the arguments.
with open('outgoing.msg', 'rb') as fp:
    msg = BytesParser(policy=policy.default).parse(fp)

# Now the header items can be accessed as a dictionary, and any non-ASCII will
# be converted to unicode:
print('To:', msg['to'])
print('From:', msg['from'])
print('Subject:', msg['subject'])

# If we want to print a preview of the message content, we can extract whatever
# the least formatted payload is and print the first three lines. Of course,
# if the message has no plain text part printing the first three lines of html
# is probably useless, but this is just a conceptual example.
simplest = msg.get_body(preferencelist=('plain', 'html'))
print()
print(''.join(simplest.get_content().splitlines(keepends=True)[:3]))

ans = input("View full message?")
if ans.lower()[0] == 'n':
    sys.exit()

# We can extract the richest alternative in order to display it:
richest = msg.get_body()
partfiles = {}
if richest['content-type'].maintype == 'text':
    if richest['content-type'].subtype == 'plain':

```

(下页继续)

(续上页)

```

    for line in richest.get_content().splitlines():
        print(line)
    sys.exit()
elif richest['content-type'].subtype == 'html':
    body = richest
else:
    print("Don't know how to display {}".format(richest.get_content_type()))
    sys.exit()
elif richest['content-type'].content_type == 'multipart/related':
    body = richest.get_body(preferencelist=('html'))
    for part in richest.iter_attachments():
        fn = part.get_filename()
        if fn:
            extension = os.path.splitext(part.get_filename())[1]
        else:
            extension = mimetypes.guess_extension(part.get_content_type())
        with tempfile.NamedTemporaryFile(suffix=extension, delete=False) as f:
            f.write(part.get_content())
            # again strip the <> to go from email form of cid to html form.
            partfiles[part['content-id'][1:-1]] = f.name
else:
    print("Don't know how to display {}".format(richest.get_content_type()))
    sys.exit()
with tempfile.NamedTemporaryFile(mode='w', delete=False) as f:
    # The magic_html_parser has to rewrite the href="cid:..." attributes to
    # point to the filenames in partfiles. It also has to do a safety-sanitize
    # of the html. It could be written using html.parser.
    f.write(magic_html_parser(body.get_content(), partfiles))
webbrowser.open(f.name)
os.remove(f.name)
for fn in partfiles.values():
    os.remove(fn)

# Of course, there are lots of email messages that could break this simple
# minded program, but it will handle the most common ones.

```

直到输出提示，上面的输出是：

```

To: Penelope Pussycat <penelope@example.com>, Fabrette Pussycat <fabrette@example.com>
From: Pepé Le Pew <pepe@example.com>
Subject: Ayons asperges pour le déjeuner

Salut!

Cela ressemble à un excellent recipie[1] déjeuner.

```

Legacy API:

19.1.9 email.message.Message: Representing an email message using the compat32 API

The `Message` class is very similar to the `EmailMessage` class, without the methods added by that class, and with the default behavior of certain other methods being slightly different. We also document here some methods that, while supported by the `EmailMessage` class, are not recommended unless you are dealing with legacy code.

The philosophy and structure of the two classes is otherwise the same.

This document describes the behavior under the default (for *Message*) policy *Compat32*. If you are going to use another policy, you should be using the *EmailMessage* class instead.

An email message consists of *headers* and a *payload*. Headers must be **RFC 5233** style names and values, where the field name and value are separated by a colon. The colon is not part of either the field name or the field value. The payload may be a simple text message, or a binary object, or a structured sequence of sub-messages each with their own set of headers and their own payload. The latter type of payload is indicated by the message having a MIME type such as *multipart/** or *message/rfc822*.

The conceptual model provided by a *Message* object is that of an ordered dictionary of headers with additional methods for accessing both specialized information from the headers, for accessing the payload, for generating a serialized version of the message, and for recursively walking over the object tree. Note that duplicate headers are supported but special methods must be used to access them.

The *Message* pseudo-dictionary is indexed by the header names, which must be ASCII values. The values of the dictionary are strings that are supposed to contain only ASCII characters; there is some special handling for non-ASCII input, but it doesn't always produce the correct results. Headers are stored and returned in case-preserving form, but field names are matched case-insensitively. There may also be a single envelope header, also known as the *Unix-From* header or the *From_* header. The *payload* is either a string or bytes, in the case of simple message objects, or a list of *Message* objects, for MIME container documents (e.g. *multipart/** and *message/rfc822*).

Here are the methods of the *Message* class:

class email.message.*Message* (*policy=compat32*)

If *policy* is specified (it must be an instance of a *policy* class) use the rules it specifies to update and serialize the representation of the message. If *policy* is not set, use the *compat32* policy, which maintains backward compatibility with the Python 3.2 version of the email package. For more information see the *policy* documentation.

在 3.3 版更改: The *policy* keyword argument was added.

as_string (*unixfrom=False, maxheaderlen=0, policy=None*)

Return the entire message flattened as a string. When optional *unixfrom* is true, the envelope header is included in the returned string. *unixfrom* defaults to *False*. For backward compatibility reasons, *maxheaderlen* defaults to 0, so if you want a different value you must override it explicitly (the value specified for *max_line_length* in the policy will be ignored by this method). The *policy* argument may be used to override the default policy obtained from the message instance. This can be used to control some of the formatting produced by the method, since the specified *policy* will be passed to the *Generator*.

Flattening the message may trigger changes to the *Message* if defaults need to be filled in to complete the transformation to a string (for example, MIME boundaries may be generated or modified).

Note that this method is provided as a convenience and may not always format the message the way you want. For example, by default it does not do the mangling of lines that begin with *From* that is required by the unix mbox format. For more flexibility, instantiate a *Generator* instance and use its *flatten()* method directly. For example:

```
from io import StringIO
from email.generator import Generator
fp = StringIO()
g = Generator(fp, mangle_from_=True, maxheaderlen=60)
g.flatten(msg)
text = fp.getvalue()
```

If the message object contains binary data that is not encoded according to RFC standards, the non-compliant data will be replaced by unicode “unknown character” code points. (See also *as_bytes()* and *BytesGenerator*.)

在 3.4 版更改: the *policy* keyword argument was added.

__str__()

Equivalent to `as_string()`. Allows `str(msg)` to produce a string containing the formatted message.

as_bytes (*unixfrom=False, policy=None*)

Return the entire message flattened as a bytes object. When optional *unixfrom* is true, the envelope header is included in the returned string. *unixfrom* defaults to `False`. The *policy* argument may be used to override the default policy obtained from the message instance. This can be used to control some of the formatting produced by the method, since the specified *policy* will be passed to the `BytesGenerator`.

Flattening the message may trigger changes to the `Message` if defaults need to be filled in to complete the transformation to a string (for example, MIME boundaries may be generated or modified).

Note that this method is provided as a convenience and may not always format the message the way you want. For example, by default it does not do the mangling of lines that begin with `From` that is required by the unix mbox format. For more flexibility, instantiate a `BytesGenerator` instance and use its `flatten()` method directly. For example:

```
from io import BytesIO
from email.generator import BytesGenerator
fp = BytesIO()
g = BytesGenerator(fp, mangle_from_=True, maxheaderlen=60)
g.flatten(msg)
text = fp.getvalue()
```

3.4 新版功能.

__bytes__()

Equivalent to `as_bytes()`. Allows `bytes(msg)` to produce a bytes object containing the formatted message.

3.4 新版功能.

is_multipart()

Return `True` if the message's payload is a list of sub-`Message` objects, otherwise return `False`. When `is_multipart()` returns `False`, the payload should be a string object (which might be a CTE encoded binary payload). (Note that `is_multipart()` returning `True` does not necessarily mean that "`msg.get_content_maintype() == 'multipart'`" will return the `True`. For example, `is_multipart` will return `True` when the `Message` is of type `message/rfc822`.)

set_unixfrom (*unixfrom*)

Set the message's envelope header to *unixfrom*, which should be a string.

get_unixfrom()

Return the message's envelope header. Defaults to `None` if the envelope header was never set.

attach (*payload*)

Add the given *payload* to the current payload, which must be `None` or a list of `Message` objects before the call. After the call, the payload will always be a list of `Message` objects. If you want to set the payload to a scalar object (e.g. a string), use `set_payload()` instead.

This is a legacy method. On the `EmailMessage` class its functionality is replaced by `set_content()` and the related `make` and `add` methods.

get_payload (*i=None, decode=False*)

Return the current payload, which will be a list of `Message` objects when `is_multipart()` is `True`, or a string when `is_multipart()` is `False`. If the payload is a list and you mutate the list object, you modify the message's payload in place.

With optional argument *i*, `get_payload()` will return the *i*-th element of the payload, counting from zero, if `is_multipart()` is `True`. An `IndexError` will be raised if *i* is less than 0 or greater than or equal

to the number of items in the payload. If the payload is a string (i.e. `is_multipart()` is False) and `i` is given, a `TypeError` is raised.

Optional `decode` is a flag indicating whether the payload should be decoded or not, according to the `Content-Transfer-Encoding` header. When True and the message is not a multipart, the payload will be decoded if this header's value is quoted-printable or base64. If some other encoding is used, or `Content-Transfer-Encoding` header is missing, the payload is returned as-is (undecoded). In all cases the returned value is binary data. If the message is a multipart and the `decode` flag is True, then None is returned. If the payload is base64 and it was not perfectly formed (missing padding, characters outside the base64 alphabet), then an appropriate defect will be added to the message's defect property (`InvalidBase64PaddingDefect` or `InvalidBase64CharactersDefect`, respectively).

When `decode` is False (the default) the body is returned as a string without decoding the `Content-Transfer-Encoding`. However, for a `Content-Transfer-Encoding` of 8bit, an attempt is made to decode the original bytes using the charset specified by the `Content-Type` header, using the replace error handler. If no charset is specified, or if the charset given is not recognized by the email package, the body is decoded using the default ASCII charset.

This is a legacy method. On the `EmailMessage` class its functionality is replaced by `get_content()` and `iter_parts()`.

set_payload (payload, charset=None)

Set the entire message object's payload to `payload`. It is the client's responsibility to ensure the payload invariants. Optional `charset` sets the message's default character set; see `set_charset()` for details.

This is a legacy method. On the `EmailMessage` class its functionality is replaced by `set_content()`.

set_charset (charset)

Set the character set of the payload to `charset`, which can either be a `Charset` instance (see `email.charset`), a string naming a character set, or None. If it is a string, it will be converted to a `Charset` instance. If `charset` is None, the `charset` parameter will be removed from the `Content-Type` header (the message will not be otherwise modified). Anything else will generate a `TypeError`.

If there is no existing `MIME-Version` header one will be added. If there is no existing `Content-Type` header, one will be added with a value of `text/plain`. Whether the `Content-Type` header already exists or not, its `charset` parameter will be set to `charset.output_charset`. If `charset.input_charset` and `charset.output_charset` differ, the payload will be re-encoded to the `output_charset`. If there is no existing `Content-Transfer-Encoding` header, then the payload will be transfer-encoded, if needed, using the specified `Charset`, and a header with the appropriate value will be added. If a `Content-Transfer-Encoding` header already exists, the payload is assumed to already be correctly encoded using that `Content-Transfer-Encoding` and is not modified.

This is a legacy method. On the `EmailMessage` class its functionality is replaced by the `charset` parameter of the `email.message.EmailMessage.set_content()` method.

get_charset ()

Return the `Charset` instance associated with the message's payload.

This is a legacy method. On the `EmailMessage` class it always returns None.

The following methods implement a mapping-like interface for accessing the message's **RFC 2822** headers. Note that there are some semantic differences between these methods and a normal mapping (i.e. dictionary) interface. For example, in a dictionary there are no duplicate keys, but here there may be duplicate message headers. Also, in dictionaries there is no guaranteed order to the keys returned by `keys()`, but in a `Message` object, headers are always returned in the order they appeared in the original message, or were added to the message later. Any header deleted and then re-added are always appended to the end of the header list.

These semantic differences are intentional and are biased toward maximal convenience.

Note that in all cases, any envelope header present in the message is not included in the mapping interface.

In a model generated from bytes, any header values that (in contravention of the RFCs) contain non-ASCII bytes will, when retrieved through this interface, be represented as *Header* objects with a charset of *unknown-8bit*.

__len__()

Return the total number of headers, including duplicates.

__contains__(name)

Return true if the message object has a field named *name*. Matching is done case-insensitively and *name* should not include the trailing colon. Used for the *in* operator, e.g.:

```
if 'message-id' in myMessage:
    print('Message-ID:', myMessage['message-id'])
```

__getitem__(name)

Return the value of the named header field. *name* should not include the colon field separator. If the header is missing, None is returned; a *KeyError* is never raised.

Note that if the named field appears more than once in the message's headers, exactly which of those field values will be returned is undefined. Use the *get_all()* method to get the values of all the extant named headers.

__setitem__(name, val)

Add a header to the message with field name *name* and value *val*. The field is appended to the end of the message's existing fields.

Note that this does *not* overwrite or delete any existing header with the same name. If you want to ensure that the new header is the only one present in the message with field name *name*, delete the field first, e.g.:

```
del msg['subject']
msg['subject'] = 'Python roolz!'
```

__delitem__(name)

Delete all occurrences of the field with name *name* from the message's headers. No exception is raised if the named field isn't present in the headers.

keys()

Return a list of all the message's header field names.

values()

Return a list of all the message's field values.

items()

Return a list of 2-tuples containing all the message's field headers and values.

get(name, failobj=None)

Return the value of the named header field. This is identical to *__getitem__()* except that optional *failobj* is returned if the named header is missing (defaults to None).

Here are some additional useful methods:

get_all(name, failobj=None)

Return a list of all the values for the field named *name*. If there are no such named headers in the message, *failobj* is returned (defaults to None).

add_header(_name, _value, **_params)

Extended header setting. This method is similar to *__setitem__()* except that additional header parameters can be provided as keyword arguments. *_name* is the header field to add and *_value* is the *primary* value for the header.

For each item in the keyword argument dictionary *_params*, the key is taken as the parameter name, with underscores converted to dashes (since dashes are illegal in Python identifiers). Normally, the parameter will

be added as `key="value"` unless the value is `None`, in which case only the key will be added. If the value contains non-ASCII characters, it can be specified as a three tuple in the format `(CHARSET, LANGUAGE, VALUE)`, where `CHARSET` is a string naming the charset to be used to encode the value, `LANGUAGE` can usually be set to `None` or the empty string (see [RFC 2231](#) for other possibilities), and `VALUE` is the string value containing non-ASCII code points. If a three tuple is not passed and the value contains non-ASCII characters, it is automatically encoded in [RFC 2231](#) format using a `CHARSET` of `utf-8` and a `LANGUAGE` of `None`.

Here's an example:

```
msg.add_header('Content-Disposition', 'attachment', filename='bud.gif')
```

This will add a header that looks like

```
Content-Disposition: attachment; filename="bud.gif"
```

An example with non-ASCII characters:

```
msg.add_header('Content-Disposition', 'attachment',
               filename=('iso-8859-1', '', 'Fußballer.ppt'))
```

Which produces

```
Content-Disposition: attachment; filename*="iso-8859-1'Fu%DFballer.ppt"
```

replace_header (*_name*, *_value*)

Replace a header. Replace the first header found in the message that matches *_name*, retaining header order and field name case. If no matching header was found, a *KeyError* is raised.

get_content_type ()

Return the message's content type. The returned string is coerced to lower case of the form *maintype/subtype*. If there was no *Content-Type* header in the message the default type as given by *get_default_type* () will be returned. Since according to [RFC 2045](#), messages always have a default type, *get_content_type* () will always return a value.

[RFC 2045](#) defines a message's default type to be *text/plain* unless it appears inside a *multipart/digest* container, in which case it would be *message/rfc822*. If the *Content-Type* header has an invalid type specification, [RFC 2045](#) mandates that the default type be *text/plain*.

get_content_maintype ()

Return the message's main content type. This is the *maintype* part of the string returned by *get_content_type* ().

get_content_subtype ()

Return the message's sub-content type. This is the *subtype* part of the string returned by *get_content_type* ().

get_default_type ()

Return the default content type. Most messages have a default content type of *text/plain*, except for messages that are subparts of *multipart/digest* containers. Such subparts have a default content type of *message/rfc822*.

set_default_type (*ctype*)

Set the default content type. *ctype* should either be *text/plain* or *message/rfc822*, although this is not enforced. The default content type is not stored in the *Content-Type* header.

get_params (*failobj=None*, *header='content-type'*, *unquote=True*)

Return the message's *Content-Type* parameters, as a list. The elements of the returned list are 2-tuples of key/value pairs, as split on the '=' sign. The left hand side of the '=' is the key, while the right hand

side is the value. If there is no '=' sign in the parameter the value is the empty string, otherwise the value is as described in `get_param()` and is unquoted if optional `unquote` is `True` (the default).

Optional `failobj` is the object to return if there is no `Content-Type` header. Optional `header` is the header to search instead of `Content-Type`.

This is a legacy method. On the `EmailMessage` class its functionality is replaced by the `params` property of the individual header objects returned by the header access methods.

get_param (*param*, *failobj*=None, *header*='content-type', *unquote*=True)

Return the value of the `Content-Type` header's parameter *param* as a string. If the message has no `Content-Type` header or if there is no such parameter, then *failobj* is returned (defaults to None).

Optional *header* if given, specifies the message header to use instead of `Content-Type`.

Parameter keys are always compared case insensitively. The return value can either be a string, or a 3-tuple if the parameter was [RFC 2231](#) encoded. When it's a 3-tuple, the elements of the value are of the form (CHARSET, LANGUAGE, VALUE). Note that both CHARSET and LANGUAGE can be None, in which case you should consider VALUE to be encoded in the `us-ascii` charset. You can usually ignore LANGUAGE.

If your application doesn't care whether the parameter was encoded as in [RFC 2231](#), you can collapse the parameter value by calling `email.utils.collapse_rfc2231_value()`, passing in the return value from `get_param()`. This will return a suitably decoded Unicode string when the value is a tuple, or the original string unquoted if it isn't. For example:

```
rawparam = msg.get_param('foo')
param = email.utils.collapse_rfc2231_value(rawparam)
```

In any case, the parameter value (either the returned string, or the VALUE item in the 3-tuple) is always unquoted, unless `unquote` is set to `False`.

This is a legacy method. On the `EmailMessage` class its functionality is replaced by the `params` property of the individual header objects returned by the header access methods.

set_param (*param*, *value*, *header*='Content-Type', *requote*=True, *charset*=None, *language*=", *replace*=False)

Set a parameter in the `Content-Type` header. If the parameter already exists in the header, its value will be replaced with *value*. If the `Content-Type` header has not yet been defined for this message, it will be set to `text/plain` and the new parameter value will be appended as per [RFC 2045](#).

Optional *header* specifies an alternative header to `Content-Type`, and all parameters will be quoted as necessary unless optional *requote* is `False` (the default is `True`).

If optional *charset* is specified, the parameter will be encoded according to [RFC 2231](#). Optional *language* specifies the RFC 2231 language, defaulting to the empty string. Both *charset* and *language* should be strings.

If *replace* is `False` (the default) the header is moved to the end of the list of headers. If *replace* is `True`, the header will be updated in place.

在 3.4 版更改: `replace` keyword was added.

del_param (*param*, *header*='content-type', *requote*=True)

Remove the given parameter completely from the `Content-Type` header. The header will be re-written in place without the parameter or its value. All values will be quoted as necessary unless *requote* is `False` (the default is `True`). Optional *header* specifies an alternative to `Content-Type`.

set_type (*type*, *header*='Content-Type', *requote*=True)

Set the main type and subtype for the `Content-Type` header. *type* must be a string in the form `maintype/subtype`, otherwise a `ValueError` is raised.

This method replaces the *Content-Type* header, keeping all the parameters in place. If *quote* is `False`, this leaves the existing header's quoting as is, otherwise the parameters will be quoted (the default).

An alternative header can be specified in the *header* argument. When the *Content-Type* header is set a *MIME-Version* header is also added.

This is a legacy method. On the `EmailMessage` class its functionality is replaced by the `make_` and `add_` methods.

get_filename (*failobj=None*)

Return the value of the *filename* parameter of the *Content-Disposition* header of the message. If the header does not have a *filename* parameter, this method falls back to looking for the *name* parameter on the *Content-Type* header. If neither is found, or the header is missing, then *failobj* is returned. The returned string will always be unquoted as per `email.utils.unquote()`.

get_boundary (*failobj=None*)

Return the value of the *boundary* parameter of the *Content-Type* header of the message, or *failobj* if either the header is missing, or has no *boundary* parameter. The returned string will always be unquoted as per `email.utils.unquote()`.

set_boundary (*boundary*)

Set the *boundary* parameter of the *Content-Type* header to *boundary*. `set_boundary()` will always quote *boundary* if necessary. A `HeaderParseError` is raised if the message object has no *Content-Type* header.

Note that using this method is subtly different than deleting the old *Content-Type* header and adding a new one with the new boundary via `add_header()`, because `set_boundary()` preserves the order of the *Content-Type* header in the list of headers. However, it does *not* preserve any continuation lines which may have been present in the original *Content-Type* header.

get_content_charset (*failobj=None*)

Return the *charset* parameter of the *Content-Type* header, coerced to lower case. If there is no *Content-Type* header, or if that header has no *charset* parameter, *failobj* is returned.

Note that this method differs from `get_charset()` which returns the `Charset` instance for the default encoding of the message body.

get_charsets (*failobj=None*)

Return a list containing the character set names in the message. If the message is a *multipart*, then the list will contain one element for each subpart in the payload, otherwise, it will be a list of length 1.

Each item in the list will be a string which is the value of the *charset* parameter in the *Content-Type* header for the represented subpart. However, if the subpart has no *Content-Type* header, no *charset* parameter, or is not of the *text* main MIME type, then that item in the returned list will be *failobj*.

get_content_disposition ()

Return the lowercased value (without parameters) of the message's *Content-Disposition* header if it has one, or `None`. The possible values for this method are *inline*, *attachment* or `None` if the message follows [RFC 2183](#).

3.5 新版功能.

walk ()

The `walk()` method is an all-purpose generator which can be used to iterate over all the parts and subparts of a message object tree, in depth-first traversal order. You will typically use `walk()` as the iterator in a `for` loop; each iteration returns the next subpart.

Here's an example that prints the MIME type of every part of a multipart message structure:

```
>>> for part in msg.walk():
...     print(part.get_content_type())
multipart/report
text/plain
message/delivery-status
text/plain
text/plain
message/rfc822
text/plain
```

`walk` iterates over the subparts of any part where `is_multipart()` returns `True`, even though `msg.get_content_maintype() == 'multipart'` may return `False`. We can see this in our example by making use of the `_structure` debug helper function:

```
>>> for part in msg.walk():
...     print(part.get_content_maintype() == 'multipart'),
...           part.is_multipart())
True True
False False
False True
False False
False False
False True
False False
>>> _structure(msg)
multipart/report
    text/plain
message/delivery-status
    text/plain
    text/plain
message/rfc822
    text/plain
```

Here the message parts are not multipart, but they do contain subparts. `is_multipart()` returns `True` and `walk` descends into the subparts.

Message objects can also optionally contain two instance attributes, which can be used when generating the plain text of a MIME message.

preamble

The format of a MIME document allows for some text between the blank line following the headers, and the first multipart boundary string. Normally, this text is never visible in a MIME-aware mail reader because it falls outside the standard MIME armor. However, when viewing the raw text of the message, or when viewing the message in a non-MIME aware reader, this text can become visible.

The *preamble* attribute contains this leading extra-armor text for MIME documents. When the *Parser* discovers some text after the headers but before the first boundary string, it assigns this text to the message's *preamble* attribute. When the *Generator* is writing out the plain text representation of a MIME message, and it finds the message has a *preamble* attribute, it will write this text in the area between the headers and the first boundary. See *email.parser* and *email.generator* for details.

Note that if the message object has no preamble, the *preamble* attribute will be `None`.

epilogue

The *epilogue* attribute acts the same way as the *preamble* attribute, except that it contains text that appears between the last boundary and the end of the message.

You do not need to set the epilogue to the empty string in order for the *Generator* to print a newline at the end of the file.

defects

The *defects* attribute contains a list of all the problems found when parsing this message. See *email.errors* for a detailed description of the possible parsing defects.

19.1.10 email.mime: Creating email and MIME objects from scratch

Source code: [Lib/email/mime/](#)

This module is part of the legacy (Compat32) email API. Its functionality is partially replaced by the *contentmanager* in the new API, but in certain applications these classes may still be useful, even in non-legacy code.

Ordinarily, you get a message object structure by passing a file or some text to a parser, which parses the text and returns the root message object. However you can also build a complete message structure from scratch, or even individual *Message* objects by hand. In fact, you can also take an existing structure and add new *Message* objects, move them around, etc. This makes a very convenient interface for slicing-and-dicing MIME messages.

You can create a new object structure by creating *Message* instances, adding attachments and all the appropriate headers manually. For MIME messages though, the *email* package provides some convenient subclasses to make things easier.

Here are the classes:

```
class email.mime.base.MIMEBase (_maintype, _subtype, *, policy=compat32, **_params)
```

Module: *email.mime.base*

This is the base class for all the MIME-specific subclasses of *Message*. Ordinarily you won't create instances specifically of *MIMEBase*, although you could. *MIMEBase* is provided primarily as a convenient base class for more specific MIME-aware subclasses.

_maintype is the *Content-Type* major type (e.g. *text* or *image*), and *_subtype* is the *Content-Type* minor type (e.g. *plain* or *gif*). *_params* is a parameter key/value dictionary and is passed directly to *Message.add_header*.

If *policy* is specified, (defaults to the *compat32* policy) it will be passed to *Message*.

The *MIMEBase* class always adds a *Content-Type* header (based on *_maintype*, *_subtype*, and *_params*), and a *MIME-Version* header (always set to 1.0).

在 3.6 版更改: Added *policy* keyword-only parameter.

```
class email.mime.nonmultipart.MIMENonMultipart
```

Module: *email.mime.nonmultipart*

A subclass of *MIMEBase*, this is an intermediate base class for MIME messages that are not *multipart*. The primary purpose of this class is to prevent the use of the *attach()* method, which only makes sense for *multipart* messages. If *attach()* is called, a *MultipartConversionError* exception is raised.

```
class email.mime.multipart.MIMEMultipart (_subtype='mixed', boundary=None, _sub-  
parts=None, *, policy=compat32, **_params)
```

Module: *email.mime.multipart*

A subclass of *MIMEBase*, this is an intermediate base class for MIME messages that are *multipart*. Optional *_subtype* defaults to *mixed*, but can be used to specify the subtype of the message. A *Content-Type* header of *multipart/_subtype* will be added to the message object. A *MIME-Version* header will also be added.

Optional *boundary* is the multipart boundary string. When *None* (the default), the boundary is calculated when needed (for example, when the message is serialized).

_subparts is a sequence of initial subparts for the payload. It must be possible to convert this sequence to a list. You can always attach new subparts to the message by using the *Message.attach* method.

Optional *policy* argument defaults to *compat32*.

Additional parameters for the *Content-Type* header are taken from the keyword arguments, or passed into the *_params* argument, which is a keyword dictionary.

在 3.6 版更改: Added *policy* keyword-only parameter.

```
class email.mime.application.MIMEApplication (_data, _subtype='octet-stream', _en-
                                             coder=email.encoders.encode_base64,
                                             *, policy=compat32, **_params)
```

Module: `email.mime.application`

A subclass of *MIMENonMultipart*, the *MIMEApplication* class is used to represent MIME message objects of major type *application*. *_data* is a string containing the raw byte data. Optional *_subtype* specifies the MIME subtype and defaults to *octet-stream*.

Optional *_encoder* is a callable (i.e. function) which will perform the actual encoding of the data for transport. This callable takes one argument, which is the *MIMEApplication* instance. It should use *get_payload()* and *set_payload()* to change the payload to encoded form. It should also add any *Content-Transfer-Encoding* or other headers to the message object as necessary. The default encoding is base64. See the *email.encoders* module for a list of the built-in encoders.

Optional *policy* argument defaults to *compat32*.

_params are passed straight through to the base class constructor.

在 3.6 版更改: Added *policy* keyword-only parameter.

```
class email.mime.audio.MIMEAudio (_audiodata, _subtype=None, _en-
                                   coder=email.encoders.encode_base64, *, policy=compat32,
                                   **_params)
```

Module: `email.mime.audio`

A subclass of *MIMENonMultipart*, the *MIMEAudio* class is used to create MIME message objects of major type *audio*. *_audiodata* is a string containing the raw audio data. If this data can be decoded by the standard Python module *sndhdr*, then the subtype will be automatically included in the *Content-Type* header. Otherwise you can explicitly specify the audio subtype via the *_subtype* argument. If the minor type could not be guessed and *_subtype* was not given, then *TypeError* is raised.

Optional *_encoder* is a callable (i.e. function) which will perform the actual encoding of the audio data for transport. This callable takes one argument, which is the *MIMEAudio* instance. It should use *get_payload()* and *set_payload()* to change the payload to encoded form. It should also add any *Content-Transfer-Encoding* or other headers to the message object as necessary. The default encoding is base64. See the *email.encoders* module for a list of the built-in encoders.

Optional *policy* argument defaults to *compat32*.

_params are passed straight through to the base class constructor.

在 3.6 版更改: Added *policy* keyword-only parameter.

```
class email.mime.image.MIMEImage (_imagedata, _subtype=None, _en-
                                   coder=email.encoders.encode_base64, *, policy=compat32,
                                   **_params)
```

Module: `email.mime.image`

A subclass of *MIMENonMultipart*, the *MIMEImage* class is used to create MIME message objects of major type *image*. *_imagedata* is a string containing the raw image data. If this data can be decoded by the standard Python module *imghdr*, then the subtype will be automatically included in the *Content-Type* header. Otherwise you can explicitly specify the image subtype via the *_subtype* argument. If the minor type could not be guessed and *_subtype* was not given, then *TypeError* is raised.

Optional `_encoder` is a callable (i.e. function) which will perform the actual encoding of the image data for transport. This callable takes one argument, which is the `MIMEImage` instance. It should use `get_payload()` and `set_payload()` to change the payload to encoded form. It should also add any *Content-Transfer-Encoding* or other headers to the message object as necessary. The default encoding is base64. See the `email.encoders` module for a list of the built-in encoders.

Optional `policy` argument defaults to `compat32`.

`_params` are passed straight through to the `MIMEBase` constructor.

在 3.6 版更改: Added `policy` keyword-only parameter.

```
class email.mime.message.MIMEMessage (_msg, _subtype='rfc822', *, policy=compat32)
Module: email.mime.message
```

A subclass of `MIMENonMultipart`, the `MIMEMessage` class is used to create MIME objects of main type *message*. `_msg` is used as the payload, and must be an instance of class `Message` (or a subclass thereof), otherwise a `TypeError` is raised.

Optional `_subtype` sets the subtype of the message; it defaults to `rfc822`.

Optional `policy` argument defaults to `compat32`.

在 3.6 版更改: Added `policy` keyword-only parameter.

```
class email.mime.text.MIMEText (_text, _subtype='plain', _charset=None, *, policy=compat32)
Module: email.mime.text
```

A subclass of `MIMENonMultipart`, the `MIMEText` class is used to create MIME objects of major type *text*. `_text` is the string for the payload. `_subtype` is the minor type and defaults to `plain`. `_charset` is the character set of the text and is passed as an argument to the `MIMENonMultipart` constructor; it defaults to `us-ascii` if the string contains only `ascii` code points, and `utf-8` otherwise. The `_charset` parameter accepts either a string or a `Charset` instance.

Unless the `_charset` argument is explicitly set to `None`, the `MIMEText` object created will have both a *Content-Type* header with a `charset` parameter, and a *Content-Transfer-Encoding* header. This means that a subsequent `set_payload` call will not result in an encoded payload, even if a `charset` is passed in the `set_payload` command. You can “reset” this behavior by deleting the *Content-Transfer-Encoding* header, after which a `set_payload` call will automatically encode the new payload (and add a new *Content-Transfer-Encoding* header).

Optional `policy` argument defaults to `compat32`.

在 3.5 版更改: `_charset` also accepts `Charset` instances.

在 3.6 版更改: Added `policy` keyword-only parameter.

19.1.11 email.header: Internationalized headers

Source code: [Lib/email/header.py](#)

This module is part of the legacy (Compat32) email API. In the current API encoding and decoding of headers is handled transparently by the dictionary-like API of the `EmailMessage` class. In addition to uses in legacy code, this module can be useful in applications that need to completely control the character sets used when encoding headers.

本段落中的剩余文本是该模块的原始文档。

RFC 2822 is the base standard that describes the format of email messages. It derives from the older **RFC 822** standard which came into widespread use at a time when most email was composed of ASCII characters only. **RFC 2822** is a specification written assuming email contains only 7-bit ASCII characters.

Of course, as email has been deployed worldwide, it has become internationalized, such that language specific character sets can now be used in email messages. The base standard still requires email messages to be transferred using only 7-bit ASCII characters, so a slew of RFCs have been written describing how to encode email containing non-ASCII characters into **RFC 2822**-compliant format. These RFCs include **RFC 2045**, **RFC 2046**, **RFC 2047**, and **RFC 2231**. The *email* package supports these standards in its *email.header* and *email.charset* modules.

If you want to include non-ASCII characters in your email headers, say in the *Subject* or *To* fields, you should use the *Header* class and assign the field in the *Message* object to an instance of *Header* instead of using a string for the header value. Import the *Header* class from the *email.header* module. For example:

```
>>> from email.message import Message
>>> from email.header import Header
>>> msg = Message()
>>> h = Header('p\xF6stal', 'iso-8859-1')
>>> msg['Subject'] = h
>>> msg.as_string()
'Subject: =?iso-8859-1?q?p=F6stal?=\n\n'
```

Notice here how we wanted the *Subject* field to contain a non-ASCII character? We did this by creating a *Header* instance and passing in the character set that the byte string was encoded in. When the subsequent *Message* instance was flattened, the *Subject* field was properly **RFC 2047** encoded. MIME-aware mail readers would show this header using the embedded ISO-8859-1 character.

Here is the *Header* class description:

```
class email.header.Header (s=None, charset=None, maxlinelen=None, header_name=None, continuation_ws=' ', errors='strict')
```

Create a MIME-compliant header that can contain strings in different character sets.

Optional *s* is the initial header value. If *None* (the default), the initial header value is not set. You can later append to the header with *append()* method calls. *s* may be an instance of *bytes* or *str*, but see the *append()* documentation for semantics.

Optional *charset* serves two purposes: it has the same meaning as the *charset* argument to the *append()* method. It also sets the default character set for all subsequent *append()* calls that omit the *charset* argument. If *charset* is not provided in the constructor (the default), the *us-ascii* character set is used both as *s*'s initial charset and as the default for subsequent *append()* calls.

The maximum line length can be specified explicitly via *maxlinelen*. For splitting the first line to a shorter value (to account for the field header which isn't included in *s*, e.g. *Subject*) pass in the name of the field in *header_name*. The default *maxlinelen* is 76, and the default value for *header_name* is *None*, meaning it is not taken into account for the first line of a long, split header.

Optional *continuation_ws* must be **RFC 2822**-compliant folding whitespace, and is usually either a space or a hard tab character. This character will be prepended to continuation lines. *continuation_ws* defaults to a single space character.

Optional *errors* is passed straight through to the *append()* method.

```
append (s, charset=None, errors='strict')
```

Append the string *s* to the MIME header.

Optional *charset*, if given, should be a *Charset* instance (see *email.charset*) or the name of a character set, which will be converted to a *Charset* instance. A value of *None* (the default) means that the *charset* given in the constructor is used.

s may be an instance of *bytes* or *str*. If it is an instance of *bytes*, then *charset* is the encoding of that byte string, and a *UnicodeError* will be raised if the string cannot be decoded with that character set.

If *s* is an instance of *str*, then *charset* is a hint specifying the character set of the characters in the string.

In either case, when producing an [RFC 2822](#)-compliant header using [RFC 2047](#) rules, the string will be encoded using the output codec of the charset. If the string cannot be encoded using the output codec, a `UnicodeError` will be raised.

Optional *errors* is passed as the *errors* argument to the decode call if *s* is a byte string.

encode (*splitchars*=', \t', *maxlinelen*=None, *linesep*='\n')

Encode a message header into an RFC-compliant format, possibly wrapping long lines and encapsulating non-ASCII parts in base64 or quoted-printable encodings.

Optional *splitchars* is a string containing characters which should be given extra weight by the splitting algorithm during normal header wrapping. This is in very rough support of [RFC 2822](#)'s 'higher level syntactic breaks': split points preceded by a splitchar are preferred during line splitting, with the characters preferred in the order in which they appear in the string. Space and tab may be included in the string to indicate whether preference should be given to one over the other as a split point when other split chars do not appear in the line being split. Splitchars does not affect [RFC 2047](#) encoded lines.

maxlinelen, if given, overrides the instance's *s* value for the maximum line length.

linesep specifies the characters used to separate the lines of the folded header. It defaults to the most useful value for Python application code (`\n`), but `\r\n` can be specified in order to produce headers with RFC-compliant line separators.

在 3.2 版更改: Added the *linesep* argument.

The [Header](#) class also provides a number of methods to support standard operators and built-in functions.

__str__ ()

Returns an approximation of the [Header](#) as a string, using an unlimited line length. All pieces are converted to unicode using the specified encoding and joined together appropriately. Any pieces with a charset of 'unknown-8bit' are decoded as ASCII using the 'replace' error handler.

在 3.2 版更改: Added handling for the 'unknown-8bit' charset.

__eq__ (*other*)

This method allows you to compare two [Header](#) instances for equality.

__ne__ (*other*)

This method allows you to compare two [Header](#) instances for inequality.

The [email.header](#) module also provides the following convenient functions.

[email.header.decode_header](#) (*header*)

Decode a message header value without converting the character set. The header value is in *header*.

This function returns a list of (*decoded_string*, *charset*) pairs containing each of the decoded parts of the header. *charset* is None for non-encoded parts of the header, otherwise a lower case string containing the name of the character set specified in the encoded string.

Here's an example:

```
>>> from email.header import decode_header
>>> decode_header('=?iso-8859-1?q?p=F6stal?='')
[(b'p\xf6stal', 'iso-8859-1')]
```

[email.header.make_header](#) (*decoded_seq*, *maxlinelen*=None, *header_name*=None, *continuation_ws*='')

Create a [Header](#) instance from a sequence of pairs as returned by [decode_header](#) ().

[decode_header](#) () takes a header value string and returns a sequence of pairs of the format (*decoded_string*, *charset*) where *charset* is the name of the character set.

This function takes one of those sequence of pairs and returns a [Header](#) instance. Optional *maxlinelen*, *header_name*, and *continuation_ws* are as in the [Header](#) constructor.

19.1.12 `email.charset`: Representing character sets

Source code: [Lib/email/charset.py](#)

This module is part of the legacy (Compat32) email API. In the new API only the aliases table is used.

本段落中的剩余文本是该模块的原始文档。

This module provides a class `Charset` for representing character sets and character set conversions in email messages, as well as a character set registry and several convenience methods for manipulating this registry. Instances of `Charset` are used in several other modules within the `email` package.

Import this class from the `email.charset` module.

class `email.charset.Charset` (*input_charset=DEFAULT_CHARSET*)

Map character sets to their email properties.

This class provides information about the requirements imposed on email for a specific character set. It also provides convenience routines for converting between character sets, given the availability of the applicable codecs. Given a character set, it will do its best to provide information on how to use that character set in an email message in an RFC-compliant way.

Certain character sets must be encoded with quoted-printable or base64 when used in email headers or bodies. Certain character sets must be converted outright, and are not allowed in email.

Optional *input_charset* is as described below; it is always coerced to lower case. After being alias normalized it is also used as a lookup into the registry of character sets to find out the header encoding, body encoding, and output conversion codec to be used for the character set. For example, if *input_charset* is `iso-8859-1`, then headers and bodies will be encoded using quoted-printable and no output conversion codec is necessary. If *input_charset* is `eur-jp`, then headers will be encoded with base64, bodies will not be encoded, but output text will be converted from the `eur-jp` character set to the `iso-2022-jp` character set.

`Charset` instances have the following data attributes:

`input_charset`

The initial character set specified. Common aliases are converted to their *official* email names (e.g. `latin_1` is converted to `iso-8859-1`). Defaults to 7-bit `us-ascii`.

`header_encoding`

If the character set must be encoded before it can be used in an email header, this attribute will be set to `Charset.QP` (for quoted-printable), `Charset.BASE64` (for base64 encoding), or `Charset.SHORTEST` for the shortest of QP or BASE64 encoding. Otherwise, it will be `None`.

`body_encoding`

Same as *header_encoding*, but describes the encoding for the mail message's body, which indeed may be different than the header encoding. `Charset.SHORTEST` is not allowed for *body_encoding*.

`output_charset`

Some character sets must be converted before they can be used in email headers or bodies. If the *input_charset* is one of them, this attribute will contain the name of the character set output will be converted to. Otherwise, it will be `None`.

`input_codec`

The name of the Python codec used to convert the *input_charset* to Unicode. If no conversion codec is necessary, this attribute will be `None`.

`output_codec`

The name of the Python codec used to convert Unicode to the *output_charset*. If no conversion codec is necessary, this attribute will have the same value as the *input_codec*.

Charset instances also have the following methods:

get_body_encoding()

Return the content transfer encoding used for body encoding.

This is either the string `quoted-printable` or `base64` depending on the encoding used, or it is a function, in which case you should call the function with a single argument, the *Message* object being encoded. The function should then set the *Content-Transfer-Encoding* header itself to whatever is appropriate.

Returns the string `quoted-printable` if *body_encoding* is `QP`, returns the string `base64` if *body_encoding* is `BASE64`, and returns the string `7bit` otherwise.

get_output_charset()

Return the output character set.

This is the *output_charset* attribute if that is not `None`, otherwise it is *input_charset*.

header_encode(string)

Header-encode the string *string*.

The type of encoding (`base64` or `quoted-printable`) will be based on the *header_encoding* attribute.

header_encode_lines(string, maxlengths)

Header-encode a *string* by converting it first to bytes.

This is similar to *header_encode()* except that the string is fit into maximum line lengths as given by the argument *maxlengths*, which must be an iterator: each element returned from this iterator will provide the next maximum line length.

body_encode(string)

Body-encode the string *string*.

The type of encoding (`base64` or `quoted-printable`) will be based on the *body_encoding* attribute.

The *Charset* class also provides a number of methods to support standard operations and built-in functions.

__str__()

Returns *input_charset* as a string coerced to lower case. *__repr__()* is an alias for *__str__()*.

__eq__(other)

This method allows you to compare two *Charset* instances for equality.

__ne__(other)

This method allows you to compare two *Charset* instances for inequality.

The *email.charset* module also provides the following functions for adding new entries to the global character set, alias, and codec registries:

email.charset.add_charset(charset, header_enc=None, body_enc=None, output_charset=None)

Add character properties to the global registry.

charset is the input character set, and must be the canonical name of a character set.

Optional *header_enc* and *body_enc* is either *Charset.QP* for quoted-printable, *Charset.BASE64* for base64 encoding, *Charset.SHORTEST* for the shortest of quoted-printable or base64 encoding, or `None` for no encoding. *SHORTEST* is only valid for *header_enc*. The default is `None` for no encoding.

Optional *output_charset* is the character set that the output should be in. Conversions will proceed from input charset, to Unicode, to the output charset when the method *Charset.convert()* is called. The default is to output in the same character set as the input.

Both *input_charset* and *output_charset* must have Unicode codec entries in the module's character set-to-codec mapping; use *add_codec()* to add codecs the module does not know about. See the *codecs* module's documentation for more information.

The global character set registry is kept in the module global dictionary *CHARSETS*.

`email.charset.add_alias(alias, canonical)`

Add a character set alias. *alias* is the alias name, e.g. *latin-1*. *canonical* is the character set's canonical name, e.g. *iso-8859-1*.

The global charset alias registry is kept in the module global dictionary *ALIASES*.

`email.charset.add_codec(charset, codecname)`

Add a codec that map characters in the given character set to and from Unicode.

charset is the canonical name of a character set. *codecname* is the name of a Python codec, as appropriate for the second argument to the *str*'s *encode()* method.

19.1.13 email.encoders: Encoders

Source code: [Lib/email/encoders.py](#)

This module is part of the legacy (Compat32) email API. In the new API the functionality is provided by the *cte* parameter of the *set_content()* method.

本段落中的剩余文本是该模块的原始文档。

When creating *Message* objects from scratch, you often need to encode the payloads for transport through compliant mail servers. This is especially true for *image/** and *text/** type messages containing binary data.

email 包在其 *encoders* 模块中提供了一些方便的编码。这些编码器实际上由 *MIMEAudio* 和 *MIMEImage* 类构造函数使用，以提供默认编码。所有编码器函数只接受一个参数，即要编码的消息对象。它们通常提取有效数据，对其进行编码，并将有效数据重置为此新编码的值。他们还应该根据需要设置 *Content-Transfer-Encoding* 标头。

请注意，这些函数对于多段消息没有意义。它们必须应用到各个单独的段上面，而不是整体。如果直接传递一个多段类型的消息，会产生一个 *TypeError* 错误。

下面是提供的编码函数：

`email.encoders.encode_quopri(msg)`

将有效数据编码为 *quoted-printable* 形式，并将 *mailheader:Content-Transfer-Encoding* 标头设置为 *“quoted-printable”*¹。当大多数实际的数据是普通的可打印数据但包含少量不可打印的字符时，这是一个很好的编码。

`email.encoders.encode_base64(msg)`

Encodes the payload into base64 form and sets the *Content-Transfer-Encoding* header to *base64*. This is a good encoding to use when most of your payload is unprintable data since it is a more compact form than *quoted-printable*. The drawback of base64 encoding is that it renders the text non-human readable.

`email.encoders.encode_7or8bit(msg)`

This doesn't actually modify the message's payload, but it does set the *Content-Transfer-Encoding* header to either *7bit* or *8bit* as appropriate, based on the payload data.

`email.encoders.encode_noop(msg)`

这样什么都不会做；它甚至不会设置 *Content-Transfer-Encoding* 标头。

¹ Note that encoding with *encode_quopri()* also encodes all tabs and space characters in the data.

19.1.14 email.utils: 其他工具

源代码: `Lib/email/utils.py`

`email.utils` 模块提供如下几个工具

`email.utils.localtime (dt=None)`

Return local time as an aware datetime object. If called without arguments, return current time. Otherwise *dt* argument should be a `datetime` instance, and it is converted to the local time zone according to the system time zone database. If *dt* is naive (that is, `dt.tzinfo` is `None`), it is assumed to be in local time. In this case, a positive or zero value for *isdst* causes `localtime` to presume initially that summer time (for example, Daylight Saving Time) is or is not (respectively) in effect for the specified time. A negative value for *isdst* causes the `localtime` to attempt to divine whether summer time is in effect for the specified time.

3.3 新版功能.

`email.utils.make_msgid (idstring=None, domain=None)`

Returns a string suitable for an [RFC 2822](#)-compliant `Message-ID` header. Optional *idstring* if given, is a string used to strengthen the uniqueness of the message id. Optional *domain* if given provides the portion of the msgid after the '@'. The default is the local hostname. It is not normally necessary to override this default, but may be useful certain cases, such as a constructing distributed system that uses a consistent domain name across multiple hosts.

在 3.2 版更改: 增加了关键字 *domain*

The remaining functions are part of the legacy (Compat32) email API. There is no need to directly use these with the new API, since the parsing and formatting they provide is done automatically by the header parsing machinery of the new API.

`email.utils.quote (str)`

Return a new string with backslashes in *str* replaced by two backslashes, and double quotes replaced by backslash-double quote.

`email.utils.unquote (str)`

Return a new string which is an *unquoted* version of *str*. If *str* ends and begins with double quotes, they are stripped off. Likewise if *str* ends and begins with angle brackets, they are stripped off.

`email.utils.parseaddr (address)`

Parse address – which should be the value of some address-containing field such as *To* or *Cc* – into its constituent *realname* and *email address* parts. Returns a tuple of that information, unless the parse fails, in which case a 2-tuple of ('', '') is returned.

`email.utils.formataddr (pair, charset='utf-8')`

The inverse of `parseaddr()`, this takes a 2-tuple of the form (*realname*, *email_address*) and returns the string value suitable for a *To* or *Cc* header. If the first element of *pair* is false, then the second element is returned unmodified.

Optional *charset* is the character set that will be used in the [RFC 2047](#) encoding of the *realname* if the *realname* contains non-ASCII characters. Can be an instance of *str* or a *Charset*. Defaults to `utf-8`.

在 3.3 版更改: Added the *charset* option.

`email.utils.getaddresses (fieldvalues)`

This method returns a list of 2-tuples of the form returned by `parseaddr()`. *fieldvalues* is a sequence of header field values as might be returned by `Message.get_all`. Here's a simple example that gets all the recipients of a message:

```

from email.utils import getaddresses

tos = msg.get_all('to', [])
ccs = msg.get_all('cc', [])
resent_tos = msg.get_all('resent-to', [])
resent_ccs = msg.get_all('resent-cc', [])
all_recipients = getaddresses(tos + ccs + resent_tos + resent_ccs)

```

`email.utils.parsedate(date)`

Attempts to parse a date according to the rules in [RFC 2822](#), however, some mailers don't follow that format as specified, so `parsedate()` tries to guess correctly in such cases. `date` is a string containing an [RFC 2822](#) date, such as "Mon, 20 Nov 1995 19:12:08 -0500". If it succeeds in parsing the date, `parsedate()` returns a 9-tuple that can be passed directly to `time.mktime()`; otherwise `None` will be returned. Note that indexes 6, 7, and 8 of the result tuple are not usable.

`email.utils.parsedate_tz(date)`

Performs the same function as `parsedate()`, but returns either `None` or a 10-tuple; the first 9 elements make up a tuple that can be passed directly to `time.mktime()`, and the tenth is the offset of the date's timezone from UTC (which is the official term for Greenwich Mean Time)¹. If the input string has no timezone, the last element of the tuple returned is `None`. Note that indexes 6, 7, and 8 of the result tuple are not usable.

`email.utils.parsedate_to_datetime(date)`

The inverse of `format_datetime()`. Performs the same function as `parsedate()`, but on success returns a `datetime`. If the input date has a timezone of -0000, the `datetime` will be a naive `datetime`, and if the date is conforming to the RFCs it will represent a time in UTC but with no indication of the actual source timezone of the message the date comes from. If the input date has any other valid timezone offset, the `datetime` will be an aware `datetime` with the corresponding a `timezone tzinfo`.

3.3 新版功能.

`email.utils.mktime_tz(tuple)`

Turn a 10-tuple as returned by `parsedate_tz()` into a UTC timestamp (seconds since the Epoch). If the timezone item in the tuple is `None`, assume local time.

`email.utils.formatdate(timeval=None, localtime=False, usegmt=False)`

Returns a date string as per [RFC 2822](#), e.g.:

```
Fri, 09 Nov 2001 01:08:47 -0000
```

Optional `timeval` if given is a floating point time value as accepted by `time.gmtime()` and `time.localtime()`, otherwise the current time is used.

Optional `localtime` is a flag that when `True`, interprets `timeval`, and returns a date relative to the local timezone instead of UTC, properly taking daylight savings time into account. The default is `False` meaning UTC is used.

Optional `usegmt` is a flag that when `True`, outputs a date string with the timezone as an ascii string GMT, rather than a numeric -0000. This is needed for some protocols (such as HTTP). This only applies when `localtime` is `False`. The default is `False`.

`email.utils.format_datetime(dt, usegmt=False)`

Like `formatdate`, but the input is a `datetime` instance. If it is a naive `datetime`, it is assumed to be "UTC with no information about the source timezone", and the conventional -0000 is used for the timezone. If it is an aware `datetime`, then the numeric timezone offset is used. If it is an aware timezone with offset zero, then `usegmt` may be set to `True`, in which case the string GMT is used instead of the numeric timezone offset. This provides a way to generate standards conformant HTTP date headers.

3.3 新版功能.

¹ Note that the sign of the timezone offset is the opposite of the sign of the `time.timezone` variable for the same timezone; the latter variable follows the POSIX standard while this module follows [RFC 2822](#).

`email.utils.decode_rfc2231(s)`

Decode the string *s* according to [RFC 2231](#).

`email.utils.encode_rfc2231(s, charset=None, language=None)`

Encode the string *s* according to [RFC 2231](#). Optional *charset* and *language*, if given is the character set name and language name to use. If neither is given, *s* is returned as-is. If *charset* is given but *language* is not, the string is encoded using the empty string for *language*.

`email.utils.collapse_rfc2231_value(value, errors='replace', fallback_charset='us-ascii')`

When a header parameter is encoded in [RFC 2231](#) format, `Message.get_param` may return a 3-tuple containing the character set, language, and value. `collapse_rfc2231_value()` turns this into a unicode string. Optional *errors* is passed to the *errors* argument of `str`'s `encode()` method; it defaults to 'replace'. Optional *fallback_charset* specifies the character set to use if the one in the [RFC 2231](#) header is not known by Python; it defaults to 'us-ascii'.

For convenience, if the *value* passed to `collapse_rfc2231_value()` is not a tuple, it should be a string and it is returned unquoted.

`email.utils.decode_params(params)`

Decode parameters list according to [RFC 2231](#). *params* is a sequence of 2-tuples containing elements of the form (content-type, string-value).

备注

19.1.15 email.iterators: Iterators

Source code: [Lib/email/iterators.py](#)

Iterating over a message object tree is fairly easy with the `Message.walk` method. The `email.iterators` module provides some useful higher level iterations over message object trees.

`email.iterators.body_line_iterator(msg, decode=False)`

This iterates over all the payloads in all the subparts of *msg*, returning the string payloads line-by-line. It skips over all the subpart headers, and it skips over any subpart with a payload that isn't a Python string. This is somewhat equivalent to reading the flat text representation of the message from a file using `readline()`, skipping over all the intervening headers.

Optional *decode* is passed through to `Message.get_payload`.

`email.iterators.typed_subpart_iterator(msg, maintype='text', subtype=None)`

This iterates over all the subparts of *msg*, returning only those subparts that match the MIME type specified by *maintype* and *subtype*.

Note that *subtype* is optional; if omitted, then subpart MIME type matching is done only with the main type. *maintype* is optional too; it defaults to `text`.

Thus, by default `typed_subpart_iterator()` returns each subpart that has a MIME type of `text/*`.

The following function has been added as a useful debugging tool. It should *not* be considered part of the supported public interface for the package.

`email.iterators._structure(msg, fp=None, level=0, include_default=False)`

Prints an indented representation of the content types of the message object structure. For example:

```
>>> msg = email.message_from_file(somefile)
>>> _structure(msg)
multipart/mixed
```

(下页继续)

(续上页)

```

text/plain
text/plain
multipart/digest
    message/rfc822
        text/plain
    message/rfc822
        text/plain
    message/rfc822
        text/plain
    message/rfc822
        text/plain
    message/rfc822
        text/plain
    message/rfc822
        text/plain
text/plain

```

Optional *fp* is a file-like object to print the output to. It must be suitable for Python's `print()` function. *level* is used internally. *include_default*, if true, prints the default type as well.

参见:

Module **smtplib** SMTP (Simple Mail Transport Protocol) client

Module **poplib** POP (Post Office Protocol) client

Module **imaplib** IMAP (Internet Message Access Protocol) client

Module **nntplib** NNTP (Net News Transport Protocol) client

Module **mailbox** Tools for creating, reading, and managing collections of messages on disk using a variety standard formats.

Module **smtpd** SMTP server framework (primarily useful for testing)

19.2 json —JSON 编码和解码器

源代码: `Lib/json/__init__.py`

JSON (JavaScript Object Notation), 由 **RFC 7159** (which obsoletes **RFC 4627**) 和 **ECMA-404** 指定, 是一个受 JavaScript 的对象字面量语法启发的轻量级数据交换格式, 尽管它不仅仅是一个严格意义上的 JavaScript 的字典¹。

`json` 提供了与标准库 `marshal` 和 `pickle` 相似的 API 接口。

对基本的 Python 对象层次结构进行编码:

```

>>> import json
>>> json.dumps(['foo', {'bar': ('baz', None, 1.0, 2)}])
'["foo", {"bar": ["baz", null, 1.0, 2]}]'
>>> print(json.dumps("\"foo\\bar\""))
\"foo\\bar\"
>>> print(json.dumps('\\u1234'))
\"u1234\"
>>> print(json.dumps('\\\\'))
'\\\\'

```

(下页继续)

¹ As noted in the errata for **RFC 7159**, JSON permits literal U+2028 (LINE SEPARATOR) and U+2029 (PARAGRAPH SEPARATOR) characters in strings, whereas JavaScript (as of ECMA Script Edition 5.1) does not.

(续上页)

```

"\\"
>>> print(json.dumps({"c": 0, "b": 0, "a": 0}, sort_keys=True))
{"a": 0, "b": 0, "c": 0}
>>> from io import StringIO
>>> io = StringIO()
>>> json.dump(['streaming API'], io)
>>> io.getvalue()
'["streaming API"]'

```

紧凑编码:

```

>>> import json
>>> json.dumps([1, 2, 3, {'4': 5, '6': 7}], separators=(',', ':'))
'[1,2,3,{"4":5,"6":7}]'

```

美化输出:

```

>>> import json
>>> print(json.dumps({'4': 5, '6': 7}, sort_keys=True, indent=4))
{
    "4": 5,
    "6": 7
}

```

JSON 解码:

```

>>> import json
>>> json.loads('["foo", {"bar":["baz", null, 1.0, 2]}]')
['foo', {'bar': ['baz', None, 1.0, 2]}]
>>> json.loads('"\"foo\\bar\"')
'foo\x08ar'
>>> from io import StringIO
>>> io = StringIO('["streaming API"]')
>>> json.load(io)
['streaming API']

```

特殊 JSON 对象解码:

```

>>> import json
>>> def as_complex(dct):
...     if '__complex__' in dct:
...         return complex(dct['real'], dct['imag'])
...     return dct
...
>>> json.loads('{"__complex__": true, "real": 1, "imag": 2}',
... object_hook=as_complex)
(1+2j)
>>> import decimal
>>> json.loads('1.1', parse_float=decimal.Decimal)
Decimal('1.1')

```

扩展 `JSONEncoder`:

```

>>> import json
>>> class ComplexEncoder(json.JSONEncoder):
...     def default(self, obj):

```

(下页继续)

(续上页)

```

...     if isinstance(obj, complex):
...         return [obj.real, obj.imag]
...         # Let the base class default method raise the TypeError
...         return json.JSONEncoder.default(self, obj)
...
>>> json.dumps(2 + 1j, cls=ComplexEncoder)
'[2.0, 1.0]'
>>> ComplexEncoder().encode(2 + 1j)
'[2.0, 1.0]'
>>> list(ComplexEncoder().iterencode(2 + 1j))
['[2.0', ', ', '1.0', ', ', ']']

```

从命令行使用 `json.tool` 来验证并美化输出：

```

$ echo '{"json":"obj"}' | python -m json.tool
{
    "json": "obj"
}
$ echo '{1.2:3.4}' | python -m json.tool
Expecting property name enclosed in double quotes: line 1 column 2 (char 1)

```

详细文档请参见 [Command Line Interface](#)。

注解：JSON 是 [YAML 1.2](#) 的一个子集。由该模块的默认设置生成的 JSON（尤其是默认的“分隔符”设置值）也是 [YAML 1.0 and 1.1](#) 的一个子集。因此该模块也能够用于序列化为 [YAML](#)。

19.2.1 基本使用

`json.dump(obj, fp, *, skipkeys=False, ensure_ascii=True, check_circular=True, allow_nan=True, cls=None, indent=None, separators=None, default=None, sort_keys=False, **kw)`

使用这个 [conversion table](#) 来序列化 `obj` 为一个 JSON 格式的流并输出到 `fp`（一个支持 `.write()` 的 *file-like object*）。

如果 `skipkeys` 是 `true`（默认为 `False`），那么那些不是基本对象（包括 `str`、`int`、`float`、`bool`、`None`）的字典的键会被跳过；否则引发一个 `TypeError`。

`json` 模块始终产生 `str` 对象而非 `bytes` 对象。因此，`fp.write()` 必须支持 `str` 输入。

如果 `ensure_ascii` 是 `true`（即默认值），输出保证将所有输入的非 ASCII 字符转义。如果 `ensure_ascii` 是 `false`，这些字符会原样输出。

如果 `check_circular` 是为假值（默认为 `True`），那么容器类型的循环引用检验会被跳过并且循环引用会引发一个 `OverflowError`（或者更糟的情况）。

如果 `allow_nan` 是 `false`（默认为 `True`），那么在对严格 JSON 规格范围外的 `float` 类型值（`nan`、`inf` 和 `-inf`）进行序列化时会引发一个 `ValueError`。如果 `allow_nan` 是 `true`，则使用它们的 JavaScript 等价形式（`NaN`、`Infinity` 和 `-Infinity`）。

如果 `indent` 是一个非负整数或者字符串，那么 JSON 数组元素和对象成员会被美化输出为该值指定的缩进等级。如果缩进等级为零、负数或者 `""`，则只会添加换行符。`None`（默认值）选择最紧凑的表达。使用一个正整数会让每一层缩进同样数量的空格。如果 `*indent*` 是一个字符串（比如 `""\t"`），那个字符串会被用于缩进每一层。

在 3.2 版更改：允许使用字符串作为 `indent` 而不再仅仅是整数。

当指定时, *separators* 应当是一个 (*item_separator*, *key_separator*) 元组。当 *indent* 为 *None* 时, 默认值取 (*' , ' , ' : ' '*), 否则取 (*' , ' , ' : ' '*)。为了得到最紧凑的 JSON 表达式, 你应该指定其为 (*' , ' , ' : ' '*) 以消除空白字符。

在 3.4 版更改: 现当 *indent* 不是 *None* 时, 采用 (*' , ' , ' : ' '*) 作为默认值。

当 *default* 被指定时, 其应该是一个函数, 每当某个对象无法被序列化时它会被调用。它应该返回该对象的一个可以被 JSON 编码的版本或者引发一个 *TypeError*。如果没有被指定, 则会直接引发 *TypeError*。

如果 *sort_keys* 是 *true* (默认为 *False*), 那么字典的输出会以键的顺序排序。

为了使用一个自定义的 *JSONEncoder* 子类 (比如: 覆盖了 *default()* 方法来序列化额外的类型), 通过 *cls* 关键字参数来指定; 否则将使用 *JSONEncoder*。

在 3.6 版更改: 所有的可选参数现在是 *keyword-only* 的了。

注解: 与 *pickle* 和 *marshal* 不同, JSON 不是一个具有框架的协议, 所以尝试多次使用同一个 *fp* 调用 *dump()* 来序列化多个对象会产生一个不合规的 JSON 文件。

`json.dumps(obj, *, skipkeys=False, ensure_ascii=True, check_circular=True, allow_nan=True, cls=None, indent=None, separators=None, default=None, sort_keys=False, **kw)`
使用这个转换表 将 *obj* 序列化为 JSON 格式的 *str*。其参数的含义与 *dump()* 中的相同。

注解: JSON 中的键-值对中的键永远是 *str* 类型的。当一个对象被转化为 JSON 时, 字典中所有的键都会被强制转换为字符串。这所造成的结果是字典被转换为 JSON 然后转换回字典时可能和原来的不相等。换句话说, 如果 *x* 具有非字符串的键, 则有 `loads(dumps(x)) != x`。

`json.load(fp, *, cls=None, object_hook=None, parse_float=None, parse_int=None, parse_constant=None, object_pairs_hook=None, **kw)`
使用这个转换表 将 *fp* (一个支持 *.read()* 并包含一个 JSON 文档的 *text file* 或者 *binary file*) 反序列化为一个 Python 对象。

object_hook 是一个可选的函数, 它会被调用于每一个解码出的对象字面量 (即一个 *dict*)。 *object_hook* 的返回值会取代原本的 *dict*。这一特性能够被用于实现自定义解码器 (如 JSON-RPC 的类型提示)。

object_pairs_hook is an optional function that will be called with the result of any object literal decoded with an ordered list of pairs. The return value of *object_pairs_hook* will be used instead of the *dict*. This feature can be used to implement custom decoders that rely on the order that the key and value pairs are decoded (for example, *collections.OrderedDict()* will remember the order of insertion). If *object_hook* is also defined, the *object_pairs_hook* takes priority.

在 3.1 版更改: Added support for *object_pairs_hook*.

parse_float, if specified, will be called with the string of every JSON float to be decoded. By default, this is equivalent to `float(num_str)`. This can be used to use another datatype or parser for JSON floats (e.g. *decimal.Decimal*).

parse_int, if specified, will be called with the string of every JSON int to be decoded. By default, this is equivalent to `int(num_str)`. This can be used to use another datatype or parser for JSON integers (e.g. *float*).

parse_constant, if specified, will be called with one of the following strings: *'-Infinity'*, *'Infinity'*, *'NaN'*. This can be used to raise an exception if invalid JSON numbers are encountered.

在 3.1 版更改: *parse_constant* doesn't get called on *'null'*, *'true'*, *'false'* anymore.

To use a custom *JSONDecoder* subclass, specify it with the *cls* kwarg; otherwise *JSONDecoder* is used. Additional keyword arguments will be passed to the constructor of the class.

If the data being deserialized is not a valid JSON document, a `JSONDecodeError` will be raised.

在 3.6 版更改: 所有的可选参数现在是`keyword-only`的了。

在 3.6 版更改: `fp` can now be a `binary file`. The input encoding should be UTF-8, UTF-16 or UTF-32.

`json.loads(s, *, encoding=None, cls=None, object_hook=None, parse_float=None, parse_int=None, parse_constant=None, object_pairs_hook=None, **kw)`
 Deserialize `s` (a `str`, `bytes` or `bytearray` instance containing a JSON document) to a Python object using this [conversion table](#).

The other arguments have the same meaning as in `load()`, except `encoding` which is ignored and deprecated.

If the data being deserialized is not a valid JSON document, a `JSONDecodeError` will be raised.

在 3.6 版更改: `s` can now be of type `bytes` or `bytearray`. The input encoding should be UTF-8, UTF-16 or UTF-32.

19.2.2 Encoders and Decoders

`class json.JSONDecoder(*, object_hook=None, parse_float=None, parse_int=None, parse_constant=None, strict=True, object_pairs_hook=None)`

Simple JSON decoder.

Performs the following translations in decoding by default:

JSON	Python
object	dict
array	list
string	str
整数	int
非整数	float
true	True
false	False
null	None

It also understands NaN, Infinity, and -Infinity as their corresponding `float` values, which is outside the JSON spec.

`object_hook`, if specified, will be called with the result of every JSON object decoded and its return value will be used in place of the given `dict`. This can be used to provide custom deserializations (e.g. to support JSON-RPC class hinting).

`object_pairs_hook`, if specified will be called with the result of every JSON object decoded with an ordered list of pairs. The return value of `object_pairs_hook` will be used instead of the `dict`. This feature can be used to implement custom decoders that rely on the order that the key and value pairs are decoded (for example, `collections.OrderedDict()` will remember the order of insertion). If `object_hook` is also defined, the `object_pairs_hook` takes priority.

在 3.1 版更改: Added support for `object_pairs_hook`.

`parse_float`, if specified, will be called with the string of every JSON float to be decoded. By default, this is equivalent to `float(num_str)`. This can be used to use another datatype or parser for JSON floats (e.g. `decimal.Decimal`).

`parse_int`, if specified, will be called with the string of every JSON int to be decoded. By default, this is equivalent to `int(num_str)`. This can be used to use another datatype or parser for JSON integers (e.g. `float`).

parse_constant, if specified, will be called with one of the following strings: `'-Infinity'`, `'Infinity'`, `'NaN'`. This can be used to raise an exception if invalid JSON numbers are encountered.

If *strict* is false (True is the default), then control characters will be allowed inside strings. Control characters in this context are those with character codes in the 0–31 range, including `'\t'` (tab), `'\n'`, `'\r'` and `'\0'`.

If the data being deserialized is not a valid JSON document, a `JSONDecodeError` will be raised.

在 3.6 版更改: All parameters are now *keyword-only*.

decode (*s*)

Return the Python representation of *s* (a *str* instance containing a JSON document).

`JSONDecodeError` will be raised if the given JSON document is not valid.

raw_decode (*s*)

Decode a JSON document from *s* (a *str* beginning with a JSON document) and return a 2-tuple of the Python representation and the index in *s* where the document ended.

This can be used to decode a JSON document from a string that may have extraneous data at the end.

class `json.JSONEncoder` (*, *skipkeys=False*, *ensure_ascii=True*, *check_circular=True*, *allow_nan=True*, *sort_keys=False*, *indent=None*, *separators=None*, *default=None*)

Extensible JSON encoder for Python data structures.

Supports the following objects and types by default:

Python	JSON
dict	object
list, tuple	array
str	string
int, float, int- & float-derived Enums	number
True	true
False	false
None	null

在 3.4 版更改: Added support for int- and float-derived Enum classes.

To extend this to recognize other objects, subclass and implement a `default()` method with another method that returns a serializable object for `o` if possible, otherwise it should call the superclass implementation (to raise `TypeError`).

If *skipkeys* is false (the default), then it is a `TypeError` to attempt encoding of keys that are not *str*, *int*, *float* or `None`. If *skipkeys* is true, such items are simply skipped.

如果 *ensure_ascii* 是 true (即默认值), 输出保证将所有输入的非 ASCII 字符转义。如果 *ensure_ascii* 是 false, 这些字符会原样输出。

If *check_circular* is true (the default), then lists, dicts, and custom encoded objects will be checked for circular references during encoding to prevent an infinite recursion (which would cause an `OverflowError`). Otherwise, no such check takes place.

If *allow_nan* is true (the default), then NaN, Infinity, and -Infinity will be encoded as such. This behavior is not JSON specification compliant, but is consistent with most JavaScript based encoders and decoders. Otherwise, it will be a `ValueError` to encode such floats.

If *sort_keys* is true (default: False), then the output of dictionaries will be sorted by key; this is useful for regression tests to ensure that JSON serializations can be compared on a day-to-day basis.

如果 *indent* 是一个非负整数或者字符串, 那么 JSON 数组元素和对象成员会被美化输出为该值指定的缩进等级。如果缩进等级为零、负数或者 "", 则只会添加换行符。None (默认值) 选择最紧凑的表达。

使用一个正整数会让每一层缩进同样数量的空格。如果 `*indent*` 是一个字符串（比如 ```\t```），那个字符串会被用于缩进每一层。

在 3.2 版更改：允许使用字符串作为 *indent* 而不再仅仅是整数。

当指定时，*separators* 应当是一个 (*item_separator*, *key_separator*) 元组。当 *indent* 为 `None` 时，默认值取 (`' , ' , ' : ' '`)，否则取 (`' , ' , ' : ' '`)。为了得到最紧凑的 JSON 表达式，你应该指定其为 (`' , ' , ' : ' '`) 以消除空白字符。

在 3.4 版更改：现当 *indent* 不是 `None` 时，采用 (`' , ' , ' : ' '`) 作为默认值。

当 *default* 被指定时，其应该是一个函数，每当某个对象无法被序列化时它会被调用。它应该返回该对象的一个可以被 JSON 编码的版本或者引发一个 `TypeError`。如果没有被指定，则会直接引发 `TypeError`。

在 3.6 版更改：All parameters are now *keyword-only*.

default (*o*)

Implement this method in a subclass such that it returns a serializable object for *o*, or calls the base implementation (to raise a `TypeError`).

For example, to support arbitrary iterators, you could implement *default* like this:

```
def default(self, o):
    try:
        iterable = iter(o)
    except TypeError:
        pass
    else:
        return list(iterable)
    # Let the base class default method raise the TypeError
    return json.JSONEncoder.default(self, o)
```

encode (*o*)

Return a JSON string representation of a Python data structure, *o*. For example:

```
>>> json.JSONEncoder().encode({"foo": ["bar", "baz"]})
'{"foo": ["bar", "baz"]}'
```

iterencode (*o*)

Encode the given object, *o*, and yield each string representation as available. For example:

```
for chunk in json.JSONEncoder().iterencode(bigobject):
    mysocket.write(chunk)
```

19.2.3 异常

exception `json.JSONDecodeError` (*msg*, *doc*, *pos*)

Subclass of `ValueError` with the following additional attributes:

msg

未格式化的错误消息。

doc

The JSON document being parsed.

pos

The start index of *doc* where parsing failed.

linenoThe line corresponding to *pos*.**colno**The column corresponding to *pos*.

3.5 新版功能.

19.2.4 Standard Compliance and Interoperability

The JSON format is specified by [RFC 7159](#) and by [ECMA-404](#). This section details this module's level of compliance with the RFC. For simplicity, [JSONEncoder](#) and [JSONDecoder](#) subclasses, and parameters other than those explicitly mentioned, are not considered.

This module does not comply with the RFC in a strict fashion, implementing some extensions that are valid JavaScript but not valid JSON. In particular:

- Infinite and NaN number values are accepted and output;
- Repeated names within an object are accepted, and only the value of the last name-value pair is used.

Since the RFC permits RFC-compliant parsers to accept input texts that are not RFC-compliant, this module's deserializer is technically RFC-compliant under default settings.

Character Encodings

The RFC requires that JSON be represented using either UTF-8, UTF-16, or UTF-32, with UTF-8 being the recommended default for maximum interoperability.

As permitted, though not required, by the RFC, this module's serializer sets *ensure_ascii=True* by default, thus escaping the output so that the resulting strings only contain ASCII characters.

Other than the *ensure_ascii* parameter, this module is defined strictly in terms of conversion between Python objects and *Unicode strings*, and thus does not otherwise directly address the issue of character encodings.

The RFC prohibits adding a byte order mark (BOM) to the start of a JSON text, and this module's serializer does not add a BOM to its output. The RFC permits, but does not require, JSON deserializers to ignore an initial BOM in their input. This module's deserializer raises a *ValueError* when an initial BOM is present.

The RFC does not explicitly forbid JSON strings which contain byte sequences that don't correspond to valid Unicode characters (e.g. unpaired UTF-16 surrogates), but it does note that they may cause interoperability problems. By default, this module accepts and outputs (when present in the original *str*) code points for such sequences.

Infinite and NaN Number Values

The RFC does not permit the representation of infinite or NaN number values. Despite that, by default, this module accepts and outputs *Infinity*, *-Infinity*, and *NaN* as if they were valid JSON number literal values:

```
>>> # Neither of these calls raises an exception, but the results are not valid JSON
>>> json.dumps(float('-inf'))
'-Infinity'
>>> json.dumps(float('nan'))
'NaN'
>>> # Same when deserializing
>>> json.loads('-Infinity')
-inf
```

(下页继续)

(续上页)

```
>>> json.loads('NaN')
nan
```

In the serializer, the *allow_nan* parameter can be used to alter this behavior. In the deserializer, the *parse_constant* parameter can be used to alter this behavior.

Repeated Names Within an Object

The RFC specifies that the names within a JSON object should be unique, but does not mandate how repeated names in JSON objects should be handled. By default, this module does not raise an exception; instead, it ignores all but the last name-value pair for a given name:

```
>>> weird_json = '{"x": 1, "x": 2, "x": 3}'
>>> json.loads(weird_json)
{'x': 3}
```

The *object_pairs_hook* parameter can be used to alter this behavior.

Top-level Non-Object, Non-Array Values

The old version of JSON specified by the obsolete [RFC 4627](#) required that the top-level value of a JSON text must be either a JSON object or array (Python *dict* or *list*), and could not be a JSON null, boolean, number, or string value. [RFC 7159](#) removed that restriction, and this module does not and has never implemented that restriction in either its serializer or its deserializer.

Regardless, for maximum interoperability, you may wish to voluntarily adhere to the restriction yourself.

Implementation Limitations

Some JSON deserializer implementations may set limits on:

- the size of accepted JSON texts
- the maximum level of nesting of JSON objects and arrays
- the range and precision of JSON numbers
- the content and maximum length of JSON strings

This module does not impose any such limits beyond those of the relevant Python datatypes themselves or the Python interpreter itself.

When serializing to JSON, beware any such limitations in applications that may consume your JSON. In particular, it is common for JSON numbers to be deserialized into IEEE 754 double precision numbers and thus subject to that representation's range and precision limitations. This is especially relevant when serializing Python *int* values of extremely large magnitude, or when serializing instances of “exotic” numerical types such as *decimal.Decimal*.

19.2.5 Command Line Interface

Source code: [Lib/json/tool.py](#)

The *json.tool* module provides a simple command line interface to validate and pretty-print JSON objects.

If the optional `infile` and `outfile` arguments are not specified, `sys.stdin` and `sys.stdout` will be used respectively:

```
$ echo '{"json": "obj"}' | python -m json.tool
{
  "json": "obj"
}
$ echo '{1.2:3.4}' | python -m json.tool
Expecting property name enclosed in double quotes: line 1 column 2 (char 1)
```

在 3.5 版更改: The output is now in the same order as the input. Use the `--sort-keys` option to sort the output of dictionaries alphabetically by key.

Command line options

`infile`

The JSON file to be validated or pretty-printed:

```
$ python -m json.tool mp_films.json
[
  {
    "title": "And Now for Something Completely Different",
    "year": 1971
  },
  {
    "title": "Monty Python and the Holy Grail",
    "year": 1975
  }
]
```

If `infile` is not specified, read from `sys.stdin`.

`outfile`

Write the output of the `infile` to the given `outfile`. Otherwise, write it to `sys.stdout`.

`--sort-keys`

Sort the output of dictionaries alphabetically by key.

3.5 新版功能.

`-h, --help`

Show the help message.

19.3 mailcap —Mailcap file handling

Source code: [Lib/mailcap.py](#)

Mailcap files are used to configure how MIME-aware applications such as mail readers and Web browsers react to files with different MIME types. (The name “mailcap” is derived from the phrase “mail capability”.) For example, a mailcap file might contain a line like `video/mpeg; xmpeg %s`. Then, if the user encounters an email message or Web document with the MIME type `video/mpeg`, `%s` will be replaced by a filename (usually one belonging to a temporary file) and the `xmpeg` program can be automatically started to view the file.

The mailcap format is documented in [RFC 1524](#), “A User Agent Configuration Mechanism For Multimedia Mail Format Information,” but is not an Internet standard. However, mailcap files are supported on most Unix systems.

`mailcap.findmatch(caps, MIMEtype, key='view', filename='/dev/null', plist=[])`

Return a 2-tuple; the first element is a string containing the command line to be executed (which can be passed to `os.system()`), and the second element is the mailcap entry for a given MIME type. If no matching MIME type can be found, `(None, None)` is returned.

`key` is the name of the field desired, which represents the type of activity to be performed; the default value is 'view', since in the most common case you simply want to view the body of the MIME-typed data. Other possible values might be 'compose' and 'edit', if you wanted to create a new body of the given MIME type or alter the existing body data. See [RFC 1524](#) for a complete list of these fields.

`filename` is the filename to be substituted for `%s` in the command line; the default value is `'/dev/null'` which is almost certainly not what you want, so usually you'll override it by specifying a filename.

`plist` can be a list containing named parameters; the default value is simply an empty list. Each entry in the list must be a string containing the parameter name, an equals sign ('='), and the parameter's value. Mailcap entries can contain named parameters like `%{foo}`, which will be replaced by the value of the parameter named 'foo'. For example, if the command line `showpartial %{id} %{number} %{total}` was in a mailcap file, and `plist` was set to `['id=1', 'number=2', 'total=3']`, the resulting command line would be `'showpartial 1 2 3'`.

In a mailcap file, the "test" field can optionally be specified to test some external condition (such as the machine architecture, or the window system in use) to determine whether or not the mailcap line applies. `findmatch()` will automatically check such conditions and skip the entry if the check fails.

`mailcap.getcaps()`

Returns a dictionary mapping MIME types to a list of mailcap file entries. This dictionary must be passed to the `findmatch()` function. An entry is stored as a list of dictionaries, but it shouldn't be necessary to know the details of this representation.

The information is derived from all of the mailcap files found on the system. Settings in the user's mailcap file `$HOME/.mailcap` will override settings in the system mailcap files `/etc/mailcap`, `/usr/etc/mailcap`, and `/usr/local/etc/mailcap`.

An example usage:

```
>>> import mailcap
>>> d = mailcap.getcaps()
>>> mailcap.findmatch(d, 'video/mpeg', filename='tmp1223')
('xmpeg tmp1223', {'view': 'xmpeg %s'})
```

19.4 mailbox — Manipulate mailboxes in various formats

Source code: [Lib/mailbox.py](#)

This module defines two classes, `Mailbox` and `Message`, for accessing and manipulating on-disk mailboxes and the messages they contain. `Mailbox` offers a dictionary-like mapping from keys to messages. `Message` extends the `email.message` module's `Message` class with format-specific state and behavior. Supported mailbox formats are Maildir, mbox, MH, Babyl, and MMDF.

参见:

Module `email` Represent and manipulate messages.

19.4.1 Mailbox objects

class mailbox.Mailbox

A mailbox, which may be inspected and modified.

The *Mailbox* class defines an interface and is not intended to be instantiated. Instead, format-specific subclasses should inherit from *Mailbox* and your code should instantiate a particular subclass.

The *Mailbox* interface is dictionary-like, with small keys corresponding to messages. Keys are issued by the *Mailbox* instance with which they will be used and are only meaningful to that *Mailbox* instance. A key continues to identify a message even if the corresponding message is modified, such as by replacing it with another message.

Messages may be added to a *Mailbox* instance using the set-like method *add()* and removed using a *del* statement or the set-like methods *remove()* and *discard()*.

Mailbox interface semantics differ from dictionary semantics in some noteworthy ways. Each time a message is requested, a new representation (typically a *Message* instance) is generated based upon the current state of the mailbox. Similarly, when a message is added to a *Mailbox* instance, the provided message representation's contents are copied. In neither case is a reference to the message representation kept by the *Mailbox* instance.

The default *Mailbox* iterator iterates over message representations, not keys as the default dictionary iterator does. Moreover, modification of a mailbox during iteration is safe and well-defined. Messages added to the mailbox after an iterator is created will not be seen by the iterator. Messages removed from the mailbox before the iterator yields them will be silently skipped, though using a key from an iterator may result in a *KeyError* exception if the corresponding message is subsequently removed.

警告: Be very cautious when modifying mailboxes that might be simultaneously changed by some other process. The safest mailbox format to use for such tasks is Maildir; try to avoid using single-file formats such as mbox for concurrent writing. If you're modifying a mailbox, you *must* lock it by calling the *lock()* and *unlock()* methods *before* reading any messages in the file or making any changes by adding or deleting a message. Failing to lock the mailbox runs the risk of losing messages or corrupting the entire mailbox.

Mailbox instances have the following methods:

add (*message*)

Add *message* to the mailbox and return the key that has been assigned to it.

Parameter *message* may be a *Message* instance, an *email.message.Message* instance, a string, a byte string, or a file-like object (which should be open in binary mode). If *message* is an instance of the appropriate format-specific *Message* subclass (e.g., if it's an *mboxMessage* instance and this is an *mbox* instance), its format-specific information is used. Otherwise, reasonable defaults for format-specific information are used.

在 3.2 版更改: Support for binary input was added.

remove (*key*)

__delitem__ (*key*)

discard (*key*)

Delete the message corresponding to *key* from the mailbox.

If no such message exists, a *KeyError* exception is raised if the method was called as *remove()* or *__delitem__()* but no exception is raised if the method was called as *discard()*. The behavior of *discard()* may be preferred if the underlying mailbox format supports concurrent modification by other processes.

__setitem__ (*key*, *message*)

Replace the message corresponding to *key* with *message*. Raise a *KeyError* exception if no message already corresponds to *key*.

As with *add()*, parameter *message* may be a *Message* instance, an *email.message.Message* instance, a string, a byte string, or a file-like object (which should be open in binary mode). If *message* is an instance of the appropriate format-specific *Message* subclass (e.g., if it's an *mbboxMessage* instance and this is an *mbbox* instance), its format-specific information is used. Otherwise, the format-specific information of the message that currently corresponds to *key* is left unchanged.

iterkeys ()

keys ()

Return an iterator over all keys if called as *iterkeys()* or return a list of keys if called as *keys()*.

intervalues ()

__iter__ ()

values ()

Return an iterator over representations of all messages if called as *intervalues()* or *__iter__()* or return a list of such representations if called as *values()*. The messages are represented as instances of the appropriate format-specific *Message* subclass unless a custom message factory was specified when the *Mailbox* instance was initialized.

注解: The behavior of *__iter__()* is unlike that of dictionaries, which iterate over keys.

iteritems ()

items ()

Return an iterator over (*key*, *message*) pairs, where *key* is a key and *message* is a message representation, if called as *iteritems()* or return a list of such pairs if called as *items()*. The messages are represented as instances of the appropriate format-specific *Message* subclass unless a custom message factory was specified when the *Mailbox* instance was initialized.

get (*key*, *default=None*)

__getitem__ (*key*)

Return a representation of the message corresponding to *key*. If no such message exists, *default* is returned if the method was called as *get()* and a *KeyError* exception is raised if the method was called as *__getitem__()*. The message is represented as an instance of the appropriate format-specific *Message* subclass unless a custom message factory was specified when the *Mailbox* instance was initialized.

get_message (*key*)

Return a representation of the message corresponding to *key* as an instance of the appropriate format-specific *Message* subclass, or raise a *KeyError* exception if no such message exists.

get_bytes (*key*)

Return a byte representation of the message corresponding to *key*, or raise a *KeyError* exception if no such message exists.

3.2 新版功能.

get_string (*key*)

Return a string representation of the message corresponding to *key*, or raise a *KeyError* exception if no such message exists. The message is processed through *email.message.Message* to convert it to a 7bit clean representation.

get_file (*key*)

Return a file-like representation of the message corresponding to *key*, or raise a *KeyError* exception if no such message exists. The file-like object behaves as if open in binary mode. This file should be closed once it is no longer needed.

在 3.2 版更改: The file object really is a binary file; previously it was incorrectly returned in text mode. Also, the file-like object now supports the context management protocol: you can use a `with` statement to automatically close it.

注解: Unlike other representations of messages, file-like representations are not necessarily independent of the *Mailbox* instance that created them or of the underlying mailbox. More specific documentation is provided by each subclass.

__contains__ (*key*)

Return True if *key* corresponds to a message, False otherwise.

__len__ ()

Return a count of messages in the mailbox.

clear ()

Delete all messages from the mailbox.

pop (*key*, *default=None*)

Return a representation of the message corresponding to *key* and delete the message. If no such message exists, return *default*. The message is represented as an instance of the appropriate format-specific *Message* subclass unless a custom message factory was specified when the *Mailbox* instance was initialized.

popitem ()

Return an arbitrary (*key*, *message*) pair, where *key* is a key and *message* is a message representation, and delete the corresponding message. If the mailbox is empty, raise a *KeyError* exception. The message is represented as an instance of the appropriate format-specific *Message* subclass unless a custom message factory was specified when the *Mailbox* instance was initialized.

update (*arg*)

Parameter *arg* should be a *key*-to-*message* mapping or an iterable of (*key*, *message*) pairs. Updates the mailbox so that, for each given *key* and *message*, the message corresponding to *key* is set to *message* as if by using `__setitem__()`. As with `__setitem__()`, each *key* must already correspond to a message in the mailbox or else a *KeyError* exception will be raised, so in general it is incorrect for *arg* to be a *Mailbox* instance.

注解: Unlike with dictionaries, keyword arguments are not supported.

flush ()

Write any pending changes to the filesystem. For some *Mailbox* subclasses, changes are always written immediately and `flush()` does nothing, but you should still make a habit of calling this method.

lock ()

Acquire an exclusive advisory lock on the mailbox so that other processes know not to modify it. An *ExternalClashError* is raised if the lock is not available. The particular locking mechanisms used depend upon the mailbox format. You should *always* lock the mailbox before making any modifications to its contents.

unlock ()

Release the lock on the mailbox, if any.

close ()

Flush the mailbox, unlock it if necessary, and close any open files. For some *Mailbox* subclasses, this method does nothing.

Maildir

class mailbox.**Maildir** (*dirname*, *factory=None*, *create=True*)

A subclass of *Mailbox* for mailboxes in Maildir format. Parameter *factory* is a callable object that accepts a file-like message representation (which behaves as if opened in binary mode) and returns a custom representation. If *factory* is *None*, *MaildirMessage* is used as the default message representation. If *create* is *True*, the mailbox is created if it does not exist.

It is for historical reasons that *dirname* is named as such rather than *path*.

Maildir is a directory-based mailbox format invented for the qmail mail transfer agent and now widely supported by other programs. Messages in a Maildir mailbox are stored in separate files within a common directory structure. This design allows Maildir mailboxes to be accessed and modified by multiple unrelated programs without data corruption, so file locking is unnecessary.

Maildir mailboxes contain three subdirectories, namely: *tmp*, *new*, and *cur*. Messages are created momentarily in the *tmp* subdirectory and then moved to the *new* subdirectory to finalize delivery. A mail user agent may subsequently move the message to the *cur* subdirectory and store information about the state of the message in a special “info” section appended to its file name.

Folders of the style introduced by the Courier mail transfer agent are also supported. Any subdirectory of the main mailbox is considered a folder if *'.'* is the first character in its name. Folder names are represented by *Maildir* without the leading *'.'*. Each folder is itself a Maildir mailbox but should not contain other folders. Instead, a logical nesting is indicated using *'.'* to delimit levels, e.g., “*Archived.2005.07*”.

注解: The Maildir specification requires the use of a colon (*':'*) in certain message file names. However, some operating systems do not permit this character in file names. If you wish to use a Maildir-like format on such an operating system, you should specify another character to use instead. The exclamation point (*'!'*) is a popular choice. For example:

```
import mailbox
mailbox.Maildir.colon = '!'
```

The *colon* attribute may also be set on a per-instance basis.

Maildir instances have all of the methods of *Mailbox* in addition to the following:

list_folders()

Return a list of the names of all folders.

get_folder(folder)

Return a *Maildir* instance representing the folder whose name is *folder*. A *NoSuchMailboxError* exception is raised if the folder does not exist.

add_folder(folder)

Create a folder whose name is *folder* and return a *Maildir* instance representing it.

remove_folder(folder)

Delete the folder whose name is *folder*. If the folder contains any messages, a *NotEmptyError* exception will be raised and the folder will not be deleted.

clean()

Delete temporary files from the mailbox that have not been accessed in the last 36 hours. The Maildir specification says that mail-reading programs should do this occasionally.

Some *Mailbox* methods implemented by *Maildir* deserve special remarks:

add(message)

__setitem__(key, message)

update (*arg*)

警告: These methods generate unique file names based upon the current process ID. When using multiple threads, undetected name clashes may occur and cause corruption of the mailbox unless threads are coordinated to avoid using these methods to manipulate the same mailbox simultaneously.

flush ()

All changes to Maildir mailboxes are immediately applied, so this method does nothing.

lock ()

unlock ()

Maildir mailboxes do not support (or require) locking, so these methods do nothing.

close ()

Maildir instances do not keep any open files and the underlying mailboxes do not support locking, so this method does nothing.

get_file (*key*)

Depending upon the host platform, it may not be possible to modify or remove the underlying message while the returned file remains open.

参见:

maildir man page from gmail The original specification of the format.

Using maildir format Notes on Maildir by its inventor. Includes an updated name-creation scheme and details on “info” semantics.

maildir man page from Courier Another specification of the format. Describes a common extension for supporting folders.

mbox

class mailbox.**mbox** (*path*, *factory=None*, *create=True*)

A subclass of *Mailbox* for mailboxes in mbox format. Parameter *factory* is a callable object that accepts a file-like message representation (which behaves as if opened in binary mode) and returns a custom representation. If *factory* is None, *mboxMessage* is used as the default message representation. If *create* is True, the mailbox is created if it does not exist.

The mbox format is the classic format for storing mail on Unix systems. All messages in an mbox mailbox are stored in a single file with the beginning of each message indicated by a line whose first five characters are “From “.

Several variations of the mbox format exist to address perceived shortcomings in the original. In the interest of compatibility, *mbox* implements the original format, which is sometimes referred to as *mboxo*. This means that the *Content-Length* header, if present, is ignored and that any occurrences of “From ” at the beginning of a line in a message body are transformed to “>From ” when storing the message, although occurrences of “>From ” are not transformed to “From ” when reading the message.

Some *Mailbox* methods implemented by *mbox* deserve special remarks:

get_file (*key*)

Using the file after calling *flush* () or *close* () on the *mbox* instance may yield unpredictable results or raise an exception.

lock ()

unlock()

Three locking mechanisms are used—dot locking and, if available, the `flock()` and `lockf()` system calls.

参见:

mbox man page from gmail A specification of the format and its variations.

mbox man page from tin Another specification of the format, with details on locking.

Configuring Netscape Mail on Unix: Why The Content-Length Format is Bad An argument for using the original mbox format rather than a variation.

“mbox” is a family of several mutually incompatible mailbox formats A history of mbox variations.

MH

class `mailbox.MH` (*path*, *factory*=None, *create*=True)

A subclass of `Mailbox` for mailboxes in MH format. Parameter *factory* is a callable object that accepts a file-like message representation (which behaves as if opened in binary mode) and returns a custom representation. If *factory* is None, `MHMessage` is used as the default message representation. If *create* is True, the mailbox is created if it does not exist.

MH is a directory-based mailbox format invented for the MH Message Handling System, a mail user agent. Each message in an MH mailbox resides in its own file. An MH mailbox may contain other MH mailboxes (called *folders*) in addition to messages. Folders may be nested indefinitely. MH mailboxes also support *sequences*, which are named lists used to logically group messages without moving them to sub-folders. Sequences are defined in a file called `.mh_sequences` in each folder.

The `MH` class manipulates MH mailboxes, but it does not attempt to emulate all of `mh`'s behaviors. In particular, it does not modify and is not affected by the `context` or `.mh_profile` files that are used by `mh` to store its state and configuration.

`MH` instances have all of the methods of `Mailbox` in addition to the following:

list_folders()

Return a list of the names of all folders.

get_folder (*folder*)

Return an `MH` instance representing the folder whose name is *folder*. A `NoSuchMailboxError` exception is raised if the folder does not exist.

add_folder (*folder*)

Create a folder whose name is *folder* and return an `MH` instance representing it.

remove_folder (*folder*)

Delete the folder whose name is *folder*. If the folder contains any messages, a `NotEmptyError` exception will be raised and the folder will not be deleted.

get_sequences()

Return a dictionary of sequence names mapped to key lists. If there are no sequences, the empty dictionary is returned.

set_sequences (*sequences*)

Re-define the sequences that exist in the mailbox based upon *sequences*, a dictionary of names mapped to key lists, like returned by `get_sequences()`.

pack()

Rename messages in the mailbox as necessary to eliminate gaps in numbering. Entries in the sequences list are updated correspondingly.

注解: Already-issued keys are invalidated by this operation and should not be subsequently used.

Some *Mailbox* methods implemented by *MH* deserve special remarks:

remove (*key*)

__delitem__ (*key*)

discard (*key*)

These methods immediately delete the message. The MH convention of marking a message for deletion by prepending a comma to its name is not used.

lock ()

unlock ()

Three locking mechanisms are used—dot locking and, if available, the `flock()` and `lockf()` system calls. For MH mailboxes, locking the mailbox means locking the `.mh_sequences` file and, only for the duration of any operations that affect them, locking individual message files.

get_file (*key*)

Depending upon the host platform, it may not be possible to remove the underlying message while the returned file remains open.

flush ()

All changes to MH mailboxes are immediately applied, so this method does nothing.

close ()

MH instances do not keep any open files, so this method is equivalent to `unlock()`.

参见:

nmh - Message Handling System Home page of **nmh**, an updated version of the original **mh**.

MH & nmh: Email for Users & Programmers A GPL-licensed book on **mh** and **nmh**, with some information on the mailbox format.

Babyl

class mailbox.**Babyl** (*path*, *factory=None*, *create=True*)

A subclass of *Mailbox* for mailboxes in Babyl format. Parameter *factory* is a callable object that accepts a file-like message representation (which behaves as if opened in binary mode) and returns a custom representation. If *factory* is `None`, *BabylMessage* is used as the default message representation. If *create* is `True`, the mailbox is created if it does not exist.

Babyl is a single-file mailbox format used by the Rmail mail user agent included with Emacs. The beginning of a message is indicated by a line containing the two characters Control-Underscore (`'\037'`) and Control-L (`'\014'`). The end of a message is indicated by the start of the next message or, in the case of the last message, a line containing a Control-Underscore (`'\037'`) character.

Messages in a Babyl mailbox have two sets of headers, original headers and so-called visible headers. Visible headers are typically a subset of the original headers that have been reformatted or abridged to be more attractive. Each message in a Babyl mailbox also has an accompanying list of *labels*, or short strings that record extra information about the message, and a list of all user-defined labels found in the mailbox is kept in the Babyl options section.

Babyl instances have all of the methods of *Mailbox* in addition to the following:

get_labels ()

Return a list of the names of all user-defined labels used in the mailbox.

注解: The actual messages are inspected to determine which labels exist in the mailbox rather than consulting the list of labels in the Babyl options section, but the Babyl section is updated whenever the mailbox is modified.

Some *Mailbox* methods implemented by *Babyl* deserve special remarks:

get_file (*key*)

In Babyl mailboxes, the headers of a message are not stored contiguously with the body of the message. To generate a file-like representation, the headers and body are copied together into an *io.BytesIO* instance, which has an API identical to that of a file. As a result, the file-like object is truly independent of the underlying mailbox but does not save memory compared to a string representation.

lock ()

unlock ()

Three locking mechanisms are used—dot locking and, if available, the *flock* () and *lockf* () system calls.

参见:

Format of Version 5 Babyl Files A specification of the Babyl format.

Reading Mail with Rmail The Rmail manual, with some information on Babyl semantics.

MMDF

class mailbox.**MMDF** (*path*, *factory=None*, *create=True*)

A subclass of *Mailbox* for mailboxes in MMDF format. Parameter *factory* is a callable object that accepts a file-like message representation (which behaves as if opened in binary mode) and returns a custom representation. If *factory* is None, *MMDFMessage* is used as the default message representation. If *create* is True, the mailbox is created if it does not exist.

MMDF is a single-file mailbox format invented for the Multichannel Memorandum Distribution Facility, a mail transfer agent. Each message is in the same form as an mbox message but is bracketed before and after by lines containing four Control-A (' \001 ') characters. As with the mbox format, the beginning of each message is indicated by a line whose first five characters are “From “, but additional occurrences of “From ” are not transformed to “>From ” when storing messages because the extra message separator lines prevent mistaking such occurrences for the starts of subsequent messages.

Some *Mailbox* methods implemented by *MMDF* deserve special remarks:

get_file (*key*)

Using the file after calling *flush* () or *close* () on the *MMDF* instance may yield unpredictable results or raise an exception.

lock ()

unlock ()

Three locking mechanisms are used—dot locking and, if available, the *flock* () and *lockf* () system calls.

参见:

mmdf man page from tin A specification of MMDF format from the documentation of tin, a newsreader.

MMDF A Wikipedia article describing the Multichannel Memorandum Distribution Facility.

19.4.2 Message objects

class mailbox.**Message** (*message=None*)

A subclass of the `email.message` module's `Message`. Subclasses of `mailbox.Message` add mailbox-format-specific state and behavior.

If *message* is omitted, the new instance is created in a default, empty state. If *message* is an `email.message.Message` instance, its contents are copied; furthermore, any format-specific information is converted insofar as possible if *message* is a `Message` instance. If *message* is a string, a byte string, or a file, it should contain an **RFC 2822**-compliant message, which is read and parsed. Files should be open in binary mode, but text mode files are accepted for backward compatibility.

The format-specific state and behaviors offered by subclasses vary, but in general it is only the properties that are not specific to a particular mailbox that are supported (although presumably the properties are specific to a particular mailbox format). For example, file offsets for single-file mailbox formats and file names for directory-based mailbox formats are not retained, because they are only applicable to the original mailbox. But state such as whether a message has been read by the user or marked as important is retained, because it applies to the message itself.

There is no requirement that `Message` instances be used to represent messages retrieved using `Mailbox` instances. In some situations, the time and memory required to generate `Message` representations might not be acceptable. For such situations, `Mailbox` instances also offer string and file-like representations, and a custom message factory may be specified when a `Mailbox` instance is initialized.

MaildirMessage

class mailbox.**MaildirMessage** (*message=None*)

A message with Maildir-specific behaviors. Parameter *message* has the same meaning as with the `Message` constructor.

Typically, a mail user agent application moves all of the messages in the `new` subdirectory to the `cur` subdirectory after the first time the user opens and closes the mailbox, recording that the messages are old whether or not they've actually been read. Each message in `cur` has an “info” section added to its file name to store information about its state. (Some mail readers may also add an “info” section to messages in `new`.) The “info” section may take one of two forms: it may contain “2,” followed by a list of standardized flags (e.g., “2,FR”) or it may contain “1,” followed by so-called experimental information. Standard flags for Maildir messages are as follows:

标志	意义	解释
D	Draft	Under composition
F	Flagged	Marked as important
P	Passed	Forwarded, resent, or bounced
R	Replied	Replied to
S	Seen	Read
T	Trashed	Marked for subsequent deletion

`MaildirMessage` instances offer the following methods:

get_subdir ()

Return either “new” (if the message should be stored in the `new` subdirectory) or “cur” (if the message should be stored in the `cur` subdirectory).

注解: A message is typically moved from `new` to `cur` after its mailbox has been accessed, whether or not the message is has been read. A message `msg` has been read if `"S" in msg.get_flags()` is `True`.

set_subdir (*subdir*)

Set the subdirectory the message should be stored in. Parameter *subdir* must be either “new” or “cur” .

get_flags ()

Return a string specifying the flags that are currently set. If the message complies with the standard Maildir format, the result is the concatenation in alphabetical order of zero or one occurrence of each of 'D', 'F', 'P', 'R', 'S', and 'T'. The empty string is returned if no flags are set or if “info” contains experimental semantics.

set_flags (*flags*)

Set the flags specified by *flags* and unset all others.

add_flag (*flag*)

Set the flag(s) specified by *flag* without changing other flags. To add more than one flag at a time, *flag* may be a string of more than one character. The current “info” is overwritten whether or not it contains experimental information rather than flags.

remove_flag (*flag*)

Unset the flag(s) specified by *flag* without changing other flags. To remove more than one flag at a time, *flag* maybe a string of more than one character. If “info” contains experimental information rather than flags, the current “info” is not modified.

get_date ()

Return the delivery date of the message as a floating-point number representing seconds since the epoch.

set_date (*date*)

Set the delivery date of the message to *date*, a floating-point number representing seconds since the epoch.

get_info ()

Return a string containing the “info” for a message. This is useful for accessing and modifying “info” that is experimental (i.e., not a list of flags).

set_info (*info*)

Set “info” to *info*, which should be a string.

When a *MaildirMessage* instance is created based upon an *mboxMessage* or *MMDFMessage* instance, the *Status* and *X-Status* headers are omitted and the following conversions take place:

Resulting state	<i>mboxMessage</i> or <i>MMDFMessage</i> state
“cur” subdirectory	O flag
F flag	F flag
R flag	A flag
S flag	R flag
T flag	D flag

When a *MaildirMessage* instance is created based upon an *MHMessage* instance, the following conversions take place:

Resulting state	<i>MHMessage</i> state
“cur” subdirectory	“unseen” sequence
“cur” subdirectory and S flag	no “unseen” sequence
F flag	“flagged” sequence
R flag	“replied” sequence

When a *MaildirMessage* instance is created based upon a *BabylMessage* instance, the following conversions take place:

Resulting state	<i>BabylMessage</i> state
“cur” subdirectory	“unseen” label
“cur” subdirectory and S flag	no “unseen” label
P flag	“forwarded” or “resent” label
R flag	“answered” label
T flag	“deleted” label

mbboxMessage

class mailbox.**mbboxMessage** (*message=None*)

A message with mbox-specific behaviors. Parameter *message* has the same meaning as with the *Message* constructor.

Messages in an mbox mailbox are stored together in a single file. The sender’s envelope address and the time of delivery are typically stored in a line beginning with “From ” that is used to indicate the start of a message, though there is considerable variation in the exact format of this data among mbox implementations. Flags that indicate the state of the message, such as whether it has been read or marked as important, are typically stored in *Status* and *X-Status* headers.

Conventional flags for mbox messages are as follows:

标志	意义	解释
R	Read	Read
O	Old	Previously detected by MUA
D	Deleted	Marked for subsequent deletion
F	Flagged	Marked as important
A	Answered	Replied to

The “R” and “O” flags are stored in the *Status* header, and the “D”, “F”, and “A” flags are stored in the *X-Status* header. The flags and headers typically appear in the order mentioned.

mbboxMessage instances offer the following methods:

get_from()

Return a string representing the “From ” line that marks the start of the message in an mbox mailbox. The leading “From ” and the trailing newline are excluded.

set_from (*from_, time_=None*)

Set the “From ” line to *from_*, which should be specified without a leading “From ” or trailing newline. For convenience, *time_* may be specified and will be formatted appropriately and appended to *from_*. If *time_* is specified, it should be a *time.struct_time* instance, a tuple suitable for passing to *time.strftime()*, or True (to use *time.gmtime()*).

get_flags()

Return a string specifying the flags that are currently set. If the message complies with the conventional format, the result is the concatenation in the following order of zero or one occurrence of each of 'R', 'O', 'D', 'F', and 'A'.

set_flags (*flags*)

Set the flags specified by *flags* and unset all others. Parameter *flags* should be the concatenation in any order of zero or more occurrences of each of 'R', 'O', 'D', 'F', and 'A'.

add_flag (*flag*)

Set the flag(s) specified by *flag* without changing other flags. To add more than one flag at a time, *flag* may be a string of more than one character.

remove_flag (*flag*)

Unset the flag(s) specified by *flag* without changing other flags. To remove more than one flag at a time, *flag* maybe a string of more than one character.

When an *mbxMessage* instance is created based upon a *MaiDirMessage* instance, a “From ” line is generated based upon the *MaiDirMessage* instance’s delivery date, and the following conversions take place:

Resulting state	<i>MaiDirMessage</i> state
R flag	S flag
O flag	“cur” subdirectory
D flag	T flag
F flag	F flag
A flag	R flag

When an *mbxMessage* instance is created based upon an *MHMessage* instance, the following conversions take place:

Resulting state	<i>MHMessage</i> state
R flag and O flag	no “unseen” sequence
O flag	“unseen” sequence
F flag	“flagged” sequence
A flag	“replied” sequence

When an *mbxMessage* instance is created based upon a *BabylMessage* instance, the following conversions take place:

Resulting state	<i>BabylMessage</i> state
R flag and O flag	no “unseen” label
O flag	“unseen” label
D flag	“deleted” label
A flag	“answered” label

When a *Message* instance is created based upon an *MMDFMessage* instance, the “From ” line is copied and all flags directly correspond:

Resulting state	<i>MMDFMessage</i> state
R flag	R flag
O flag	O flag
D flag	D flag
F flag	F flag
A flag	A flag

MHMessage

class mailbox.**MHMessage** (*message=None*)

A message with MH-specific behaviors. Parameter *message* has the same meaning as with the *Message* constructor.

MH messages do not support marks or flags in the traditional sense, but they do support sequences, which are logical groupings of arbitrary messages. Some mail reading programs (although not the standard **mh** and **nmh**) use sequences in much the same way flags are used with other formats, as follows:

Sequence	解释
unseen	Not read, but previously detected by MUA
replied	Replied to
flagged	Marked as important

MHMessage instances offer the following methods:

get_sequences ()

Return a list of the names of sequences that include this message.

set_sequences (*sequences*)

Set the list of sequences that include this message.

add_sequence (*sequence*)

Add *sequence* to the list of sequences that include this message.

remove_sequence (*sequence*)

Remove *sequence* from the list of sequences that include this message.

When an *MHMessage* instance is created based upon a *MaildirMessage* instance, the following conversions take place:

Resulting state	<i>MaildirMessage</i> state
“unseen” sequence	no S flag
“replied” sequence	R flag
“flagged” sequence	F flag

When an *MHMessage* instance is created based upon an *mboxMessage* or *MMDFMessage* instance, the *Status* and *X-Status* headers are omitted and the following conversions take place:

Resulting state	<i>mboxMessage</i> or <i>MMDFMessage</i> state
“unseen” sequence	no R flag
“replied” sequence	A flag
“flagged” sequence	F flag

When an *MHMessage* instance is created based upon a *BabylMessage* instance, the following conversions take place:

Resulting state	<i>BabylMessage</i> state
“unseen” sequence	“unseen” label
“replied” sequence	“answered” label

BabylMessage

class mailbox.**BabylMessage** (*message=None*)

A message with Babyl-specific behaviors. Parameter *message* has the same meaning as with the *Message* constructor.

Certain message labels, called *attributes*, are defined by convention to have special meanings. The attributes are as follows:

Label	解释
unseen	Not read, but previously detected by MUA
deleted	Marked for subsequent deletion
filed	Copied to another file or mailbox
answered	Replied to
forwarded	Forwarded
edited	Modified by the user
resent	Resent

By default, Rmail displays only visible headers. The *BabylMessage* class, though, uses the original headers because they are more complete. Visible headers may be accessed explicitly if desired.

BabylMessage instances offer the following methods:

get_labels()

Return a list of labels on the message.

set_labels(labels)

Set the list of labels on the message to *labels*.

add_label(label)

Add *label* to the list of labels on the message.

remove_label(label)

Remove *label* from the list of labels on the message.

get_visible()

Return an *Message* instance whose headers are the message's visible headers and whose body is empty.

set_visible(visible)

Set the message's visible headers to be the same as the headers in *message*. Parameter *visible* should be a *Message* instance, an *email.message.Message* instance, a string, or a file-like object (which should be open in text mode).

update_visible()

When a *BabylMessage* instance's original headers are modified, the visible headers are not automatically modified to correspond. This method updates the visible headers as follows: each visible header with a corresponding original header is set to the value of the original header, each visible header without a corresponding original header is removed, and any of *Date*, *From*, *Reply-To*, *To*, *CC*, and *Subject* that are present in the original headers but not the visible headers are added to the visible headers.

When a *BabylMessage* instance is created based upon a *MaildirMessage* instance, the following conversions take place:

Resulting state	<i>MaildirMessage</i> state
"unseen" label	no S flag
"deleted" label	T flag
"answered" label	R flag
"forwarded" label	P flag

When a *BabylMessage* instance is created based upon an *mbxMessage* or *MMDfMessage* instance, the *Status* and *X-Status* headers are omitted and the following conversions take place:

Resulting state	<i>mbxMessage</i> or <i>MMDfMessage</i> state
"unseen" label	no R flag
"deleted" label	D flag
"answered" label	A flag

When a *BabylMessage* instance is created based upon an *MHMessage* instance, the following conversions take place:

Resulting state	<i>MHMessage</i> state
“unseen” label	“unseen” sequence
“answered” label	“replied” sequence

MMDFMessage

class mailbox.**MMDFMessage** (*message=None*)

A message with MMDF-specific behaviors. Parameter *message* has the same meaning as with the *Message* constructor.

As with message in an mbox mailbox, MMDF messages are stored with the sender’s address and the delivery date in an initial line beginning with “From “. Likewise, flags that indicate the state of the message are typically stored in *Status* and *X-Status* headers.

Conventional flags for MMDF messages are identical to those of mbox message and are as follows:

标志	意义	解释
R	Read	Read
O	Old	Previously detected by MUA
D	Deleted	Marked for subsequent deletion
F	Flagged	Marked as important
A	Answered	Replied to

The “R” and “O” flags are stored in the *Status* header, and the “D”, “F”, and “A” flags are stored in the *X-Status* header. The flags and headers typically appear in the order mentioned.

MMDFMessage instances offer the following methods, which are identical to those offered by *mboxMessage*:

get_from ()

Return a string representing the “From ” line that marks the start of the message in an mbox mailbox. The leading “From ” and the trailing newline are excluded.

set_from (*from_*, *time_=None*)

Set the “From ” line to *from_*, which should be specified without a leading “From ” or trailing newline. For convenience, *time_* may be specified and will be formatted appropriately and appended to *from_*. If *time_* is specified, it should be a *time.struct_time* instance, a tuple suitable for passing to *time.strftime()*, or True (to use *time.gmtime()*).

get_flags ()

Return a string specifying the flags that are currently set. If the message complies with the conventional format, the result is the concatenation in the following order of zero or one occurrence of each of 'R', 'O', 'D', 'F', and 'A'.

set_flags (*flags*)

Set the flags specified by *flags* and unset all others. Parameter *flags* should be the concatenation in any order of zero or more occurrences of each of 'R', 'O', 'D', 'F', and 'A'.

add_flag (*flag*)

Set the flag(s) specified by *flag* without changing other flags. To add more than one flag at a time, *flag* may be a string of more than one character.

remove_flag (*flag*)

Unset the flag(s) specified by *flag* without changing other flags. To remove more than one flag at a time, *flag* maybe a string of more than one character.

When an *MMDFMessage* instance is created based upon a *MaildirMessage* instance, a “From ” line is generated based upon the *MaildirMessage* instance’s delivery date, and the following conversions take place:

Resulting state	<i>MaildirMessage</i> state
R flag	S flag
O flag	“cur” subdirectory
D flag	T flag
F flag	F flag
A flag	R flag

When an *MMDFMessage* instance is created based upon an *MHMessage* instance, the following conversions take place:

Resulting state	<i>MHMessage</i> state
R flag and O flag	no “unseen” sequence
O flag	“unseen” sequence
F flag	“flagged” sequence
A flag	“replied” sequence

When an *MMDFMessage* instance is created based upon a *BabylMessage* instance, the following conversions take place:

Resulting state	<i>BabylMessage</i> state
R flag and O flag	no “unseen” label
O flag	“unseen” label
D flag	“deleted” label
A flag	“answered” label

When an *MMDFMessage* instance is created based upon an *mbxMessage* instance, the “From ” line is copied and all flags directly correspond:

Resulting state	<i>mbxMessage</i> state
R flag	R flag
O flag	O flag
D flag	D flag
F flag	F flag
A flag	A flag

19.4.3 异常

The following exception classes are defined in the *mailbox* module:

exception `mailbox.Error`

The based class for all other module-specific exceptions.

exception `mailbox.NoSuchMailboxError`

Raised when a mailbox is expected but is not found, such as when instantiating a *Mailbox* subclass with a path that does not exist (and with the *create* parameter set to *False*), or when opening a folder that does not exist.

exception `mailbox.NotEmptyError`

Raised when a mailbox is not empty but is expected to be, such as when deleting a folder that contains messages.

exception mailbox.ExternalClashError

Raised when some mailbox-related condition beyond the control of the program causes it to be unable to proceed, such as when failing to acquire a lock that another program already holds a lock, or when a uniquely-generated file name already exists.

exception mailbox.FormatError

Raised when the data in a file cannot be parsed, such as when an *MH* instance attempts to read a corrupted .mh_sequences file.

19.4.4 示例

A simple example of printing the subjects of all messages in a mailbox that seem interesting:

```
import mailbox
for message in mailbox.mbox('~/.mbox'):
    subject = message['subject']          # Could possibly be None.
    if subject and 'python' in subject.lower():
        print(subject)
```

To copy all mail from a Babyl mailbox to an MH mailbox, converting all of the format-specific information that can be converted:

```
import mailbox
destination = mailbox.MH('~/.Mail')
destination.lock()
for message in mailbox.Babyl('~/.RMAIL'):
    destination.add(mailbox.MHMessage(message))
destination.flush()
destination.unlock()
```

This example sorts mail from several mailing lists into different mailboxes, being careful to avoid mail corruption due to concurrent modification by other programs, mail loss due to interruption of the program, or premature termination due to malformed messages in the mailbox:

```
import mailbox
import email.errors

list_names = ('python-list', 'python-dev', 'python-bugs')

boxes = {name: mailbox.mbox('~/.email/%s' % name) for name in list_names}
inbox = mailbox.Maildir('~/.Maildir', factory=None)

for key in inbox.iterkeys():
    try:
        message = inbox[key]
    except email.errors.MessageParseError:
        continue          # The message is malformed. Just leave it.

    for name in list_names:
        list_id = message['list-id']
        if list_id and name in list_id:
            # Get mailbox to use
            box = boxes[name]

            # Write copy to disk before removing original.
            # If there's a crash, you might duplicate a message, but
```

(下页继续)

(续上页)

```

        # that's better than losing a message completely.
        box.lock()
        box.add(message)
        box.flush()
        box.unlock()

        # Remove original message
        inbox.lock()
        inbox.discard(key)
        inbox.flush()
        inbox.unlock()
        break          # Found destination, so stop looking.

for box in boxes.itervalues():
    box.close()

```

19.5 mimetypes — Map filenames to MIME types

Source code: [Lib/mimetypes.py](#)

The *mimetypes* module converts between a filename or URL and the MIME type associated with the filename extension. Conversions are provided from filename to MIME type and from MIME type to filename extension; encodings are not supported for the latter conversion.

The module provides one class and a number of convenience functions. The functions are the normal interface to this module, but some applications may be interested in the class as well.

The functions described below provide the primary interface for this module. If the module has not been initialized, they will call *init()* if they rely on the information *init()* sets up.

mimetypes.guess_type(url, strict=True)

Guess the type of a file based on its filename or URL, given by *url*. The return value is a tuple (*type*, *encoding*) where *type* is None if the type can't be guessed (missing or unknown suffix) or a string of the form 'type/subtype', usable for a MIME *content-type* header.

encoding is None for no encoding or the name of the program used to encode (e.g. **compress** or **gzip**). The encoding is suitable for use as a *Content-Encoding* header, **not** as a *Content-Transfer-Encoding* header. The mappings are table driven. Encoding suffixes are case sensitive; type suffixes are first tried case sensitively, then case insensitively.

The optional *strict* argument is a flag specifying whether the list of known MIME types is limited to only the official types registered with IANA. When *strict* is True (the default), only the IANA types are supported; when *strict* is False, some additional non-standard but commonly used MIME types are also recognized.

mimetypes.guess_all_extensions(type, strict=True)

Guess the extensions for a file based on its MIME type, given by *type*. The return value is a list of strings giving all possible filename extensions, including the leading dot ('.'). The extensions are not guaranteed to have been associated with any particular data stream, but would be mapped to the MIME type *type* by *guess_type()*.

The optional *strict* argument has the same meaning as with the *guess_type()* function.

mimetypes.guess_extension(type, strict=True)

Guess the extension for a file based on its MIME type, given by *type*. The return value is a string giving a filename extension, including the leading dot ('.'). The extension is not guaranteed to have been associated with any

particular data stream, but would be mapped to the MIME type *type* by `guess_type()`. If no extension can be guessed for *type*, `None` is returned.

The optional *strict* argument has the same meaning as with the `guess_type()` function.

Some additional functions and data items are available for controlling the behavior of the module.

`mimetypes.init(files=None)`

Initialize the internal data structures. If given, *files* must be a sequence of file names which should be used to augment the default type map. If omitted, the file names to use are taken from `knownfiles`; on Windows, the current registry settings are loaded. Each file named in *files* or `knownfiles` takes precedence over those named before it. Calling `init()` repeatedly is allowed.

Specifying an empty list for *files* will prevent the system defaults from being applied: only the well-known values will be present from a built-in list.

在 3.2 版更改: Previously, Windows registry settings were ignored.

`mimetypes.read_mime_types(filename)`

Load the type map given in the file *filename*, if it exists. The type map is returned as a dictionary mapping filename extensions, including the leading dot ('.'), to strings of the form 'type/subtype'. If the file *filename* does not exist or cannot be read, `None` is returned.

`mimetypes.add_type(type, ext, strict=True)`

Add a mapping from the MIME type *type* to the extension *ext*. When the extension is already known, the new type will replace the old one. When the type is already known the extension will be added to the list of known extensions.

When *strict* is `True` (the default), the mapping will be added to the official MIME types, otherwise to the non-standard ones.

`mimetypes.inited`

Flag indicating whether or not the global data structures have been initialized. This is set to `True` by `init()`.

`mimetypes.knownfiles`

List of type map file names commonly installed. These files are typically named `mime.types` and are installed in different locations by different packages.

`mimetypes.suffix_map`

Dictionary mapping suffixes to suffixes. This is used to allow recognition of encoded files for which the encoding and the type are indicated by the same extension. For example, the `.tgz` extension is mapped to `.tar.gz` to allow the encoding and type to be recognized separately.

`mimetypes.encodings_map`

Dictionary mapping filename extensions to encoding types.

`mimetypes.types_map`

Dictionary mapping filename extensions to MIME types.

`mimetypes.common_types`

Dictionary mapping filename extensions to non-standard, but commonly found MIME types.

An example usage of the module:

```
>>> import mimetypes
>>> mimetypes.init()
>>> mimetypes.knownfiles
['/etc/mime.types', '/etc/httpd/mime.types', ... ]
>>> mimetypes.suffix_map['.tgz']
'.tar.gz'
>>> mimetypes.encodings_map['.gz']
```

(下页继续)

(续上页)

```
'gzip'
>>> mimetypes.types_map['.tgz']
'application/x-tar-gz'
```

19.5.1 MimeTypes Objects

The *MimeTypes* class may be useful for applications which may want more than one MIME-type database; it provides an interface similar to the one of the *mimetypes* module.

class *mimetypes.MimeTypes* (*filenames=()*, *strict=True*)

This class represents a MIME-types database. By default, it provides access to the same database as the rest of this module. The initial database is a copy of that provided by the module, and may be extended by loading additional *mime.types*-style files into the database using the *read()* or *readfp()* methods. The mapping dictionaries may also be cleared before loading additional data if the default data is not desired.

The optional *filenames* parameter can be used to cause additional files to be loaded “on top” of the default database.

suffix_map

Dictionary mapping suffixes to suffixes. This is used to allow recognition of encoded files for which the encoding and the type are indicated by the same extension. For example, the *.tgz* extension is mapped to *.tar.gz* to allow the encoding and type to be recognized separately. This is initially a copy of the global *suffix_map* defined in the module.

encodings_map

Dictionary mapping filename extensions to encoding types. This is initially a copy of the global *encodings_map* defined in the module.

types_map

Tuple containing two dictionaries, mapping filename extensions to MIME types: the first dictionary is for the non-standards types and the second one is for the standard types. They are initialized by *common_types* and *types_map*.

types_map_inv

Tuple containing two dictionaries, mapping MIME types to a list of filename extensions: the first dictionary is for the non-standards types and the second one is for the standard types. They are initialized by *common_types* and *types_map*.

guess_extension (*type*, *strict=True*)

Similar to the *guess_extension()* function, using the tables stored as part of the object.

guess_type (*url*, *strict=True*)

Similar to the *guess_type()* function, using the tables stored as part of the object.

guess_all_extensions (*type*, *strict=True*)

Similar to the *guess_all_extensions()* function, using the tables stored as part of the object.

read (*filename*, *strict=True*)

Load MIME information from a file named *filename*. This uses *readfp()* to parse the file.

If *strict* is *True*, information will be added to list of standard types, else to the list of non-standard types.

readfp (*fp*, *strict=True*)

Load MIME type information from an open file *fp*. The file must have the format of the standard *mime.types* files.

If *strict* is *True*, information will be added to the list of standard types, else to the list of non-standard types.

read_windows_registry (*strict=True*)

Load MIME type information from the Windows registry. Availability: Windows.

If *strict* is `True`, information will be added to the list of standard types, else to the list of non-standard types.

3.2 新版功能.

19.6 base64 — Base16, Base32, Base64, Base85 数据编码

源代码: [Lib/base64.py](#)

此模块提供了将二进制数据编码为可打印的 ASCII 字符以及将这些编码解码回二进制数据的函数。它为 **RFC 3548** 指定的 Base16, Base32 和 Base64 编码以及已被广泛接受的 Ascii85 和 Base85 编码提供了编码和解码函数。

RFC 3548 编码的目的是使得二进制数据可以作为电子邮件的内容正确地发送, 用作 URL 的一部分, 或者作为 HTTP POST 请求的一部分。其中的编码算法和 **uuencode** 程序是不同的。

此模块提供了两个接口。新的接口提供了从类字节对象到 ASCII 字节 *bytes* 的编码, 以及将 ASCII 的类字节对象或字符串解码到 *bytes* 的操作。此模块支持定义在 **RFC 3548** 中的所有 base-64 字母表 (普通的、URL 安全的和文件系统安全的)。

旧的接口不提供从字符串的解码操作, 但提供了操作文件对象的编码和解码函数。旧接口只支持标准的 Base64 字母表, 并且按照 **RFC 2045** 的规范每 76 个字符增加一个换行符。注意: 如果你需要支持 **RFC 2045**, 那么使用 *email* 模块可能更加合适。

在 3.3 版更改: 新的接口提供的解码函数现在已经支持只包含 ASCII 的 Unicode 字符串。

在 3.4 版更改: 所有类字节对象 现在已经被所有编码和解码函数接受。添加了对 Ascii85/Base85 的支持。

新的接口提供:

base64.b64encode (*s*, *altchars=None*)

对 *bytes-like object* *s* 进行 Base64 编码, 并返回编码后的 *bytes*。

可选项 *altchars* 必须是一个长 2 字节的 *bytes-like object*, 它指定了用于替换 + 和 / 的字符。这允许应用程序生成 URL 或文件系统安全的 Base64 字符串。默认值是 `None`, 使用标准 Base64 字母表。

base64.b64decode (*s*, *altchars=None*, *validate=False*)

解码 Base64 编码过的 *bytes-like object* 或 ASCII 字符串 *s* 并返回解码过的 *bytes*。

可选项 *altchars* 必须是一个长 2 字节的 *bytes-like object*, 它指定了用于替换 + 和 / 的字符。

如果 *s* 被不正确地填写, 一个 *binascii.Error* 错误将被抛出。

如果 *validate* 值为 `False` (默认情况), 则在填充检查前, 将丢弃既不在标准 base-64 字母表之中也不在备用字母表中的字符。如果 *validate* 为 `True`, 这些非 base64 字符将导致 *binascii.Error*。

base64.standard_b64encode (*s*)

编码 *bytes-like object* *s*, 使用标准 Base64 字母表并返回编码过的 *bytes*。

base64.standard_b64decode (*s*)

解码 *bytes-like object* 或 ASCII 字符串 *s*, 使用标准 Base64 字母表并返回编码过的 *bytes*。

base64.urlsafe_b64encode (*s*)

编码 *bytes-like object* *s*, 使用 URL 与文件系统安全的字母表, 使用 - 以及 _ 代替标准 Base64 字母表中的 + 和 /。返回编码过的 *bytes*。结果中可能包含 =。

`base64.urlsafe_b64decode(s)`

解码 *bytes-like object* 或 ASCII 字符串 *s*，使用 URL 与文件系统安全的字母表，使用 `-` 以及 `_` 代替标准 Base64 字母表中的 `+` 和 `/`。返回解码过的 *bytes*

`base64.b32encode(s)`

用 Base32 编码 *bytes-like object* *s* 并返回编码过的 *bytes*

`base64.b32decode(s, casefold=False, map01=None)`

解码 Base32 编码过的 *bytes-like object* 或 ASCII 字符串 *s* 并返回解码过的 *bytes*。

可选的 *casefold* 是一个指定小写字母是否可接受为输入的标志。为了安全考虑，默认值为 `False`。

RFC 3548 允许将字母 0(zero) 映射为字母 O(oh)，并可以选择是否将字母 1(one) 映射为 I(eye) 或 L(el)。可选参数 *map01* 不是 `None` 时，它的值指定 1 的映射目标 (I 或 l)，当 *map01* 非 `None` 时，0 总是被映射为 O。为了安全考虑，默认值被设为 `None`，如果是这样，0 和 1 不允许被作为输入。

如果 *s* 被错误地填写或输入中存在字母表之外的字符，将抛出 *binascii.Error*。

`base64.b16encode(s)`

用 Base16 编码 *bytes-like object* *s* 并返回编码过的 *bytes*

`base64.b16decode(s, casefold=False)`

解码 Base16 编码过的 *bytes-like object* 或 ASCII 字符串 *s* 并返回解码过的 *bytes*。

可选的 *casefold* 是一个指定小写字母是否可接受为输入的标志。为了安全考虑，默认值为 `False`。

如果 *s* 被错误地填写或输入中存在字母表之外的字符，将抛出 *binascii.Error*。

`base64.a85encode(b, *, foldspaces=False, wrapcol=0, pad=False, adobe=False)`

用 Ascii85 编码 *bytes-like object* *s* 并返回编码过的 *bytes*

foldspaces 是一个可选的标志，使用特殊的短序列 ‘y’ 代替 ‘btoa’ 提供的 4 个连续空格 (ASCII 0x20)。这个特性不被 “标准” Ascii85 编码支持。

wrapcol 控制了输出是否包含换行符 (b'\n')。如果该值非零，则每一行只有该值所限制的字符长度。

pad 控制在编码之前输入是否填充为 4 的倍数。请注意，btoa 实现总是填充。

adobe controls whether the encoded byte sequence is framed with `<~` and `~>`, which is used by the Adobe implementation.

3.4 新版功能.

`base64.a85decode(b, *, foldspaces=False, adobe=False, ignorechars=b'\t\n\r\v')`

解码 Ascii85 编码过的 *bytes-like object* 或 ASCII 字符串 *s* 并返回解码过的 *bytes*。

foldspaces is a flag that specifies whether the ‘y’ short sequence should be accepted as shorthand for 4 consecutive spaces (ASCII 0x20). This feature is not supported by the “standard” Ascii85 encoding.

adobe controls whether the input sequence is in Adobe Ascii85 format (i.e. is framed with `<~` and `~>`).

ignorechars should be a *bytes-like object* or ASCII string containing characters to ignore from the input. This should only contain whitespace characters, and by default contains all whitespace characters in ASCII.

3.4 新版功能.

`base64.b85encode(b, pad=False)`

Encode the *bytes-like object* *b* using base85 (as used in e.g. git-style binary diffs) and return the encoded *bytes*.

If *pad* is true, the input is padded with `b'\0'` so its length is a multiple of 4 bytes before encoding.

3.4 新版功能.

`base64.b85decode(b)`

Decode the base85-encoded *bytes-like object* or ASCII string *b* and return the decoded *bytes*. Padding is implicitly removed, if necessary.

3.4 新版功能.

The legacy interface:

`base64.decode(input, output)`

Decode the contents of the binary *input* file and write the resulting binary data to the *output* file. *input* and *output* must be *file objects*. *input* will be read until `input.readline()` returns an empty bytes object.

`base64.decodebytes(s)`

Decode the *bytes-like object* *s*, which must contain one or more lines of base64 encoded data, and return the decoded *bytes*.

3.1 新版功能.

`base64.decodestring(s)`

Deprecated alias of `decodebytes()`.

3.1 版后已移除.

`base64.encode(input, output)`

Encode the contents of the binary *input* file and write the resulting base64 encoded data to the *output* file. *input* and *output* must be *file objects*. *input* will be read until `input.read()` returns an empty bytes object. `encode()` inserts a newline character (`b'\n'`) after every 76 bytes of the output, as well as ensuring that the output always ends with a newline, as per **RFC 2045** (MIME).

`base64.encodebytes(s)`

Encode the *bytes-like object* *s*, which can contain arbitrary binary data, and return *bytes* containing the base64-encoded data, with newlines (`b'\n'`) inserted after every 76 bytes of output, and ensuring that there is a trailing newline, as per **RFC 2045** (MIME).

3.1 新版功能.

`base64.encodestring(s)`

Deprecated alias of `encodebytes()`.

3.1 版后已移除.

An example usage of the module:

```
>>> import base64
>>> encoded = base64.b64encode(b'data to be encoded')
>>> encoded
b'ZGF0YSB0byBiZSB1bmNvZGVk'
>>> data = base64.b64decode(encoded)
>>> data
b'data to be encoded'
```

参见:

模块 `binascii` 支持模块, 包含 ASCII 到二进制和二进制到 ASCII 转换。

RFC 1521 - MIME (Multipurpose Internet Mail Extensions) Part One: Mechanisms for Specifying and Describing the Format of
Section 5.2, “Base64 Content-Transfer-Encoding,” provides the definition of the base64 encoding.

19.7 binhex 一对 binhex4 文件进行编码和解码

源代码: `Lib/binhex.py`

此模块以 binhex4 格式对文件进行编码和解码，该格式允许 Macintosh 文件以 ASCII 格式表示。仅处理数据分支。

`binhex` 模块定义了以下功能：

`binhex.binhex(input, output)`

将带有文件名 输入的二进制文件转换为 binhex 文件 输出。输出参数可以是文件名或类文件对象 (`write()` 和 `close()` 方法的任何对象)。

`binhex.hexbin(input, output)`

解码 binhex 文件输入。输入可以是支持 `read()` 和 `close()` 方法的文件名或类文件对象。生成的文件将写入名为 `output` 的文件，除非参数为 `None`，在这种情况下，从 binhex 文件中读取输出文件名。

还定义了以下异常：

exception `binhex.Error`

当无法使用 binhex 格式编码某些内容时（例如，文件名太长而无法放入文件名字段中），或者输入未正确编码的 binhex 数据时，会引发异常。

参见：

模块 `binascii` 支持模块，包含 ASCII 到二进制和二进制到 ASCII 转换。

19.7.1 注释

还有一个替代的、功能更强大的编码器和解码器接口，详细信息请参见源代码。

如果您在非 Macintosh 平台上编码或解码文本文件，它们仍将使用旧的 Macintosh 换行符约定（回车符作为行尾）。

19.8 binascii — 二进制和 ASCII 码互转

`binascii` 模块包含很多在二进制和二进制表示的各种 ASCII 码之间转换的方法。通常情况不会直接使用这些函数，而是使用像 `uu`，`base64`，或 `binhex` 这样的封装模块。为了执行效率高，`binascii` 模块含有许多用 C 写的低级函数，这些底层函数被一些高级模块所使用。

注解： `a2b_*` 函数接受只含有 ASCII 码的 Unicode 字符串。其他函数只接受字节类对象（例如 `bytes`，`bytearray` 和其他支持缓冲区协议的对象）。

在 3.3 版更改：ASCII-only unicode strings are now accepted by the `a2b_*` functions.

`binascii` 模块定义了以下函数：

`binascii.a2b_uu(string)`

将单行 `uu` 编码数据转换成二进制数据并返回。`uu` 编码每行的数据通常包含 45 个（二进制）字节，最后一行除外。每行数据后面可能跟有空格。

`binascii.b2a_uu(data)`

Convert binary data to a line of ASCII characters, the return value is the converted line, including a newline char. The length of `data` should be at most 45.

`binascii.a2b_base64(string)`

将 base64 数据块转换成二进制并以二进制数据形式返回。一次可以传递多行数据。

`binascii.b2a_base64 (data, *, newline=True)`

将二进制数据转换为一行用 base64 编码的 ASCII 字符串。返回值是转换后的行数据，如果 `newline` 为 `true`，则返回值包括换行符。该函数的输出符合：rfc: 3548。

在 3.6 版更改：增加 `newline` 形参。

`binascii.a2b_qp (data, header=False)`

将一个引号可打印的数据块转换成二进制数据并返回。一次可以转换多行。如果可选参数 `header` 存在且为 `true`，则数据中的下划线将被解码成空格。

`binascii.b2a_qp (data, quotetabs=False, istext=True, header=False)`

将二进制数据转换为一行或多行带引号可打印编码的 ASCII 字符串。返回值是转换后的行数据。如果可选参数 `quotetabs` 存在且为真值，则对所有制表符和空格进行编码。如果可选参数 `istext` 存在且为真值，则不对新行进行编码，但将对尾随空格进行编码。如果可选参数 `header` 存在且为 `true`，则空格将被编码为下划线 **RFC 1522**。如果可选参数 `header` 存在且为假值，则也会对换行符进行编码；不进行换行转换编码可能会破坏二进制数据流。

`binascii.a2b_hqx (string)`

将 binhex4 格式的 ASCII 数据不进行 RLE 解压缩直接转换为二进制数据。该字符串应包含完整数量的二进制字节，或者（在 binhex4 数据最后部分）剩余位为零。

`binascii.rledecode_hqx (data)`

根据 binhex4 标准对数据执行 RLE 解压缩。该算法在一个字节的数据后使用 `0x90` 作为重复指示符，然后计数。计数 0 指定字节值 `0x90`。该例程返回解压缩的数据，输入数据以孤立的重复指示符结束的情况下，将引发 `Incomplete` 异常。

在 3.2 版更改：仅接受 `bytestring` 或 `bytearray` 对象作为输入。

`binascii.rlecode_hqx (data)`

在 `data` 上执行 binhex4 游程编码压缩并返回结果。

`binascii.b2a_hqx (data)`

执行 hexbin4 类型二进制到 ASCII 码的转换并返回结果字符串。输入数据应经过 RLE 编码，且数据长度可被 3 整除（除了最后一个片段）。

`binascii.crc_hqx (data, value)`

以 `value` 作为初始 CRC 计算 `data` 的 16 位 CRC 值，返回其结果。这里使用 CRC-CCITT 生成多项式 $x^{16} + x^{12} + x^5 + 1$ ，通常表示为 `0x1021`。该 CRC 被用于 binhex4 格式。

`binascii.crc32 (data[, value])`

计算 CRC-32，从 `value` 的初始 CRC 开始计算 `data` 的 32 位校验和。默认初始 CRC 为零。该算法与 ZIP 文件校验和一致。由于该算法被设计用作校验和算法，因此不适合用作通用散列算法。使用方法如下：

```
print(binascii.crc32(b"hello world"))
# Or, in two pieces:
crc = binascii.crc32(b"hello")
crc = binascii.crc32(b" world", crc)
print('crc32 = {:#010x}'.format(crc))
```

在 3.0 版更改：校验结果始终是无符号类型的。要在所有 Python 版本和平台上生成相同的数值，请使用 `crc32(data) & 0xffffffff`。

`binascii.b2a_hex (data)`

`binascii.hexlify (data)`

返回二进制数据 `data` 的十六进制表示形式。`data` 的每个字节都被转换为相应的 2 位十六进制表示形式。因此返回的字节对象的长度是 `data` 的两倍。

`binascii.a2b_hex (hexstr)`

`binascii.unhexlify (hexstr)`

返回由十六进制字符串 `hexstr` 表示的二进制数据。此函数功能与 `b2a_hex()` 相反。`hexstr` 必须包含偶数个十六进制数字（可以是大写或小写），否则会引发 `Error` 异常。

exception `binascii.Error`

通常是因为编程错误引发的异常。

exception `binascii.Incomplete`

数据不完整引发的异常。通常不是编程错误导致的，可以通过读取更多的数据并再次尝试来处理该异常。

参见：

模块 `base64` 支持在 16, 32, 64, 85 进制中进行符合 RFC 协议的 base64 样式编码。

模块 `binhex` 支持在 Macintosh 上使用的 binhex 格式。

模块 `uu` 支持在 Unix 上使用的 UU 编码。

模块 `quopri` 支持在 MIME 版本电子邮件中使用引号可打印编码。

19.9 quopri —Encode and decode MIME quoted-printable data

Source code: [Lib/quopri.py](#)

This module performs quoted-printable transport encoding and decoding, as defined in [RFC 1521](#): “MIME (Multipurpose Internet Mail Extensions) Part One: Mechanisms for Specifying and Describing the Format of Internet Message Bodies” . The quoted-printable encoding is designed for data where there are relatively few nonprintable characters; the base64 encoding scheme available via the `base64` module is more compact if there are many such characters, as when sending a graphics file.

`quopri.decode(input, output, header=False)`

Decode the contents of the *input* file and write the resulting decoded binary data to the *output* file. *input* and *output* must be *binary file objects*. If the optional argument *header* is present and true, underscore will be decoded as space. This is used to decode “Q”-encoded headers as described in [RFC 1522](#): “MIME (Multipurpose Internet Mail Extensions) Part Two: Message Header Extensions for Non-ASCII Text” .

`quopri.encode(input, output, quotetabs, header=False)`

Encode the contents of the *input* file and write the resulting quoted-printable data to the *output* file. *input* and *output* must be *binary file objects*. *quotetabs*, a non-optional flag which controls whether to encode embedded spaces and tabs; when true it encodes such embedded whitespace, and when false it leaves them unencoded. Note that spaces and tabs appearing at the end of lines are always encoded, as per [RFC 1521](#). *header* is a flag which controls if spaces are encoded as underscores as per [RFC 1522](#).

`quopri.decodestring(s, header=False)`

Like `decode()`, except that it accepts a source *bytes* and returns the corresponding decoded *bytes*.

`quopri.encodestring(s, quotetabs=False, header=False)`

Like `encode()`, except that it accepts a source *bytes* and returns the corresponding encoded *bytes*. By default, it sends a `False` value to *quotetabs* parameter of the `encode()` function.

参见：

模块 `base64` Encode and decode MIME base64 data

19.10 uu —Encode and decode uuencode files

Source code: [Lib/uu.py](#)

This module encodes and decodes files in uuencode format, allowing arbitrary binary data to be transferred over ASCII-only connections. Wherever a file argument is expected, the methods accept a file-like object. For backwards compatibility, a string containing a pathname is also accepted, and the corresponding file will be opened for reading and writing; the pathname '-' is understood to mean the standard input or output. However, this interface is deprecated; it's better for the caller to open the file itself, and be sure that, when required, the mode is 'rb' or 'wb' on Windows.

This code was contributed by Lance Ellinghouse, and modified by Jack Jansen.

The `uu` module defines the following functions:

`uu.encode(in_file, out_file, name=None, mode=None)`

Uuencode file *in_file* into file *out_file*. The uuencoded file will have the header specifying *name* and *mode* as the defaults for the results of decoding the file. The default defaults are taken from *in_file*, or '-' and 0o666 respectively.

`uu.decode(in_file, out_file=None, mode=None, quiet=False)`

This call decodes uuencoded file *in_file* placing the result on file *out_file*. If *out_file* is a pathname, *mode* is used to set the permission bits if the file must be created. Defaults for *out_file* and *mode* are taken from the uuencode header. However, if the file specified in the header already exists, a `uu.Error` is raised.

`decode()` may print a warning to standard error if the input was produced by an incorrect uuencoder and Python could recover from that error. Setting *quiet* to a true value silences this warning.

exception `uu.Error`

Subclass of `Exception`, this can be raised by `uu.decode()` under various situations, such as described above, but also including a badly formatted header, or truncated input file.

参见:

模块 `binascii` 支持模块, 包含 ASCII 到二进制和二进制到 ASCII 转换。

结构化标记处理工具

Python 支持各种模块，以处理各种形式的结构化数据标记。这包括使用标准通用标记语言（SGML）和超文本标记语言（HTML）的模块，以及使用可扩展标记语言（XML）的几个接口。

20.1 html —超文本标记语言支持

源码： `Lib/html/__init__.py`

该模块定义了操作 HTML 的工具。

`html.escape(s, quote=True)`

将字符串 *s* 中的字符 “&”、< 和 > 转换为安全的 HTML 序列。如果需要在 HTML 中显示可能包含此类字符的文本，请使用此选项。如果可选的标志 *quote* 为真值，则字符 (") 和 (') 也被转换；这有助于包含在由引号分隔的 HTML 属性中，如 ``。

3.2 新版功能.

`html.unescape(s)`

将字符串 *s* 中的所有命名和数字字符引用 (例如 `>`, `>`, `>`) 转换为相应的 Unicode 字符。此函数使用 HTML 5 标准为有效和无效字符引用定义的规则，以及 [HTML 5 命名字符引用列表](#)。

3.4 新版功能.

html 包中的子模块是：

- `html.parser` —— 具有宽松解析模式的 HTML / XHTML 解析器
- `html.entities` – HTML 实体定义

20.2 `html.parser` — 简单的 HTML 和 XHTML 解析器

源代码: `Lib/html/parser.py`

这个模块定义了一个 `HTMLParser` 类, 为 HTML (超文本标记语言) 和 XHTML 文本文件解析提供基础。

class `html.parser.HTMLParser` (*, `convert_charrefs=True`)

创建一个能解析无效标记的解析器实例。

如果 `convert_charrefs` 为 `True` (默认值), 则所有字符引用 (`script/style` 元素中的除外) 都会自动转换为相应的 Unicode 字符。

一个 `HTMLParser` 类的实例用来接受 HTML 数据, 并在标记开始、标记结束、文本、注释和其他元素标记出现的时候调用对应的方法。要实现具体的行为, 请使用 `HTMLParser` 的子类并重载其方法。

这个解析器不检查结束标记是否与开始标记匹配, 也不会因外层元素完毕而隐式关闭了的元素引发结束标记处理。

在 3.4 版更改: `convert_charrefs` 关键字参数被添加。

在 3.5 版更改: `convert_charrefs` 参数的默认值现在为 `True`。

20.2.1 HTML 解析器的示例程序

下面是简单的 HTML 解析器的一个基本示例, 使用 `HTMLParser` 类, 当遇到开始标记、结束标记以及数据的时候将内容打印出来。

```
from html.parser import HTMLParser

class MyHTMLParser(HTMLParser):
    def handle_starttag(self, tag, attrs):
        print("Encountered a start tag:", tag)

    def handle_endtag(self, tag):
        print("Encountered an end tag :", tag)

    def handle_data(self, data):
        print("Encountered some data  :", data)

parser = MyHTMLParser()
parser.feed('<html><head><title>Test</title></head>'
          '<body><h1>Parse me!</h1></body></html>')
```

输出是:

```
Encountered a start tag: html
Encountered a start tag: head
Encountered a start tag: title
Encountered some data  : Test
Encountered an end tag : title
Encountered an end tag : head
Encountered a start tag: body
Encountered a start tag: h1
Encountered some data  : Parse me!
Encountered an end tag : h1
Encountered an end tag : body
Encountered an end tag : html
```

20.2.2 HTMLParser 方法

HTMLParser 实例有下列方法：

HTMLParser.feed(*data*)

填充一些文本到解析器中。如果包含完整的元素，则被处理；如果数据不完整，将被缓冲直到更多的数据被填充，或者 *close()* 被调用。*data* 必须为 *str* 类型。

HTMLParser.close()

如同后面跟着一个文件结束标记一样，强制处理所有缓冲数据。这个方法能被派生类重新定义，用于在输入的末尾定义附加处理，但是重定义的版本应当始终调用基类 *HTMLParser* 的 *close()* 方法。

HTMLParser.reset()

重置实例。丢失所有未处理的数据。在实例化阶段被隐式调用。

HTMLParser.getpos()

返回当前行号和偏移值。

HTMLParser.get_starttag_text()

返回最近打开的开始标记中的文本。结构化处理时通常应该不需要这个，但在处理“已部署”的 HTML 或是在以最小改变来重新生成输入时可能会有用处（例如可以保留属性间的空格等）。

下列方法将在遇到数据或者标记元素的时候被调用。他们需要在子类中重载。基类的实现中没有任何实际操作（除了 *handle_startendtag()*）：

HTMLParser.handle_starttag(*tag*, *attrs*)

这个方法在标签开始的时候被调用（例如：<div id="main">）。

tag 参数是小写的标记名。*attrs* 参数是一个 (name, value) 形式的列表，包含了所有在标记的 <> 括号中找到的属性。*name* 转换为小写，*value* 的引号被去除，字符和实体引用都会被替换。

实例中，对于标签 ，这个方法将以下列形式被调用 *handle_starttag('a', [('href', 'https://www.cwi.nl/')])*。

html.entities 中的所有实体引用，会被替换为属性值。

HTMLParser.handle_endtag(*tag*)

此方法被用来处理元素的结束标记（例如：</div>）。

tag 参数是小写的标签名。

HTMLParser.handle_startendtag(*tag*, *attrs*)

类似于 *handle_starttag()*，只是在解析器遇到 XHTML 样式的空标记时被调用（）。这个方法能被需要这种特殊词法信息的子类重载；默认实现仅简单调用 *handle_starttag()* 和 *handle_endtag()*。

HTMLParser.handle_data(*data*)

这个方法被用来处理任意数据（例如：文本节点和 <script>...</script> 以及 <style>...</style> 中的内容）。

HTMLParser.handle_entityref(*name*)

这个方法被用于处理 &name; 形式的命名字符引用（例如 >），其中 *name* 是通用的实体引用（例如：'gt'）。如果 *convert_charrefs* 为 True，该方法永远不会被调用。

HTMLParser.handle_charref(*name*)

这个方法被用来处理 &#NNN; 和 &#xNNN; 形式的十进制和十六进制字符引用。例如，> 等效的十进制形式为 >，而十六进制形式为 >；在这种情况下，方法将收到 '62' 或 'x3E'。如果 *convert_charrefs* 为 True，则该方法永远不会被调用。

HTMLParser.handle_comment(*data*)

这个方法在遇到注释的时候被调用（例如：<!--comment-->）。

例如，<!-- comment --> 这个注释会用 'comment' 作为参数调用此方法。

Internet Explorer 条件注释 (condcoms) 的内容也被发送到这个方法, 因此, 对于 `<!--[if IE 9]>IE9-specific content<![endif]-->`, 这个方法将接收到 `'[if IE 9]>IE9-specific content<![endif]'`。

`HTMLParser.handle_decl(decl)`

这个方法用来处理 HTML doctype 申明 (例如 `<!DOCTYPE html>`)。

`decl` 形参为 `<![...>` 标记中的所有内容 (例如: `'DOCTYPE html'`)。

`HTMLParser.handle_pi(data)`

此方法在遇到处理指令的时候被调用。`data` 形参将包含整个处理指令。例如, 对于处理指令 `<?proc color='red'>`, 这个方法将以 `handle_pi("proc color='red'")` 形式被调用。它旨在被派生类重载; 基类实现中无任何实际操作。

注解: `HTMLParser` 类使用 SGML 语法规则处理指令。使用 `'?'` 结尾的 XHTML 处理指令将导致 `'?'` 包含在 `data` 中。

`HTMLParser.unknown_decl(data)`

当解析器读到无法识别的声明时, 此方法被调用。

`data` 形参为 `<![...>` 标记中的所有内容。某些时候对派生类的重载很有用。基类实现中无任何实际操作。

20.2.3 示例

下面的类实现了一个解析器, 用于更多示例的演示:

```
from html.parser import HTMLParser
from html.entities import name2codepoint

class MyHTMLParser(HTMLParser):
    def handle_starttag(self, tag, attrs):
        print("Start tag:", tag)
        for attr in attrs:
            print("    attr:", attr)

    def handle_endtag(self, tag):
        print("End tag  :", tag)

    def handle_data(self, data):
        print("Data      :", data)

    def handle_comment(self, data):
        print("Comment   :", data)

    def handle_entityref(self, name):
        c = chr(name2codepoint[name])
        print("Named ent:", c)

    def handle_charref(self, name):
        if name.startswith('x'):
            c = chr(int(name[1:], 16))
        else:
            c = chr(int(name))
        print("Num ent  :", c)
```

(下页继续)

(续上页)

```
def handle_decl(self, data):
    print("Decl      :", data)

parser = MyHTMLParser()
```

解析一个文档类型声明:

```
>>> parser.feed('<!DOCTYPE HTML PUBLIC "-//W3C//DTD HTML 4.01//EN" '
...             '"http://www.w3.org/TR/html4/strict.dtd">')
Decl      : DOCTYPE HTML PUBLIC "-//W3C//DTD HTML 4.01//EN" "http://www.w3.org/TR/
↳html4/strict.dtd"
```

解析一个具有一些属性和标题的元素:

```
>>> parser.feed('')
Start tag: img
  attr: ('src', 'python-logo.png')
  attr: ('alt', 'The Python logo')
>>>
>>> parser.feed('<h1>Python</h1>')
Start tag: h1
Data      : Python
End tag   : h1
```

script 和 style 元素中的内容原样返回, 无需进一步解析:

```
>>> parser.feed('<style type="text/css">#python { color: green }</style>')
Start tag: style
  attr: ('type', 'text/css')
Data      : #python { color: green }
End tag   : style

>>> parser.feed('<script type="text/javascript">'
...             'alert("<strong>hello!</strong>");</script>')
Start tag: script
  attr: ('type', 'text/javascript')
Data      : alert("<strong>hello!</strong>");
End tag   : script
```

解析注释:

```
>>> parser.feed('<!-- a comment -->'
...             '<!--[if IE 9]>IE-specific content<![endif]-->')
Comment   : a comment
Comment   : [if IE 9]>IE-specific content<![endif]
```

解析命名或数字形式的字符引用, 并把他们转换到正确的字符(注意: 这3种转义都是 '>'):

```
>>> parser.feed('&gt;&#62;&#x3E;')
Named ent: >
Num ent   : >
Num ent   : >
```

填充不完整的块给 `feed()` 执行, `handle_data()` 可能会多次调用(除非 `convert_charrefs` 被设置为 `True`):

```
>>> for chunk in ['<sp', 'an>buff', 'ered ', 'text</s', 'pan>']:
...     parser.feed(chunk)
...
Start tag: span
Data      : buff
Data      : ered
Data      : text
End tag   : span
```

解析无效的 HTML (例如: 未引用的属性) 也能正常运行:

```
>>> parser.feed('<p><a class=link href=#main>tag soup</p ></a>')
Start tag: p
Start tag: a
  attr: ('class', 'link')
  attr: ('href', '#main')
Data    : tag soup
End tag : p
End tag : a
```

20.3 `html.entities` —HTML 一般实体的定义

源码: [Lib/html/entities.py](#)

该模块定义了四个词典, `html5`、`name2codepoint`、`codepoint2name`、以及`entitydefs`。

`html.entities.html5`

将 HTML5 命名字符引用¹ 映射到等效的 Unicode 字符的字典, 例如 `html5['gt;'] == '>'`。请注意, 尾随的分号包含在名称中 (例如 `'gt;'`), 但是即使没有分号, 一些名称也会被标准接受, 在这种情况下, 名称出现时带有和不带有 `';`。另见 `html.unescape()`。

3.3 新版功能.

`html.entities.entitydefs`

将 XHTML 1.0 实体定义映射到 ISO Latin-1 中的替换文本的字典。

`html.entities.name2codepoint`

将 HTML 实体名称映射到 Unicode 代码点的字典。

`html.entities.codepoint2name`

将 Unicode 代码点映射到 HTML 实体名称的字典。

20.4 XML 处理模块

源码: [Lib/xml/](#)

用于处理 XML 的 Python 接口分组在 `xml` 包中。

¹ 参见 <https://www.w3.org/TR/html5/syntax.html#named-character-references>

警告： XML 模块对于错误或恶意构造的数据是不安全的。如果需要解析不受信任或未经身份验证的数据，请参阅[XML 漏洞](#)和[defusedxml](#)和[defusedexpat](#)软件包部分。

值得注意的是`xml`包中的模块要求至少有一个 SAX 兼容的 XML 解析器可用。在 Python 中包含 Expat 解析器，因此`xml.parsers.expat`模块将始终可用。

`xml.dom`和`xml.sax`包的文档是 DOM 和 SAX 接口的 Python 绑定的定义。

XML 处理子模块包括：

- `xml.etree.ElementTree`: ElementTree API，一个简单而轻量级的 XML 处理器
- `xml.dom`: DOM API 定义
- `xml.dom.minidom`: 最小的 DOM 实现
- `xml.dom.pulldom`: 支持构建部分 DOM 树
- `xml.sax`: SAX2 基类和便利函数
- `xml.parsers.expat`: Expat 解析器绑定

20.4.1 XML 漏洞

XML 处理模块对于恶意构造的数据是不安全的。攻击者可能滥用 XML 功能来执行拒绝服务攻击、访问本地文件、生成与其它计算机的网络连接或绕过防火墙。

下表概述了已知的攻击以及各种模块是否容易受到攻击。

种类	sax	etree	minidom	pulldom	xmlrpc
billion laughs	易受攻击	易受攻击	易受攻击	易受攻击	易受攻击
quadratic blowup	易受攻击	易受攻击	易受攻击	易受攻击	易受攻击
external entity expansion	安全 (4)	安全 (1)	安全 (2)	安全 (4)	安全 (3)
DTD retrieval	安全 (4)	安全	安全	安全 (4)	安全
decompression bomb	安全	安全	安全	安全	易受攻击

1. `xml.etree.ElementTree` 不会扩展外部实体并在实体发生时引发 `ParserError`。
2. `xml.dom.minidom` 不会扩展外部实体，只是简单地返回未扩展的实体。
3. `xmlrpclib` 不扩展外部实体并省略它们。
4. Since Python 3.8.0, external general entities are no longer processed by default since Python.

billion laughs / exponential entity expansion (狂笑/递归实体扩展) [Billion Laughs](#) 攻击 – 也称为递归实体扩展 – 使用多级嵌套实体。每个实体多次引用另一个实体，最终实体定义包含一个小字符串。指数级扩展导致几千 GB 的文本，并消耗大量内存和 CPU 时间。

quadratic blowup entity expansion (二次爆炸实体扩展) 二次爆炸攻击类似于 [Billion Laughs](#) 攻击，它也滥用实体扩展。它不是嵌套实体，而是一遍又一遍地重复一个具有几千个字符的大型实体。攻击不如递归情况有效，但它避免触发禁止深度嵌套实体的解析器对策。

external entity expansion 实体声明可以包含的不仅仅是替换文本。它们还可以指向外部资源或本地文件。XML 解析器访问资源并将内容嵌入到 XML 文档中。

DTD retrieval Python 的一些 XML 库`xml.dom.pulldom`从远程或本地位置检索文档类型定义。该功能与外部实体扩展问题具有相似的含义。

decompression bomb Decompression bombs（解压炸弹，又名 **ZIP bomb**）适用于所有可以解析压缩 XML 流（例如 gzip 压缩的 HTTP 流或 LZMA 压缩的文件）的 XML 库。对于攻击者来说，它可以将传输的数据量减少三个量级或更多。

PyPI 上 `defusedxml` 的文档包含有关所有已知攻击向量的更多信息以及示例和参考。

20.4.2 `defusedxml` 和 `defusedexpat` 软件包

`defusedxml` 是一个纯 Python 软件包，它修改了所有标准库 XML 解析器的子类，可以防止任何潜在的恶意操作。对于解析不受信任的 XML 数据的任何服务器代码，建议使用此程序包。该软件包还提供了有关更多 XML 漏洞（如 XPath 注入）的示例漏洞和扩展文档。

`defusedexpat` 提供了一个修改过的 `libexpat` 和一个打过补丁的 `pyexpat` 模块，它有针对性扩展 DoS 攻击的对策。`defusedexpat` 模块仍然允许合理且可配置的实体扩展量。这些修改可能包含在 Python 的某些未来版本中，但不会包含在 Python 的任何修复版本中，因为它们会破坏向后兼容性。

20.5 `xml.etree.ElementTree` —The `ElementTree` XML API

Source code: `Lib/xml/etree/ElementTree.py`

The `xml.etree.ElementTree` module implements a simple and efficient API for parsing and creating XML data.

在 3.3 版更改: This module will use a fast implementation whenever available. The `xml.etree.cElementTree` module is deprecated.

警告: The `xml.etree.ElementTree` module is not secure against maliciously constructed data. If you need to parse untrusted or unauthenticated data see [XML 漏洞](#).

20.5.1 教程

This is a short tutorial for using `xml.etree.ElementTree` (ET in short). The goal is to demonstrate some of the building blocks and basic concepts of the module.

XML tree and elements

XML is an inherently hierarchical data format, and the most natural way to represent it is with a tree. ET has two classes for this purpose - `ElementTree` represents the whole XML document as a tree, and `Element` represents a single node in this tree. Interactions with the whole document (reading and writing to/from files) are usually done on the `ElementTree` level. Interactions with a single XML element and its sub-elements are done on the `Element` level.

Parsing XML

We'll be using the following XML document as the sample data for this section:

```
<?xml version="1.0"?>
<data>
  <country name="Liechtenstein">
    <rank>1</rank>
    <year>2008</year>
    <gdppc>141100</gdppc>
    <neighbor name="Austria" direction="E"/>
    <neighbor name="Switzerland" direction="W"/>
  </country>
  <country name="Singapore">
    <rank>4</rank>
    <year>2011</year>
    <gdppc>59900</gdppc>
    <neighbor name="Malaysia" direction="N"/>
  </country>
  <country name="Panama">
    <rank>68</rank>
    <year>2011</year>
    <gdppc>13600</gdppc>
    <neighbor name="Costa Rica" direction="W"/>
    <neighbor name="Colombia" direction="E"/>
  </country>
</data>
```

We can import this data by reading from a file:

```
import xml.etree.ElementTree as ET
tree = ET.parse('country_data.xml')
root = tree.getroot()
```

Or directly from a string:

```
root = ET.fromstring(country_data_as_string)
```

fromstring() parses XML from a string directly into an *Element*, which is the root element of the parsed tree. Other parsing functions may create an *ElementTree*. Check the documentation to be sure.

As an *Element*, root has a tag and a dictionary of attributes:

```
>>> root.tag
'data'
>>> root.attrib
{}
```

It also has children nodes over which we can iterate:

```
>>> for child in root:
...     print(child.tag, child.attrib)
...
country {'name': 'Liechtenstein'}
country {'name': 'Singapore'}
country {'name': 'Panama'}
```

Children are nested, and we can access specific child nodes by index:

```
>>> root[0][1].text
'2008'
```

注解: Not all elements of the XML input will end up as elements of the parsed tree. Currently, this module skips over any XML comments, processing instructions, and document type declarations in the input. Nevertheless, trees built using this module's API rather than parsing from XML text can have comments and processing instructions in them; they will be included when generating XML output. A document type declaration may be accessed by passing a custom *TreeBuilder* instance to the *XMLParser* constructor.

Pull API for non-blocking parsing

Most parsing functions provided by this module require the whole document to be read at once before returning any result. It is possible to use an *XMLParser* and feed data into it incrementally, but it is a push API that calls methods on a callback target, which is too low-level and inconvenient for most needs. Sometimes what the user really wants is to be able to parse XML incrementally, without blocking operations, while enjoying the convenience of fully constructed *Element* objects.

The most powerful tool for doing this is *XMLPullParser*. It does not require a blocking read to obtain the XML data, and is instead fed with data incrementally with *XMLPullParser.feed()* calls. To get the parsed XML elements, call *XMLPullParser.read_events()*. Here is an example:

```
>>> parser = ET.XMLPullParser(['start', 'end'])
>>> parser.feed('<mytag>sometext')
>>> list(parser.read_events())
[('start', <Element 'mytag' at 0x7fa66db2be58>)]
>>> parser.feed(' more text</mytag>')
>>> for event, elem in parser.read_events():
...     print(event)
...     print(elem.tag, 'text=', elem.text)
...
end
```

The obvious use case is applications that operate in a non-blocking fashion where the XML data is being received from a socket or read incrementally from some storage device. In such cases, blocking reads are unacceptable.

Because it's so flexible, *XMLPullParser* can be inconvenient to use for simpler use-cases. If you don't mind your application blocking on reading XML data but would still like to have incremental parsing capabilities, take a look at *iterparse()*. It can be useful when you're reading a large XML document and don't want to hold it wholly in memory.

Finding interesting elements

Element has some useful methods that help iterate recursively over all the sub-tree below it (its children, their children, and so on). For example, *Element.iter()*:

```
>>> for neighbor in root.iter('neighbor'):
...     print(neighbor.attrib)
...
{'name': 'Austria', 'direction': 'E'}
{'name': 'Switzerland', 'direction': 'W'}
{'name': 'Malaysia', 'direction': 'N'}
{'name': 'Costa Rica', 'direction': 'W'}
{'name': 'Colombia', 'direction': 'E'}
```

`Element.findall()` finds only elements with a tag which are direct children of the current element. `Element.find()` finds the *first* child with a particular tag, and `Element.text` accesses the element's text content. `Element.get()` accesses the element's attributes:

```
>>> for country in root.findall('country'):
...     rank = country.find('rank').text
...     name = country.get('name')
...     print(name, rank)
...
Liechtenstein 1
Singapore 4
Panama 68
```

More sophisticated specification of which elements to look for is possible by using *XPath*.

Modifying an XML File

`ElementTree` provides a simple way to build XML documents and write them to files. The `ElementTree.write()` method serves this purpose.

Once created, an `Element` object may be manipulated by directly changing its fields (such as `Element.text`), adding and modifying attributes (`Element.set()` method), as well as adding new children (for example with `Element.append()`).

Let's say we want to add one to each country's rank, and add an `updated` attribute to the rank element:

```
>>> for rank in root.iter('rank'):
...     new_rank = int(rank.text) + 1
...     rank.text = str(new_rank)
...     rank.set('updated', 'yes')
...
>>> tree.write('output.xml')
```

Our XML now looks like this:

```
<?xml version="1.0"?>
<data>
  <country name="Liechtenstein">
    <rank updated="yes">2</rank>
    <year>2008</year>
    <gdppc>141100</gdppc>
    <neighbor name="Austria" direction="E"/>
    <neighbor name="Switzerland" direction="W"/>
  </country>
  <country name="Singapore">
    <rank updated="yes">5</rank>
    <year>2011</year>
    <gdppc>59900</gdppc>
    <neighbor name="Malaysia" direction="N"/>
  </country>
  <country name="Panama">
    <rank updated="yes">69</rank>
    <year>2011</year>
    <gdppc>13600</gdppc>
    <neighbor name="Costa Rica" direction="W"/>
    <neighbor name="Colombia" direction="E"/>
  </country>
</data>
```

(下页继续)

(续上页)

```
</country>
</data>
```

We can remove elements using `Element.remove()`. Let's say we want to remove all countries with a rank higher than 50:

```
>>> for country in root.findall('country'):
...     rank = int(country.find('rank').text)
...     if rank > 50:
...         root.remove(country)
...
>>> tree.write('output.xml')
```

Our XML now looks like this:

```
<?xml version="1.0"?>
<data>
  <country name="Liechtenstein">
    <rank updated="yes">2</rank>
    <year>2008</year>
    <gdppc>141100</gdppc>
    <neighbor name="Austria" direction="E"/>
    <neighbor name="Switzerland" direction="W"/>
  </country>
  <country name="Singapore">
    <rank updated="yes">5</rank>
    <year>2011</year>
    <gdppc>59900</gdppc>
    <neighbor name="Malaysia" direction="N"/>
  </country>
</data>
```

Building XML documents

The `SubElement()` function also provides a convenient way to create new sub-elements for a given element:

```
>>> a = ET.Element('a')
>>> b = ET.SubElement(a, 'b')
>>> c = ET.SubElement(a, 'c')
>>> d = ET.SubElement(c, 'd')
>>> ET.dump(a)
<a><b /><c><d /></c></a>
```

Parsing XML with Namespaces

If the XML input has `namespaces`, tags and attributes with prefixes in the form `prefix:sometag` get expanded to `{uri}sometag` where the *prefix* is replaced by the full *URI*. Also, if there is a `default namespace`, that full URI gets prepended to all of the non-prefixed tags.

Here is an XML example that incorporates two namespaces, one with the prefix “fictional” and the other serving as the default namespace:

```
<?xml version="1.0"?>
<actors xmlns:fictional="http://characters.example.com"
        xmlns="http://people.example.com">
  <actor>
    <name>John Cleese</name>
    <fictional:character>Lancelot</fictional:character>
    <fictional:character>Archie Leach</fictional:character>
  </actor>
  <actor>
    <name>Eric Idle</name>
    <fictional:character>Sir Robin</fictional:character>
    <fictional:character>Gunther</fictional:character>
    <fictional:character>Commander Clement</fictional:character>
  </actor>
</actors>
```

One way to search and explore this XML example is to manually add the URI to every tag or attribute in the xpath of a `find()` or `findall()`:

```
root = fromstring(xml_text)
for actor in root.findall('{http://people.example.com}actor'):
    name = actor.find('{http://people.example.com}name')
    print(name.text)
    for char in actor.findall('{http://characters.example.com}character'):
        print(' |-->', char.text)
```

A better way to search the namespaced XML example is to create a dictionary with your own prefixes and use those in the search functions:

```
ns = {'real_person': 'http://people.example.com',
      'role': 'http://characters.example.com'}

for actor in root.findall('real_person:actor', ns):
    name = actor.find('real_person:name', ns)
    print(name.text)
    for char in actor.findall('role:character', ns):
        print(' |-->', char.text)
```

These two approaches both output:

```
John Cleese
 |--> Lancelot
 |--> Archie Leach
Eric Idle
 |--> Sir Robin
 |--> Gunther
 |--> Commander Clement
```

Additional resources

See <http://effbot.org/zone/element-index.htm> for tutorials and links to other docs.

20.5.2 XPath support

This module provides limited support for [XPath expressions](#) for locating elements in a tree. The goal is to support a small subset of the abbreviated syntax; a full XPath engine is outside the scope of the module.

示例

Here’s an example that demonstrates some of the XPath capabilities of the module. We’ll be using the `countrydata` XML document from the [Parsing XML](#) section:

```
import xml.etree.ElementTree as ET

root = ET.fromstring(countrydata)

# Top-level elements
root.findall(".")

# All 'neighbor' grand-children of 'country' children of the top-level
# elements
root.findall("./country/neighbor")

# Nodes with name='Singapore' that have a 'year' child
root.findall("./year/..[@name='Singapore']")

# 'year' nodes that are children of nodes with name='Singapore'
root.findall("./*[@name='Singapore']/year")

# All 'neighbor' nodes that are the second child of their parent
root.findall("./neighbor[2]")
```

Supported XPath syntax

Syntax	意义
tag	Selects all child elements with the given tag. For example, <code>spam</code> selects all child elements named <code>spam</code> , and <code>spam/egg</code> selects all grandchildren named <code>egg</code> in all children named <code>spam</code> .
*	Selects all child elements. For example, <code>*/egg</code> selects all grandchildren named <code>egg</code> .
.	Selects the current node. This is mostly useful at the beginning of the path, to indicate that it’s a relative path.
//	Selects all subelements, on all levels beneath the current element. For example, <code>./egg</code> selects all <code>egg</code> elements in the entire tree.
..	Selects the parent element. Returns <code>None</code> if the path attempts to reach the ancestors of the start element (the element <code>find</code> was called on).
[@attrib]	Selects all elements that have the given attribute.
[@attrib='value']	Selects all elements for which the given attribute has the given value. The value cannot contain quotes.
[tag]	Selects all elements that have a child named <code>tag</code> . Only immediate children are supported.
[tag='text']	Selects all elements that have a child named <code>tag</code> whose complete text content, including descendants, equals the given <code>text</code> .
[position]	Selects all elements that are located at the given position. The position can be either an integer (1 is the first position), the expression <code>last()</code> (for the last position), or a position relative to the last position (e.g. <code>last()-1</code>).

Predicates (expressions within square brackets) must be preceded by a tag name, an asterisk, or another predicate. `position` predicates must be preceded by a tag name.

20.5.3 参考引用

函数

`xml.etree.ElementTree.Comment (text=None)`

Comment element factory. This factory function creates a special element that will be serialized as an XML comment by the standard serializer. The comment string can be either a bytestring or a Unicode string. *text* is a string containing the comment string. Returns an element instance representing a comment.

Note that *XMLParser* skips over comments in the input instead of creating comment objects for them. An *ElementTree* will only contain comment nodes if they have been inserted into the tree using one of the *Element* methods.

`xml.etree.ElementTree.dump (elem)`

Writes an element tree or element structure to `sys.stdout`. This function should be used for debugging only.

The exact output format is implementation dependent. In this version, it's written as an ordinary XML file.

elem is an element tree or an individual element.

`xml.etree.ElementTree.fromstring (text)`

Parses an XML section from a string constant. Same as *XML()*. *text* is a string containing XML data. Returns an *Element* instance.

`xml.etree.ElementTree.fromstringlist (sequence, parser=None)`

Parses an XML document from a sequence of string fragments. *sequence* is a list or other sequence containing XML data fragments. *parser* is an optional parser instance. If not given, the standard *XMLParser* parser is used. Returns an *Element* instance.

3.2 新版功能.

`xml.etree.ElementTree.iselement (element)`

Checks if an object appears to be a valid element object. *element* is an element instance. Returns a true value if this is an element object.

`xml.etree.ElementTree.iterparse (source, events=None, parser=None)`

Parses an XML section into an element tree incrementally, and reports what's going on to the user. *source* is a filename or *file object* containing XML data. *events* is a sequence of events to report back. The supported events are the strings "start", "end", "start-ns" and "end-ns" (the "ns" events are used to get detailed namespace information). If *events* is omitted, only "end" events are reported. *parser* is an optional parser instance. If not given, the standard *XMLParser* parser is used. *parser* must be a subclass of *XMLParser* and can only use the default *TreeBuilder* as a target. Returns an *iterator* providing (event, elem) pairs.

Note that while *iterparse()* builds the tree incrementally, it issues blocking reads on *source* (or the file it names). As such, it's unsuitable for applications where blocking reads can't be made. For fully non-blocking parsing, see *XMLPullParser*.

注解: *iterparse()* only guarantees that it has seen the ">" character of a starting tag when it emits a "start" event, so the attributes are defined, but the contents of the text and tail attributes are undefined at that point. The same applies to the element children; they may or may not be present.

If you need a fully populated element, look for "end" events instead.

3.4 版后已移除: The *parser* argument.

`xml.etree.ElementTree.parse(source, parser=None)`

Parses an XML section into an element tree. *source* is a filename or file object containing XML data. *parser* is an optional parser instance. If not given, the standard *XMLParser* parser is used. Returns an *ElementTree* instance.

`xml.etree.ElementTree.ProcessingInstruction(target, text=None)`

PI element factory. This factory function creates a special element that will be serialized as an XML processing instruction. *target* is a string containing the PI target. *text* is a string containing the PI contents, if given. Returns an element instance, representing a processing instruction.

Note that *XMLParser* skips over processing instructions in the input instead of creating comment objects for them. An *ElementTree* will only contain processing instruction nodes if they have been inserted into the tree using one of the *Element* methods.

`xml.etree.ElementTree.register_namespace(prefix, uri)`

Registers a namespace prefix. The registry is global, and any existing mapping for either the given prefix or the namespace URI will be removed. *prefix* is a namespace prefix. *uri* is a namespace uri. Tags and attributes in this namespace will be serialized with the given prefix, if at all possible.

3.2 新版功能.

`xml.etree.ElementTree.SubElement(parent, tag, attrib={}, **extra)`

Subelement factory. This function creates an element instance, and appends it to an existing element.

The element name, attribute names, and attribute values can be either bytestrings or Unicode strings. *parent* is the parent element. *tag* is the subelement name. *attrib* is an optional dictionary, containing element attributes. *extra* contains additional attributes, given as keyword arguments. Returns an element instance.

`xml.etree.ElementTree.tostring(element, encoding="us-ascii", method="xml", *, short_empty_elements=True)`

Generates a string representation of an XML element, including all subelements. *element* is an *Element* instance. *encoding*¹ is the output encoding (default is US-ASCII). Use *encoding="unicode"* to generate a Unicode string (otherwise, a bytestring is generated). *method* is either "xml", "html" or "text" (default is "xml"). *short_empty_elements* has the same meaning as in *ElementTree.write()*. Returns an (optionally) encoded string containing the XML data.

3.4 新版功能: The *short_empty_elements* parameter.

`xml.etree.ElementTree.tostringlist(element, encoding="us-ascii", method="xml", *, short_empty_elements=True)`

Generates a string representation of an XML element, including all subelements. *element* is an *Element* instance. *encoding*¹ is the output encoding (default is US-ASCII). Use *encoding="unicode"* to generate a Unicode string (otherwise, a bytestring is generated). *method* is either "xml", "html" or "text" (default is "xml"). *short_empty_elements* has the same meaning as in *ElementTree.write()*. Returns a list of (optionally) encoded strings containing the XML data. It does not guarantee any specific sequence, except that `b"".join(tostringlist(element)) == tostring(element)`.

3.2 新版功能.

3.4 新版功能: The *short_empty_elements* parameter.

`xml.etree.ElementTree.XML(text, parser=None)`

Parses an XML section from a string constant. This function can be used to embed “XML literals” in Python code. *text* is a string containing XML data. *parser* is an optional parser instance. If not given, the standard *XMLParser* parser is used. Returns an *Element* instance.

`xml.etree.ElementTree.XMLID(text, parser=None)`

Parses an XML section from a string constant, and also returns a dictionary which maps from element id:s to

¹ The encoding string included in XML output should conform to the appropriate standards. For example, “UTF-8” is valid, but “UTF8” is not. See <https://www.w3.org/TR/2006/REC-xml11-20060816/#NT-EncodingDecl> and <https://www.iana.org/assignments/character-sets/character-sets.xhtml>.

elements. *text* is a string containing XML data. *parser* is an optional parser instance. If not given, the standard `XMLParser` parser is used. Returns a tuple containing an `Element` instance and a dictionary.

Element Objects

class `xml.etree.ElementTree.Element` (*tag*, *attrib*={}, ****extra**)

Element class. This class defines the Element interface, and provides a reference implementation of this interface.

The element name, attribute names, and attribute values can be either bytestrings or Unicode strings. *tag* is the element name. *attrib* is an optional dictionary, containing element attributes. *extra* contains additional attributes, given as keyword arguments.

tag

一个标识此元素意味着何种数据的字符串 (换句话说, 元素类型)。

text

tail

These attributes can be used to hold additional data associated with the element. Their values are usually strings but may be any application-specific object. If the element is created from an XML file, the *text* attribute holds either the text between the element's start tag and its first child or end tag, or `None`, and the *tail* attribute holds either the text between the element's end tag and the next tag, or `None`. For the XML data

```
<a><b>1<c>2<d/>3</c></b>4</a>
```

the *a* element has `None` for both *text* and *tail* attributes, the *b* element has *text* "1" and *tail* "4", the *c* element has *text* "2" and *tail* `None`, and the *d* element has *text* `None` and *tail* "3".

To collect the inner text of an element, see `itertext()`, for example `"".join(element.itertext())`.

Applications may store arbitrary objects in these attributes.

attrib

A dictionary containing the element's attributes. Note that while the *attrib* value is always a real mutable Python dictionary, an `ElementTree` implementation may choose to use another internal representation, and create the dictionary only if someone asks for it. To take advantage of such implementations, use the dictionary methods below whenever possible.

The following dictionary-like methods work on the element attributes.

clear()

Resets an element. This function removes all subelements, clears all attributes, and sets the text and tail attributes to `None`.

get (*key*, *default*=`None`)

Gets the element attribute named *key*.

Returns the attribute value, or *default* if the attribute was not found.

items()

Returns the element attributes as a sequence of (name, value) pairs. The attributes are returned in an arbitrary order.

keys()

Returns the elements attribute names as a list. The names are returned in an arbitrary order.

set (*key*, *value*)

Set the attribute *key* on the element to *value*.

The following methods work on the element's children (subelements).

append (*subelement*)

Adds the element *subelement* to the end of this element's internal list of subelements. Raises *TypeError* if *subelement* is not an *Element*.

extend (*subelements*)

Appends *subelements* from a sequence object with zero or more elements. Raises *TypeError* if a subelement is not an *Element*.

3.2 新版功能.

find (*match*, *namespaces=None*)

Finds the first subelement matching *match*. *match* may be a tag name or a *path*. Returns an element instance or *None*. *namespaces* is an optional mapping from namespace prefix to full name.

findall (*match*, *namespaces=None*)

Finds all matching subelements, by tag name or *path*. Returns a list containing all matching elements in document order. *namespaces* is an optional mapping from namespace prefix to full name.

findtext (*match*, *default=None*, *namespaces=None*)

Finds text for the first subelement matching *match*. *match* may be a tag name or a *path*. Returns the text content of the first matching element, or *default* if no element was found. Note that if the matching element has no text content an empty string is returned. *namespaces* is an optional mapping from namespace prefix to full name.

getchildren ()

3.2 版后已移除: Use `list(elem)` or iteration.

getiterator (*tag=None*)

3.2 版后已移除: Use method *Element.iter*() instead.

insert (*index*, *subelement*)

Inserts *subelement* at the given position in this element. Raises *TypeError* if *subelement* is not an *Element*.

iter (*tag=None*)

Creates a tree *iterator* with the current element as the root. The iterator iterates over this element and all elements below it, in document (depth first) order. If *tag* is not *None* or *'*'*, only elements whose tag equals *tag* are returned from the iterator. If the tree structure is modified during iteration, the result is undefined.

3.2 新版功能.

iterfind (*match*, *namespaces=None*)

Finds all matching subelements, by tag name or *path*. Returns an iterable yielding all matching elements in document order. *namespaces* is an optional mapping from namespace prefix to full name.

3.2 新版功能.

itertext ()

Creates a text iterator. The iterator loops over this element and all subelements, in document order, and returns all inner text.

3.2 新版功能.

makeelement (*tag*, *attrib*)

Creates a new element object of the same type as this element. Do not call this method, use the *SubElement* () factory function instead.

remove (*subelement*)

Removes *subelement* from the element. Unlike the *find** methods this method compares elements based on the instance identity, not on tag value or contents.

Element objects also support the following sequence type methods for working with subelements: `__delitem__()`, `__getitem__()`, `__setitem__()`, `__len__()`.

Caution: Elements with no subelements will test as `False`. This behavior will change in future versions. Use specific `len(elem)` or `elem is None` test instead.

```
element = root.find('foo')

if not element: # careful!
    print("element not found, or element has no subelements")

if element is None:
    print("element not found")
```

ElementTree Objects

class `xml.etree.ElementTree.ElementTree` (*element=None, file=None*)

ElementTree wrapper class. This class represents an entire element hierarchy, and adds some extra support for serialization to and from standard XML.

element is the root element. The tree is initialized with the contents of the XML *file* if given.

_setroot (*element*)

Replaces the root element for this tree. This discards the current contents of the tree, and replaces it with the given element. Use with care. *element* is an element instance.

find (*match, namespaces=None*)

Same as `Element.find()`, starting at the root of the tree.

findall (*match, namespaces=None*)

Same as `Element.findall()`, starting at the root of the tree.

findtext (*match, default=None, namespaces=None*)

Same as `Element.findtext()`, starting at the root of the tree.

getiterator (*tag=None*)

3.2 版后已移除: Use method `ElementTree.iter()` instead.

getroot ()

Returns the root element for this tree.

iter (*tag=None*)

Creates and returns a tree iterator for the root element. The iterator loops over all elements in this tree, in section order. *tag* is the tag to look for (default is to return all elements).

iterfind (*match, namespaces=None*)

Same as `Element.iterfind()`, starting at the root of the tree.

3.2 新版功能.

parse (*source, parser=None*)

Loads an external XML section into this element tree. *source* is a file name or *file object*. *parser* is an optional parser instance. If not given, the standard `XMLParser` parser is used. Returns the section root element.

write (*file, encoding="us-ascii", xml_declaration=None, default_namespace=None, method="xml", *, short_empty_elements=True*)

Writes the element tree to a file, as XML. *file* is a file name, or a *file object* opened for writing. *encoding*¹ is the output encoding (default is US-ASCII). *xml_declaration* controls if an XML declaration should be added to the file. Use `False` for never, `True` for always, `None` for only if not US-ASCII or UTF-8 or Unicode (default is `None`). *default_namespace* sets the default XML namespace (for “xmlns”). *method* is either “xml”,

"html" or "text" (default is "xml"). The keyword-only *short_empty_elements* parameter controls the formatting of elements that contain no content. If `True` (the default), they are emitted as a single self-closed tag, otherwise they are emitted as a pair of start/end tags.

The output is either a string (*str*) or binary (*bytes*). This is controlled by the *encoding* argument. If *encoding* is "unicode", the output is a string; otherwise, it's binary. Note that this may conflict with the type of *file* if it's an open *file object*; make sure you do not try to write a string to a binary stream and vice versa.

3.4 新版功能: The *short_empty_elements* parameter.

This is the XML file that is going to be manipulated:

```
<html>
  <head>
    <title>Example page</title>
  </head>
  <body>
    <p>Moved to <a href="http://example.org/">example.org</a>
    or <a href="http://example.com/">example.com</a>.</p>
  </body>
</html>
```

Example of changing the attribute "target" of every link in first paragraph:

```
>>> from xml.etree.ElementTree import ElementTree
>>> tree = ElementTree()
>>> tree.parse("index.xhtml")
<Element 'html' at 0xb77e6fac>
>>> p = tree.find("body/p")      # Finds first occurrence of tag p in body
>>> p
<Element 'p' at 0xb77ec26c>
>>> links = list(p.iter("a"))    # Returns list of all links
>>> links
[<Element 'a' at 0xb77ec2ac>, <Element 'a' at 0xb77ec1cc>]
>>> for i in links:              # Iterates through all found links
...     i.attrib["target"] = "blank"
>>> tree.write("output.xhtml")
```

QName Objects

class `xml.etree.ElementTree.QName` (*text_or_uri*, *tag=None*)

QName wrapper. This can be used to wrap a QName attribute value, in order to get proper namespace handling on output. *text_or_uri* is a string containing the QName value, in the form {uri}local, or, if the *tag* argument is given, the URI part of a QName. If *tag* is given, the first argument is interpreted as a URI, and this argument is interpreted as a local name. *QName* instances are opaque.

TreeBuilder Objects

class `xml.etree.ElementTree.TreeBuilder` (*element_factory=None*)

Generic element structure builder. This builder converts a sequence of start, data, and end method calls to a well-formed element structure. You can use this class to build an element structure using a custom XML parser, or a parser for some other XML-like format. *element_factory*, when given, must be a callable accepting two positional arguments: a tag and a dict of attributes. It is expected to return a new element instance.

close()

Flushes the builder buffers, and returns the toplevel document element. Returns an *Element* instance.

data(data)

Adds text to the current element. *data* is a string. This should be either a bytestring, or a Unicode string.

end(tag)

Closes the current element. *tag* is the element name. Returns the closed element.

start(tag, attrs)

Opens a new element. *tag* is the element name. *attrs* is a dictionary containing element attributes. Returns the opened element.

In addition, a custom *TreeBuilder* object can provide the following method:

doctype(name, pubid, system)

Handles a doctype declaration. *name* is the doctype name. *pubid* is the public identifier. *system* is the system identifier. This method does not exist on the default *TreeBuilder* class.

3.2 新版功能.

XMLParser Objects

class xml.etree.ElementTree.XMLParser(html=0, target=None, encoding=None)

This class is the low-level building block of the module. It uses *xml.parsers.expat* for efficient, event-based parsing of XML. It can be fed XML data incrementally with the *feed()* method, and parsing events are translated to a push API - by invoking callbacks on the *target* object. If *target* is omitted, the standard *TreeBuilder* is used. The *html* argument was historically used for backwards compatibility and is now deprecated. If *encoding*¹ is given, the value overrides the encoding specified in the XML file.

3.4 版后已移除: The *html* argument. The remaining arguments should be passed via keyword to prepare for the removal of the *html* argument.

close()

Finishes feeding data to the parser. Returns the result of calling the *close()* method of the *target* passed during construction; by default, this is the toplevel document element.

doctype(name, pubid, system)

3.2 版后已移除: Define the *TreeBuilder.doctype()* method on a custom *TreeBuilder* target.

feed(data)

Feeds data to the parser. *data* is encoded data.

XMLParser.feed() calls *target*'s *start(tag, attrs_dict)* method for each opening tag, its *end(tag)* method for each closing tag, and data is processed by method *data(data)*. *XMLParser.close()* calls *target*'s method *close()*. *XMLParser* can be used not only for building a tree structure. This is an example of counting the maximum depth of an XML file:

```
>>> from xml.etree.ElementTree import XMLParser
>>> class MaxDepth:                                # The target object of the parser
...     maxDepth = 0
...     depth = 0
...     def start(self, tag, attrib):               # Called for each opening tag.
...         self.depth += 1
...         if self.depth > self.maxDepth:
...             self.maxDepth = self.depth
...     def end(self, tag):                           # Called for each closing tag.
...         self.depth -= 1
...     def data(self, data):
```

(下页继续)

(续上页)

```

...     pass # We do not need to do anything with data.
...     def close(self): # Called when all data has been parsed.
...         return self.maxDepth
...
>>> target = MaxDepth()
>>> parser = XMLParser(target=target)
>>> exampleXml = """
... <a>
...     <b>
...     </b>
...     <b>
...         <c>
...         <d>
...         </d>
...     </c>
... </b>
... </a>"""
>>> parser.feed(exampleXml)
>>> parser.close()
4

```

XMLPullParser Objects

class xml.etree.ElementTree.XMLPullParser(*events=None*)

A pull parser suitable for non-blocking applications. Its input-side API is similar to that of *XMLParser*, but instead of pushing calls to a callback target, *XMLPullParser* collects an internal list of parsing events and lets the user read from it. *events* is a sequence of events to report back. The supported events are the strings "start", "end", "start-ns" and "end-ns" (the "ns" events are used to get detailed namespace information). If *events* is omitted, only "end" events are reported.

feed(*data*)

Feed the given bytes data to the parser.

close()

Signal the parser that the data stream is terminated. Unlike *XMLParser.close()*, this method always returns *None*. Any events not yet retrieved when the parser is closed can still be read with *read_events()*.

read_events()

Return an iterator over the events which have been encountered in the data fed to the parser. The iterator yields (*event*, *elem*) pairs, where *event* is a string representing the type of event (e.g. "end") and *elem* is the encountered *Element* object.

Events provided in a previous call to *read_events()* will not be yielded again. Events are consumed from the internal queue only when they are retrieved from the iterator, so multiple readers iterating in parallel over iterators obtained from *read_events()* will have unpredictable results.

注解: *XMLPullParser* only guarantees that it has seen the ">" character of a starting tag when it emits a "start" event, so the attributes are defined, but the contents of the text and tail attributes are undefined at that point. The same applies to the element children; they may or may not be present.

If you need a fully populated element, look for "end" events instead.

3.4 新版功能.

异常**class** `xml.etree.ElementTree.ParseError`

XML parse error, raised by the various parsing methods in this module when parsing fails. The string representation of an instance of this exception will contain a user-friendly error message. In addition, it will have the following attributes available:

code

A numeric error code from the expat parser. See the documentation of `xml.parsers.expat` for the list of error codes and their meanings.

position

A tuple of *line*, *column* numbers, specifying where the error occurred.

20.6 `xml.dom` —The Document Object Model API

Source code: `Lib/xml/dom/__init__.py`

The Document Object Model, or “DOM,” is a cross-language API from the World Wide Web Consortium (W3C) for accessing and modifying XML documents. A DOM implementation presents an XML document as a tree structure, or allows client code to build such a structure from scratch. It then gives access to the structure through a set of objects which provided well-known interfaces.

The DOM is extremely useful for random-access applications. SAX only allows you a view of one bit of the document at a time. If you are looking at one SAX element, you have no access to another. If you are looking at a text node, you have no access to a containing element. When you write a SAX application, you need to keep track of your program’s position in the document somewhere in your own code. SAX does not do it for you. Also, if you need to look ahead in the XML document, you are just out of luck.

Some applications are simply impossible in an event driven model with no access to a tree. Of course you could build some sort of tree yourself in SAX events, but the DOM allows you to avoid writing that code. The DOM is a standard tree representation for XML data.

The Document Object Model is being defined by the W3C in stages, or “levels” in their terminology. The Python mapping of the API is substantially based on the DOM Level 2 recommendation.

DOM applications typically start by parsing some XML into a DOM. How this is accomplished is not covered at all by DOM Level 1, and Level 2 provides only limited improvements: There is a `DOMImplementation` object class which provides access to `Document` creation methods, but no way to access an XML reader/parser/Document builder in an implementation-independent way. There is also no well-defined way to access these methods without an existing `Document` object. In Python, each DOM implementation will provide a function `getDOMImplementation()`. DOM Level 3 adds a Load/Store specification, which defines an interface to the reader, but this is not yet available in the Python standard library.

Once you have a DOM document object, you can access the parts of your XML document through its properties and methods. These properties are defined in the DOM specification; this portion of the reference manual describes the interpretation of the specification in Python.

The specification provided by the W3C defines the DOM API for Java, ECMAScript, and OMG IDL. The Python mapping defined here is based in large part on the IDL version of the specification, but strict compliance is not required (though implementations are free to support the strict mapping from IDL). See section [Conformance](#) for a detailed discussion of mapping requirements.

参见:

Document Object Model (DOM) Level 2 Specification The W3C recommendation upon which the Python DOM API is based.

Document Object Model (DOM) Level 1 Specification The W3C recommendation for the DOM supported by `xml.dom.minidom`.

Python Language Mapping Specification This specifies the mapping from OMG IDL to Python.

20.6.1 模块内容

The `xml.dom` contains the following functions:

`xml.dom.registerDOMImplementation(name, factory)`

Register the *factory* function with the name *name*. The factory function should return an object which implements the `DOMImplementation` interface. The factory function can return the same object every time, or a new one for each call, as appropriate for the specific implementation (e.g. if that implementation supports some customization).

`xml.dom.getDOMImplementation(name=None, features=())`

Return a suitable DOM implementation. The *name* is either well-known, the module name of a DOM implementation, or `None`. If it is not `None`, imports the corresponding module and returns a `DOMImplementation` object if the import succeeds. If no name is given, and if the environment variable `PYTHON_DOM` is set, this variable is used to find the implementation.

If name is not given, this examines the available implementations to find one with the required feature set. If no implementation can be found, raise an `ImportError`. The features list must be a sequence of (*feature*, *version*) pairs which are passed to the `hasFeature()` method on available `DOMImplementation` objects.

Some convenience constants are also provided:

`xml.dom.EMPTY_NAMESPACE`

The value used to indicate that no namespace is associated with a node in the DOM. This is typically found as the *namespaceURI* of a node, or used as the *namespaceURI* parameter to a namespaces-specific method.

`xml.dom.XML_NAMESPACE`

The namespace URI associated with the reserved prefix `xml`, as defined by [Namespaces in XML](#) (section 4).

`xml.dom.XMLNS_NAMESPACE`

The namespace URI for namespace declarations, as defined by [Document Object Model \(DOM\) Level 2 Core Specification](#) (section 1.1.8).

`xml.dom.XHTML_NAMESPACE`

The URI of the XHTML namespace as defined by [XHTML 1.0: The Extensible HyperText Markup Language](#) (section 3.1.1).

In addition, `xml.dom` contains a base `Node` class and the DOM exception classes. The `Node` class provided by this module does not implement any of the methods or attributes defined by the DOM specification; concrete DOM implementations must provide those. The `Node` class provided as part of this module does provide the constants used for the *nodeType* attribute on concrete `Node` objects; they are located within the class rather than at the module level to conform with the DOM specifications.

20.6.2 Objects in the DOM

The definitive documentation for the DOM is the DOM specification from the W3C.

Note that DOM attributes may also be manipulated as nodes instead of as simple strings. It is fairly rare that you must do this, however, so this usage is not yet documented.

Interface	Section	目的
DOMImplementation	<i>DOMImplementation Objects</i>	Interface to the underlying implementation.
Node	<i>Node Objects</i>	Base interface for most objects in a document.
NodeList	<i>NodeList Objects</i>	Interface for a sequence of nodes.
DocumentType	<i>DocumentType Objects</i>	Information about the declarations needed to process a document.
Document	<i>Document Objects</i>	Object which represents an entire document.
Element	<i>Element Objects</i>	Element nodes in the document hierarchy.
Attr	<i>Attr Objects</i>	Attribute value nodes on element nodes.
Comment	<i>Comment Objects</i>	Representation of comments in the source document.
Text	<i>Text and CDATASection Objects</i>	Nodes containing textual content from the document.
ProcessingInstruction	<i>ProcessingInstruction Objects</i>	Processing instruction representation.

An additional section describes the exceptions defined for working with the DOM in Python.

DOMImplementation Objects

The `DOMImplementation` interface provides a way for applications to determine the availability of particular features in the DOM they are using. DOM Level 2 added the ability to create new `Document` and `DocumentType` objects using the `DOMImplementation` as well.

`DOMImplementation.hasFeature` (*feature*, *version*)

Return true if the feature identified by the pair of strings *feature* and *version* is implemented.

`DOMImplementation.createDocument` (*namespaceUri*, *qualifiedName*, *doctype*)

Return a new `Document` object (the root of the DOM), with a child `Element` object having the given *namespaceUri* and *qualifiedName*. The *doctype* must be a `DocumentType` object created by `createDocumentType()`, or `None`. In the Python DOM API, the first two arguments can also be `None` in order to indicate that no `Element` child is to be created.

`DOMImplementation.createDocumentType` (*qualifiedName*, *publicId*, *systemId*)

Return a new `DocumentType` object that encapsulates the given *qualifiedName*, *publicId*, and *systemId* strings, representing the information contained in an XML document type declaration.

Node Objects

All of the components of an XML document are subclasses of `Node`.

`Node.nodeType`

An integer representing the node type. Symbolic constants for the types are on the `Node` object: `ELEMENT_NODE`, `ATTRIBUTE_NODE`, `TEXT_NODE`, `CDATA_SECTION_NODE`, `ENTITY_NODE`, `PROCESSING_INSTRUCTION_NODE`, `COMMENT_NODE`, `DOCUMENT_NODE`, `DOCUMENT_TYPE_NODE`, `NOTATION_NODE`. This is a read-only attribute.

`Node.parentNode`

The parent of the current node, or `None` for the document node. The value is always a `Node` object or `None`. For `Element` nodes, this will be the parent element, except for the root element, in which case it will be the `Document` object. For `Attr` nodes, this is always `None`. This is a read-only attribute.

Node.attributes

A `NamedNodeMap` of attribute objects. Only elements have actual values for this; others provide `None` for this attribute. This is a read-only attribute.

Node.previousSibling

The node that immediately precedes this one with the same parent. For instance the element with an end-tag that comes just before the *self* element's start-tag. Of course, XML documents are made up of more than just elements so the previous sibling could be text, a comment, or something else. If this node is the first child of the parent, this attribute will be `None`. This is a read-only attribute.

Node.nextSibling

The node that immediately follows this one with the same parent. See also *previousSibling*. If this is the last child of the parent, this attribute will be `None`. This is a read-only attribute.

Node.childNodes

A list of nodes contained within this node. This is a read-only attribute.

Node.firstChild

The first child of the node, if there are any, or `None`. This is a read-only attribute.

Node.lastChild

The last child of the node, if there are any, or `None`. This is a read-only attribute.

Node.localName

The part of the `tagName` following the colon if there is one, else the entire `tagName`. The value is a string.

Node.prefix

The part of the `tagName` preceding the colon if there is one, else the empty string. The value is a string, or `None`.

Node.namespaceURI

The namespace associated with the element name. This will be a string or `None`. This is a read-only attribute.

Node.nodeName

This has a different meaning for each node type; see the DOM specification for details. You can always get the information you would get here from another property such as the `tagName` property for elements or the `name` property for attributes. For all node types, the value of this attribute will be either a string or `None`. This is a read-only attribute.

Node.nodeValue

This has a different meaning for each node type; see the DOM specification for details. The situation is similar to that with *nodeName*. The value is a string or `None`.

Node.hasAttributes()

Returns true if the node has any attributes.

Node.hasChildNodes()

Returns true if the node has any child nodes.

Node.isSameNode(*other*)

Returns true if *other* refers to the same node as this node. This is especially useful for DOM implementations which use any sort of proxy architecture (because more than one object can refer to the same node).

注解: This is based on a proposed DOM Level 3 API which is still in the “working draft” stage, but this particular interface appears uncontroversial. Changes from the W3C will not necessarily affect this method in the Python DOM interface (though any new W3C API for this would also be supported).

Node.appendChild(*newChild*)

Add a new child node to this node at the end of the list of children, returning *newChild*. If the node was already in the tree, it is removed first.

`Node.insertBefore(newChild, refChild)`

Insert a new child node before an existing child. It must be the case that *refChild* is a child of this node; if not, *ValueError* is raised. *newChild* is returned. If *refChild* is `None`, it inserts *newChild* at the end of the children's list.

`Node.removeChild(oldChild)`

Remove a child node. *oldChild* must be a child of this node; if not, *ValueError* is raised. *oldChild* is returned on success. If *oldChild* will not be used further, its `unlink()` method should be called.

`Node.replaceChild(newChild, oldChild)`

Replace an existing node with a new node. It must be the case that *oldChild* is a child of this node; if not, *ValueError* is raised.

`Node.normalize()`

Join adjacent text nodes so that all stretches of text are stored as single `Text` instances. This simplifies processing text from a DOM tree for many applications.

`Node.cloneNode(deep)`

Clone this node. Setting *deep* means to clone all child nodes as well. This returns the clone.

NodeList Objects

A `NodeList` represents a sequence of nodes. These objects are used in two ways in the DOM Core recommendation: an `Element` object provides one as its list of child nodes, and the `getElementsByTagName()` and `getElementsByTagNameNS()` methods of `Node` return objects with this interface to represent query results.

The DOM Level 2 recommendation defines one method and one attribute for these objects:

`NodeList.item(i)`

Return the *i*'th item from the sequence, if there is one, or `None`. The index *i* is not allowed to be less than zero or greater than or equal to the length of the sequence.

`NodeList.length`

The number of nodes in the sequence.

In addition, the Python DOM interface requires that some additional support is provided to allow `NodeList` objects to be used as Python sequences. All `NodeList` implementations must include support for `__len__()` and `__getitem__()`; this allows iteration over the `NodeList` in `for` statements and proper support for the `len()` built-in function.

If a DOM implementation supports modification of the document, the `NodeList` implementation must also support the `__setitem__()` and `__delitem__()` methods.

DocumentType Objects

Information about the notations and entities declared by a document (including the external subset if the parser uses it and can provide the information) is available from a `DocumentType` object. The `DocumentType` for a document is available from the `Document` object's `doctype` attribute; if there is no `DOCTYPE` declaration for the document, the document's `doctype` attribute will be set to `None` instead of an instance of this interface.

`DocumentType` is a specialization of `Node`, and adds the following attributes:

`DocumentType.publicId`

The public identifier for the external subset of the document type definition. This will be a string or `None`.

`DocumentType.systemId`

The system identifier for the external subset of the document type definition. This will be a URI as a string, or `None`.

`DocumentType.internalSubset`

A string giving the complete internal subset from the document. This does not include the brackets which enclose the subset. If the document has no internal subset, this should be `None`.

`DocumentType.name`

The name of the root element as given in the DOCTYPE declaration, if present.

`DocumentType.entities`

This is a `NamedNodeMap` giving the definitions of external entities. For entity names defined more than once, only the first definition is provided (others are ignored as required by the XML recommendation). This may be `None` if the information is not provided by the parser, or if no entities are defined.

`DocumentType.notations`

This is a `NamedNodeMap` giving the definitions of notations. For notation names defined more than once, only the first definition is provided (others are ignored as required by the XML recommendation). This may be `None` if the information is not provided by the parser, or if no notations are defined.

Document Objects

A `Document` represents an entire XML document, including its constituent elements, attributes, processing instructions, comments etc. Remember that it inherits properties from `Node`.

`Document.documentElement`

The one and only root element of the document.

`Document.createElement(tagName)`

Create and return a new element node. The element is not inserted into the document when it is created. You need to explicitly insert it with one of the other methods such as `insertBefore()` or `appendChild()`.

`Document.createElementNS(namespaceURI, tagName)`

Create and return a new element with a namespace. The *tagName* may have a prefix. The element is not inserted into the document when it is created. You need to explicitly insert it with one of the other methods such as `insertBefore()` or `appendChild()`.

`Document.createTextNode(data)`

Create and return a text node containing the data passed as a parameter. As with the other creation methods, this one does not insert the node into the tree.

`Document.createComment(data)`

Create and return a comment node containing the data passed as a parameter. As with the other creation methods, this one does not insert the node into the tree.

`Document.createProcessingInstruction(target, data)`

Create and return a processing instruction node containing the *target* and *data* passed as parameters. As with the other creation methods, this one does not insert the node into the tree.

`Document.createAttribute(name)`

Create and return an attribute node. This method does not associate the attribute node with any particular element. You must use `setAttributeNode()` on the appropriate `Element` object to use the newly created attribute instance.

`Document.createAttributeNS(namespaceURI, qualifiedName)`

Create and return an attribute node with a namespace. The *tagName* may have a prefix. This method does not associate the attribute node with any particular element. You must use `setAttributeNode()` on the appropriate `Element` object to use the newly created attribute instance.

`Document.getElementsByTagName(tagName)`

Search for all descendants (direct children, children's children, etc.) with a particular element type name.

`Document.getElementsByTagNameNS(namespaceURI, localName)`

Search for all descendants (direct children, children's children, etc.) with a particular namespace URI and local-name. The localname is the part of the namespace after the prefix.

Element Objects

`Element` is a subclass of `Node`, so inherits all the attributes of that class.

`Element.tagName`

The element type name. In a namespace-using document it may have colons in it. The value is a string.

`Element.getElementsByTagName(tagName)`

Same as equivalent method in the `Document` class.

`Element.getElementsByTagNameNS(namespaceURI, localName)`

Same as equivalent method in the `Document` class.

`Element.hasAttribute(name)`

Returns true if the element has an attribute named by *name*.

`Element.hasAttributeNS(namespaceURI, localName)`

Returns true if the element has an attribute named by *namespaceURI* and *localName*.

`Element.getAttribute(name)`

Return the value of the attribute named by *name* as a string. If no such attribute exists, an empty string is returned, as if the attribute had no value.

`Element.getAttributeNode(attrname)`

Return the `Attr` node for the attribute named by *attrname*.

`Element.getAttributeNS(namespaceURI, localName)`

Return the value of the attribute named by *namespaceURI* and *localName* as a string. If no such attribute exists, an empty string is returned, as if the attribute had no value.

`Element.getAttributeNodeNS(namespaceURI, localName)`

Return an attribute value as a node, given a *namespaceURI* and *localName*.

`Element.removeAttribute(name)`

Remove an attribute by name. If there is no matching attribute, a `NotFoundError` is raised.

`Element.removeAttributeNode(oldAttr)`

Remove and return *oldAttr* from the attribute list, if present. If *oldAttr* is not present, `NotFoundError` is raised.

`Element.removeAttributeNS(namespaceURI, localName)`

Remove an attribute by name. Note that it uses a *localName*, not a *qname*. No exception is raised if there is no matching attribute.

`Element.setAttribute(name, value)`

Set an attribute value from a string.

`Element.setAttributeNode(newAttr)`

Add a new attribute node to the element, replacing an existing attribute if necessary if the *name* attribute matches. If a replacement occurs, the old attribute node will be returned. If *newAttr* is already in use, `InuseAttributeErr` will be raised.

`Element.setAttributeNodeNS(newAttr)`

Add a new attribute node to the element, replacing an existing attribute if necessary if the *namespaceURI* and *localName* attributes match. If a replacement occurs, the old attribute node will be returned. If *newAttr* is already in use, `InuseAttributeErr` will be raised.

`Element.setAttributeNS(namespaceURI, qname, value)`

Set an attribute value from a string, given a *namespaceURI* and a *qname*. Note that a *qname* is the whole attribute name. This is different than above.

Attr Objects

`Attr` inherits from `Node`, so inherits all its attributes.

`Attr.name`

The attribute name. In a namespace-using document it may include a colon.

`Attr.localName`

The part of the name following the colon if there is one, else the entire name. This is a read-only attribute.

`Attr.prefix`

The part of the name preceding the colon if there is one, else the empty string.

`Attr.value`

The text value of the attribute. This is a synonym for the `nodeValue` attribute.

NamedNodeMap Objects

`NamedNodeMap` does *not* inherit from `Node`.

`NamedNodeMap.length`

The length of the attribute list.

`NamedNodeMap.item(index)`

Return an attribute with a particular index. The order you get the attributes in is arbitrary but will be consistent for the life of a DOM. Each item is an attribute node. Get its value with the `value` attribute.

There are also experimental methods that give this class more mapping behavior. You can use them or you can use the standardized `getAttribute*()` family of methods on the `Element` objects.

Comment Objects

`Comment` represents a comment in the XML document. It is a subclass of `Node`, but cannot have child nodes.

`Comment.data`

The content of the comment as a string. The attribute contains all characters between the leading `<!--` and trailing `-->`, but does not include them.

Text and CDATASection Objects

The `Text` interface represents text in the XML document. If the parser and DOM implementation support the DOM's XML extension, portions of the text enclosed in CDATA marked sections are stored in `CDATASection` objects. These two interfaces are identical, but provide different values for the `nodeType` attribute.

These interfaces extend the `Node` interface. They cannot have child nodes.

`Text.data`

The content of the text node as a string.

注解: The use of a `CDATASection` node does not indicate that the node represents a complete CDATA marked section, only that the content of the node was part of a CDATA section. A single CDATA section may be represented by

more than one node in the document tree. There is no way to determine whether two adjacent `CDATASection` nodes represent different CDATA marked sections.

ProcessingInstruction Objects

Represents a processing instruction in the XML document; this inherits from the `Node` interface and cannot have child nodes.

`ProcessingInstruction.target`

The content of the processing instruction up to the first whitespace character. This is a read-only attribute.

`ProcessingInstruction.data`

The content of the processing instruction following the first whitespace character.

异常

The DOM Level 2 recommendation defines a single exception, *DOMException*, and a number of constants that allow applications to determine what sort of error occurred. *DOMException* instances carry a *code* attribute that provides the appropriate value for the specific exception.

The Python DOM interface provides the constants, but also expands the set of exceptions so that a specific exception exists for each of the exception codes defined by the DOM. The implementations must raise the appropriate specific exception, each of which carries the appropriate value for the *code* attribute.

exception `xml.dom.DOMException`

Base exception class used for all specific DOM exceptions. This exception class cannot be directly instantiated.

exception `xml.dom.DomstringSizeErr`

Raised when a specified range of text does not fit into a string. This is not known to be used in the Python DOM implementations, but may be received from DOM implementations not written in Python.

exception `xml.dom.HierarchyRequestErr`

Raised when an attempt is made to insert a node where the node type is not allowed.

exception `xml.dom.IndexSizeErr`

Raised when an index or size parameter to a method is negative or exceeds the allowed values.

exception `xml.dom.InuseAttributeErr`

Raised when an attempt is made to insert an `Attr` node that is already present elsewhere in the document.

exception `xml.dom.InvalidAccessErr`

Raised if a parameter or an operation is not supported on the underlying object.

exception `xml.dom.InvalidCharacterErr`

This exception is raised when a string parameter contains a character that is not permitted in the context it's being used in by the XML 1.0 recommendation. For example, attempting to create an `Element` node with a space in the element type name will cause this error to be raised.

exception `xml.dom.InvalidModificationErr`

Raised when an attempt is made to modify the type of a node.

exception `xml.dom.InvalidStateErr`

Raised when an attempt is made to use an object that is not defined or is no longer usable.

exception `xml.dom.NamespaceErr`

If an attempt is made to change any object in a way that is not permitted with regard to the [Namespaces in XML](#) recommendation, this exception is raised.

exception xml.dom.NotFoundErr

Exception when a node does not exist in the referenced context. For example, `NamedNodeMap.removeNamedItem()` will raise this if the node passed in does not exist in the map.

exception xml.dom.NotSupportedErr

Raised when the implementation does not support the requested type of object or operation.

exception xml.dom.NoDataAllowedErr

This is raised if data is specified for a node which does not support data.

exception xml.dom.NoModificationAllowedErr

Raised on attempts to modify an object where modifications are not allowed (such as for read-only nodes).

exception xml.dom.SyntaxErr

Raised when an invalid or illegal string is specified.

exception xml.dom.WrongDocumentErr

Raised when a node is inserted in a different document than it currently belongs to, and the implementation does not support migrating the node from one document to the other.

The exception codes defined in the DOM recommendation map to the exceptions described above according to this table:

常数	Exception
DOMSTRING_SIZE_ERR	<i>DomstringSizeErr</i>
HIERARCHY_REQUEST_ERR	<i>HierarchyRequestErr</i>
INDEX_SIZE_ERR	<i>IndexSizeErr</i>
INUSE_ATTRIBUTE_ERR	<i>InuseAttributeErr</i>
INVALID_ACCESS_ERR	<i>InvalidAccessErr</i>
INVALID_CHARACTER_ERR	<i>InvalidCharacterErr</i>
INVALID_MODIFICATION_ERR	<i>InvalidModificationErr</i>
INVALID_STATE_ERR	<i>InvalidStateErr</i>
NAMESPACE_ERR	<i>NamespaceErr</i>
NOT_FOUND_ERR	<i>NotFoundErr</i>
NOT_SUPPORTED_ERR	<i>NotSupportedErr</i>
NO_DATA_ALLOWED_ERR	<i>NoDataAllowedErr</i>
NO_MODIFICATION_ALLOWED_ERR	<i>NoModificationAllowedErr</i>
SYNTAX_ERR	<i>SyntaxErr</i>
WRONG_DOCUMENT_ERR	<i>WrongDocumentErr</i>

20.6.3 Conformance

This section describes the conformance requirements and relationships between the Python DOM API, the W3C DOM recommendations, and the OMG IDL mapping for Python.

Type Mapping

The IDL types used in the DOM specification are mapped to Python types according to the following table.

IDL Type	Python Type
boolean	bool or int
int	int
long int	int
unsigned int	int
DOMString	str or bytes
null	None

Accessor Methods

The mapping from OMG IDL to Python defines accessor functions for IDL attribute declarations in much the way the Java mapping does. Mapping the IDL declarations

```
readonly attribute string someValue;
    attribute string anotherValue;
```

yields three accessor functions: a “get” method for `someValue` (`_get_someValue()`), and “get” and “set” methods for `anotherValue` (`_get_anotherValue()` and `_set_anotherValue()`). The mapping, in particular, does not require that the IDL attributes are accessible as normal Python attributes: `object.someValue` is *not* required to work, and may raise an `AttributeError`.

The Python DOM API, however, *does* require that normal attribute access work. This means that the typical surrogates generated by Python IDL compilers are not likely to work, and wrapper objects may be needed on the client if the DOM objects are accessed via CORBA. While this does require some additional consideration for CORBA DOM clients, the implementers with experience using DOM over CORBA from Python do not consider this a problem. Attributes that are declared `readonly` may not restrict write access in all DOM implementations.

In the Python DOM API, accessor functions are not required. If provided, they should take the form defined by the Python IDL mapping, but these methods are considered unnecessary since the attributes are accessible directly from Python. “Set” accessors should never be provided for `readonly` attributes.

The IDL definitions do not fully embody the requirements of the W3C DOM API, such as the notion of certain objects, such as the return value of `getElementsByTagName()`, being “live”. The Python DOM API does not require implementations to enforce such requirements.

20.7 xml.dom.minidom — Minimal DOM implementation

Source code: [Lib/xml/dom/minidom.py](#)

`xml.dom.minidom` is a minimal implementation of the Document Object Model interface, with an API similar to that in other languages. It is intended to be simpler than the full DOM and also significantly smaller. Users who are not already proficient with the DOM should consider using the `xml.etree.ElementTree` module for their XML processing instead.

警告: The `xml.dom.minidom` module is not secure against maliciously constructed data. If you need to parse untrusted or unauthenticated data see [XML 漏洞](#).

DOM applications typically start by parsing some XML into a DOM. With `xml.dom.minidom`, this is done through the parse functions:

```

from xml.dom.minidom import parse, parseString

dom1 = parse('c:\\temp\\mydata.xml')  # parse an XML file by name

datasource = open('c:\\temp\\mydata.xml')
dom2 = parse(datasource)  # parse an open file

dom3 = parseString('<myxml>Some data<empty/> some more data</myxml>')

```

The `parse()` function can take either a filename or an open file object.

`xml.dom.minidom.parse(filename_or_file, parser=None, bufsize=None)`

Return a Document from the given input. *filename_or_file* may be either a file name, or a file-like object. *parser*, if given, must be a SAX2 parser object. This function will change the document handler of the parser and activate namespace support; other parser configuration (like setting an entity resolver) must have been done in advance.

If you have XML in a string, you can use the `parseString()` function instead:

`xml.dom.minidom.parseString(string, parser=None)`

Return a Document that represents the *string*. This method creates an `io.StringIO` object for the string and passes that on to `parse()`.

Both functions return a Document object representing the content of the document.

What the `parse()` and `parseString()` functions do is connect an XML parser with a “DOM builder” that can accept parse events from any SAX parser and convert them into a DOM tree. The name of the functions are perhaps misleading, but are easy to grasp when learning the interfaces. The parsing of the document will be completed before these functions return; it’s simply that these functions do not provide a parser implementation themselves.

You can also create a Document by calling a method on a “DOM Implementation” object. You can get this object either by calling the `getDOMImplementation()` function in the `xml.dom` package or the `xml.dom.minidom` module. Once you have a Document, you can add child nodes to it to populate the DOM:

```

from xml.dom.minidom import getDOMImplementation

impl = getDOMImplementation()

newdoc = impl.createDocument(None, "some_tag", None)
top_element = newdoc.documentElement
text = newdoc.createTextNode('Some textual content.')
top_element.appendChild(text)

```

Once you have a DOM document object, you can access the parts of your XML document through its properties and methods. These properties are defined in the DOM specification. The main property of the document object is the `documentElement` property. It gives you the main element in the XML document: the one that holds all others. Here is an example program:

```

dom3 = parseString("<myxml>Some data</myxml>")
assert dom3.documentElement.tagName == "myxml"

```

When you are finished with a DOM tree, you may optionally call the `unlink()` method to encourage early cleanup of the now-unneeded objects. `unlink()` is an `xml.dom.minidom`-specific extension to the DOM API that renders the node and its descendants are essentially useless. Otherwise, Python’s garbage collector will eventually take care of the objects in the tree.

参见:

Document Object Model (DOM) Level 1 Specification The W3C recommendation for the DOM supported by `xml.dom.minidom`.

20.7.1 DOM Objects

The definition of the DOM API for Python is given as part of the `xml.dom` module documentation. This section lists the differences between the API and `xml.dom.minidom`.

Node.**unlink**()

Break internal references within the DOM so that it will be garbage collected on versions of Python without cyclic GC. Even when cyclic GC is available, using this can make large amounts of memory available sooner, so calling this on DOM objects as soon as they are no longer needed is good practice. This only needs to be called on the `Document` object, but may be called on child nodes to discard children of that node.

You can avoid calling this method explicitly by using the `with` statement. The following code will automatically unlink `dom` when the `with` block is exited:

```
with xml.dom.minidom.parse(datasource) as dom:
    ... # Work with dom.
```

Node.**writexml**(*writer*, *indent=""*, *addindent=""*, *newl=""*)

Write XML to the writer object. The writer should have a `write()` method which matches that of the file object interface. The *indent* parameter is the indentation of the current node. The *addindent* parameter is the incremental indentation to use for subnodes of the current one. The *newl* parameter specifies the string to use to terminate newlines.

For the `Document` node, an additional keyword argument *encoding* can be used to specify the encoding field of the XML header.

Node.**toxml**(*encoding=None*)

Return a string or byte string containing the XML represented by the DOM node.

With an explicit *encoding*¹ argument, the result is a byte string in the specified encoding. With no *encoding* argument, the result is a Unicode string, and the XML declaration in the resulting string does not specify an encoding. Encoding this string in an encoding other than UTF-8 is likely incorrect, since UTF-8 is the default encoding of XML.

Node.**toprettyxml**(*indent="\t"*, *newl="\n"*, *encoding=None*)

Return a pretty-printed version of the document. *indent* specifies the indentation string and defaults to a tabulator; *newl* specifies the string emitted at the end of each line and defaults to `\n`.

The *encoding* argument behaves like the corresponding argument of `toxml()`.

20.7.2 DOM Example

This example program is a fairly realistic example of a simple program. In this particular case, we do not take much advantage of the flexibility of the DOM.

```
import xml.dom.minidom

document = """\
<slideshow>
<title>Demo slideshow</title>
<slide><title>Slide title</title>
<point>This is a demo</point>
<point>Of a program for processing slides</point>
</slide>
```

(下页继续)

¹ The encoding name included in the XML output should conform to the appropriate standards. For example, “UTF-8” is valid, but “UTF8” is not valid in an XML document’s declaration, even though Python accepts it as an encoding name. See <https://www.w3.org/TR/2006/REC-xml11-20060816/#NT-EncodingDecl> and <https://www.iana.org/assignments/character-sets/character-sets.xhtml>.

```

<slide><title>Another demo slide</title>
<point>It is important</point>
<point>To have more than</point>
<point>one slide</point>
</slide>
</slideshow>
"""

dom = xml.dom.minidom.parseString(document)

def getText(nodelist):
    rc = []
    for node in nodelist:
        if node.nodeType == node.TEXT_NODE:
            rc.append(node.data)
    return ''.join(rc)

def handleSlideshow(slideshow):
    print("<html>")
    handleSlideshowTitle(slideshow.getElementsByTagName("title")[0])
    slides = slideshow.getElementsByTagName("slide")
    handleToc(slides)
    handleSlides(slides)
    print("</html>")

def handleSlides(slides):
    for slide in slides:
        handleSlide(slide)

def handleSlide(slide):
    handleSlideTitle(slide.getElementsByTagName("title")[0])
    handlePoints(slide.getElementsByTagName("point"))

def handleSlideshowTitle(title):
    print("<title>%s</title>" % getText(title.childNodes))

def handleSlideTitle(title):
    print("<h2>%s</h2>" % getText(title.childNodes))

def handlePoints(points):
    print("<ul>")
    for point in points:
        handlePoint(point)
    print("</ul>")

def handlePoint(point):
    print("<li>%s</li>" % getText(point.childNodes))

def handleToc(slides):
    for slide in slides:
        title = slide.getElementsByTagName("title")[0]
        print("<p>%s</p>" % getText(title.childNodes))

handleSlideshow(dom)

```

20.7.3 minidom and the DOM standard

The `xml.dom.minidom` module is essentially a DOM 1.0-compatible DOM with some DOM 2 features (primarily namespace features).

Usage of the DOM interface in Python is straight-forward. The following mapping rules apply:

- Interfaces are accessed through instance objects. Applications should not instantiate the classes themselves; they should use the creator functions available on the `Document` object. Derived interfaces support all operations (and attributes) from the base interfaces, plus any new operations.
- Operations are used as methods. Since the DOM uses only `in` parameters, the arguments are passed in normal order (from left to right). There are no optional arguments. `void` operations return `None`.
- IDL attributes map to instance attributes. For compatibility with the OMG IDL language mapping for Python, an attribute `foo` can also be accessed through accessor methods `_get_foo()` and `_set_foo()`. `readonly` attributes must not be changed; this is not enforced at runtime.
- The types `short` `int`, `unsigned int`, `unsigned long` `long`, and `boolean` all map to Python integer objects.
- The type `DOMString` maps to Python strings. `xml.dom.minidom` supports either bytes or strings, but will normally produce strings. Values of type `DOMString` may also be `None` where allowed to have the IDL null value by the DOM specification from the W3C.
- `const` declarations map to variables in their respective scope (e.g. `xml.dom.minidom.Node.PROCESSING_INSTRUCTION_NODE`); they must not be changed.
- `DOMException` is currently not supported in `xml.dom.minidom`. Instead, `xml.dom.minidom` uses standard Python exceptions such as `TypeError` and `AttributeError`.
- `NodeList` objects are implemented using Python's built-in list type. These objects provide the interface defined in the DOM specification, but with earlier versions of Python they do not support the official API. They are, however, much more “Pythonic” than the interface defined in the W3C recommendations.

The following interfaces have no implementation in `xml.dom.minidom`:

- `DOMTimeStamp`
- `DocumentType`
- `DOMImplementation`
- `CharacterData`
- `CDATASection`
- `Notation`
- `Entity`
- `EntityReference`
- `DocumentFragment`

Most of these reflect information in the XML document that is not of general utility to most DOM users.

20.8 xml.dom.pulldom — Support for building partial DOM trees

Source code: `Lib/xml/dom/pulldom.py`

The `xml.dom.pulldom` module provides a “pull parser” which can also be asked to produce DOM-accessible fragments of the document where necessary. The basic concept involves pulling “events” from a stream of incoming XML and processing them. In contrast to SAX which also employs an event-driven processing model together with callbacks, the user of a pull parser is responsible for explicitly pulling events from the stream, looping over those events until either processing is finished or an error condition occurs.

警告: The `xml.dom.pulldom` module is not secure against maliciously constructed data. If you need to parse untrusted or unauthenticated data see [XML 漏洞](#).

在 3.6.7 版更改: The SAX parser no longer processes general external entities by default to increase security by default. To enable processing of external entities, pass a custom parser instance in:

```
from xml.dom.pulldom import parse
from xml.sax import make_parser
from xml.sax.handler import feature_external_ges

parser = make_parser()
parser.setFeature(feature_external_ges, True)
parse(filename, parser=parser)
```

示例:

```
from xml.dom import pulldom

doc = pulldom.parse('sales_items.xml')
for event, node in doc:
    if event == pulldom.START_ELEMENT and node.tagName == 'item':
        if int(node.getAttribute('price')) > 50:
            doc.expandNode(node)
            print(node.toxml())
```

event is a constant and can be one of:

- `START_ELEMENT`
- `END_ELEMENT`
- `COMMENT`
- `START_DOCUMENT`
- `END_DOCUMENT`
- `CHARACTERS`
- `PROCESSING_INSTRUCTION`
- `IGNORABLE_WHITESPACE`

node is an object of type `xml.dom.minidom.Document`, `xml.dom.minidom.Element` or `xml.dom.minidom.Text`.

Since the document is treated as a “flat” stream of events, the document “tree” is implicitly traversed and the desired elements are found regardless of their depth in the tree. In other words, one does not need to consider hierarchical issues such as recursive searching of the document nodes, although if the context of elements were important, one would either need to maintain some context-related state (i.e. remembering where one is in the document at any given point) or to make use of the `DOMEventStream.expandNode()` method and switch to DOM-related processing.

class `xml.dom.pulldom.PullDom` (*documentFactory=None*)
Subclass of `xml.sax.handler.ContentHandler`.

class `xml.dom.pulldom.SAX2DOM` (*documentFactory=None*)
 Subclass of `xml.sax.handler.ContentHandler`.

`xml.dom.pulldom.parse` (*stream_or_string*, *parser=None*, *bufsize=None*)
 Return a *DOMEventStream* from the given input. *stream_or_string* may be either a file name, or a file-like object. *parser*, if given, must be an *XMLReader* object. This function will change the document handler of the parser and activate namespace support; other parser configuration (like setting an entity resolver) must have been done in advance.

If you have XML in a string, you can use the *parseString()* function instead:

`xml.dom.pulldom.parseString` (*string*, *parser=None*)
 Return a *DOMEventStream* that represents the (Unicode) *string*.

`xml.dom.pulldom.default_bufsize`
 Default value for the *bufsize* parameter to *parse()*.

The value of this variable can be changed before calling *parse()* and the new value will take effect.

20.8.1 DOMEventStream Objects

class `xml.dom.pulldom.DOMEventStream` (*stream*, *parser*, *bufsize*)

getEvent ()
 Return a tuple containing *event* and the current *node* as `xml.dom.minidom.Document` if *event* equals `START_DOCUMENT`, `xml.dom.minidom.Element` if *event* equals `START_ELEMENT` or `END_ELEMENT` or `xml.dom.minidom.Text` if *event* equals `CHARACTERS`. The current node does not contain information about its children, unless *expandNode()* is called.

expandNode (*node*)
 Expands all children of *node* into *node*. Example:

```
from xml.dom import pulldom

xml = '<html><title>Foo</title> <p>Some text <div>and more</div></p> </html>'
doc = pulldom.parseString(xml)
for event, node in doc:
    if event == pulldom.START_ELEMENT and node.tagName == 'p':
        # Following statement only prints '<p/>'
        print(node.toxml())
        doc.expandNode(node)
        # Following statement prints node with all its children '<p>Some text
        <div>and more</div></p>'
        print(node.toxml())
```

reset ()

20.9 xml.sax —Support for SAX2 parsers

Source code: `Lib/xml/sax/__init__.py`

The *xml.sax* package provides a number of modules which implement the Simple API for XML (SAX) interface for Python. The package itself provides the SAX exceptions and the convenience functions which will be most used by users of the SAX API.

警告: The `xml.sax` module is not secure against maliciously constructed data. If you need to parse untrusted or unauthenticated data see [XML 漏洞](#).

在 3.6.7 版更改: The SAX parser no longer processes general external entities by default to increase security. Before, the parser created network connections to fetch remote files or loaded local files from the file system for DTD and entities. The feature can be enabled again with method `setFeature()` on the parser object and argument `feature_external_ges`.

The convenience functions are:

`xml.sax.make_parser(parser_list=[])`

Create and return a SAX `XMLReader` object. The first parser found will be used. If `parser_list` is provided, it must be a list of strings which name modules that have a function named `create_parser()`. Modules listed in `parser_list` will be used before modules in the default list of parsers.

`xml.sax.parse(filename_or_stream, handler, error_handler=handler.ErrorHandler())`

Create a SAX parser and use it to parse a document. The document, passed in as `filename_or_stream`, can be a filename or a file object. The `handler` parameter needs to be a SAX `ContentHandler` instance. If `error_handler` is given, it must be a SAX `ErrorHandler` instance; if omitted, `SAXParseException` will be raised on all errors. There is no return value; all work must be done by the `handler` passed in.

`xml.sax.parseString(string, handler, error_handler=handler.ErrorHandler())`

Similar to `parse()`, but parses from a buffer `string` received as a parameter. `string` must be a `str` instance or a `bytes-like object`.

在 3.5 版更改: Added support of `str` instances.

A typical SAX application uses three kinds of objects: readers, handlers and input sources. “Reader” in this context is another term for parser, i.e. some piece of code that reads the bytes or characters from the input source, and produces a sequence of events. The events then get distributed to the handler objects, i.e. the reader invokes a method on the handler. A SAX application must therefore obtain a reader object, create or open the input sources, create the handlers, and connect these objects all together. As the final step of preparation, the reader is called to parse the input. During parsing, methods on the handler objects are called based on structural and syntactic events from the input data.

For these objects, only the interfaces are relevant; they are normally not instantiated by the application itself. Since Python does not have an explicit notion of interface, they are formally introduced as classes, but applications may use implementations which do not inherit from the provided classes. The `InputSource`, `Locator`, `Attributes`, `AttributesNS`, and `XMLReader` interfaces are defined in the module `xml.sax.xmlreader`. The handler interfaces are defined in `xml.sax.handler`. For convenience, `InputSource` (which is often instantiated directly) and the handler classes are also available from `xml.sax`. These interfaces are described below.

In addition to these classes, `xml.sax` provides the following exception classes.

exception `xml.sax.SAXException(msg, exception=None)`

Encapsulate an XML error or warning. This class can contain basic error or warning information from either the XML parser or the application: it can be subclassed to provide additional functionality or to add localization. Note that although the handlers defined in the `ErrorHandler` interface receive instances of this exception, it is not required to actually raise the exception — it is also useful as a container for information.

When instantiated, `msg` should be a human-readable description of the error. The optional `exception` parameter, if given, should be `None` or an exception that was caught by the parsing code and is being passed along as information.

This is the base class for the other SAX exception classes.

exception `xml.sax.SAXParseException(msg, exception, locator)`

Subclass of `SAXException` raised on parse errors. Instances of this class are passed to the methods of the SAX `ErrorHandler` interface to provide information about the parse error. This class supports the SAX `Locator` interface as well as the `SAXException` interface.

exception `xml.sax.SAXNotRecognizedException(msg, exception=None)`

Subclass of `SAXException` raised when a SAX `XMLReader` is confronted with an unrecognized feature or property. SAX applications and extensions may use this class for similar purposes.

exception `xml.sax.SAXNotSupportedException(msg, exception=None)`

Subclass of `SAXException` raised when a SAX `XMLReader` is asked to enable a feature that is not supported, or to set a property to a value that the implementation does not support. SAX applications and extensions may use this class for similar purposes.

参见:

SAX: The Simple API for XML This site is the focal point for the definition of the SAX API. It provides a Java implementation and online documentation. Links to implementations and historical information are also available.

Module `xml.sax.handler` Definitions of the interfaces for application-provided objects.

Module `xml.sax.saxutils` Convenience functions for use in SAX applications.

Module `xml.sax.xmlreader` Definitions of the interfaces for parser-provided objects.

20.9.1 SAXException Objects

The `SAXException` exception class supports the following methods:

`SAXException.getMessage()`

Return a human-readable message describing the error condition.

`SAXException.getException()`

Return an encapsulated exception object, or `None`.

20.10 xml.sax.handler —Base classes for SAX handlers

Source code: `Lib/xml/sax/handler.py`

The SAX API defines four kinds of handlers: content handlers, DTD handlers, error handlers, and entity resolvers. Applications normally only need to implement those interfaces whose events they are interested in; they can implement the interfaces in a single object or in multiple objects. Handler implementations should inherit from the base classes provided in the module `xml.sax.handler`, so that all methods get default implementations.

class `xml.sax.handler.ContentHandler`

This is the main callback interface in SAX, and the one most important to applications. The order of events in this interface mirrors the order of the information in the document.

class `xml.sax.handler.DTDHandler`

Handle DTD events.

This interface specifies only those DTD events required for basic parsing (unparsed entities and attributes).

class `xml.sax.handler.EntityResolver`

Basic interface for resolving entities. If you create an object implementing this interface, then register the object with your Parser, the parser will call the method in your object to resolve all external entities.

class `xml.sax.handler.ErrorHandler`

Interface used by the parser to present error and warning messages to the application. The methods of this object control whether errors are immediately converted to exceptions or are handled in some other way.

In addition to these classes, `xml.sax.handler` provides symbolic constants for the feature and property names.

`xml.sax.handler.feature_namespaces`

value: "http://xml.org/sax/features/namespaces"
true: Perform Namespace processing.
false: Optionally do not perform Namespace processing (implies namespace-prefixes; default).
access: (parsing) read-only; (not parsing) read/write

`xml.sax.handler.feature_namespace_prefixes`

value: "http://xml.org/sax/features/namespace-prefixes"
true: Report the original prefixed names and attributes used for Namespace declarations.
false: Do not report attributes used for Namespace declarations, and optionally do not report original prefixed names (default).
access: (parsing) read-only; (not parsing) read/write

`xml.sax.handler.feature_string_interning`

value: "http://xml.org/sax/features/string-interning"
true: All element names, prefixes, attribute names, Namespace URIs, and local names are interned using the built-in intern function.
false: Names are not necessarily interned, although they may be (default).
access: (parsing) read-only; (not parsing) read/write

`xml.sax.handler.feature_validation`

value: "http://xml.org/sax/features/validation"
true: Report all validation errors (implies external-general-entities and external-parameter-entities).
false: Do not report validation errors.
access: (parsing) read-only; (not parsing) read/write

`xml.sax.handler.feature_external_ges`

value: "http://xml.org/sax/features/external-general-entities"
true: Include all external general (text) entities.
false: Do not include external general entities.
access: (parsing) read-only; (not parsing) read/write

`xml.sax.handler.feature_external_pes`

value: "http://xml.org/sax/features/external-parameter-entities"
true: Include all external parameter entities, including the external DTD subset.
false: Do not include any external parameter entities, even the external DTD subset.
access: (parsing) read-only; (not parsing) read/write

`xml.sax.handler.all_features`

List of all features.

`xml.sax.handler.property_lexical_handler`

value: "http://xml.org/sax/properties/lexical-handler"
data type: `xml.sax.sax2lib.LexicalHandler` (not supported in Python 2)

description: An optional extension handler for lexical events like comments.

access: read/write

`xml.sax.handler.property_declaration_handler`

value: "http://xml.org/sax/properties/declaration-handler"

data type: `xml.sax.sax2lib.DeclHandler` (not supported in Python 2)

description: An optional extension handler for DTD-related events other than notations and unparsed entities.

access: read/write

`xml.sax.handler.property_dom_node`

value: "http://xml.org/sax/properties/dom-node"

data type: `org.w3c.dom.Node` (not supported in Python 2)

description: When parsing, the current DOM node being visited if this is a DOM iterator; when not parsing, the root DOM node for iteration.

access: (parsing) read-only; (not parsing) read/write

`xml.sax.handler.property_xml_string`

value: "http://xml.org/sax/properties/xml-string"

data type: `String`

description: The literal string of characters that was the source for the current event.

access: read-only

`xml.sax.handler.all_properties`

List of all known property names.

20.10.1 ContentHandler Objects

Users are expected to subclass `ContentHandler` to support their application. The following methods are called by the parser on the appropriate events in the input document:

`ContentHandler.setDocumentLocator(locator)`

Called by the parser to give the application a locator for locating the origin of document events.

SAX parsers are strongly encouraged (though not absolutely required) to supply a locator: if it does so, it must supply the locator to the application by invoking this method before invoking any of the other methods in the `DocumentHandler` interface.

The locator allows the application to determine the end position of any document-related event, even if the parser is not reporting an error. Typically, the application will use this information for reporting its own errors (such as character content that does not match an application's business rules). The information returned by the locator is probably not sufficient for use with a search engine.

Note that the locator will return correct information only during the invocation of the events in this interface. The application should not attempt to use it at any other time.

`ContentHandler.startDocument()`

Receive notification of the beginning of a document.

The SAX parser will invoke this method only once, before any other methods in this interface or in `DTDHandler` (except for `setDocumentLocator()`).

`ContentHandler.endDocument()`

Receive notification of the end of a document.

The SAX parser will invoke this method only once, and it will be the last method invoked during the parse. The parser shall not invoke this method until it has either abandoned parsing (because of an unrecoverable error) or reached the end of input.

`ContentHandler.startPrefixMapping(prefix, uri)`

Begin the scope of a prefix-URI Namespace mapping.

The information from this event is not necessary for normal Namespace processing: the SAX XML reader will automatically replace prefixes for element and attribute names when the `feature_namespaces` feature is enabled (the default).

There are cases, however, when applications need to use prefixes in character data or in attribute values, where they cannot safely be expanded automatically; the `startPrefixMapping()` and `endPrefixMapping()` events supply the information to the application to expand prefixes in those contexts itself, if necessary.

Note that `startPrefixMapping()` and `endPrefixMapping()` events are not guaranteed to be properly nested relative to each-other: all `startPrefixMapping()` events will occur before the corresponding `startElement()` event, and all `endPrefixMapping()` events will occur after the corresponding `endElement()` event, but their order is not guaranteed.

`ContentHandler.endPrefixMapping(prefix)`

End the scope of a prefix-URI mapping.

See `startPrefixMapping()` for details. This event will always occur after the corresponding `endElement()` event, but the order of `endPrefixMapping()` events is not otherwise guaranteed.

`ContentHandler.startElement(name, attrs)`

Signals the start of an element in non-namespace mode.

The `name` parameter contains the raw XML 1.0 name of the element type as a string and the `attrs` parameter holds an object of the `Attributes` interface (see [The Attributes Interface](#)) containing the attributes of the element. The object passed as `attrs` may be re-used by the parser; holding on to a reference to it is not a reliable way to keep a copy of the attributes. To keep a copy of the attributes, use the `copy()` method of the `attrs` object.

`ContentHandler.endElement(name)`

Signals the end of an element in non-namespace mode.

The `name` parameter contains the name of the element type, just as with the `startElement()` event.

`ContentHandler.startElementNS(name, qname, attrs)`

Signals the start of an element in namespace mode.

The `name` parameter contains the name of the element type as a `(uri, localname)` tuple, the `qname` parameter contains the raw XML 1.0 name used in the source document, and the `attrs` parameter holds an instance of the `AttributesNS` interface (see [The AttributesNS Interface](#)) containing the attributes of the element. If no namespace is associated with the element, the `uri` component of `name` will be `None`. The object passed as `attrs` may be re-used by the parser; holding on to a reference to it is not a reliable way to keep a copy of the attributes. To keep a copy of the attributes, use the `copy()` method of the `attrs` object.

Parsers may set the `qname` parameter to `None`, unless the `feature_namespace_prefixes` feature is activated.

`ContentHandler.endElementNS(name, qname)`

Signals the end of an element in namespace mode.

The `name` parameter contains the name of the element type, just as with the `startElementNS()` method, likewise the `qname` parameter.

`ContentHandler.characters` (*content*)

Receive notification of character data.

The Parser will call this method to report each chunk of character data. SAX parsers may return all contiguous character data in a single chunk, or they may split it into several chunks; however, all of the characters in any single event must come from the same external entity so that the Locator provides useful information.

content may be a string or bytes instance; the `expat` reader module always produces strings.

注解: The earlier SAX 1 interface provided by the Python XML Special Interest Group used a more Java-like interface for this method. Since most parsers used from Python did not take advantage of the older interface, the simpler signature was chosen to replace it. To convert old code to the new interface, use *content* instead of slicing content with the old *offset* and *length* parameters.

`ContentHandler.ignorableWhitespace` (*whitespace*)

Receive notification of ignorable whitespace in element content.

Validating Parsers must use this method to report each chunk of ignorable whitespace (see the W3C XML 1.0 recommendation, section 2.10): non-validating parsers may also use this method if they are capable of parsing and using content models.

SAX parsers may return all contiguous whitespace in a single chunk, or they may split it into several chunks; however, all of the characters in any single event must come from the same external entity, so that the Locator provides useful information.

`ContentHandler.processingInstruction` (*target*, *data*)

Receive notification of a processing instruction.

The Parser will invoke this method once for each processing instruction found: note that processing instructions may occur before or after the main document element.

A SAX parser should never report an XML declaration (XML 1.0, section 2.8) or a text declaration (XML 1.0, section 4.3.1) using this method.

`ContentHandler.skippedEntity` (*name*)

Receive notification of a skipped entity.

The Parser will invoke this method once for each entity skipped. Non-validating processors may skip entities if they have not seen the declarations (because, for example, the entity was declared in an external DTD subset). All processors may skip external entities, depending on the values of the `feature_external_ges` and the `feature_external_pes` properties.

20.10.2 DTDHandler Objects

`DTDHandler` instances provide the following methods:

`DTDHandler.notationDecl` (*name*, *publicId*, *systemId*)

Handle a notation declaration event.

`DTDHandler.unparsedEntityDecl` (*name*, *publicId*, *systemId*, *ndata*)

Handle an unparsed entity declaration event.

20.10.3 EntityResolver Objects

`EntityResolver.resolveEntity` (*publicId*, *systemId*)

Resolve the system identifier of an entity and return either the system identifier to read from as a string, or an `InputSource` to read from. The default implementation returns *systemId*.

20.10.4 ErrorHandler Objects

Objects with this interface are used to receive error and warning information from the `XMLReader`. If you create an object that implements this interface, then register the object with your `XMLReader`, the parser will call the methods in your object to report all warnings and errors. There are three levels of errors available: warnings, (possibly) recoverable errors, and unrecoverable errors. All methods take a `SAXParseException` as the only parameter. Errors and warnings may be converted to an exception by raising the passed-in exception object.

`ErrorHandler.error(exception)`

Called when the parser encounters a recoverable error. If this method does not raise an exception, parsing may continue, but further document information should not be expected by the application. Allowing the parser to continue may allow additional errors to be discovered in the input document.

`ErrorHandler.fatalError(exception)`

Called when the parser encounters an error it cannot recover from; parsing is expected to terminate when this method returns.

`ErrorHandler.warning(exception)`

Called when the parser presents minor warning information to the application. Parsing is expected to continue when this method returns, and document information will continue to be passed to the application. Raising an exception in this method will cause parsing to end.

20.11 xml.sax.saxutils —SAX Utilities

Source code: `Lib/xml/sax/saxutils.py`

The module `xml.sax.saxutils` contains a number of classes and functions that are commonly useful when creating SAX applications, either in direct use, or as base classes.

`xml.sax.saxutils.escape(data, entities={})`

Escape '&', '<', and '>' in a string of data.

You can escape other strings of data by passing a dictionary as the optional *entities* parameter. The keys and values must all be strings; each key will be replaced with its corresponding value. The characters '&', '<' and '>' are always escaped, even if *entities* is provided.

`xml.sax.saxutils.unescape(data, entities={})`

Unescape '&'; '<'; and '>' in a string of data.

You can unescape other strings of data by passing a dictionary as the optional *entities* parameter. The keys and values must all be strings; each key will be replaced with its corresponding value. '&', '<', and '>' are always unescaped, even if *entities* is provided.

`xml.sax.saxutils.quoteattr(data, entities={})`

Similar to `escape()`, but also prepares *data* to be used as an attribute value. The return value is a quoted version of *data* with any additional required replacements. `quoteattr()` will select a quote character based on the content of *data*, attempting to avoid encoding any quote characters in the string. If both single- and double-quote characters are already in *data*, the double-quote characters will be encoded and *data* will be wrapped in double-quotes. The resulting string can be used directly as an attribute value:

```
>>> print("<element attr=%s>" % quoteattr("ab ' cd \" ef"))
<element attr="ab ' cd &quot; ef">
```

This function is useful when generating attribute values for HTML or any SGML using the reference concrete syntax.

```
class xml.sax.saxutils.XMLGenerator (out=None, encoding='iso-8859-1',
                                     short_empty_elements=False)
```

This class implements the *ContentHandler* interface by writing SAX events back into an XML document. In other words, using an *XMLGenerator* as the content handler will reproduce the original document being parsed. *out* should be a file-like object which will default to *sys.stdout*. *encoding* is the encoding of the output stream which defaults to 'iso-8859-1'. *short_empty_elements* controls the formatting of elements that contain no content: if *False* (the default) they are emitted as a pair of start/end tags, if set to *True* they are emitted as a single self-closed tag.

3.2 新版功能: The *short_empty_elements* parameter.

```
class xml.sax.saxutils.XMLFilterBase (base)
```

This class is designed to sit between an *XMLReader* and the client application's event handlers. By default, it does nothing but pass requests up to the reader and events on to the handlers unmodified, but subclasses can override specific methods to modify the event stream or the configuration requests as they pass through.

```
xml.sax.saxutils.prepare_input_source (source, base="")
```

This function takes an input source and an optional base URL and returns a fully resolved *InputSource* object ready for reading. The input source can be given as a string, a file-like object, or an *InputSource* object; parsers will use this function to implement the polymorphic *source* argument to their *parse()* method.

20.12 xml.sax.xmlreader —Interface for XML parsers

Source code: [Lib/xml/sax/xmlreader.py](#)

SAX parsers implement the *XMLReader* interface. They are implemented in a Python module, which must provide a function *create_parser()*. This function is invoked by *xml.sax.make_parser()* with no arguments to create a new parser object.

```
class xml.sax.xmlreader.XMLReader
```

Base class which can be inherited by SAX parsers.

```
class xml.sax.xmlreader.IncrementalParser
```

In some cases, it is desirable not to parse an input source at once, but to feed chunks of the document as they get available. Note that the reader will normally not read the entire file, but read it in chunks as well; still *parse()* won't return until the entire document is processed. So these interfaces should be used if the blocking behaviour of *parse()* is not desirable.

When the parser is instantiated it is ready to begin accepting data from the feed method immediately. After parsing has been finished with a call to close the reset method must be called to make the parser ready to accept new data, either from feed or using the parse method.

Note that these methods must *not* be called during parsing, that is, after *parse* has been called and before it returns.

By default, the class also implements the parse method of the XMLReader interface using the feed, close and reset methods of the IncrementalParser interface as a convenience to SAX 2.0 driver writers.

```
class xml.sax.xmlreader.Locator
```

Interface for associating a SAX event with a document location. A locator object will return valid results only during calls to DocumentHandler methods; at any other time, the results are unpredictable. If information is not available, methods may return *None*.

```
class xml.sax.xmlreader.InputSource (system_id=None)
```

Encapsulation of the information needed by the *XMLReader* to read entities.

This class may include information about the public identifier, system identifier, byte stream (possibly with character encoding information) and/or the character stream of an entity.

Applications will create objects of this class for use in the `XMLReader.parse()` method and for returning from `EntityResolver.resolveEntity`.

An `InputSource` belongs to the application, the `XMLReader` is not allowed to modify `InputSource` objects passed to it from the application, although it may make copies and modify those.

class `xml.sax.xmlreader.AttributesImpl (attrs)`

This is an implementation of the `Attributes` interface (see section *The Attributes Interface*). This is a dictionary-like object which represents the element attributes in a `startElement()` call. In addition to the most useful dictionary operations, it supports a number of other methods as described by the interface. Objects of this class should be instantiated by readers; `attrs` must be a dictionary-like object containing a mapping from attribute names to attribute values.

class `xml.sax.xmlreader.AttributesNSImpl (attrs, qnames)`

Namespace-aware variant of `AttributesImpl`, which will be passed to `startElementNS()`. It is derived from `AttributesImpl`, but understands attribute names as two-tuples of `namespaceURI` and `localname`. In addition, it provides a number of methods expecting qualified names as they appear in the original document. This class implements the `AttributesNS` interface (see section *The AttributesNS Interface*).

20.12.1 XMLReader Objects

The `XMLReader` interface supports the following methods:

`XMLReader.parse (source)`

Process an input source, producing SAX events. The `source` object can be a system identifier (a string identifying the input source – typically a file name or a URL), a file-like object, or an `InputSource` object. When `parse()` returns, the input is completely processed, and the parser object can be discarded or reset.

在 3.5 版更改: Added support of character streams.

`XMLReader.getContentHandler ()`

Return the current `ContentHandler`.

`XMLReader.setContentHandler (handler)`

Set the current `ContentHandler`. If no `ContentHandler` is set, content events will be discarded.

`XMLReader.getDTDHandler ()`

Return the current `DTDHandler`.

`XMLReader.setDTDHandler (handler)`

Set the current `DTDHandler`. If no `DTDHandler` is set, DTD events will be discarded.

`XMLReader.getEntityResolver ()`

Return the current `EntityResolver`.

`XMLReader.setEntityResolver (handler)`

Set the current `EntityResolver`. If no `EntityResolver` is set, attempts to resolve an external entity will result in opening the system identifier for the entity, and fail if it is not available.

`XMLReader.getErrorHandler ()`

Return the current `ErrorHandler`.

`XMLReader.setErrorHandler (handler)`

Set the current error handler. If no `ErrorHandler` is set, errors will be raised as exceptions, and warnings will be printed.

`XMLReader.setLocale (locale)`

Allow an application to set the locale for errors and warnings.

SAX parsers are not required to provide localization for errors and warnings; if they cannot support the requested locale, however, they must raise a SAX exception. Applications may request a locale change in the middle of a parse.

`XMLReader.getFeature(featurename)`

Return the current setting for feature *featurename*. If the feature is not recognized, `SAXNotRecognizedException` is raised. The well-known featurenames are listed in the module `xml.sax.handler`.

`XMLReader.setFeature(featurename, value)`

Set the *featurename* to *value*. If the feature is not recognized, `SAXNotRecognizedException` is raised. If the feature or its setting is not supported by the parser, `SAXNotSupportedException` is raised.

`XMLReader.getProperty(propertyname)`

Return the current setting for property *propertyname*. If the property is not recognized, a `SAXNotRecognizedException` is raised. The well-known propertynames are listed in the module `xml.sax.handler`.

`XMLReader.setProperty(propertyname, value)`

Set the *propertyname* to *value*. If the property is not recognized, `SAXNotRecognizedException` is raised. If the property or its setting is not supported by the parser, `SAXNotSupportedException` is raised.

20.12.2 IncrementalParser Objects

Instances of `IncrementalParser` offer the following additional methods:

`IncrementalParser.feed(data)`

Process a chunk of *data*.

`IncrementalParser.close()`

Assume the end of the document. That will check well-formedness conditions that can be checked only at the end, invoke handlers, and may clean up resources allocated during parsing.

`IncrementalParser.reset()`

This method is called after `close` has been called to reset the parser so that it is ready to parse new documents. The results of calling `parse` or `feed` after `close` without calling `reset` are undefined.

20.12.3 Locator Objects

Instances of `Locator` provide these methods:

`Locator.getColumnNumber()`

Return the column number where the current event begins.

`Locator.getLineNumber()`

Return the line number where the current event begins.

`Locator.getPublicId()`

Return the public identifier for the current event.

`Locator.getSystemId()`

Return the system identifier for the current event.

20.12.4 InputSource Objects

`InputSource.setPublicId(id)`

Sets the public identifier of this `InputSource`.

`InputSource.getPublicId()`

Returns the public identifier of this *InputSource*.

`InputSource.setSystemId(id)`

Sets the system identifier of this *InputSource*.

`InputSource.getSystemId()`

Returns the system identifier of this *InputSource*.

`InputSource.setEncoding(encoding)`

Sets the character encoding of this *InputSource*.

The encoding must be a string acceptable for an XML encoding declaration (see section 4.3.3 of the XML recommendation).

The encoding attribute of the *InputSource* is ignored if the *InputSource* also contains a character stream.

`InputSource.getEncoding()`

Get the character encoding of this *InputSource*.

`InputSource.setByteStream(bytefile)`

Set the byte stream (a *binary file*) for this input source.

The SAX parser will ignore this if there is also a character stream specified, but it will use a byte stream in preference to opening a URI connection itself.

If the application knows the character encoding of the byte stream, it should set it with the `setEncoding` method.

`InputSource.getByteStream()`

Get the byte stream for this input source.

The `getEncoding` method will return the character encoding for this byte stream, or `None` if unknown.

`InputSource.setCharacterStream(charfile)`

Set the character stream (a *text file*) for this input source.

If there is a character stream specified, the SAX parser will ignore any byte stream and will not attempt to open a URI connection to the system identifier.

`InputSource.getCharacterStream()`

Get the character stream for this input source.

20.12.5 The Attributes Interface

Attributes objects implement a portion of the *mapping protocol*, including the methods `copy()`, `get()`, `__contains__()`, `items()`, `keys()`, and `values()`. The following methods are also provided:

`Attributes.getLength()`

Return the number of attributes.

`Attributes.getNames()`

Return the names of the attributes.

`Attributes.getType(name)`

Returns the type of the attribute *name*, which is normally 'CDATA'.

`Attributes.getValue(name)`

Return the value of attribute *name*.

20.12.6 The `AttributesNS` Interface

This interface is a subtype of the `Attributes` interface (see section [The `Attributes` Interface](#)). All methods supported by that interface are also available on `AttributesNS` objects.

The following methods are also available:

`AttributesNS.get`**ValueByQName** (*name*)

Return the value for a qualified name.

`AttributesNS.get`**NameByQName** (*name*)

Return the (namespace, localname) pair for a qualified *name*.

`AttributesNS.get`**QNameByName** (*name*)

Return the qualified name for a (namespace, localname) pair.

`AttributesNS.get`**QNames** ()

Return the qualified names of all attributes.

20.13 `xml.parsers.expat` —Fast XML parsing using Expat

警告: The `pyexpat` module is not secure against maliciously constructed data. If you need to parse untrusted or unauthenticated data see [XML 漏洞](#).

The `xml.parsers.expat` module is a Python interface to the Expat non-validating XML parser. The module provides a single extension type, `xmlparser`, that represents the current state of an XML parser. After an `xmlparser` object has been created, various attributes of the object can be set to handler functions. When an XML document is then fed to the parser, the handler functions are called for the character data and markup in the XML document.

This module uses the `pyexpat` module to provide access to the Expat parser. Direct use of the `pyexpat` module is deprecated.

This module provides one exception and one type object:

exception `xml.parsers.expat.ExpatError`

The exception raised when Expat reports an error. See section [ExpatError Exceptions](#) for more information on interpreting Expat errors.

exception `xml.parsers.expat.error`

Alias for [ExpatError](#).

`xml.parsers.expat.XMLParserType`

The type of the return values from the [ParserCreate](#) () function.

The `xml.parsers.expat` module contains two functions:

`xml.parsers.expat.ErrorString` (*errno*)

Returns an explanatory string for a given error number *errno*.

`xml.parsers.expat.ParserCreate` (*encoding=None, namespace_separator=None*)

Creates and returns a new `xmlparser` object. *encoding*, if specified, must be a string naming the encoding used by the XML data. Expat doesn't support as many encodings as Python does, and its repertoire of encodings can't

t be extended; it supports UTF-8, UTF-16, ISO-8859-1 (Latin1), and ASCII. If *encoding*¹ is given it will override the implicit or explicit encoding of the document.

Expat can optionally do XML namespace processing for you, enabled by providing a value for *namespace_separator*. The value must be a one-character string; a *ValueError* will be raised if the string has an illegal length (None is considered the same as omission). When namespace processing is enabled, element type names and attribute names that belong to a namespace will be expanded. The element name passed to the element handlers *StartElementHandler* and *EndElementHandler* will be the concatenation of the namespace URI, the namespace separator character, and the local part of the name. If the namespace separator is a zero byte (`chr(0)`) then the namespace URI and the local part will be concatenated without any separator.

For example, if *namespace_separator* is set to a space character (' ') and the following document is parsed:

```
<?xml version="1.0"?>
<root xmlns      = "http://default-namespace.org/"
      xmlns:py   = "http://www.python.org/ns/">
  <py:elem1 />
  <elem2 xmlns="" />
</root>
```

StartElementHandler will receive the following strings for each element:

```
http://default-namespace.org/ root
http://www.python.org/ns/ elem1
elem2
```

Due to limitations in the Expat library used by *pyexpat*, the *xmlparser* instance returned can only be used to parse a single XML document. Call *ParserCreate* for each document to provide unique parser instances.

参见:

The Expat XML Parser Home page of the Expat project.

20.13.1 XMLParser Objects

xmlparser objects have the following methods:

xmlparser.Parse (*data* [, *isfinal*])

Parses the contents of the string *data*, calling the appropriate handler functions to process the parsed data. *isfinal* must be true on the final call to this method; it allows the parsing of a single file in fragments, not the submission of multiple files. *data* can be the empty string at any time.

xmlparser.ParseFile (*file*)

Parse XML data reading from the object *file*. *file* only needs to provide the *read(nbytes)* method, returning the empty string when there's no more data.

xmlparser.SetBase (*base*)

Sets the base to be used for resolving relative URIs in system identifiers in declarations. Resolving relative identifiers is left to the application: this value will be passed through as the *base* argument to the *ExternalEntityRefHandler()*, *NotationDeclHandler()*, and *UnparsedEntityDeclHandler()* functions.

xmlparser.GetBase ()

Returns a string containing the base set by a previous call to *SetBase()*, or None if *SetBase()* hasn't been called.

¹ The encoding string included in XML output should conform to the appropriate standards. For example, "UTF-8" is valid, but "UTF8" is not. See <https://www.w3.org/TR/2006/REC-xml11-20060816/#NT-EncodingDecl> and <https://www.iana.org/assignments/character-sets/character-sets.xhtml>.

`xmlparser.GetInputContext()`

Returns the input data that generated the current event as a string. The data is in the encoding of the entity which contains the text. When called while an event handler is not active, the return value is `None`.

`xmlparser.ExternalEntityParserCreate(context[, encoding])`

Create a “child” parser which can be used to parse an external parsed entity referred to by content parsed by the parent parser. The `context` parameter should be the string passed to the `ExternalEntityRefHandler()` handler function, described below. The child parser is created with the `ordered_attributes` and `specified_attributes` set to the values of this parser.

`xmlparser.SetParamEntityParsing(flag)`

Control parsing of parameter entities (including the external DTD subset). Possible `flag` values are `XML_PARAM_ENTITY_PARSING_NEVER`, `XML_PARAM_ENTITY_PARSING_UNLESS_STANDALONE` and `XML_PARAM_ENTITY_PARSING_ALWAYS`. Return true if setting the flag was successful.

`xmlparser.UseForeignDTD([flag])`

Calling this with a true value for `flag` (the default) will cause Expat to call the `ExternalEntityRefHandler` with `None` for all arguments to allow an alternate DTD to be loaded. If the document does not contain a document type declaration, the `ExternalEntityRefHandler` will still be called, but the `StartDoctypeDeclHandler` and `EndDoctypeDeclHandler` will not be called.

Passing a false value for `flag` will cancel a previous call that passed a true value, but otherwise has no effect.

This method can only be called before the `Parse()` or `ParseFile()` methods are called; calling it after either of those have been called causes `ExpatError` to be raised with the `code` attribute set to `errors.codes[errors.XML_ERROR_CANT_CHANGE_FEATURE_ONCE_PARSING]`.

`xmlparser` objects have the following attributes:

`xmlparser.buffer_size`

The size of the buffer used when `buffer_text` is true. A new buffer size can be set by assigning a new integer value to this attribute. When the size is changed, the buffer will be flushed.

`xmlparser.buffer_text`

Setting this to true causes the `xmlparser` object to buffer textual content returned by Expat to avoid multiple calls to the `CharacterDataHandler()` callback whenever possible. This can improve performance substantially since Expat normally breaks character data into chunks at every line ending. This attribute is false by default, and may be changed at any time.

`xmlparser.buffer_used`

If `buffer_text` is enabled, the number of bytes stored in the buffer. These bytes represent UTF-8 encoded text. This attribute has no meaningful interpretation when `buffer_text` is false.

`xmlparser.ordered_attributes`

Setting this attribute to a non-zero integer causes the attributes to be reported as a list rather than a dictionary. The attributes are presented in the order found in the document text. For each attribute, two list entries are presented: the attribute name and the attribute value. (Older versions of this module also used this format.) By default, this attribute is false; it may be changed at any time.

`xmlparser.specified_attributes`

If set to a non-zero integer, the parser will report only those attributes which were specified in the document instance and not those which were derived from attribute declarations. Applications which set this need to be especially careful to use what additional information is available from the declarations as needed to comply with the standards for the behavior of XML processors. By default, this attribute is false; it may be changed at any time.

The following attributes contain values relating to the most recent error encountered by an `xmlparser` object, and will only have correct values once a call to `Parse()` or `ParseFile()` has raised an `xml.parsers.expat.ExpatError` exception.

`xmlparser.ErrorByteIndex`

Byte index at which an error occurred.

`xmlparser.ErrorCode`

Numeric code specifying the problem. This value can be passed to the `ErrorString()` function, or compared to one of the constants defined in the `errors` object.

`xmlparser.ErrorColumnNumber`

Column number at which an error occurred.

`xmlparser.ErrorLineNumber`

Line number at which an error occurred.

The following attributes contain values relating to the current parse location in an `xmlparser` object. During a callback reporting a parse event they indicate the location of the first of the sequence of characters that generated the event. When called outside of a callback, the position indicated will be just past the last parse event (regardless of whether there was an associated callback).

`xmlparser.CurrentByteIndex`

Current byte index in the parser input.

`xmlparser.CurrentColumnNumber`

Current column number in the parser input.

`xmlparser.CurrentLineNumber`

Current line number in the parser input.

Here is the list of handlers that can be set. To set a handler on an `xmlparser` object *o*, use `o.handlername = func`. *handlername* must be taken from the following list, and *func* must be a callable object accepting the correct number of arguments. The arguments are all strings, unless otherwise stated.

`xmlparser.XmlDeclHandler` (*version, encoding, standalone*)

Called when the XML declaration is parsed. The XML declaration is the (optional) declaration of the applicable version of the XML recommendation, the encoding of the document text, and an optional “standalone” declaration. *version* and *encoding* will be strings, and *standalone* will be 1 if the document is declared standalone, 0 if it is declared not to be standalone, or -1 if the standalone clause was omitted. This is only available with Expat version 1.95.0 or newer.

`xmlparser.StartDoctypeDeclHandler` (*doctypeName, systemId, publicId, has_internal_subset*)

Called when Expat begins parsing the document type declaration (`<!DOCTYPE . . .`). The *doctypeName* is provided exactly as presented. The *systemId* and *publicId* parameters give the system and public identifiers if specified, or None if omitted. *has_internal_subset* will be true if the document contains an internal document declaration subset. This requires Expat version 1.2 or newer.

`xmlparser.EndDoctypeDeclHandler` ()

Called when Expat is done parsing the document type declaration. This requires Expat version 1.2 or newer.

`xmlparser.ElementDeclHandler` (*name, model*)

Called once for each element type declaration. *name* is the name of the element type, and *model* is a representation of the content model.

`xmlparser.AttnlistDeclHandler` (*elname, attname, type, default, required*)

Called for each declared attribute for an element type. If an attribute list declaration declares three attributes, this handler is called three times, once for each attribute. *elname* is the name of the element to which the declaration applies and *attname* is the name of the attribute declared. The attribute type is a string passed as *type*; the possible values are 'CDATA', 'ID', 'IDREF', ... *default* gives the default value for the attribute used when the attribute is not specified by the document instance, or None if there is no default value (#IMPLIED values). If the attribute is required to be given in the document instance, *required* will be true. This requires Expat version 1.95.0 or newer.

`xmlparser.StartElementHandler` (*name, attributes*)

Called for the start of every element. *name* is a string containing the element name, and *attributes* is the element

attributes. If `ordered_attributes` is true, this is a list (see `ordered_attributes` for a full description). Otherwise it's a dictionary mapping names to values.

`xmlparser.EndElementHandler` (*name*)

Called for the end of every element.

`xmlparser.ProcessingInstructionHandler` (*target, data*)

Called for every processing instruction.

`xmlparser.CharacterDataHandler` (*data*)

Called for character data. This will be called for normal character data, CDATA marked content, and ignorable whitespace. Applications which must distinguish these cases can use the `StartCdataSectionHandler`, `EndCdataSectionHandler`, and `ElementDeclHandler` callbacks to collect the required information.

`xmlparser.UnparsedEntityDeclHandler` (*entityName, base, systemId, publicId, notationName*)

Called for unparsed (NDATA) entity declarations. This is only present for version 1.2 of the Expat library; for more recent versions, use `EntityDeclHandler` instead. (The underlying function in the Expat library has been declared obsolete.)

`xmlparser.EntityDeclHandler` (*entityName, is_parameter_entity, value, base, systemId, publicId, notationName*)

Called for all entity declarations. For parameter and internal entities, *value* will be a string giving the declared contents of the entity; this will be `None` for external entities. The *notationName* parameter will be `None` for parsed entities, and the name of the notation for unparsed entities. *is_parameter_entity* will be true if the entity is a parameter entity or false for general entities (most applications only need to be concerned with general entities). This is only available starting with version 1.95.0 of the Expat library.

`xmlparser.NotationDeclHandler` (*notationName, base, systemId, publicId*)

Called for notation declarations. *notationName*, *base*, and *systemId*, and *publicId* are strings if given. If the public identifier is omitted, *publicId* will be `None`.

`xmlparser.StartNamespaceDeclHandler` (*prefix, uri*)

Called when an element contains a namespace declaration. Namespace declarations are processed before the `StartElementHandler` is called for the element on which declarations are placed.

`xmlparser.EndNamespaceDeclHandler` (*prefix*)

Called when the closing tag is reached for an element that contained a namespace declaration. This is called once for each namespace declaration on the element in the reverse of the order for which the `StartNamespaceDeclHandler` was called to indicate the start of each namespace declaration's scope. Calls to this handler are made after the corresponding `EndElementHandler` for the end of the element.

`xmlparser.CommentHandler` (*data*)

Called for comments. *data* is the text of the comment, excluding the leading '`<!--`' and trailing '`-->`'.

`xmlparser.StartCdataSectionHandler` ()

Called at the start of a CDATA section. This and `EndCdataSectionHandler` are needed to be able to identify the syntactical start and end for CDATA sections.

`xmlparser.EndCdataSectionHandler` ()

Called at the end of a CDATA section.

`xmlparser.DefaultHandler` (*data*)

Called for any characters in the XML document for which no applicable handler has been specified. This means characters that are part of a construct which could be reported, but for which no handler has been supplied.

`xmlparser.DefaultHandlerExpand` (*data*)

This is the same as the `DefaultHandler()`, but doesn't inhibit expansion of internal entities. The entity reference will not be passed to the default handler.

`xmlparser.NotStandaloneHandler` ()

Called if the XML document hasn't been declared as being a standalone document. This happens when there is

an external subset or a reference to a parameter entity, but the XML declaration does not set standalone to `yes` in an XML declaration. If this handler returns 0, then the parser will raise an `XML_ERROR_NOT_STANDALONE` error. If this handler is not set, no exception is raised by the parser for this condition.

`xmlparser.ExternalEntityRefHandler` (*context*, *base*, *systemId*, *publicId*)

Called for references to external entities. *base* is the current base, as set by a previous call to `SetBase()`. The public and system identifiers, *systemId* and *publicId*, are strings if given; if the public identifier is not given, *publicId* will be `None`. The *context* value is opaque and should only be used as described below.

For external entities to be parsed, this handler must be implemented. It is responsible for creating the sub-parser using `ExternalEntityParserCreate(context)`, initializing it with the appropriate callbacks, and parsing the entity. This handler should return an integer; if it returns 0, the parser will raise an `XML_ERROR_EXTERNAL_ENTITY_HANDLING` error, otherwise parsing will continue.

If this handler is not provided, external entities are reported by the `DefaultHandler` callback, if provided.

20.13.2 ExpatError Exceptions

`ExpatError` exceptions have a number of interesting attributes:

`ExpatError.code`

Expat's internal error number for the specific error. The `errors.messages` dictionary maps these error numbers to Expat's error messages. For example:

```
from xml.parsers.expat import ParserCreate, ExpatError, errors

p = ParserCreate()
try:
    p.Parse(some_xml_document)
except ExpatError as err:
    print("Error:", errors.messages[err.code])
```

The `errors` module also provides error message constants and a dictionary `codes` mapping these messages back to the error codes, see below.

`ExpatError.lineno`

Line number on which the error was detected. The first line is numbered 1.

`ExpatError.offset`

Character offset into the line where the error occurred. The first column is numbered 0.

20.13.3 示例

The following program defines three handlers that just print out their arguments.

```
import xml.parsers.expat

# 3 handler functions
def start_element(name, attrs):
    print('Start element:', name, attrs)
def end_element(name):
    print('End element:', name)
def char_data(data):
    print('Character data:', repr(data))

p = xml.parsers.expat.ParserCreate()
```

(下页继续)

(续上页)

```

p.StartElementHandler = start_element
p.EndElementHandler = end_element
p.CharacterDataHandler = char_data

p.Parse("""<?xml version="1.0"?>
<parent id="top"><child1 name="paul">Text goes here</child1>
<child2 name="fred">More text</child2>
</parent>""", 1)

```

The output from this program is:

```

Start element: parent {'id': 'top'}
Start element: child1 {'name': 'paul'}
Character data: 'Text goes here'
End element: child1
Character data: '\n'
Start element: child2 {'name': 'fred'}
Character data: 'More text'
End element: child2
Character data: '\n'
End element: parent

```

20.13.4 Content Model Descriptions

Content models are described using nested tuples. Each tuple contains four values: the type, the quantifier, the name, and a tuple of children. Children are simply additional content model descriptions.

The values of the first two fields are constants defined in the `xml.parsers.expat.model` module. These constants can be collected in two groups: the model type group and the quantifier group.

The constants in the model type group are:

`xml.parsers.expat.model.XML_CTYPE_ANY`

The element named by the model name was declared to have a content model of ANY.

`xml.parsers.expat.model.XML_CTYPE_CHOICE`

The named element allows a choice from a number of options; this is used for content models such as (A | B | C).

`xml.parsers.expat.model.XML_CTYPE_EMPTY`

Elements which are declared to be EMPTY have this model type.

`xml.parsers.expat.model.XML_CTYPE_MIXED`

`xml.parsers.expat.model.XML_CTYPE_NAME`

`xml.parsers.expat.model.XML_CTYPE_SEQ`

Models which represent a series of models which follow one after the other are indicated with this model type. This is used for models such as (A, B, C).

The constants in the quantifier group are:

`xml.parsers.expat.model.XML_CQUANT_NONE`

No modifier is given, so it can appear exactly once, as for A.

`xml.parsers.expat.model.XML_CQUANT_OPT`

The model is optional: it can appear once or not at all, as for A?.

`xml.parsers.expat.model.XML_CQUANT_PLUS`

The model must occur one or more times (like A+).

`xml.parsers.expat.model.XML_CQUANT_REP`

The model must occur zero or more times, as for A*.

20.13.5 Expat error constants

The following constants are provided in the `xml.parsers.expat.errors` module. These constants are useful in interpreting some of the attributes of the `ExpatError` exception objects raised when an error has occurred. Since for backwards compatibility reasons, the constants' value is the error *message* and not the numeric error *code*, you do this by comparing its `code` attribute with `errors.codes[errors.XML_ERROR_CONSTANT_NAME]`.

The `errors` module has the following attributes:

`xml.parsers.expat.errors.codes`

A dictionary mapping numeric error codes to their string descriptions.

3.2 新版功能.

`xml.parsers.expat.errors.messages`

A dictionary mapping string descriptions to their error codes.

3.2 新版功能.

`xml.parsers.expat.errors.XML_ERROR_ASYNC_ENTITY`

`xml.parsers.expat.errors.XML_ERROR_ATTRIBUTE_EXTERNAL_ENTITY_REF`

An entity reference in an attribute value referred to an external entity instead of an internal entity.

`xml.parsers.expat.errors.XML_ERROR_BAD_CHAR_REF`

A character reference referred to a character which is illegal in XML (for example, character 0, or '�').

`xml.parsers.expat.errors.XML_ERROR_BINARY_ENTITY_REF`

An entity reference referred to an entity which was declared with a notation, so cannot be parsed.

`xml.parsers.expat.errors.XML_ERROR_DUPLICATE_ATTRIBUTE`

An attribute was used more than once in a start tag.

`xml.parsers.expat.errors.XML_ERROR_INCORRECT_ENCODING`

`xml.parsers.expat.errors.XML_ERROR_INVALID_TOKEN`

Raised when an input byte could not properly be assigned to a character; for example, a NUL byte (value 0) in a UTF-8 input stream.

`xml.parsers.expat.errors.XML_ERROR_JUNK_AFTER_DOC_ELEMENT`

Something other than whitespace occurred after the document element.

`xml.parsers.expat.errors.XML_ERROR_MISPLACED_XML_PI`

An XML declaration was found somewhere other than the start of the input data.

`xml.parsers.expat.errors.XML_ERROR_NO_ELEMENTS`

The document contains no elements (XML requires all documents to contain exactly one top-level element)..

`xml.parsers.expat.errors.XML_ERROR_NO_MEMORY`

Expat was not able to allocate memory internally.

`xml.parsers.expat.errors.XML_ERROR_PARAM_ENTITY_REF`

A parameter entity reference was found where it was not allowed.

`xml.parsers.expat.errors.XML_ERROR_PARTIAL_CHAR`

An incomplete character was found in the input.

`xml.parsers.expat.errors.XML_ERROR_RECURSIVE_ENTITY_REF`
 An entity reference contained another reference to the same entity; possibly via a different name, and possibly indirectly.

`xml.parsers.expat.errors.XML_ERROR_SYNTAX`
 Some unspecified syntax error was encountered.

`xml.parsers.expat.errors.XML_ERROR_TAG_MISMATCH`
 An end tag did not match the innermost open start tag.

`xml.parsers.expat.errors.XML_ERROR_UNCLOSED_TOKEN`
 Some token (such as a start tag) was not closed before the end of the stream or the next token was encountered.

`xml.parsers.expat.errors.XML_ERROR_UNDEFINED_ENTITY`
 A reference was made to an entity which was not defined.

`xml.parsers.expat.errors.XML_ERROR_UNKNOWN_ENCODING`
 The document encoding is not supported by Expat.

`xml.parsers.expat.errors.XML_ERROR_UNCLOSED_CDATA_SECTION`
 A CDATA marked section was not closed.

`xml.parsers.expat.errors.XML_ERROR_EXTERNAL_ENTITY_HANDLING`

`xml.parsers.expat.errors.XML_ERROR_NOT_STANDALONE`
 The parser determined that the document was not “standalone” though it declared itself to be in the XML declaration, and the `NotStandaloneHandler` was set and returned 0.

`xml.parsers.expat.errors.XML_ERROR_UNEXPECTED_STATE`

`xml.parsers.expat.errors.XML_ERROR_ENTITY_DECLARED_IN_PE`

`xml.parsers.expat.errors.XML_ERROR_FEATURE_REQUIRES_XML_DTD`
 An operation was requested that requires DTD support to be compiled in, but Expat was configured without DTD support. This should never be reported by a standard build of the `xml.parsers.expat` module.

`xml.parsers.expat.errors.XML_ERROR_CANT_CHANGE_FEATURE_ONCE_PARSING`
 A behavioral change was requested after parsing started that can only be changed before parsing has started. This is (currently) only raised by `UseForeignDTD()`.

`xml.parsers.expat.errors.XML_ERROR_UNBOUND_PREFIX`
 An undeclared prefix was found when namespace processing was enabled.

`xml.parsers.expat.errors.XML_ERROR_UNDECLARING_PREFIX`
 The document attempted to remove the namespace declaration associated with a prefix.

`xml.parsers.expat.errors.XML_ERROR_INCOMPLETE_PE`
 A parameter entity contained incomplete markup.

`xml.parsers.expat.errors.XML_ERROR_XML_DECL`
 The document contained no document element at all.

`xml.parsers.expat.errors.XML_ERROR_TEXT_DECL`
 There was an error parsing a text declaration in an external entity.

`xml.parsers.expat.errors.XML_ERROR_PUBLICID`
 Characters were found in the public id that are not allowed.

`xml.parsers.expat.errors.XML_ERROR_SUSPENDED`
 The requested operation was made on a suspended parser, but isn't allowed. This includes attempts to provide additional input or to stop the parser.

`xml.parsers.expat.errors.XML_ERROR_NOT_SUSPENDED`
 An attempt to resume the parser was made when the parser had not been suspended.

`xml.parsers.expat.errors.XML_ERROR_ABORTED`

This should not be reported to Python applications.

`xml.parsers.expat.errors.XML_ERROR_FINISHED`

The requested operation was made on a parser which was finished parsing input, but isn't allowed. This includes attempts to provide additional input or to stop the parser.

`xml.parsers.expat.errors.XML_ERROR_SUSPEND_PE`

本章介绍的模块实现了互联网协议并支持相关技术。它们都是用 Python 实现的。这些模块中的大多数都需要存在依赖于系统的模块 `socket`，目前大多数流行平台都支持它。这是一个概述：

21.1 webbrowser 方便的 Web 浏览器控制器

源码： `Lib/webbrowser.py`

`webbrowser` 模块提供了一个高级接口，允许向用户显示基于 Web 的文档。在大多数情况下，只需从该模块调用 `open()` 函数就可以了。

在 Unix 下，图形浏览器在 X11 下是首选，但如果图形浏览器不可用或 X11 显示不可用，则将使用文本模式浏览器。如果使用文本模式浏览器，则调用进程将阻塞，直到用户退出浏览器。

如果存在环境变量 `BROWSER`，则将其解释为 `os.pathsep` 分隔的浏览器列表，以便在平台默认值之前尝试。当列表部分的值包含字符串 `%s` 时，它被解释为一个文字浏览器命令行，用于替换 `%s` 的参数 `URL`；如果该部分不包含 `%s`，则它只被解释为要启动的浏览器的名称。¹

对于非 Unix 平台，或者当 Unix 上有远程浏览器时，控制过程不会等待用户完成浏览器，而是允许远程浏览器在显示界面上维护自己的窗口。如果 Unix 上没有远程浏览器，控制进程将启动一个新的浏览器并等待。

脚本 `webbrowser` 可以用作模块的命令行界面。它接受一个 `URL` 作为参数。还接受以下可选参数：`-n` 如果可能，在新的浏览器窗口中打开 `URL`；`-t` 在新的浏览器页面（“标签”）中打开 `URL`。这些选择当然是相互排斥的。用法示例：

```
python -m webbrowser -t "http://www.python.org"
```

定义了以下异常：

exception webbrowser.Error
发生浏览器控件错误时引发异常。

¹ 这里命名的不带完整路径的可执行文件将在 `PATH` 环境变量给出的目录中搜索。

定义了以下函数：

`webbrowser.open(url, new=0, autoraise=True)`

使用默认浏览器显示 `url`。如果 `new` 为 0，则尽可能在同一浏览器窗口中打开 `url`。如果 `new` 为 1，则尽可能打开新的浏览器窗口。如果 `new` 为 2，则尽可能打开新的浏览器页面（“标签”）。如果 `autoraise` 为 “True”，则会尽可能置前窗口（请注意，在许多窗口管理器下，无论此变量的设置如何，都会置前窗口）。

请注意，在某些平台上，尝试使用此函数打开文件名，可能会起作用并启动操作系统的关联程序。但是，这种方式不被支持也不可移植。

`webbrowser.open_new(url)`

如果可能，在默认浏览器的新窗口中打开 `url`，否则，在唯一的浏览器窗口中打开 `url`。

`webbrowser.open_new_tab(url)`

如果可能，在默认浏览器的新页面（“标签”）中打开 `url`，否则等效于 `open_new()`。

`webbrowser.get(using=None)`

返回浏览器类型为 `using` 指定的控制器对象。如果 `using` 为 `None`，则返回适用于调用者环境的默认浏览器的控制器。

`webbrowser.register(name, constructor, instance=None)`

注册 `name` 浏览器类型。注册浏览器类型后，`get()` 函数可以返回该浏览器类型的控制器。如果没有提供 `instance`，或者为 `None`，`constructor` 将在没有参数的情况下被调用，以在需要时创建实例。如果提供了 `instance`，则永远不会调用 `constructor`，并且可能是 `None`。

This entry point is only useful if you plan to either set the BROWSER variable or call `get()` with a nonempty argument matching the name of a handler you declare.

预定义了许多浏览器类型。此表给出了可以传递给 `get()` 函数的类型名称以及控制器类的相应实例化，这些都在此模块中定义。

类型名	类名	注释
'mozilla'	Mozilla('mozilla')	
'firefox'	Mozilla('mozilla')	
'netscape'	Mozilla('netscape')	
'galeon'	Galeon('galeon')	
'epiphany'	Galeon('epiphany')	
'skipstone'	BackgroundBrowser('skipstone')	
'kfmclient'	Konqueror()	(1)
'konqueror'	Konqueror()	(1)
'kfm'	Konqueror()	(1)
'mosaic'	BackgroundBrowser('mosaic')	
'opera'	Opera()	
'grail'	Grail()	
'links'	GenericBrowser('links')	
'elinks'	Elinks('elinks')	
'lynx'	GenericBrowser('lynx')	
'w3m'	GenericBrowser('w3m')	
'windows-default'	WindowsDefault	(2)
'macosx'	MacOSX('default')	(3)
'safari'	MacOSX('safari')	(3)
'google-chrome'	Chrome('google-chrome')	
'chrome'	Chrome('chrome')	
'chromium'	Chromium('chromium')	
'chromium-browser'	Chromium('chromium-browser')	

注释:

- (1) “Konqueror” 是 Unix 的 KDE 桌面环境的文件管理器，只有在 KDE 运行时才有意义。一些可靠地检测 KDE 的方法会很好；仅检查 `KDEDIR` 变量是不够的。另请注意，KDE 2 的 **konqueror** 命令，会使用名称 “kfm” — 此实现选择运行的 Konqueror 的最佳策略。
- (2) 仅限 Windows 平台。
- (3) 仅限 Mac OS X 平台。

3.3 新版功能: 添加了对 Chrome/Chromium 的支持。

以下是一些简单的例子:

```
url = 'http://docs.python.org/'

# Open URL in a new tab, if a browser window is already open.
webbrowser.open_new_tab(url)

# Open URL in new window, raising the window if possible.
webbrowser.open_new(url)
```

21.1.1 浏览器控制器对象

浏览器控制器提供三个与模块级便捷函数相同的方法:

`controller.open(url, new=0, autoraise=True)`

使用此控制器处理的浏览器显示 `url`。如果 `new` 为 1，则尽可能打开新的浏览器窗口。如果 `new` 为 2，则尽可能打开新的浏览器页面（“标签”）。

`controller.open_new(url)`

如果可能，在此控制器处理的浏览器的新窗口中打开 `url`，否则，在唯一的浏览器窗口中打开 `url`。别名 `open_new()`。

`controller.open_new_tab(url)`

如果可能，在此控制器处理的浏览器的新页面（“标签”）中打开 `url`，否则等效于 `open_new()`。

21.2 cgi — Common Gateway Interface support

Source code: [Lib/cgi.py](#)

Support module for Common Gateway Interface (CGI) scripts.

This module defines a number of utilities for use by CGI scripts written in Python.

21.2.1 概述

A CGI script is invoked by an HTTP server, usually to process user input submitted through an HTML `<FORM>` or `<ISINDEX>` element.

Most often, CGI scripts live in the server’s special `cgi-bin` directory. The HTTP server places all sorts of information about the request (such as the client’s hostname, the requested URL, the query string, and lots of other goodies) in the script’s shell environment, executes the script, and sends the script’s output back to the client.

The script's input is connected to the client too, and sometimes the form data is read this way; at other times the form data is passed via the “query string” part of the URL. This module is intended to take care of the different cases and provide a simpler interface to the Python script. It also provides a number of utilities that help in debugging scripts, and the latest addition is support for file uploads from a form (if your browser supports it).

The output of a CGI script should consist of two sections, separated by a blank line. The first section contains a number of headers, telling the client what kind of data is following. Python code to generate a minimal header section looks like this:

```
print("Content-Type: text/html")    # HTML is following
print()                          # blank line, end of headers
```

The second section is usually HTML, which allows the client software to display nicely formatted text with header, in-line images, etc. Here's Python code that prints a simple piece of HTML:

```
print("<TITLE>CGI script output</TITLE>")
print("<H1>This is my first CGI script</H1>")
print("Hello, world!")
```

21.2.2 Using the cgi module

Begin by writing `import cgi`.

When you write a new script, consider adding these lines:

```
import cgitb
cgitb.enable()
```

This activates a special exception handler that will display detailed reports in the Web browser if any errors occur. If you'd rather not show the guts of your program to users of your script, you can have the reports saved to files instead, with code like this:

```
import cgitb
cgitb.enable(display=0, logdir="/path/to/logdir")
```

It's very helpful to use this feature during script development. The reports produced by `cgitb` provide information that can save you a lot of time in tracking down bugs. You can always remove the `cgitb` line later when you have tested your script and are confident that it works correctly.

To get at submitted form data, use the `FieldStorage` class. If the form contains non-ASCII characters, use the `encoding` keyword parameter set to the value of the encoding defined for the document. It is usually contained in the META tag in the HEAD section of the HTML document or by the `Content-Type` header). This reads the form contents from the standard input or the environment (depending on the value of various environment variables set according to the CGI standard). Since it may consume standard input, it should be instantiated only once.

The `FieldStorage` instance can be indexed like a Python dictionary. It allows membership testing with the `in` operator, and also supports the standard dictionary method `keys()` and the built-in function `len()`. Form fields containing empty strings are ignored and do not appear in the dictionary; to keep such values, provide a true value for the optional `keep_blank_values` keyword parameter when creating the `FieldStorage` instance.

For instance, the following code (which assumes that the `Content-Type` header and blank line have already been printed) checks that the fields `name` and `addr` are both set to a non-empty string:

```
form = cgi.FieldStorage()
if "name" not in form or "addr" not in form:
    print("<H1>Error</H1>")
```

(下页继续)

(续上页)

```

    print("Please fill in the name and addr fields.")
    return
print("<p>name:", form["name"].value)
print("<p>addr:", form["addr"].value)
...further form processing here...

```

Here the fields, accessed through `form[key]`, are themselves instances of `FieldStorage` (or `MiniFieldStorage`, depending on the form encoding). The `value` attribute of the instance yields the string value of the field. The `getvalue()` method returns this string value directly; it also accepts an optional second argument as a default to return if the requested key is not present.

If the submitted form data contains more than one field with the same name, the object retrieved by `form[key]` is not a `FieldStorage` or `MiniFieldStorage` instance but a list of such instances. Similarly, in this situation, `form.getvalue(key)` would return a list of strings. If you expect this possibility (when your HTML form contains multiple fields with the same name), use the `getlist()` method, which always returns a list of values (so that you do not need to special-case the single item case). For example, this code concatenates any number of username fields, separated by commas:

```

value = form.getlist("username")
usernames = ",".join(value)

```

If a field represents an uploaded file, accessing the value via the `value` attribute or the `getvalue()` method reads the entire file in memory as bytes. This may not be what you want. You can test for an uploaded file by testing either the `filename` attribute or the `file` attribute. You can then read the data from the `file` attribute before it is automatically closed as part of the garbage collection of the `FieldStorage` instance (the `read()` and `readline()` methods will return bytes):

```

fileitem = form["userfile"]
if fileitem.file:
    # It's an uploaded file; count lines
    linecount = 0
    while True:
        line = fileitem.file.readline()
        if not line: break
        linecount = linecount + 1

```

`FieldStorage` objects also support being used in a `with` statement, which will automatically close them when done.

If an error is encountered when obtaining the contents of an uploaded file (for example, when the user interrupts the form submission by clicking on a Back or Cancel button) the `done` attribute of the object for the field will be set to the value `-1`.

The file upload draft standard entertains the possibility of uploading multiple files from one field (using a recursive `multipart/*` encoding). When this occurs, the item will be a dictionary-like `FieldStorage` item. This can be determined by testing its `type` attribute, which should be `multipart/form-data` (or perhaps another MIME type matching `multipart/*`). In this case, it can be iterated over recursively just like the top-level form object.

When a form is submitted in the “old” format (as the query string or as a single data part of type `application/x-www-form-urlencoded`), the items will actually be instances of the class `MiniFieldStorage`. In this case, the `list`, `file`, and `filename` attributes are always `None`.

A form submitted via POST that also has a query string will contain both `FieldStorage` and `MiniFieldStorage` items.

在 3.4 版更改: The `file` attribute is automatically closed upon the garbage collection of the creating `FieldStorage` instance.

在 3.5 版更改: Added support for the context management protocol to the `FieldStorage` class.

21.2.3 Higher Level Interface

The previous section explains how to read CGI form data using the `FieldStorage` class. This section describes a higher level interface which was added to this class to allow one to do it in a more readable and intuitive way. The interface doesn't make the techniques described in previous sections obsolete—they are still useful to process file uploads efficiently, for example.

The interface consists of two simple methods. Using the methods you can process form data in a generic way, without the need to worry whether only one or more values were posted under one name.

In the previous section, you learned to write following code anytime you expected a user to post more than one value under one name:

```
item = form.getvalue("item")
if isinstance(item, list):
    # The user is requesting more than one item.
else:
    # The user is requesting only one item.
```

This situation is common for example when a form contains a group of multiple checkboxes with the same name:

```
<input type="checkbox" name="item" value="1" />
<input type="checkbox" name="item" value="2" />
```

In most situations, however, there's only one form control with a particular name in a form and then you expect and need only one value associated with this name. So you write a script containing for example this code:

```
user = form.getvalue("user").upper()
```

The problem with the code is that you should never expect that a client will provide valid input to your scripts. For example, if a curious user appends another `user=foo` pair to the query string, then the script would crash, because in this situation the `getvalue("user")` method call returns a list instead of a string. Calling the `upper()` method on a list is not valid (since lists do not have a method of this name) and results in an `AttributeError` exception.

Therefore, the appropriate way to read form data values was to always use the code which checks whether the obtained value is a single value or a list of values. That's annoying and leads to less readable scripts.

A more convenient approach is to use the methods `getfirst()` and `getlist()` provided by this higher level interface.

`FieldStorage.getfirst(name, default=None)`

This method always returns only one value associated with form field *name*. The method returns only the first value in case that more values were posted under such name. Please note that the order in which the values are received may vary from browser to browser and should not be counted on.¹ If no such form field or value exists then the method returns the value specified by the optional parameter *default*. This parameter defaults to `None` if not specified.

`FieldStorage.getlist(name)`

This method always returns a list of values associated with form field *name*. The method returns an empty list if no such form field or value exists for *name*. It returns a list consisting of one item if only one such value exists.

Using these methods you can write nice compact code:

```
import cgi
form = cgi.FieldStorage()
user = form.getfirst("user", "").upper()    # This way it's safe.
```

(下页继续)

¹ Note that some recent versions of the HTML specification do state what order the field values should be supplied in, but knowing whether a request was received from a conforming browser, or even from a browser at all, is tedious and error-prone.

(续上页)

```
for item in form.getlist("item"):
    do_something(item)
```

21.2.4 函数

These are useful if you want more control, or if you want to employ some of the algorithms implemented in this module in other circumstances.

cgi.parse (*fp=None, environ=os.environ, keep_blank_values=False, strict_parsing=False*)

Parse a query in the environment or from a file (the file defaults to `sys.stdin`). The *keep_blank_values* and *strict_parsing* parameters are passed to `urllib.parse.parse_qs()` unchanged.

cgi.parse_qs (*qs, keep_blank_values=False, strict_parsing=False*)

This function is deprecated in this module. Use `urllib.parse.parse_qs()` instead. It is maintained here only for backward compatibility.

cgi.parse_qs1 (*qs, keep_blank_values=False, strict_parsing=False*)

This function is deprecated in this module. Use `urllib.parse.parse_qs1()` instead. It is maintained here only for backward compatibility.

cgi.parse_multipart (*fp, pdict*)

Parse input of type *multipart/form-data* (for file uploads). Arguments are *fp* for the input file and *pdict* for a dictionary containing other parameters in the *Content-Type* header.

Returns a dictionary just like `urllib.parse.parse_qs()` keys are the field names, each value is a list of values for that field. This is easy to use but not much good if you are expecting megabytes to be uploaded—in that case, use the `FieldStorage` class instead which is much more flexible.

Note that this does not parse nested multipart parts—use `FieldStorage` for that.

cgi.parse_header (*string*)

Parse a MIME header (such as *Content-Type*) into a main value and a dictionary of parameters.

cgi.test ()

Robust test CGI script, usable as main program. Writes minimal HTTP headers and formats all information provided to the script in HTML form.

cgi.print_environ ()

Format the shell environment in HTML.

cgi.print_form (*form*)

Format a form in HTML.

cgi.print_directory ()

Format the current directory in HTML.

cgi.print_environ_usage ()

Print a list of useful (used by CGI) environment variables in HTML.

cgi.escape (*s, quote=False*)

Convert the characters '&', '<' and '>' in string *s* to HTML-safe sequences. Use this if you need to display text that might contain such characters in HTML. If the optional flag *quote* is true, the quotation mark character (") is also translated; this helps for inclusion in an HTML attribute value delimited by double quotes, as in ``. Note that single quotes are never translated.

3.2 版后已移除: This function is unsafe because *quote* is false by default, and therefore deprecated. Use `html.escape()` instead.

21.2.5 Caring about security

There's one important rule: if you invoke an external program (via the `os.system()` or `os.popen()` functions, or others with similar functionality), make very sure you don't pass arbitrary strings received from the client to the shell. This is a well-known security hole whereby clever hackers anywhere on the Web can exploit a gullible CGI script to invoke arbitrary shell commands. Even parts of the URL or field names cannot be trusted, since the request doesn't have to come from your form!

To be on the safe side, if you must pass a string gotten from a form to a shell command, you should make sure the string contains only alphanumeric characters, dashes, underscores, and periods.

21.2.6 Installing your CGI script on a Unix system

Read the documentation for your HTTP server and check with your local system administrator to find the directory where CGI scripts should be installed; usually this is in a directory `cgi-bin` in the server tree.

Make sure that your script is readable and executable by “others”; the Unix file mode should be `00755` octal (use `chmod 0755 filename`). Make sure that the first line of the script contains `#!` starting in column 1 followed by the pathname of the Python interpreter, for instance:

```
#!/usr/local/bin/python
```

Make sure the Python interpreter exists and is executable by “others”.

Make sure that any files your script needs to read or write are readable or writable, respectively, by “others”—their mode should be `00644` for readable and `00666` for writable. This is because, for security reasons, the HTTP server executes your script as user “nobody”, without any special privileges. It can only read (write, execute) files that everybody can read (write, execute). The current directory at execution time is also different (it is usually the server's `cgi-bin` directory) and the set of environment variables is also different from what you get when you log in. In particular, don't count on the shell's search path for executables (`PATH`) or the Python module search path (`PYTHONPATH`) to be set to anything interesting.

If you need to load modules from a directory which is not on Python's default module search path, you can change the path in your script, before importing other modules. For example:

```
import sys
sys.path.insert(0, "/usr/home/joe/lib/python")
sys.path.insert(0, "/usr/local/lib/python")
```

(This way, the directory inserted last will be searched first!)

Instructions for non-Unix systems will vary; check your HTTP server's documentation (it will usually have a section on CGI scripts).

21.2.7 Testing your CGI script

Unfortunately, a CGI script will generally not run when you try it from the command line, and a script that works perfectly from the command line may fail mysteriously when run from the server. There's one reason why you should still test your script from the command line: if it contains a syntax error, the Python interpreter won't execute it at all, and the HTTP server will most likely send a cryptic error to the client.

Assuming your script has no syntax errors, yet it does not work, you have no choice but to read the next section.

21.2.8 Debugging CGI scripts

First of all, check for trivial installation errors — reading the section above on installing your CGI script carefully can save you a lot of time. If you wonder whether you have understood the installation procedure correctly, try installing a copy of this module file (`cgi.py`) as a CGI script. When invoked as a script, the file will dump its environment and the contents of the form in HTML form. Give it the right mode etc, and send it a request. If it's installed in the standard `cgi-bin` directory, it should be possible to send it a request by entering a URL into your browser of the form:

```
http://yourhostname/cgi-bin/cgi.py?name=Joe+Blow&addr=At+Home
```

If this gives an error of type 404, the server cannot find the script — perhaps you need to install it in a different directory. If it gives another error, there's an installation problem that you should fix before trying to go any further. If you get a nicely formatted listing of the environment and form content (in this example, the fields should be listed as “addr” with value “At Home” and “name” with value “Joe Blow”), the `cgi.py` script has been installed correctly. If you follow the same procedure for your own script, you should now be able to debug it.

The next step could be to call the `cgi` module's `test()` function from your script: replace its main code with the single statement

```
cgi.test()
```

This should produce the same results as those gotten from installing the `cgi.py` file itself.

When an ordinary Python script raises an unhandled exception (for whatever reason: of a typo in a module name, a file that can't be opened, etc.), the Python interpreter prints a nice traceback and exits. While the Python interpreter will still do this when your CGI script raises an exception, most likely the traceback will end up in one of the HTTP server's log files, or be discarded altogether.

Fortunately, once you have managed to get your script to execute *some* code, you can easily send tracebacks to the Web browser using the `cgitb` module. If you haven't done so already, just add the lines:

```
import cgitb
cgitb.enable()
```

to the top of your script. Then try running it again; when a problem occurs, you should see a detailed report that will likely make apparent the cause of the crash.

If you suspect that there may be a problem in importing the `cgitb` module, you can use an even more robust approach (which only uses built-in modules):

```
import sys
sys.stderr = sys.stdout
print("Content-Type: text/plain")
print()
...your code here...
```

This relies on the Python interpreter to print the traceback. The content type of the output is set to plain text, which disables all HTML processing. If your script works, the raw HTML will be displayed by your client. If it raises an exception, most likely after the first two lines have been printed, a traceback will be displayed. Because no HTML interpretation is going on, the traceback will be readable.

21.2.9 Common problems and solutions

- Most HTTP servers buffer the output from CGI scripts until the script is completed. This means that it is not possible to display a progress report on the client's display while the script is running.
- Check the installation instructions above.

- Check the HTTP server's log files. (`tail -f logfile` in a separate window may be useful!)
- Always check a script for syntax errors first, by doing something like `python script.py`.
- If your script does not have any syntax errors, try adding `import cgi; cgi.enable()` to the top of the script.
- When invoking external programs, make sure they can be found. Usually, this means using absolute path names —`PATH` is usually not set to a very useful value in a CGI script.
- When reading or writing external files, make sure they can be read or written by the `userid` under which your CGI script will be running: this is typically the `userid` under which the web server is running, or some explicitly specified `userid` for a web server's `suexec` feature.
- Don't try to give a CGI script a `set-uid` mode. This doesn't work on most systems, and is a security liability as well.

21.3 `cgitb` —Traceback manager for CGI scripts

Source code: [Lib/cgi.py](#)

The `cgitb` module provides a special exception handler for Python scripts. (Its name is a bit misleading. It was originally designed to display extensive traceback information in HTML for CGI scripts. It was later generalized to also display this information in plain text.) After this module is activated, if an uncaught exception occurs, a detailed, formatted report will be displayed. The report includes a traceback showing excerpts of the source code for each level, as well as the values of the arguments and local variables to currently running functions, to help you debug the problem. Optionally, you can save this information to a file instead of sending it to the browser.

To enable this feature, simply add this to the top of your CGI script:

```
import cgi
cgi.enable()
```

The options to the `enable()` function control whether the report is displayed in the browser and whether the report is logged to a file for later analysis.

`cgitb.enable (display=1, logdir=None, context=5, format="html")`

This function causes the `cgitb` module to take over the interpreter's default handling for exceptions by setting the value of `sys.excepthook`.

The optional argument `display` defaults to 1 and can be set to 0 to suppress sending the traceback to the browser. If the argument `logdir` is present, the traceback reports are written to files. The value of `logdir` should be a directory where these files will be placed. The optional argument `context` is the number of lines of context to display around the current line of source code in the traceback; this defaults to 5. If the optional argument `format` is "html", the output is formatted as HTML. Any other value forces plain text output. The default value is "html".

`cgitb.handler (info=None)`

This function handles an exception using the default settings (that is, show a report in the browser, but don't log to a file). This can be used when you've caught an exception and want to report it using `cgitb`. The optional `info` argument should be a 3-tuple containing an exception type, exception value, and traceback object, exactly like the tuple returned by `sys.exc_info()`. If the `info` argument is not supplied, the current exception is obtained from `sys.exc_info()`.

21.4 wsgiref —WSGI Utilities and Reference Implementation

The Web Server Gateway Interface (WSGI) is a standard interface between web server software and web applications written in Python. Having a standard interface makes it easy to use an application that supports WSGI with a number of different web servers.

Only authors of web servers and programming frameworks need to know every detail and corner case of the WSGI design. You don't need to understand every detail of WSGI just to install a WSGI application or to write a web application using an existing framework.

`wsgiref` is a reference implementation of the WSGI specification that can be used to add WSGI support to a web server or framework. It provides utilities for manipulating WSGI environment variables and response headers, base classes for implementing WSGI servers, a demo HTTP server that serves WSGI applications, and a validation tool that checks WSGI servers and applications for conformance to the WSGI specification ([PEP 3333](#)).

See <https://wsgi.readthedocs.org/> for more information about WSGI, and links to tutorials and other resources.

21.4.1 wsgiref.util – WSGI environment utilities

This module provides a variety of utility functions for working with WSGI environments. A WSGI environment is a dictionary containing HTTP request variables as described in [PEP 3333](#). All of the functions taking an *environ* parameter expect a WSGI-compliant dictionary to be supplied; please see [PEP 3333](#) for a detailed specification.

`wsgiref.util.guess_scheme(environ)`

Return a guess for whether `wsgi.url_scheme` should be “http” or “https”, by checking for a HTTPS environment variable in the *environ* dictionary. The return value is a string.

This function is useful when creating a gateway that wraps CGI or a CGI-like protocol such as FastCGI. Typically, servers providing such protocols will include a HTTPS variable with a value of “1”, “yes”, or “on” when a request is received via SSL. So, this function returns “https” if such a value is found, and “http” otherwise.

`wsgiref.util.request_uri(environ, include_query=True)`

Return the full request URI, optionally including the query string, using the algorithm found in the “URL Reconstruction” section of [PEP 3333](#). If *include_query* is false, the query string is not included in the resulting URI.

`wsgiref.util.application_uri(environ)`

Similar to `request_uri()`, except that the `PATH_INFO` and `QUERY_STRING` variables are ignored. The result is the base URI of the application object addressed by the request.

`wsgiref.util.shift_path_info(environ)`

Shift a single name from `PATH_INFO` to `SCRIPT_NAME` and return the name. The *environ* dictionary is *modified* in-place; use a copy if you need to keep the original `PATH_INFO` or `SCRIPT_NAME` intact.

If there are no remaining path segments in `PATH_INFO`, None is returned.

Typically, this routine is used to process each portion of a request URI path, for example to treat the path as a series of dictionary keys. This routine modifies the passed-in environment to make it suitable for invoking another WSGI application that is located at the target URI. For example, if there is a WSGI application at `/foo`, and the request URI path is `/foo/bar/baz`, and the WSGI application at `/foo` calls `shift_path_info()`, it will receive the string “bar”, and the environment will be updated to be suitable for passing to a WSGI application at `/foo/bar`. That is, `SCRIPT_NAME` will change from `/foo` to `/foo/bar`, and `PATH_INFO` will change from `/bar/baz` to `/baz`.

When `PATH_INFO` is just a `"/`, this routine returns an empty string and appends a trailing slash to `SCRIPT_NAME`, even though empty path segments are normally ignored, and `SCRIPT_NAME` doesn't normally end in a slash. This is intentional behavior, to ensure that an application can tell the difference between URIs ending in `/x` from ones ending in `/x/` when using this routine to do object traversal.

`wsgiref.util.setup_testing_defaults(environ)`

Update *environ* with trivial defaults for testing purposes.

This routine adds various parameters required for WSGI, including `HTTP_HOST`, `SERVER_NAME`, `SERVER_PORT`, `REQUEST_METHOD`, `SCRIPT_NAME`, `PATH_INFO`, and all of the [PEP 3333](#)-defined `wsgi.*` variables. It only supplies default values, and does not replace any existing settings for these variables.

This routine is intended to make it easier for unit tests of WSGI servers and applications to set up dummy environments. It should NOT be used by actual WSGI servers or applications, since the data is fake!

Example usage:

```
from wsgiref.util import setup_testing_defaults
from wsgiref.simple_server import make_server

# A relatively simple WSGI application. It's going to print out the
# environment dictionary after being updated by setup_testing_defaults
def simple_app(environ, start_response):
    setup_testing_defaults(environ)

    status = '200 OK'
    headers = [('Content-type', 'text/plain; charset=utf-8')]

    start_response(status, headers)

    ret = [("%s: %s\n" % (key, value)).encode("utf-8")
            for key, value in environ.items()]
    return ret

with make_server('', 8000, simple_app) as httpd:
    print("Serving on port 8000...")
    httpd.serve_forever()
```

In addition to the environment functions above, the `wsgiref.util` module also provides these miscellaneous utilities:

`wsgiref.util.is_hop_by_hop(header_name)`

Return true if `'header_name'` is an HTTP/1.1 “Hop-by-Hop” header, as defined by [RFC 2616](#).

class `wsgiref.util.FileWrapper(filelike, blksize=8192)`

A wrapper to convert a file-like object to an *iterator*. The resulting objects support both `__getitem__()` and `__iter__()` iteration styles, for compatibility with Python 2.1 and Jython. As the object is iterated over, the optional `blksize` parameter will be repeatedly passed to the *filelike* object's `read()` method to obtain bytestrings to yield. When `read()` returns an empty bytestring, iteration is ended and is not resumable.

If *filelike* has a `close()` method, the returned object will also have a `close()` method, and it will invoke the *filelike* object's `close()` method when called.

Example usage:

```
from io import StringIO
from wsgiref.util import FileWrapper

# We're using a StringIO-buffer for as the file-like object
filelike = StringIO("This is an example file-like object"*10)
```

(下页继续)

(续上页)

```
wrapper = FileWrapper(filelike, blksize=5)

for chunk in wrapper:
    print(chunk)
```

21.4.2 wsgiref.headers – WSGI response header tools

This module provides a single class, *Headers*, for convenient manipulation of WSGI response headers using a mapping-like interface.

class wsgiref.headers.*Headers* ([*headers*])

Create a mapping-like object wrapping *headers*, which must be a list of header name/value tuples as described in [PEP 3333](#). The default value of *headers* is an empty list.

Headers objects support typical mapping operations including `__getitem__()`, `get()`, `__setitem__()`, `setdefault()`, `__delitem__()` and `__contains__()`. For each of these methods, the key is the header name (treated case-insensitively), and the value is the first value associated with that header name. Setting a header deletes any existing values for that header, then adds a new value at the end of the wrapped header list. Headers' existing order is generally maintained, with new headers added to the end of the wrapped list.

Unlike a dictionary, *Headers* objects do not raise an error when you try to get or delete a key that isn't in the wrapped header list. Getting a nonexistent header just returns `None`, and deleting a nonexistent header does nothing.

Headers objects also support `keys()`, `values()`, and `items()` methods. The lists returned by `keys()` and `items()` can include the same key more than once if there is a multi-valued header. The `len()` of a *Headers* object is the same as the length of its `items()`, which is the same as the length of the wrapped header list. In fact, the `items()` method just returns a copy of the wrapped header list.

Calling `bytes()` on a *Headers* object returns a formatted bytestring suitable for transmission as HTTP response headers. Each header is placed on a line with its value, separated by a colon and a space. Each line is terminated by a carriage return and line feed, and the bytestring is terminated with a blank line.

In addition to their mapping interface and formatting features, *Headers* objects also have the following methods for querying and adding multi-valued headers, and for adding headers with MIME parameters:

get_all (*name*)

Return a list of all the values for the named header.

The returned list will be sorted in the order they appeared in the original header list or were added to this instance, and may contain duplicates. Any fields deleted and re-inserted are always appended to the header list. If no fields exist with the given name, returns an empty list.

add_header (*name*, *value*, ***_params*)

Add a (possibly multi-valued) header, with optional MIME parameters specified via keyword arguments.

name is the header field to add. Keyword arguments can be used to set MIME parameters for the header field. Each parameter must be a string or `None`. Underscores in parameter names are converted to dashes, since dashes are illegal in Python identifiers, but many MIME parameter names include dashes. If the parameter value is a string, it is added to the header value parameters in the form `name="value"`. If it is `None`, only the parameter name is added. (This is used for MIME parameters without a value.) Example usage:

```
h.add_header('content-disposition', 'attachment', filename='bud.gif')
```

The above will add a header that looks like this:

Content-Disposition: attachment; filename="bud.gif"

在 3.5 版更改: *headers* parameter is optional.

21.4.3 wsgiref.simple_server – a simple WSGI HTTP server

This module implements a simple HTTP server (based on [http.server](#)) that serves WSGI applications. Each server instance serves a single WSGI application on a given host and port. If you want to serve multiple applications on a single host and port, you should create a WSGI application that parses `PATH_INFO` to select which application to invoke for each request. (E.g., using the `shift_path_info()` function from [wsgiref.util](#).)

`wsgiref.simple_server.make_server(host, port, app, server_class=WSGIServer, handler_class=WSGIRequestHandler)`

Create a new WSGI server listening on *host* and *port*, accepting connections for *app*. The return value is an instance of the supplied *server_class*, and will process requests using the specified *handler_class*. *app* must be a WSGI application object, as defined by [PEP 3333](#).

Example usage:

```
from wsgiref.simple_server import make_server, demo_app

with make_server('', 8000, demo_app) as httpd:
    print("Serving HTTP on port 8000...")

    # Respond to requests until process is killed
    httpd.serve_forever()

    # Alternative: serve one request, then exit
    httpd.handle_request()
```

`wsgiref.simple_server.demo_app(environ, start_response)`

This function is a small but complete WSGI application that returns a text page containing the message “Hello world!” and a list of the key/value pairs provided in the *environ* parameter. It’s useful for verifying that a WSGI server (such as [wsgiref.simple_server](#)) is able to run a simple WSGI application correctly.

class `wsgiref.simple_server.WSGIServer(server_address, RequestHandlerClass)`

Create a [WSGIServer](#) instance. *server_address* should be a (host, port) tuple, and *RequestHandlerClass* should be the subclass of [http.server.BaseHTTPRequestHandler](#) that will be used to process requests.

You do not normally need to call this constructor, as the `make_server()` function can handle all the details for you.

[WSGIServer](#) is a subclass of [http.server.HTTPServer](#), so all of its methods (such as `serve_forever()` and `handle_request()`) are available. [WSGIServer](#) also provides these WSGI-specific methods:

set_app(application)

Sets the callable *application* as the WSGI application that will receive requests.

get_app()

Returns the currently-set application callable.

Normally, however, you do not need to use these additional methods, as `set_app()` is normally called by `make_server()`, and the `get_app()` exists mainly for the benefit of request handler instances.

class `wsgiref.simple_server.WSGIRequestHandler(request, client_address, server)`

Create an HTTP handler for the given *request* (i.e. a socket), *client_address* (a (host, port) tuple), and *server* ([WSGIServer](#) instance).

You do not need to create instances of this class directly; they are automatically created as needed by `WSGIServer` objects. You can, however, subclass this class and supply it as a `handler_class` to the `make_server()` function. Some possibly relevant methods for overriding in subclasses:

get_environ()

Returns a dictionary containing the WSGI environment for a request. The default implementation copies the contents of the `WSGIServer` object's `base_environ` dictionary attribute and then adds various headers derived from the HTTP request. Each call to this method should return a new dictionary containing all of the relevant CGI environment variables as specified in [PEP 3333](#).

get_stderr()

Return the object that should be used as the `wsgi.errors` stream. The default implementation just returns `sys.stderr`.

handle()

Process the HTTP request. The default implementation creates a handler instance using a `wsgiref.handlers` class to implement the actual WSGI application interface.

21.4.4 `wsgiref.validate` —WSGI conformance checker

When creating new WSGI application objects, frameworks, servers, or middleware, it can be useful to validate the new code's conformance using `wsgiref.validate`. This module provides a function that creates WSGI application objects that validate communications between a WSGI server or gateway and a WSGI application object, to check both sides for protocol conformance.

Note that this utility does not guarantee complete [PEP 3333](#) compliance; an absence of errors from this module does not necessarily mean that errors do not exist. However, if this module does produce an error, then it is virtually certain that either the server or application is not 100% compliant.

This module is based on the `paste.lint` module from Ian Bicking's "Python Paste" library.

`wsgiref.validate.validator(application)`

Wrap *application* and return a new WSGI application object. The returned application will forward all requests to the original *application*, and will check that both the *application* and the server invoking it are conforming to the WSGI specification and to [RFC 2616](#).

Any detected nonconformance results in an `AssertionError` being raised; note, however, that how these errors are handled is server-dependent. For example, `wsgiref.simple_server` and other servers based on `wsgiref.handlers` (that don't override the error handling methods to do something else) will simply output a message that an error has occurred, and dump the traceback to `sys.stderr` or some other error stream.

This wrapper may also generate output using the `warnings` module to indicate behaviors that are questionable but which may not actually be prohibited by [PEP 3333](#). Unless they are suppressed using Python command-line options or the `warnings` API, any such warnings will be written to `sys.stderr` (not `wsgi.errors`, unless they happen to be the same object).

Example usage:

```
from wsgiref.validate import validator
from wsgiref.simple_server import make_server

# Our callable object which is intentionally not compliant to the
# standard, so the validator is going to break
def simple_app(environ, start_response):
    status = '200 OK' # HTTP Status
    headers = [('Content-type', 'text/plain')] # HTTP Headers
    start_response(status, headers)
```

(下页继续)

(续上页)

```

# This is going to break because we need to return a list, and
# the validator is going to inform us
return b"Hello World"

# This is the application wrapped in a validator
validator_app = validator(simple_app)

with make_server('', 8000, validator_app) as httpd:
    print("Listening on port 8000...")
    httpd.serve_forever()

```

21.4.5 wsgiref.handlers – server/gateway base classes

This module provides base handler classes for implementing WSGI servers and gateways. These base classes handle most of the work of communicating with a WSGI application, as long as they are given a CGI-like environment, along with input, output, and error streams.

class wsgiref.handlers.CGIHandler

CGI-based invocation via `sys.stdin`, `sys.stdout`, `sys.stderr` and `os.environ`. This is useful when you have a WSGI application and want to run it as a CGI script. Simply invoke `CGIHandler().run(app)`, where `app` is the WSGI application object you wish to invoke.

This class is a subclass of `BaseCGIHandler` that sets `wsgi.run_once` to `true`, `wsgi.multithread` to `false`, and `wsgi.multiprocess` to `true`, and always uses `sys` and `os` to obtain the necessary CGI streams and environment.

class wsgiref.handlers.IISCGIHandler

A specialized alternative to `CGIHandler`, for use when deploying on Microsoft's IIS web server, without having set the config `allowPathInfo` option (IIS \geq 7) or metabase `allowPathInfoForScriptMappings` (IIS $<$ 7).

By default, IIS gives a `PATH_INFO` that duplicates the `SCRIPT_NAME` at the front, causing problems for WSGI applications that wish to implement routing. This handler strips any such duplicated path.

IIS can be configured to pass the correct `PATH_INFO`, but this causes another bug where `PATH_TRANSLATED` is wrong. Luckily this variable is rarely used and is not guaranteed by WSGI. On IIS $<$ 7, though, the setting can only be made on a vhost level, affecting all other script mappings, many of which break when exposed to the `PATH_TRANSLATED` bug. For this reason IIS $<$ 7 is almost never deployed with the fix. (Even IIS7 rarely uses it because there is still no UI for it.)

There is no way for CGI code to tell whether the option was set, so a separate handler class is provided. It is used in the same way as `CGIHandler`, i.e., by calling `IISCGIHandler().run(app)`, where `app` is the WSGI application object you wish to invoke.

3.2 新版功能.

class wsgiref.handlers.BaseCGIHandler(*stdin, stdout, stderr, environ, multithread=True, multiprocess=False*)

Similar to `CGIHandler`, but instead of using the `sys` and `os` modules, the CGI environment and I/O streams are specified explicitly. The `multithread` and `multiprocess` values are used to set the `wsgi.multithread` and `wsgi.multiprocess` flags for any applications run by the handler instance.

This class is a subclass of `SimpleHandler` intended for use with software other than HTTP “origin servers”. If you are writing a gateway protocol implementation (such as CGI, FastCGI, SCGI, etc.) that uses a `Status:` header to send an HTTP status, you probably want to subclass this instead of `SimpleHandler`.

class wsgiref.handlers.SimpleHandler(*stdin, stdout, stderr, environ, multithread=True, multiprocess=False*)

Similar to `BaseCGIHandler`, but designed for use with HTTP origin servers. If you are writing an HTTP server implementation, you will probably want to subclass this instead of `BaseCGIHandler`.

This class is a subclass of `BaseHandler`. It overrides the `__init__()`, `get_stdin()`, `get_stderr()`, `add_cgi_vars()`, `_write()`, and `_flush()` methods to support explicitly setting the environment and streams via the constructor. The supplied environment and streams are stored in the `stdin`, `stdout`, `stderr`, and `environ` attributes.

The `write()` method of `stdout` should write each chunk in full, like `io.BufferedIOBase`.

class `wsgiref.handlers.BaseHandler`

This is an abstract base class for running WSGI applications. Each instance will handle a single HTTP request, although in principle you could create a subclass that was reusable for multiple requests.

`BaseHandler` instances have only one method intended for external use:

run (*app*)

Run the specified WSGI application, *app*.

All of the other `BaseHandler` methods are invoked by this method in the process of running the application, and thus exist primarily to allow customizing the process.

The following methods **MUST** be overridden in a subclass:

_write (*data*)

Buffer the bytes *data* for transmission to the client. It's okay if this method actually transmits the data; `BaseHandler` just separates write and flush operations for greater efficiency when the underlying system actually has such a distinction.

_flush ()

Force buffered data to be transmitted to the client. It's okay if this method is a no-op (i.e., if `_write()` actually sends the data).

get_stdin ()

Return an input stream object suitable for use as the `wsgi.input` of the request currently being processed.

get_stderr ()

Return an output stream object suitable for use as the `wsgi.errors` of the request currently being processed.

add_cgi_vars ()

Insert CGI variables for the current request into the `environ` attribute.

Here are some other methods and attributes you may wish to override. This list is only a summary, however, and does not include every method that can be overridden. You should consult the docstrings and source code for additional information before attempting to create a customized `BaseHandler` subclass.

Attributes and methods for customizing the WSGI environment:

wsgi_multithread

The value to be used for the `wsgi.multithread` environment variable. It defaults to true in `BaseHandler`, but may have a different default (or be set by the constructor) in the other subclasses.

wsgi_multiprocess

The value to be used for the `wsgi.multiprocess` environment variable. It defaults to true in `BaseHandler`, but may have a different default (or be set by the constructor) in the other subclasses.

wsgi_run_once

The value to be used for the `wsgi.run_once` environment variable. It defaults to false in `BaseHandler`, but `CGIHandler` sets it to true by default.

os_environ

The default environment variables to be included in every request's WSGI environment. By default, this is

a copy of `os.environ` at the time that `wsgiref.handlers` was imported, but subclasses can either create their own at the class or instance level. Note that the dictionary should be considered read-only, since the default value is shared between multiple classes and instances.

server_software

If the `origin_server` attribute is set, this attribute's value is used to set the default `SERVER_SOFTWARE` WSGI environment variable, and also to set a default `Server:` header in HTTP responses. It is ignored for handlers (such as `BaseCGIHandler` and `CGIHandler`) that are not HTTP origin servers.

在 3.3 版更改: The term “Python” is replaced with implementation specific term like “CPython”, “Jython” etc.

get_scheme()

Return the URL scheme being used for the current request. The default implementation uses the `guess_scheme()` function from `wsgiref.util` to guess whether the scheme should be “http” or “https”, based on the current request's `environ` variables.

setup_environ()

Set the `environ` attribute to a fully-populated WSGI environment. The default implementation uses all of the above methods and attributes, plus the `get_stdin()`, `get_stderr()`, and `add_cgi_vars()` methods and the `wsgi_file_wrapper` attribute. It also inserts a `SERVER_SOFTWARE` key if not present, as long as the `origin_server` attribute is a true value and the `server_software` attribute is set.

Methods and attributes for customizing exception handling:

log_exception(exc_info)

Log the `exc_info` tuple in the server log. `exc_info` is a (type, value, traceback) tuple. The default implementation simply writes the traceback to the request's `wsgi.errors` stream and flushes it. Subclasses can override this method to change the format or retarget the output, mail the traceback to an administrator, or whatever other action may be deemed suitable.

traceback_limit

The maximum number of frames to include in tracebacks output by the default `log_exception()` method. If `None`, all frames are included.

error_output(environ, start_response)

This method is a WSGI application to generate an error page for the user. It is only invoked if an error occurs before headers are sent to the client.

This method can access the current error information using `sys.exc_info()`, and should pass that information to `start_response` when calling it (as described in the “Error Handling” section of [PEP 3333](#)).

The default implementation just uses the `error_status`, `error_headers`, and `error_body` attributes to generate an output page. Subclasses can override this to produce more dynamic error output.

Note, however, that it's not recommended from a security perspective to spit out diagnostics to any old user; ideally, you should have to do something special to enable diagnostic output, which is why the default implementation doesn't include any.

error_status

The HTTP status used for error responses. This should be a status string as defined in [PEP 3333](#); it defaults to a 500 code and message.

error_headers

The HTTP headers used for error responses. This should be a list of WSGI response headers ((name, value) tuples), as described in [PEP 3333](#). The default list just sets the content type to `text/plain`.

error_body

The error response body. This should be an HTTP response body bytestring. It defaults to the plain text, “A server error occurred. Please contact the administrator.”

Methods and attributes for **PEP 333**’s “Optional Platform-Specific File Handling” feature:

wsgi_file_wrapper

A `wsgi.file_wrapper` factory, or None. The default value of this attribute is the `wsgiref.util.FileWrapper` class.

sendfile()

Override to implement platform-specific file transmission. This method is called only if the application’s return value is an instance of the class specified by the `wsgi_file_wrapper` attribute. It should return a true value if it was able to successfully transmit the file, so that the default transmission code will not be executed. The default implementation of this method just returns a false value.

Miscellaneous methods and attributes:

origin_server

This attribute should be set to a true value if the handler’s `_write()` and `_flush()` are being used to communicate directly to the client, rather than via a CGI-like gateway protocol that wants the HTTP status in a special `Status:` header.

This attribute’s default value is true in `BaseHandler`, but false in `BaseCGIHandler` and `CGIHandler`.

http_version

If `origin_server` is true, this string attribute is used to set the HTTP version of the response set to the client. It defaults to “1.0”.

`wsgiref.handlers.read_environ()`

Transcode CGI variables from `os.environ` to PEP 333 “bytes in unicode” strings, returning a new dictionary. This function is used by `CGIHandler` and `IISCGIHandler` in place of directly using `os.environ`, which is not necessarily WSGI-compliant on all platforms and web servers using Python 3 – specifically, ones where the OS’s actual environment is Unicode (i.e. Windows), or ones where the environment is bytes, but the system encoding used by Python to decode it is anything other than ISO-8859-1 (e.g. Unix systems using UTF-8).

If you are implementing a CGI-based handler of your own, you probably want to use this routine instead of just copying values out of `os.environ` directly.

3.2 新版功能.

21.4.6 示例

This is a working “Hello World” WSGI application:

```
from wsgiref.simple_server import make_server

# Every WSGI application must have an application object - a callable
# object that accepts two arguments. For that purpose, we're going to
# use a function (note that you're not limited to a function, you can
# use a class for example). The first argument passed to the function
# is a dictionary containing CGI-style environment variables and the
# second variable is the callable object (see PEP 333).
def hello_world_app(environ, start_response):
    status = '200 OK' # HTTP Status
    headers = [('Content-type', 'text/plain; charset=utf-8')] # HTTP Headers
    start_response(status, headers)
```

(下页继续)

(续上页)

```
# The returned object is going to be printed
return [b"Hello World"]

with make_server('', 8000, hello_world_app) as httpd:
    print("Serving on port 8000...")

    # Serve until process is killed
    httpd.serve_forever()
```

21.5 urllib — URL 处理模块

源代码: [Lib/urllib/](#)

`urllib` 是一个收集了多个使用 URL 的模块的软件包:

- `urllib.request` 打开和阅读 URLs
- `urllib.error` 包含 `urllib.request` 抛出的异常
- `urllib.parse` 用于处理 URL
- `urllib.robotparser` 用于解析 `robots.txt` 文件

21.6 urllib.request — 用于打开 URL 的可扩展库

源码: [Lib/urllib/request.py](#)

`urllib.request` 模块定义了适用于在各种复杂情况下打开 URL (主要为 HTTP) 的函数和类—例如基本和精简验证, 重定向、cookies 及其它。

参见:

`Requests` 包 推荐用于更高层级的 HTTP 客户端接口。

The `urllib.request` module defines the following functions:

`urllib.request.urlopen` (*url*, *data=None* [, *timeout*], *, *cafile=None*, *capath=None*, *cadefault=False*, *context=None*)

Open the URL *url*, which can be either a string or a *Request* object.

data must be an object specifying additional data to be sent to the server, or *None* if no such data is needed. See *Request* for details.

`urllib.request` module uses HTTP/1.1 and includes `Connection:close` header in its HTTP requests.

The optional *timeout* parameter specifies a timeout in seconds for blocking operations like the connection attempt (if not specified, the global default timeout setting will be used). This actually only works for HTTP, HTTPS and FTP connections.

If *context* is specified, it must be a `ssl.SSLContext` instance describing the various SSL options. See *HTTPConnection* for more details.

The optional *cafile* and *capath* parameters specify a set of trusted CA certificates for HTTPS requests. *cafile* should point to a single file containing a bundle of CA certificates, whereas *capath* should point to a directory of hashed certificate files. More information can be found in `ssl.SSLContext.load_verify_locations()`.

The *cadefault* parameter is ignored.

This function always returns an object which can work as a *context manager* and has methods such as

- `geturl()` —return the URL of the resource retrieved, commonly used to determine if a redirect was followed
- `info()` —return the meta-information of the page, such as headers, in the form of an *email.message_from_string()* instance (see [Quick Reference to HTTP Headers](#))
- `getcode()` —return the HTTP status code of the response.

For HTTP and HTTPS URLs, this function returns a `http.client.HTTPResponse` object slightly modified. In addition to the three new methods above, the *msg* attribute contains the same information as the *reason* attribute —the reason phrase returned by server —instead of the response headers as it is specified in the documentation for *HTTPResponse*.

For FTP, file, and data URLs and requests explicitly handled by legacy *URLopener* and *FancyURLopener* classes, this function returns a `urllib.response.addinfourl` object.

Raises *URLError* on protocol errors.

Note that *None* may be returned if no handler handles the request (though the default installed global *OpenerDirector* uses *UnknownHandler* to ensure this never happens).

In addition, if proxy settings are detected (for example, when a `*_proxy` environment variable like `http_proxy` is set), *ProxyHandler* is default installed and makes sure the requests are handled through the proxy.

The legacy `urllib.urlopen` function from Python 2.6 and earlier has been discontinued; `urllib.request.urlopen()` corresponds to the old `urllib2.urlopen`. Proxy handling, which was done by passing a dictionary parameter to `urllib.urlopen`, can be obtained by using *ProxyHandler* objects.

在 3.2 版更改: *cafile* 和 *capath* were added.

在 3.2 版更改: HTTPS virtual hosts are now supported if possible (that is, if `ssl.HAS_SNI` is true).

3.2 新版功能: *data* can be an iterable object.

在 3.3 版更改: *cadefault* was added.

在 3.4.3 版更改: *context* was added.

3.6 版后已移除: *cafile*, *capath* and *cadefault* are deprecated in favor of *context*. Please use `ssl.SSLContext.load_cert_chain()` instead, or let `ssl.create_default_context()` select the system's trusted CA certificates for you.

`urllib.request.install_opener(opener)`

Install an *OpenerDirector* instance as the default global opener. Installing an opener is only necessary if you want `urlopen` to use that opener; otherwise, simply call `OpenerDirector.open()` instead of `urlopen()`. The code does not check for a real *OpenerDirector*, and any class with the appropriate interface will work.

`urllib.request.build_opener([handler, ...])`

Return an *OpenerDirector* instance, which chains the handlers in the order given. *handlers* can be either instances of *BaseHandler*, or subclasses of *BaseHandler* (in which case it must be possible to call the constructor without any parameters). Instances of the following classes will be in front of the *handlers*, unless the *handlers* contain them, instances of them or subclasses of them: *ProxyHandler* (if proxy settings are detected), *UnknownHandler*, *HTTPHandler*, *HTTPDefaultErrorHandler*, *HTTPRedirectHandler*, *FTPHandler*, *FileHandler*, *HTTPErrorProcessor*.

If the Python installation has SSL support (i.e., if the `ssl` module can be imported), `HTTPSHandler` will also be added.

A `BaseHandler` subclass may also change its `handler_order` attribute to modify its position in the handlers list.

`urllib.request.pathname2url(path)`

Convert the pathname *path* from the local syntax for a path to the form used in the path component of a URL. This does not produce a complete URL. The return value will already be quoted using the `quote()` function.

`urllib.request.url2pathname(path)`

Convert the path component *path* from a percent-encoded URL to the local syntax for a path. This does not accept a complete URL. This function uses `unquote()` to decode *path*.

`urllib.request.getproxies()`

This helper function returns a dictionary of scheme to proxy server URL mappings. It scans the environment for variables named `<scheme>_proxy`, in a case insensitive approach, for all operating systems first, and when it cannot find it, looks for proxy information from Mac OSX System Configuration for Mac OS X and Windows Systems Registry for Windows. If both lowercase and uppercase environment variables exist (and disagree), lowercase is preferred.

注解: If the environment variable `REQUEST_METHOD` is set, which usually indicates your script is running in a CGI environment, the environment variable `HTTP_PROXY` (uppercase `_PROXY`) will be ignored. This is because that variable can be injected by a client using the “Proxy:” HTTP header. If you need to use an HTTP proxy in a CGI environment, either use `ProxyHandler` explicitly, or make sure the variable name is in lowercase (or at least the `_proxy` suffix).

The following classes are provided:

class `urllib.request.Request(url, data=None, headers={}, origin_req_host=None, unverifiable=False, method=None)`

This class is an abstraction of a URL request.

url should be a string containing a valid URL.

data must be an object specifying additional data to send to the server, or `None` if no such data is needed. Currently HTTP requests are the only ones that use *data*. The supported object types include bytes, file-like objects, and iterables. If no `Content-Length` nor `Transfer-Encoding` header field has been provided, `HTTPHandler` will set these headers according to the type of *data*. `Content-Length` will be used to send bytes objects, while `Transfer-Encoding: chunked` as specified in [RFC 7230](#), Section 3.3.1 will be used to send files and other iterables.

For an HTTP POST request method, *data* should be a buffer in the standard `application/x-www-form-urlencoded` format. The `urllib.parse.urlencode()` function takes a mapping or sequence of 2-tuples and returns an ASCII string in this format. It should be encoded to bytes before being used as the *data* parameter.

headers should be a dictionary, and will be treated as if `add_header()` was called with each key and value as arguments. This is often used to “spoof” the User-Agent header value, which is used by a browser to identify itself – some HTTP servers only allow requests coming from common browsers as opposed to scripts. For example, Mozilla Firefox may identify itself as “Mozilla/5.0 (X11; U; Linux i686) Gecko/20071127 Firefox/2.0.0.11”, while `urllib`’s default user agent string is “Python-urllib/2.6” (on Python 2.6).

An appropriate `Content-Type` header should be included if the *data* argument is present. If this header has not been provided and *data* is not `None`, `Content-Type: application/x-www-form-urlencoded` will be added as a default.

The final two arguments are only of interest for correct handling of third-party HTTP cookies:

origin_req_host should be the request-host of the origin transaction, as defined by [RFC 2965](#). It defaults to `http.cookiejar.request_host(self)`. This is the host name or IP address of the original request that was initiated by the user. For example, if the request is for an image in an HTML document, this should be the request-host of the request for the page containing the image.

unverifiable should indicate whether the request is unverifiable, as defined by [RFC 2965](#). It defaults to `False`. An unverifiable request is one whose URL the user did not have the option to approve. For example, if the request is for an image in an HTML document, and the user had no option to approve the automatic fetching of the image, this should be `true`.

method should be a string that indicates the HTTP request method that will be used (e.g. `'HEAD'`). If provided, its value is stored in the *method* attribute and is used by *get_method()*. The default is `'GET'` if *data* is `None` or `'POST'` otherwise. Subclasses may indicate a different default method by setting the *method* attribute in the class itself.

注解： The request will not work as expected if the data object is unable to deliver its content more than once (e.g. a file or an iterable that can produce the content only once) and the request is retried for HTTP redirects or authentication. The *data* is sent to the HTTP server right away after the headers. There is no support for a 100-continue expectation in the library.

在 3.3 版更改: *Request.method* argument is added to the Request class.

在 3.4 版更改: Default *Request.method* may be indicated at the class level.

在 3.6 版更改: Do not raise an error if the `Content-Length` has not been provided and *data* is neither `None` nor a bytes object. Fall back to use chunked transfer encoding instead.

class `urllib.request.OpenerDirector`

The *OpenerDirector* class opens URLs via *BaseHandlers* chained together. It manages the chaining of handlers, and recovery from errors.

class `urllib.request.BaseHandler`

This is the base class for all registered handlers —and handles only the simple mechanics of registration.

class `urllib.request.HTTPDefaultErrorHandler`

A class which defines a default handler for HTTP error responses; all responses are turned into *HTTPError* exceptions.

class `urllib.request.HTTPRedirectHandler`

A class to handle redirections.

class `urllib.request.HTTPCookieProcessor` (*cookiejar=None*)

A class to handle HTTP Cookies.

class `urllib.request.ProxyHandler` (*proxies=None*)

Cause requests to go through a proxy. If *proxies* is given, it must be a dictionary mapping protocol names to URLs of proxies. The default is to read the list of proxies from the environment variables `<protocol>_proxy`. If no proxy environment variables are set, then in a Windows environment proxy settings are obtained from the registry's Internet Settings section, and in a Mac OS X environment proxy information is retrieved from the OS X System Configuration Framework.

To disable autodetected proxy pass an empty dictionary.

The `no_proxy` environment variable can be used to specify hosts which shouldn't be reached via proxy; if set, it should be a comma-separated list of hostname suffixes, optionally with `:port` appended, for example `cern.ch, ncsa.uiuc.edu, some.host:8080`.

注解: HTTP_PROXY will be ignored if a variable REQUEST_METHOD is set; see the documentation on [getproxies\(\)](#).

class urllib.request.HTTPPasswordMgr

Keep a database of (realm, uri) -> (user, password) mappings.

class urllib.request.HTTPPasswordMgrWithDefaultRealm

Keep a database of (realm, uri) -> (user, password) mappings. A realm of None is considered a catch-all realm, which is searched if no other realm fits.

class urllib.request.HTTPPasswordMgrWithPriorAuth

A variant of [HTTPPasswordMgrWithDefaultRealm](#) that also has a database of uri -> is_authenticated mappings. Can be used by a BasicAuth handler to determine when to send authentication credentials immediately instead of waiting for a 401 response first.

3.5 新版功能.

class urllib.request.AbstractBasicAuthHandler (password_mgr=None)

This is a mixin class that helps with HTTP authentication, both to the remote host and to a proxy. *password_mgr*, if given, should be something that is compatible with [HTTPPasswordMgr](#); refer to section [HTTPPasswordMgr Objects](#) for information on the interface that must be supported. If *password_mgr* also provides *is_authenticated* and *update_authenticated* methods (see [HTTPPasswordMgrWithPriorAuth Objects](#)), then the handler will use the *is_authenticated* result for a given URI to determine whether or not to send authentication credentials with the request. If *is_authenticated* returns True for the URI, credentials are sent. If *is_authenticated* is False, credentials are not sent, and then if a 401 response is received the request is re-sent with the authentication credentials. If authentication succeeds, *update_authenticated* is called to set *is_authenticated* True for the URI, so that subsequent requests to the URI or any of its super-URIs will automatically include the authentication credentials.

3.5 新版功能: Added *is_authenticated* support.

class urllib.request.HTTPBasicAuthHandler (password_mgr=None)

Handle authentication with the remote host. *password_mgr*, if given, should be something that is compatible with [HTTPPasswordMgr](#); refer to section [HTTPPasswordMgr Objects](#) for information on the interface that must be supported. HTTPBasicAuthHandler will raise a [ValueError](#) when presented with a wrong Authentication scheme.

class urllib.request.ProxyBasicAuthHandler (password_mgr=None)

Handle authentication with the proxy. *password_mgr*, if given, should be something that is compatible with [HTTPPasswordMgr](#); refer to section [HTTPPasswordMgr Objects](#) for information on the interface that must be supported.

class urllib.request.AbstractDigestAuthHandler (password_mgr=None)

This is a mixin class that helps with HTTP authentication, both to the remote host and to a proxy. *password_mgr*, if given, should be something that is compatible with [HTTPPasswordMgr](#); refer to section [HTTPPasswordMgr Objects](#) for information on the interface that must be supported.

class urllib.request.HTTPDigestAuthHandler (password_mgr=None)

Handle authentication with the remote host. *password_mgr*, if given, should be something that is compatible with [HTTPPasswordMgr](#); refer to section [HTTPPasswordMgr Objects](#) for information on the interface that must be supported. When both Digest Authentication Handler and Basic Authentication Handler are both added, Digest Authentication is always tried first. If the Digest Authentication returns a 40x response again, it is sent to Basic Authentication handler to Handle. This Handler method will raise a [ValueError](#) when presented with an authentication scheme other than Digest or Basic.

在 3.3 版更改: Raise [ValueError](#) on unsupported Authentication Scheme.

class urllib.request.**ProxyDigestAuthHandler** (*password_mgr=None*)
 Handle authentication with the proxy. *password_mgr*, if given, should be something that is compatible with *HTTPPasswordMgr*; refer to section *HTTPPasswordMgr Objects* for information on the interface that must be supported.

class urllib.request.**HTTPHandler**
 A class to handle opening of HTTP URLs.

class urllib.request.**HTTPSHandler** (*debuglevel=0, context=None, check_hostname=None*)
 A class to handle opening of HTTPS URLs. *context* and *check_hostname* have the same meaning as in *http.client.HTTPSConnection*.
 在 3.2 版更改: *context* and *check_hostname* were added.

class urllib.request.**FileHandler**
 Open local files.

class urllib.request.**DataHandler**
 Open data URLs.
 3.4 新版功能.

class urllib.request.**FTPHandler**
 Open FTP URLs.

class urllib.request.**CacheFTPHandler**
 Open FTP URLs, keeping a cache of open FTP connections to minimize delays.

class urllib.request.**UnknownHandler**
 A catch-all class to handle unknown URLs.

class urllib.request.**HTTPErrorProcessor**
 Process HTTP error responses.

21.6.1 Request Objects

The following methods describe *Request*'s public interface, and so all may be overridden in subclasses. It also defines several public attributes that can be used by clients to inspect the parsed request.

Request.**full_url**

The original URL passed to the constructor.

在 3.4 版更改.

Request.full_url is a property with setter, getter and a deleter. Getting *full_url* returns the original request URL with the fragment, if it was present.

Request.**type**

The URI scheme.

Request.**host**

The URI authority, typically a host, but may also contain a port separated by a colon.

Request.**origin_req_host**

The original host for the request, without port.

Request.**selector**

The URI path. If the *Request* uses a proxy, then selector will be the full URL that is passed to the proxy.

Request.**data**

The entity body for the request, or None if not specified.

在 3.4 版更改: Changing value of `Request.data` now deletes “Content-Length” header if it was previously set or calculated.

`Request.unverifiable`

boolean, indicates whether the request is unverifiable as defined by [RFC 2965](#).

`Request.method`

The HTTP request method to use. By default its value is `None`, which means that `get_method()` will do its normal computation of the method to be used. Its value can be set (thus overriding the default computation in `get_method()`) either by providing a default value by setting it at the class level in a `Request` subclass, or by passing a value in to the `Request` constructor via the `method` argument.

3.3 新版功能.

在 3.4 版更改: A default value can now be set in subclasses; previously it could only be set via the constructor argument.

`Request.get_method()`

Return a string indicating the HTTP request method. If `Request.method` is not `None`, return its value, otherwise return 'GET' if `Request.data` is `None`, or 'POST' if it's not. This is only meaningful for HTTP requests.

在 3.3 版更改: `get_method` now looks at the value of `Request.method`.

`Request.add_header(key, val)`

Add another header to the request. Headers are currently ignored by all handlers except HTTP handlers, where they are added to the list of headers sent to the server. Note that there cannot be more than one header with the same name, and later calls will overwrite previous calls in case the `key` collides. Currently, this is no loss of HTTP functionality, since all headers which have meaning when used more than once have a (header-specific) way of gaining the same functionality using only one header.

`Request.add_unredirected_header(key, header)`

Add a header that will not be added to a redirected request.

`Request.has_header(header)`

Return whether the instance has the named header (checks both regular and unredirected).

`Request.remove_header(header)`

Remove named header from the request instance (both from regular and unredirected headers).

3.4 新版功能.

`Request.get_full_url()`

Return the URL given in the constructor.

在 3.4 版更改.

Returns `Request.full_url`

`Request.set_proxy(host, type)`

Prepare the request by connecting to a proxy server. The `host` and `type` will replace those of the instance, and the instance's selector will be the original URL given in the constructor.

`Request.get_header(header_name, default=None)`

Return the value of the given header. If the header is not present, return the default value.

`Request.header_items()`

Return a list of tuples (header_name, header_value) of the Request headers.

在 3.4 版更改: The request methods `add_data`, `has_data`, `get_data`, `get_type`, `get_host`, `get_selector`, `get_origin_req_host` and `is_unverifiable` that were deprecated since 3.3 have been removed.

21.6.2 OpenerDirector Objects

OpenerDirector instances have the following methods:

OpenerDirector.add_handler(handler)

handler should be an instance of *BaseHandler*. The following methods are searched, and added to the possible chains (note that HTTP errors are a special case).

- *protocol_open()* —signal that the handler knows how to open *protocol* URLs.
- *http_error_type()* —signal that the handler knows how to handle HTTP errors with HTTP error code *type*.
- *protocol_error()* —signal that the handler knows how to handle errors from (non-http) *protocol*.
- *protocol_request()* —signal that the handler knows how to pre-process *protocol* requests.
- *protocol_response()* —signal that the handler knows how to post-process *protocol* responses.

OpenerDirector.open(url, data=None[, timeout])

Open the given *url* (which can be a request object or a string), optionally passing the given *data*. Arguments, return values and exceptions raised are the same as those of *urlopen()* (which simply calls the *open()* method on the currently installed global *OpenerDirector*). The optional *timeout* parameter specifies a timeout in seconds for blocking operations like the connection attempt (if not specified, the global default timeout setting will be used). The timeout feature actually works only for HTTP, HTTPS and FTP connections).

*OpenerDirector.error(proto, *args)*

Handle an error of the given protocol. This will call the registered error handlers for the given protocol with the given arguments (which are protocol specific). The HTTP protocol is a special case which uses the HTTP response code to determine the specific error handler; refer to the *http_error_*()* methods of the handler classes.

Return values and exceptions raised are the same as those of *urlopen()*.

OpenerDirector objects open URLs in three stages:

The order in which these methods are called within each stage is determined by sorting the handler instances.

1. Every handler with a method named like *protocol_request()* has that method called to pre-process the request.
2. Handlers with a method named like *protocol_open()* are called to handle the request. This stage ends when a handler either returns a non-*None* value (ie. a response), or raises an exception (usually *URLLError*). Exceptions are allowed to propagate.

In fact, the above algorithm is first tried for methods named *default_open()*. If all such methods return *None*, the algorithm is repeated for methods named like *protocol_open()*. If all such methods return *None*, the algorithm is repeated for methods named *unknown_open()*.

Note that the implementation of these methods may involve calls of the parent *OpenerDirector* instance's *open()* and *error()* methods.

3. Every handler with a method named like *protocol_response()* has that method called to post-process the response.

21.6.3 BaseHandler Objects

BaseHandler objects provide a couple of methods that are directly useful, and others that are meant to be used by derived classes. These are intended for direct use:

BaseHandler.add_parent(director)

Add a director as parent.

`BaseHandler.close()`

Remove any parents.

The following attribute and methods should only be used by classes derived from *BaseHandler*.

注解: The convention has been adopted that subclasses defining `protocol_request()` or `protocol_response()` methods are named **Processor*; all others are named **Handler*.

`BaseHandler.parent`

A valid *OpenerDirector*, which can be used to open using a different protocol, or handle errors.

`BaseHandler.default_open(req)`

This method is *not* defined in *BaseHandler*, but subclasses should define it if they want to catch all URLs.

This method, if implemented, will be called by the parent *OpenerDirector*. It should return a file-like object as described in the return value of the `open()` of *OpenerDirector*, or *None*. It should raise *URLError*, unless a truly exceptional thing happens (for example, *MemoryError* should not be mapped to *URLError*).

This method will be called before any protocol-specific open method.

`BaseHandler.protocol_open(req)`

This method is *not* defined in *BaseHandler*, but subclasses should define it if they want to handle URLs with the given protocol.

This method, if defined, will be called by the parent *OpenerDirector*. Return values should be the same as for `default_open()`.

`BaseHandler.unknown_open(req)`

This method is *not* defined in *BaseHandler*, but subclasses should define it if they want to catch all URLs with no specific registered handler to open it.

This method, if implemented, will be called by the *parent OpenerDirector*. Return values should be the same as for `default_open()`.

`BaseHandler.http_error_default(req, fp, code, msg, hdrs)`

This method is *not* defined in *BaseHandler*, but subclasses should override it if they intend to provide a catch-all for otherwise unhandled HTTP errors. It will be called automatically by the *OpenerDirector* getting the error, and should not normally be called in other circumstances.

req will be a *Request* object, *fp* will be a file-like object with the HTTP error body, *code* will be the three-digit code of the error, *msg* will be the user-visible explanation of the code and *hdrs* will be a mapping object with the headers of the error.

Return values and exceptions raised should be the same as those of `urlopen()`.

`BaseHandler.http_error_nnn(req, fp, code, msg, hdrs)`

nnn should be a three-digit HTTP error code. This method is also not defined in *BaseHandler*, but will be called, if it exists, on an instance of a subclass, when an HTTP error with code *nnn* occurs.

Subclasses should override this method to handle specific HTTP errors.

Arguments, return values and exceptions raised should be the same as for `http_error_default()`.

`BaseHandler.protocol_request(req)`

This method is *not* defined in *BaseHandler*, but subclasses should define it if they want to pre-process requests of the given protocol.

This method, if defined, will be called by the parent *OpenerDirector*. *req* will be a *Request* object. The return value should be a *Request* object.

`BaseHandler.protocol_response (req, response)`

This method is *not* defined in `BaseHandler`, but subclasses should define it if they want to post-process responses of the given protocol.

This method, if defined, will be called by the parent `OpenerDirector`. `req` will be a `Request` object. `response` will be an object implementing the same interface as the return value of `urlopen()`. The return value should implement the same interface as the return value of `urlopen()`.

21.6.4 HTTPRedirectHandler Objects

注解: Some HTTP redirections require action from this module's client code. If this is the case, `HTTPError` is raised. See [RFC 2616](#) for details of the precise meanings of the various redirection codes.

An `HTTPError` exception raised as a security consideration if the `HTTPRedirectHandler` is presented with a redirected URL which is not an HTTP, HTTPS or FTP URL.

`HTTPRedirectHandler.redirect_request (req, fp, code, msg, hdrs, newurl)`

Return a `Request` or `None` in response to a redirect. This is called by the default implementations of the `http_error_30*`() methods when a redirection is received from the server. If a redirection should take place, return a new `Request` to allow `http_error_30*`() to perform the redirect to `newurl`. Otherwise, raise `HTTPError` if no other handler should try to handle this URL, or return `None` if you can't but another handler might.

注解: The default implementation of this method does not strictly follow [RFC 2616](#), which says that 301 and 302 responses to POST requests must not be automatically redirected without confirmation by the user. In reality, browsers do allow automatic redirection of these responses, changing the POST to a GET, and the default implementation reproduces this behavior.

`HTTPRedirectHandler.http_error_301 (req, fp, code, msg, hdrs)`

Redirect to the Location: or URI: URL. This method is called by the parent `OpenerDirector` when getting an HTTP 'moved permanently' response.

`HTTPRedirectHandler.http_error_302 (req, fp, code, msg, hdrs)`

The same as `http_error_301()`, but called for the 'found' response.

`HTTPRedirectHandler.http_error_303 (req, fp, code, msg, hdrs)`

The same as `http_error_301()`, but called for the 'see other' response.

`HTTPRedirectHandler.http_error_307 (req, fp, code, msg, hdrs)`

The same as `http_error_301()`, but called for the 'temporary redirect' response.

21.6.5 HTTPCookieProcessor Objects

`HTTPCookieProcessor` instances have one attribute:

`HTTPCookieProcessor.cookiejar`

The `http.cookiejar.CookieJar` in which cookies are stored.

21.6.6 ProxyHandler Objects

`ProxyHandler.protocol_open (request)`

The `ProxyHandler` will have a method `protocol_open()` for every `protocol` which has a proxy in the

proxies dictionary given in the constructor. The method will modify requests to go through the proxy, by calling `request.set_proxy()`, and call the next handler in the chain to actually execute the protocol.

21.6.7 HTTPPasswordMgr Objects

These methods are available on `HTTPPasswordMgr` and `HTTPPasswordMgrWithDefaultRealm` objects.

`HTTPPasswordMgr.add_password(realm, uri, user, passwd)`

uri can be either a single URI, or a sequence of URIs. *realm*, *user* and *passwd* must be strings. This causes (*user*, *passwd*) to be used as authentication tokens when authentication for *realm* and a super-URI of any of the given URIs is given.

`HTTPPasswordMgr.find_user_password(realm, authuri)`

Get user/password for given realm and URI, if any. This method will return (`None`, `None`) if there is no matching user/password.

For `HTTPPasswordMgrWithDefaultRealm` objects, the realm `None` will be searched if the given *realm* has no matching user/password.

21.6.8 HTTPPasswordMgrWithPriorAuth Objects

This password manager extends `HTTPPasswordMgrWithDefaultRealm` to support tracking URIs for which authentication credentials should always be sent.

`HTTPPasswordMgrWithPriorAuth.add_password(realm, uri, user, passwd, is_authenticated=False)`

realm, *uri*, *user*, *passwd* are as for `HTTPPasswordMgr.add_password()`. *is_authenticated* sets the initial value of the *is_authenticated* flag for the given URI or list of URIs. If *is_authenticated* is specified as `True`, *realm* is ignored.

`HTTPPasswordMgrWithPriorAuth.find_user_password(realm, authuri)`

Same as for `HTTPPasswordMgrWithDefaultRealm` objects

`HTTPPasswordMgrWithPriorAuth.update_authenticated(self, uri, is_authenticated=False)`

Update the *is_authenticated* flag for the given *uri* or list of URIs.

`HTTPPasswordMgrWithPriorAuth.is_authenticated(self, authuri)`

Returns the current state of the *is_authenticated* flag for the given URI.

21.6.9 AbstractBasicAuthHandler Objects

`AbstractBasicAuthHandler.http_error_auth_reqed(authreq, host, req, headers)`

Handle an authentication request by getting a user/password pair, and re-trying the request. *authreq* should be the name of the header where the information about the realm is included in the request, *host* specifies the URL and path to authenticate for, *req* should be the (failed) `Request` object, and *headers* should be the error headers.

host is either an authority (e.g. "python.org") or a URL containing an authority component (e.g. "http://python.org/"). In either case, the authority must not contain a userinfo component (so, "python.org" and "python.org:80" are fine, "joe:password@python.org" is not).

21.6.10 HTTPBasicAuthHandler Objects

`HTTPBasicAuthHandler.http_error_401(req, fp, code, msg, hdrs)`

Retry the request with authentication information, if available.

21.6.11 ProxyBasicAuthHandler Objects

`ProxyBasicAuthHandler.http_error_407` (*req, fp, code, msg, hdrs*)
 Retry the request with authentication information, if available.

21.6.12 AbstractDigestAuthHandler Objects

`AbstractDigestAuthHandler.http_error_auth_reged` (*authreq, host, req, headers*)
authreq should be the name of the header where the information about the realm is included in the request, *host* should be the host to authenticate to, *req* should be the (failed) *Request* object, and *headers* should be the error headers.

21.6.13 HTTPDigestAuthHandler Objects

`HTTPDigestAuthHandler.http_error_401` (*req, fp, code, msg, hdrs*)
 Retry the request with authentication information, if available.

21.6.14 ProxyDigestAuthHandler Objects

`ProxyDigestAuthHandler.http_error_407` (*req, fp, code, msg, hdrs*)
 Retry the request with authentication information, if available.

21.6.15 HTTPHandler Objects

`HTTPHandler.http_open` (*req*)
 Send an HTTP request, which can be either GET or POST, depending on `req.has_data()`.

21.6.16 HTTPSHandler Objects

`HTTPSHandler.https_open` (*req*)
 Send an HTTPS request, which can be either GET or POST, depending on `req.has_data()`.

21.6.17 FileHandler Objects

`FileHandler.file_open` (*req*)
 Open the file locally, if there is no host name, or the host name is 'localhost'.
 在 3.2 版更改: This method is applicable only for local hostnames. When a remote hostname is given, an *URLError* is raised.

21.6.18 DataHandler Objects

`DataHandler.data_open` (*req*)
 Read a data URL. This kind of URL contains the content encoded in the URL itself. The data URL syntax is specified in [RFC 2397](#). This implementation ignores white spaces in base64 encoded data URLs so the URL may be wrapped in whatever source file it comes from. But even though some browsers don't mind about a missing padding at the end of a base64 encoded data URL, this implementation will raise an *ValueError* in that case.

21.6.19 FTPHandler Objects

`FTPHandler.ftp_open(req)`

Open the FTP file indicated by *req*. The login is always done with empty username and password.

21.6.20 CacheFTPHandler Objects

CacheFTPHandler objects are *FTPHandler* objects with the following additional methods:

`CacheFTPHandler.setTimeout(t)`

Set timeout of connections to *t* seconds.

`CacheFTPHandler.setMaxConns(m)`

Set maximum number of cached connections to *m*.

21.6.21 UnknownHandler Objects

`UnknownHandler.unknown_open()`

Raise a *URLError* exception.

21.6.22 HTTPErrorProcessor Objects

`HTTPErrorProcessor.http_response(request, response)`

Process HTTP error responses.

For 200 error codes, the response object is returned immediately.

For non-200 error codes, this simply passes the job on to the `protocol_error_code()` handler methods, via `OpenerDirector.error()`. Eventually, *HTTPDefaultErrorHandler* will raise an *HTTPError* if no other handler handles the error.

`HTTPErrorProcessor.https_response(request, response)`

Process HTTPS error responses.

The behavior is same as `http_response()`.

21.6.23 示例

In addition to the examples below, more examples are given in `urllib-howto`.

This example gets the `python.org` main page and displays the first 300 bytes of it.

```
>>> import urllib.request
>>> with urllib.request.urlopen('http://www.python.org/') as f:
...     print(f.read(300))
...
b'<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Transitional//EN"
"http://www.w3.org/TR/xhtml1/DTD/xhtml1-transitional.dtd">\n\n\n<html
xmlns="http://www.w3.org/1999/xhtml" xml:lang="en" lang="en">\n\n<head>\n
<meta http-equiv="content-type" content="text/html; charset=utf-8" />\n
<title>Python Programming '
```

Note that `urlopen` returns a bytes object. This is because there is no way for `urlopen` to automatically determine the encoding of the byte stream it receives from the HTTP server. In general, a program will decode the returned bytes object to string once it determines or guesses the appropriate encoding.

The following W3C document, <https://www.w3.org/International/O-charset>, lists the various ways in which an (X)HTML or an XML document could have specified its encoding information.

As the `python.org` website uses *utf-8* encoding as specified in its meta tag, we will use the same for decoding the bytes object.

```
>>> with urllib.request.urlopen('http://www.python.org/') as f:
...     print(f.read(100).decode('utf-8'))
...
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Transitional//EN"
"http://www.w3.org/TR/xhtml1/DTD/xhtml"
```

It is also possible to achieve the same result without using the *context manager* approach.

```
>>> import urllib.request
>>> f = urllib.request.urlopen('http://www.python.org/')
>>> print(f.read(100).decode('utf-8'))
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Transitional//EN"
"http://www.w3.org/TR/xhtml1/DTD/xhtml"
```

In the following example, we are sending a data-stream to the stdin of a CGI and reading the data it returns to us. Note that this example will only work when the Python installation supports SSL.

```
>>> import urllib.request
>>> req = urllib.request.Request(url='https://localhost/cgi-bin/test.cgi',
...                             data=b'This data is passed to stdin of the CGI')
>>> with urllib.request.urlopen(req) as f:
...     print(f.read().decode('utf-8'))
...
Got Data: "This data is passed to stdin of the CGI"
```

The code for the sample CGI used in the above example is:

```
#!/usr/bin/env python
import sys
data = sys.stdin.read()
print('Content-type: text/plain\n\nGot Data: "%s"' % data)
```

Here is an example of doing a PUT request using *Request*:

```
import urllib.request
DATA = b'some data'
req = urllib.request.Request(url='http://localhost:8080', data=DATA, method='PUT')
with urllib.request.urlopen(req) as f:
    pass
print(f.status)
print(f.reason)
```

Use of Basic HTTP Authentication:

```
import urllib.request
# Create an OpenerDirector with support for Basic HTTP Authentication...
auth_handler = urllib.request.HTTPBasicAuthHandler()
auth_handler.add_password(realm='PDQ Application',
```

(下页继续)

(续上页)

```

        uri='https://mahler:8092/site-updates.py',
        user='klem',
        passwd='kadidd!ehopper')
opener = urllib.request.build_opener(auth_handler)
# ...and install it globally so it can be used with urlopen.
urllib.request.install_opener(opener)
urllib.request.urlopen('http://www.example.com/login.html')

```

`build_opener()` provides many handlers by default, including a *ProxyHandler*. By default, *ProxyHandler* uses the environment variables named `<scheme>_proxy`, where `<scheme>` is the URL scheme involved. For example, the `http_proxy` environment variable is read to obtain the HTTP proxy's URL.

This example replaces the default *ProxyHandler* with one that uses programmatically-supplied proxy URLs, and adds proxy authorization support with *ProxyBasicAuthHandler*.

```

proxy_handler = urllib.request.ProxyHandler({'http': 'http://www.example.com:3128/'})
proxy_auth_handler = urllib.request.ProxyBasicAuthHandler()
proxy_auth_handler.add_password('realm', 'host', 'username', 'password')

opener = urllib.request.build_opener(proxy_handler, proxy_auth_handler)
# This time, rather than install the OpenerDirector, we use it directly:
opener.open('http://www.example.com/login.html')

```

Adding HTTP headers:

Use the *headers* argument to the *Request* constructor, or:

```

import urllib.request
req = urllib.request.Request('http://www.example.com/')
req.add_header('Referer', 'http://www.python.org/')
# Customize the default User-Agent header value:
req.add_header('User-Agent', 'urllib-example/0.1 (Contact: . . .)')
r = urllib.request.urlopen(req)

```

OpenerDirector automatically adds a *User-Agent* header to every *Request*. To change this:

```

import urllib.request
opener = urllib.request.build_opener()
opener.addheaders = [('User-agent', 'Mozilla/5.0')]
opener.open('http://www.example.com/')

```

Also, remember that a few standard headers (*Content-Length*, *Content-Type* and *Host*) are added when the *Request* is passed to `urlopen()` (or `OpenerDirector.open()`).

Here is an example session that uses the GET method to retrieve a URL containing parameters:

```

>>> import urllib.request
>>> import urllib.parse
>>> params = urllib.parse.urlencode({'spam': 1, 'eggs': 2, 'bacon': 0})
>>> url = "http://www.musi-cal.com/cgi-bin/query?%s" % params
>>> with urllib.request.urlopen(url) as f:
...     print(f.read().decode('utf-8'))
...

```

The following example uses the POST method instead. Note that params output from `urlencode` is encoded to bytes before it is sent to `urlopen` as data:

```
>>> import urllib.request
>>> import urllib.parse
>>> data = urllib.parse.urlencode({'spam': 1, 'eggs': 2, 'bacon': 0})
>>> data = data.encode('ascii')
>>> with urllib.request.urlopen("http://requestb.in/xrbl82xr", data) as f:
...     print(f.read().decode('utf-8'))
...
...
```

The following example uses an explicitly specified HTTP proxy, overriding environment settings:

```
>>> import urllib.request
>>> proxies = {'http': 'http://proxy.example.com:8080/'}
>>> opener = urllib.request.FancyURLopener(proxies)
>>> with opener.open("http://www.python.org") as f:
...     f.read().decode('utf-8')
...
...
```

The following example uses no proxies at all, overriding environment settings:

```
>>> import urllib.request
>>> opener = urllib.request.FancyURLopener({})
>>> with opener.open("http://www.python.org/") as f:
...     f.read().decode('utf-8')
...
...
```

21.6.24 Legacy interface

The following functions and classes are ported from the Python 2 module `urllib` (as opposed to `urllib2`). They might become deprecated at some point in the future.

`urllib.request.urlretrieve(url, filename=None, reporthook=None, data=None)`

Copy a network object denoted by a URL to a local file. If the URL points to a local file, the object will not be copied unless `filename` is supplied. Return a tuple (`filename`, `headers`) where `filename` is the local file name under which the object can be found, and `headers` is whatever the `info()` method of the object returned by `urlopen()` returned (for a remote object). Exceptions are the same as for `urlopen()`.

The second argument, if present, specifies the file location to copy to (if absent, the location will be a tempfile with a generated name). The third argument, if present, is a callable that will be called once on establishment of the network connection and once after each block read thereafter. The callable will be passed three arguments; a count of blocks transferred so far, a block size in bytes, and the total size of the file. The third argument may be `-1` on older FTP servers which do not return a file size in response to a retrieval request.

The following example illustrates the most common usage scenario:

```
>>> import urllib.request
>>> local_filename, headers = urllib.request.urlretrieve('http://python.org/')
>>> html = open(local_filename)
>>> html.close()
```

If the `url` uses the `http:` scheme identifier, the optional `data` argument may be given to specify a POST request (normally the request type is GET). The `data` argument must be a bytes object in standard `application/x-www-form-urlencoded` format; see the `urllib.parse.urlencode()` function.

`urlretrieve()` will raise `ContentTooShortError` when it detects that the amount of data available was less than the expected amount (which is the size reported by a `Content-Length` header). This can occur, for example, when the download is interrupted.

The *Content-Length* is treated as a lower bound: if there's more data to read, `urlretrieve` reads more data, but if less data is available, it raises the exception.

You can still retrieve the downloaded data in this case, it is stored in the `content` attribute of the exception instance.

If no *Content-Length* header was supplied, `urlretrieve` can not check the size of the data it has downloaded, and just returns it. In this case you just have to assume that the download was successful.

`urllib.request.urlcleanup()`

Cleans up temporary files that may have been left behind by previous calls to `urlretrieve()`.

class `urllib.request.URLOpener` (*proxies=None*, ***x509*)

3.3 版后已移除。

Base class for opening and reading URLs. Unless you need to support opening objects using schemes other than `http:`, `ftp:`, or `file:`, you probably want to use *FancyURLOpener*.

By default, the *URLOpener* class sends a *User-Agent* header of `urllib/VVV`, where *VVV* is the *urllib* version number. Applications can define their own *User-Agent* header by subclassing *URLOpener* or *FancyURLOpener* and setting the class attribute *version* to an appropriate string value in the subclass definition.

The optional *proxies* parameter should be a dictionary mapping scheme names to proxy URLs, where an empty dictionary turns proxies off completely. Its default value is `None`, in which case environmental proxy settings will be used if present, as discussed in the definition of `urlopen()`, above.

Additional keyword parameters, collected in *x509*, may be used for authentication of the client when using the `https:` scheme. The keywords *key_file* and *cert_file* are supported to provide an SSL key and certificate; both are needed to support client authentication.

URLOpener objects will raise an *OSError* exception if the server returns an error code.

open (*fullurl*, *data=None*)

Open *fullurl* using the appropriate protocol. This method sets up cache and proxy information, then calls the appropriate open method with its input arguments. If the scheme is not recognized, `open_unknown()` is called. The *data* argument has the same meaning as the *data* argument of `urlopen()`.

open_unknown (*fullurl*, *data=None*)

Overridable interface to open unknown URL types.

retrieve (*url*, *filename=None*, *reporthook=None*, *data=None*)

Retrieves the contents of *url* and places it in *filename*. The return value is a tuple consisting of a local filename and either an *email.message.Message* object containing the response headers (for remote URLs) or `None` (for local URLs). The caller must then open and read the contents of *filename*. If *filename* is not given and the URL refers to a local file, the input filename is returned. If the URL is non-local and *filename* is not given, the filename is the output of `tempfile.mktemp()` with a suffix that matches the suffix of the last path component of the input URL. If *reporthook* is given, it must be a function accepting three numeric parameters: A chunk number, the maximum size chunks are read in and the total size of the download (-1 if unknown). It will be called once at the start and after each chunk of data is read from the network. *reporthook* is ignored for local URLs.

If the *url* uses the `http:` scheme identifier, the optional *data* argument may be given to specify a POST request (normally the request type is GET). The *data* argument must in standard *application/x-www-form-urlencoded* format; see the `urllib.parse.urlencode()` function.

version

Variable that specifies the user agent of the opener object. To get *urllib* to tell servers that it is a particular user agent, set this in a subclass as a class variable or in the constructor before calling the base constructor.

`class urllib.request.FancyURLopener(...)`

3.3 版后已移除。

FancyURLopener subclasses *URLopener* providing default handling for the following HTTP response codes: 301, 302, 303, 307 and 401. For the 30x response codes listed above, the *Location* header is used to fetch the actual URL. For 401 response codes (authentication required), basic HTTP authentication is performed. For the 30x response codes, recursion is bounded by the value of the *maxtries* attribute, which defaults to 10.

For all other response codes, the method `http_error_default()` is called which you can override in subclasses to handle the error appropriately.

注解： According to the letter of **RFC 2616**, 301 and 302 responses to POST requests must not be automatically redirected without confirmation by the user. In reality, browsers do allow automatic redirection of these responses, changing the POST to a GET, and *urllib* reproduces this behaviour.

The parameters to the constructor are the same as those for *URLopener*.

注解： When performing basic authentication, a *FancyURLopener* instance calls its `prompt_user_passwd()` method. The default implementation asks the users for the required information on the controlling terminal. A subclass may override this method to support more appropriate behavior if needed.

The *FancyURLopener* class offers one additional method that should be overloaded to provide the appropriate behavior:

prompt_user_passwd (*host*, *realm*)

Return information needed to authenticate the user at the given host in the specified security realm. The return value should be a tuple, (*user*, *password*), which can be used for basic authentication.

The implementation prompts for this information on the terminal; an application should override this method to use an appropriate interaction model in the local environment.

21.6.25 urllib.request Restrictions

- Currently, only the following protocols are supported: HTTP (versions 0.9 and 1.0), FTP, local files, and data URLs.

在 3.4 版更改: Added support for data URLs.

- The caching feature of `urlretrieve()` has been disabled until someone finds the time to hack proper processing of Expiration time headers.
- There should be a function to query whether a particular URL is in the cache.
- For backward compatibility, if a URL appears to point to a local file but the file can't be opened, the URL is re-interpreted using the FTP protocol. This can sometimes cause confusing error messages.
- The `urlopen()` and `urlretrieve()` functions can cause arbitrarily long delays while waiting for a network connection to be set up. This means that it is difficult to build an interactive Web client using these functions without using threads.
- The data returned by `urlopen()` or `urlretrieve()` is the raw data returned by the server. This may be binary data (such as an image), plain text or (for example) HTML. The HTTP protocol provides type information

in the reply header, which can be inspected by looking at the *Content-Type* header. If the returned data is HTML, you can use the module `html.parser` to parse it.

- The code handling the FTP protocol cannot differentiate between a file and a directory. This can lead to unexpected behavior when attempting to read a URL that points to a file that is not accessible. If the URL ends in a `/`, it is assumed to refer to a directory and will be handled accordingly. But if an attempt to read a file leads to a 550 error (meaning the URL cannot be found or is not accessible, often for permission reasons), then the path is treated as a directory in order to handle the case when a directory is specified by a URL but the trailing `/` has been left off. This can cause misleading results when you try to fetch a file whose read permissions make it inaccessible; the FTP code will try to read it, fail with a 550 error, and then perform a directory listing for the unreadable file. If fine-grained control is needed, consider using the `ftplib` module, subclassing `FancyURLopener`, or changing `_urlopener` to meet your needs.

21.7 `urllib.response` —Response classes used by urllib

The `urllib.response` module defines functions and classes which define a minimal file like interface, including `read()` and `readline()`. The typical response object is an `addinfourl` instance, which defines an `info()` method and that returns headers and a `geturl()` method that returns the url. Functions defined by this module are used internally by the `urllib.request` module.

21.8 `urllib.parse` —Parse URLs into components

Source code: [Lib/urllib/parse.py](#)

This module defines a standard interface to break Uniform Resource Locator (URL) strings up in components (addressing scheme, network location, path etc.), to combine the components back into a URL string, and to convert a “relative URL” to an absolute URL given a “base URL.”

The module has been designed to match the Internet RFC on Relative Uniform Resource Locators. It supports the following URL schemes: `file`, `ftp`, `gopher`, `hdl`, `http`, `https`, `imap`, `mailto`, `mms`, `news`, `ntp`, `prospero`, `rsync`, `rtsp`, `rtspu`, `sftp`, `shttp`, `sip`, `sips`, `snews`, `svn`, `svn+ssh`, `telnet`, `wais`, `ws`, `wss`.

The `urllib.parse` module defines functions that fall into two broad categories: URL parsing and URL quoting. These are covered in detail in the following sections.

21.8.1 URL Parsing

The URL parsing functions focus on splitting a URL string into its components, or on combining URL components into a URL string.

`urllib.parse.urlparse(urlstring, scheme='', allow_fragments=True)`

Parse a URL into six components, returning a 6-tuple. This corresponds to the general structure of a URL: `scheme://netloc/path;parameters?query#fragment`. Each tuple item is a string, possibly empty. The components are not broken up in smaller parts (for example, the network location is a single string), and `%` escapes are not expanded. The delimiters as shown above are not part of the result, except for a leading slash in the `path` component, which is retained if present. For example:

```
>>> from urllib.parse import urlparse
>>> o = urlparse('http://www.cwi.nl:80/%7Eguido/Python.html')
>>> o      # doctest: +NORMALIZE_WHITESPACE
```

(下页继续)

(续上页)

```

ParseResult(scheme='http', netloc='www.cwi.nl:80', path='/%7Eguido/Python.html',
            params='', query='', fragment='')
>>> o.scheme
'http'
>>> o.port
80
>>> o.geturl()
'http://www.cwi.nl:80/%7Eguido/Python.html'

```

Following the syntax specifications in [RFC 1808](#), `urlparse` recognizes a netloc only if it is properly introduced by `'//'`. Otherwise the input is presumed to be a relative URL and thus to start with a path component.

```

>>> from urllib.parse import urlparse
>>> urlparse('//www.cwi.nl:80/%7Eguido/Python.html')
ParseResult(scheme='', netloc='www.cwi.nl:80', path='/%7Eguido/Python.html',
            params='', query='', fragment='')
>>> urlparse('www.cwi.nl/%7Eguido/Python.html')
ParseResult(scheme='', netloc='', path='www.cwi.nl/%7Eguido/Python.html',
            params='', query='', fragment='')
>>> urlparse('help/Python.html')
ParseResult(scheme='', netloc='', path='help/Python.html', params='',
            query='', fragment='')

```

The *scheme* argument gives the default addressing scheme, to be used only if the URL does not specify one. It should be the same type (text or bytes) as *urlstring*, except that the default value `''` is always allowed, and is automatically converted to `b''` if appropriate.

If the *allow_fragments* argument is false, fragment identifiers are not recognized. Instead, they are parsed as part of the path, parameters or query component, and *fragment* is set to the empty string in the return value.

The return value is actually an instance of a subclass of *tuple*. This class has the following additional read-only convenience attributes:

属性	索引	值	Value if not present
<code>scheme</code>	0	URL scheme specifier	<i>scheme</i> parameter
<code>netloc</code>	1	Network location part	empty string
<code>path</code>	2	Hierarchical path	empty string
<code>params</code>	3	Parameters for last path element	empty string
<code>query</code>	4	Query component	empty string
<code>fragment</code>	5	Fragment identifier	empty string
<code>username</code>		User name	<i>None</i>
<code>password</code>		Password	<i>None</i>
<code>hostname</code>		Host name (lower case)	<i>None</i>
<code>port</code>		Port number as integer, if present	<i>None</i>

Reading the `port` attribute will raise a *ValueError* if an invalid port is specified in the URL. See section [Structured Parse Results](#) for more information on the result object.

Unmatched square brackets in the `netloc` attribute will raise a *ValueError*.

Characters in the `netloc` attribute that decompose under NFKC normalization (as used by the IDNA encoding) into any of `/`, `?`, `#`, `@`, or `:` will raise a *ValueError*. If the URL is decomposed before parsing, no error will be raised.

在 3.2 版更改: Added IPv6 URL parsing capabilities.

在 3.3 版更改: The fragment is now parsed for all URL schemes (unless *allow_fragment* is false), in accordance with [RFC 3986](#). Previously, a whitelist of schemes that support fragments existed.

在 3.6 版更改: Out-of-range port numbers now raise *ValueError*, instead of returning *None*.

在 3.6.9 版更改: Characters that affect netloc parsing under NFKC normalization will now raise *ValueError*.

`urllib.parse.parse_qs(qs, keep_blank_values=False, strict_parsing=False, encoding='utf-8', errors='replace', max_num_fields=None)`

Parse a query string given as a string argument (data of type *application/x-www-form-urlencoded*). Data are returned as a dictionary. The dictionary keys are the unique query variable names and the values are lists of values for each name.

The optional argument *keep_blank_values* is a flag indicating whether blank values in percent-encoded queries should be treated as blank strings. A true value indicates that blanks should be retained as blank strings. The default false value indicates that blank values are to be ignored and treated as if they were not included.

The optional argument *strict_parsing* is a flag indicating what to do with parsing errors. If false (the default), errors are silently ignored. If true, errors raise a *ValueError* exception.

The optional *encoding* and *errors* parameters specify how to decode percent-encoded sequences into Unicode characters, as accepted by the *bytes.decode()* method.

The optional argument *max_num_fields* is the maximum number of fields to read. If set, then throws a *ValueError* if there are more than *max_num_fields* fields read.

Use the *urllib.parse.urlencode()* function (with the *doseq* parameter set to True) to convert such dictionaries into query strings.

在 3.2 版更改: Add *encoding* and *errors* parameters.

在 3.6.8 版更改: Added *max_num_fields* parameter.

`urllib.parse.parse_qsl(qs, keep_blank_values=False, strict_parsing=False, encoding='utf-8', errors='replace', max_num_fields=None)`

Parse a query string given as a string argument (data of type *application/x-www-form-urlencoded*). Data are returned as a list of name, value pairs.

The optional argument *keep_blank_values* is a flag indicating whether blank values in percent-encoded queries should be treated as blank strings. A true value indicates that blanks should be retained as blank strings. The default false value indicates that blank values are to be ignored and treated as if they were not included.

The optional argument *strict_parsing* is a flag indicating what to do with parsing errors. If false (the default), errors are silently ignored. If true, errors raise a *ValueError* exception.

The optional *encoding* and *errors* parameters specify how to decode percent-encoded sequences into Unicode characters, as accepted by the *bytes.decode()* method.

The optional argument *max_num_fields* is the maximum number of fields to read. If set, then throws a *ValueError* if there are more than *max_num_fields* fields read.

Use the *urllib.parse.urlencode()* function to convert such lists of pairs into query strings.

在 3.2 版更改: Add *encoding* and *errors* parameters.

在 3.6.8 版更改: Added *max_num_fields* parameter.

`urllib.parse.urlunparse(parts)`

Construct a URL from a tuple as returned by *urlparse()*. The *parts* argument can be any six-item iterable. This may result in a slightly different, but equivalent URL, if the URL that was parsed originally had unnecessary delimiters (for example, a ? with an empty query; the RFC states that these are equivalent).

`urllib.parse.urlsplit(urlstring, scheme="", allow_fragments=True)`

This is similar to *urlparse()*, but does not split the params from the URL. This should generally be used

instead of `urlparse()` if the more recent URL syntax allowing parameters to be applied to each segment of the *path* portion of the URL (see [RFC 2396](#)) is wanted. A separate function is needed to separate the path segments and parameters. This function returns a 5-tuple: (addressing scheme, network location, path, query, fragment identifier).

The return value is actually an instance of a subclass of `tuple`. This class has the following additional read-only convenience attributes:

属性	索引	值	Value if not present
<code>scheme</code>	0	URL scheme specifier	<i>scheme</i> parameter
<code>netloc</code>	1	Network location part	empty string
<code>path</code>	2	Hierarchical path	empty string
<code>query</code>	3	Query component	empty string
<code>fragment</code>	4	Fragment identifier	empty string
<code>username</code>		User name	<i>None</i>
<code>password</code>		Password	<i>None</i>
<code>hostname</code>		Host name (lower case)	<i>None</i>
<code>port</code>		Port number as integer, if present	<i>None</i>

Reading the `port` attribute will raise a `ValueError` if an invalid port is specified in the URL. See section [Structured Parse Results](#) for more information on the result object.

Unmatched square brackets in the `netloc` attribute will raise a `ValueError`.

Characters in the `netloc` attribute that decompose under NFKC normalization (as used by the IDNA encoding) into any of `/`, `?`, `#`, `@`, or `:` will raise a `ValueError`. If the URL is decomposed before parsing, no error will be raised.

在 3.6 版更改: Out-of-range port numbers now raise `ValueError`, instead of returning *None*.

在 3.6.9 版更改: Characters that affect `netloc` parsing under NFKC normalization will now raise `ValueError`.

`urllib.parse.urlunsplit(parts)`

Combine the elements of a tuple as returned by `urlsplit()` into a complete URL as a string. The *parts* argument can be any five-item iterable. This may result in a slightly different, but equivalent URL, if the URL that was parsed originally had unnecessary delimiters (for example, a `?` with an empty query; the RFC states that these are equivalent).

`urllib.parse.urljoin(base, url, allow_fragments=True)`

Construct a full (“absolute”) URL by combining a “base URL” (*base*) with another URL (*url*). Informally, this uses components of the base URL, in particular the addressing scheme, the network location and (part of) the path, to provide missing components in the relative URL. For example:

```
>>> from urllib.parse import urljoin
>>> urljoin('http://www.cwi.nl/%7Eguido/Python.html', 'FAQ.html')
'http://www.cwi.nl/%7Eguido/FAQ.html'
```

The `allow_fragments` argument has the same meaning and default as for `urlparse()`.

注解: If *url* is an absolute URL (that is, starting with `//` or `scheme://`), the *url*’s host name and/or scheme will be present in the result. For example:

```
>>> urljoin('http://www.cwi.nl/%7Eguido/Python.html',
...         '//www.python.org/%7Eguido')
'http://www.python.org/%7Eguido'
```

If you do not want that behavior, preprocess the *url* with `urlsplit()` and `urlunsplit()`, removing possible *scheme* and *netloc* parts.

在 3.5 版更改: Behaviour updated to match the semantics defined in [RFC 3986](#).

`urllib.parse.urldefrag(url)`

If *url* contains a fragment identifier, return a modified version of *url* with no fragment identifier, and the fragment identifier as a separate string. If there is no fragment identifier in *url*, return *url* unmodified and an empty string.

The return value is actually an instance of a subclass of `tuple`. This class has the following additional read-only convenience attributes:

属性	索引	值	Value if not present
<code>url</code>	0	URL with no fragment	empty string
<code>fragment</code>	1	Fragment identifier	empty string

See section [Structured Parse Results](#) for more information on the result object.

在 3.2 版更改: Result is a structured object rather than a simple 2-tuple.

21.8.2 Parsing ASCII Encoded Bytes

The URL parsing functions were originally designed to operate on character strings only. In practice, it is useful to be able to manipulate properly quoted and encoded URLs as sequences of ASCII bytes. Accordingly, the URL parsing functions in this module all operate on *bytes* and *bytearray* objects in addition to *str* objects.

If *str* data is passed in, the result will also contain only *str* data. If *bytes* or *bytearray* data is passed in, the result will contain only *bytes* data.

Attempting to mix *str* data with *bytes* or *bytearray* in a single function call will result in a `TypeError` being raised, while attempting to pass in non-ASCII byte values will trigger `UnicodeDecodeError`.

To support easier conversion of result objects between *str* and *bytes*, all return values from URL parsing functions provide either an `encode()` method (when the result contains *str* data) or a `decode()` method (when the result contains *bytes* data). The signatures of these methods match those of the corresponding *str* and *bytes* methods (except that the default encoding is 'ascii' rather than 'utf-8'). Each produces a value of a corresponding type that contains either *bytes* data (for `encode()` methods) or *str* data (for `decode()` methods).

Applications that need to operate on potentially improperly quoted URLs that may contain non-ASCII data will need to do their own decoding from bytes to characters before invoking the URL parsing methods.

The behaviour described in this section applies only to the URL parsing functions. The URL quoting functions use their own rules when producing or consuming byte sequences as detailed in the documentation of the individual URL quoting functions.

在 3.2 版更改: URL parsing functions now accept ASCII encoded byte sequences

21.8.3 Structured Parse Results

The result objects from the `urlparse()`, `urlsplit()` and `urldefrag()` functions are subclasses of the `tuple` type. These subclasses add the attributes listed in the documentation for those functions, the encoding and decoding support described in the previous section, as well as an additional method:

`urllib.parse.SplitResult.geturl()`

Return the re-combined version of the original URL as a string. This may differ from the original URL in that the scheme may be normalized to lower case and empty components may be dropped. Specifically, empty parameters, queries, and fragment identifiers will be removed.

For `urldefrag()` results, only empty fragment identifiers will be removed. For `urlsplit()` and `urlparse()` results, all noted changes will be made to the URL returned by this method.

The result of this method remains unchanged if passed back through the original parsing function:

```
>>> from urllib.parse import urlsplit
>>> url = 'HTTP://www.Python.org/doc/#'
>>> r1 = urlsplit(url)
>>> r1.geturl()
'http://www.Python.org/doc/'
>>> r2 = urlsplit(r1.geturl())
>>> r2.geturl()
'http://www.Python.org/doc/'
```

The following classes provide the implementations of the structured parse results when operating on *str* objects:

class `urllib.parse.DefragResult` (*url, fragment*)

Concrete class for `urldefrag()` results containing *str* data. The `encode()` method returns a `DefragResultBytes` instance.

3.2 新版功能.

class `urllib.parse.ParseResult` (*scheme, netloc, path, params, query, fragment*)

Concrete class for `urlparse()` results containing *str* data. The `encode()` method returns a `ParseResultBytes` instance.

class `urllib.parse.SplitResult` (*scheme, netloc, path, query, fragment*)

Concrete class for `urlsplit()` results containing *str* data. The `encode()` method returns a `SplitResultBytes` instance.

The following classes provide the implementations of the parse results when operating on *bytes* or *bytearray* objects:

class `urllib.parse.DefragResultBytes` (*url, fragment*)

Concrete class for `urldefrag()` results containing *bytes* data. The `decode()` method returns a `DefragResult` instance.

3.2 新版功能.

class `urllib.parse.ParseResultBytes` (*scheme, netloc, path, params, query, fragment*)

Concrete class for `urlparse()` results containing *bytes* data. The `decode()` method returns a `ParseResult` instance.

3.2 新版功能.

class `urllib.parse.SplitResultBytes` (*scheme, netloc, path, query, fragment*)

Concrete class for `urlsplit()` results containing *bytes* data. The `decode()` method returns a `SplitResult` instance.

3.2 新版功能.

21.8.4 URL Quoting

The URL quoting functions focus on taking program data and making it safe for use as URL components by quoting special characters and appropriately encoding non-ASCII text. They also support reversing these operations to recreate the original data from the contents of a URL component if that task isn't already covered by the URL parsing functions above.

`urllib.parse.quote` (*string, safe='/', encoding=None, errors=None*)

Replace special characters in *string* using the `%xx` escape. Letters, digits, and the characters `'_'.-` are never

quoted. By default, this function is intended for quoting the path section of URL. The optional *safe* parameter specifies additional ASCII characters that should not be quoted —its default value is `'/'`.

string may be either a *str* or a *bytes*.

The optional *encoding* and *errors* parameters specify how to deal with non-ASCII characters, as accepted by the *str.encode()* method. *encoding* defaults to `'utf-8'`. *errors* defaults to `'strict'`, meaning unsupported characters raise a *UnicodeEncodeError*. *encoding* and *errors* must not be supplied if *string* is a *bytes*, or a *TypeError* is raised.

Note that `quote(string, safe, encoding, errors)` is equivalent to `quote_from_bytes(string.encode(encoding, errors), safe)`.

Example: `quote('/El Niño/')` yields `'/El%20Ni%C3%B1o/'`.

`urllib.parse.quote_plus(string, safe='', encoding=None, errors=None)`

Like *quote()*, but also replace spaces by plus signs, as required for quoting HTML form values when building up a query string to go into a URL. Plus signs in the original string are escaped unless they are included in *safe*. It also does not have *safe* default to `'/'`.

Example: `quote_plus('/El Niño/')` yields `'%2FEl+Ni%C3%B1o%2F'`.

`urllib.parse.quote_from_bytes(bytes, safe='/')`

Like *quote()*, but accepts a *bytes* object rather than a *str*, and does not perform string-to-bytes encoding.

Example: `quote_from_bytes(b'a&\xef')` yields `'a%26%EF'`.

`urllib.parse.unquote(string, encoding='utf-8', errors='replace')`

Replace `%xx` escapes by their single-character equivalent. The optional *encoding* and *errors* parameters specify how to decode percent-encoded sequences into Unicode characters, as accepted by the *bytes.decode()* method.

string must be a *str*.

encoding defaults to `'utf-8'`. *errors* defaults to `'replace'`, meaning invalid sequences are replaced by a placeholder character.

Example: `unquote('/El%20Ni%C3%B1o/')` yields `'/El Niño/'`.

`urllib.parse.unquote_plus(string, encoding='utf-8', errors='replace')`

Like *unquote()*, but also replace plus signs by spaces, as required for unquoting HTML form values.

string must be a *str*.

Example: `unquote_plus('/El+Ni%C3%B1o/')` yields `'/El Niño/'`.

`urllib.parse.unquote_to_bytes(string)`

Replace `%xx` escapes by their single-octet equivalent, and return a *bytes* object.

string may be either a *str* or a *bytes*.

If it is a *str*, unescaped non-ASCII characters in *string* are encoded into UTF-8 bytes.

Example: `unquote_to_bytes('a%26%EF')` yields `b'a&\xef'`.

`urllib.parse.urlencode(query, doseq=False, safe='', encoding=None, errors=None, quote_via=quote_plus)`

Convert a mapping object or a sequence of two-element tuples, which may contain *str* or *bytes* objects, to a percent-encoded ASCII text string. If the resultant string is to be used as a *data* for POST operation with the *urlopen()* function, then it should be encoded to bytes, otherwise it would result in a *TypeError*.

The resulting string is a series of *key=value* pairs separated by `'&'` characters, where both *key* and *value* are quoted using the *quote_via* function. By default, *quote_plus()* is used to quote the values, which means spaces are quoted as a `'+'` character and `'/'` characters are encoded as `%2F`, which follows the standard for GET requests (`application/x-www-form-urlencoded`). An alternate function that can be passed as

`quote_via` is `quote()`, which will encode spaces as `%20` and not encode `'/'` characters. For maximum control of what is quoted, use `quote` and specify a value for `safe`.

When a sequence of two-element tuples is used as the `query` argument, the first element of each tuple is a key and the second is a value. The value element in itself can be a sequence and in that case, if the optional parameter `doseq` is evaluates to `True`, individual `key=value` pairs separated by `'&'` are generated for each element of the value sequence for the key. The order of parameters in the encoded string will match the order of parameter tuples in the sequence.

The `safe`, `encoding`, and `errors` parameters are passed down to `quote_via` (the `encoding` and `errors` parameters are only passed when a query element is a `str`).

To reverse this encoding process, `parse_qs()` and `parse_qsl()` are provided in this module to parse query strings into Python data structures.

Refer to [urllib examples](#) to find out how `urlencode` method can be used for generating query string for a URL or data for POST.

在 3.2 版更改: Query parameter supports bytes and string objects.

3.5 新版功能: `quote_via` parameter.

参见:

RFC 3986 - Uniform Resource Identifiers This is the current standard (STD66). Any changes to `urllib.parse` module should conform to this. Certain deviations could be observed, which are mostly for backward compatibility purposes and for certain de-facto parsing requirements as commonly observed in major browsers.

RFC 2732 - Format for Literal IPv6 Addresses in URL's. This specifies the parsing requirements of IPv6 URLs.

RFC 2396 - Uniform Resource Identifiers (URI): Generic Syntax Document describing the generic syntactic requirements for both Uniform Resource Names (URNs) and Uniform Resource Locators (URLs).

RFC 2368 - The mailto URL scheme. Parsing requirements for mailto URL schemes.

RFC 1808 - Relative Uniform Resource Locators This Request For Comments includes the rules for joining an absolute and a relative URL, including a fair number of “Abnormal Examples” which govern the treatment of border cases.

RFC 1738 - Uniform Resource Locators (URL) This specifies the formal syntax and semantics of absolute URLs.

21.9 urllib.error — Exception classes raised by urllib.request

Source code: [Lib/urllib/error.py](#)

The `urllib.error` module defines the exception classes for exceptions raised by `urllib.request`. The base exception class is `URLError`.

The following exceptions are raised by `urllib.error` as appropriate:

exception `urllib.error.URLError`

The handlers raise this exception (or derived exceptions) when they run into a problem. It is a subclass of `OSError`.

reason

The reason for this error. It can be a message string or another exception instance.

在 3.3 版更改: `URLError` has been made a subclass of `OSError` instead of `IOError`.

exception `urllib.error.HTTPError`

Though being an exception (a subclass of `URLError`), an `HTTPError` can also function as a non-exceptional file-like return value (the same thing that `urlopen()` returns). This is useful when handling exotic HTTP errors, such as requests for authentication.

code

An HTTP status code as defined in **RFC 2616**. This numeric value corresponds to a value found in the dictionary of codes as found in `http.server.BaseHTTPRequestHandler.responses`.

reason

This is usually a string explaining the reason for this error.

headers

The HTTP response headers for the HTTP request that caused the `HTTPError`.

3.4 新版功能.

exception `urllib.error.ContentTooShortError` (*msg, content*)

This exception is raised when the `urlretrieve()` function detects that the amount of the downloaded data is less than the expected amount (given by the *Content-Length* header). The `content` attribute stores the downloaded (and supposedly truncated) data.

21.10 urllib.robotparser — Parser for robots.txt

Source code: [Lib/urllib/robotparser.py](#)

This module provides a single class, `RobotFileParser`, which answers questions about whether or not a particular user agent can fetch a URL on the Web site that published the `robots.txt` file. For more details on the structure of `robots.txt` files, see <http://www.robotstxt.org/orig.html>.

class `urllib.robotparser.RobotFileParser` (*url=""*)

This class provides methods to read, parse and answer questions about the `robots.txt` file at *url*.

set_url (*url*)

Sets the URL referring to a `robots.txt` file.

read ()

Reads the `robots.txt` URL and feeds it to the parser.

parse (*lines*)

Parses the *lines* argument.

can_fetch (*useragent, url*)

Returns `True` if the *useragent* is allowed to fetch the *url* according to the rules contained in the parsed `robots.txt` file.

mtime ()

Returns the time the `robots.txt` file was last fetched. This is useful for long-running web spiders that need to check for new `robots.txt` files periodically.

modified ()

Sets the time the `robots.txt` file was last fetched to the current time.

crawl_delay (*useragent*)

Returns the value of the *Crawl-delay* parameter from `robots.txt` for the *useragent* in question. If there is no such parameter or it doesn't apply to the *useragent* specified or the `robots.txt` entry for this parameter has invalid syntax, return `None`.

3.6 新版功能.

request_rate (*useragent*)

Returns the contents of the Request-rate parameter from robots.txt as a *named tuple* RequestRate(requests, seconds). If there is no such parameter or it doesn't apply to the *user-agent* specified or the robots.txt entry for this parameter has invalid syntax, return None.

3.6 新版功能.

The following example demonstrates basic use of the *RobotFileParser* class:

```
>>> import urllib.robotparser
>>> rp = urllib.robotparser.RobotFileParser()
>>> rp.set_url("http://www.musi-cal.com/robots.txt")
>>> rp.read()
>>> rrate = rp.request_rate("")
>>> rrate.requests
3
>>> rrate.seconds
20
>>> rp.crawl_delay("")
6
>>> rp.can_fetch("", "http://www.musi-cal.com/cgi-bin/search?city=San+Francisco")
False
>>> rp.can_fetch("", "http://www.musi-cal.com/")
True
```

21.11 http —HTTP 模块

源代码: [Lib/http/__init__.py](#)

http 是一个包，它收集了多个用于处理超文本传输协议的模块：

- *http.client* 是一个低层级的 HTTP 协议客户端；对于高层级的 URL 访问请使用 *urllib.request*
- *http.server* 包含基于 *socketserver* 的基本 HTTP 服务类
- *http.cookies* 包含一些有用来实现通过 *cookies* 进行状态管理的工具
- *http.cookiejar* 提供了 *cookies* 的持久化

http 也是一个通过 *http.HTTPStatus* 枚举定义了一些 HTTP 状态码以及相关消息的模块

class http.HTTPStatus

3.5 新版功能.

enum.IntEnum 的子类，它定义了组 HTTP 状态码，原理短语以及用英语书写的长描述文本。

用法:

```
>>> from http import HTTPStatus
>>> HTTPStatus.OK
<HTTPStatus.OK: 200>
>>> HTTPStatus.OK == 200
True
>>> http.HTTPStatus.OK.value
200
```

(下页继续)

(续上页)

```
>>> HTTPStatus.OK.phrase
'OK'
>>> HTTPStatus.OK.description
'Request fulfilled, document follows'
>>> list(HTTPStatus)
[<HTTPStatus.CONTINUE: 100>, <HTTPStatus.SWITCHING_PROTOCOLS: 101>, ...]
```

21.11.1 HTTP 状态码

已支持并且已在http.Status.IANA 注册 的状态码有：

状态码	映射名	详情
100	CONTINUE	HTTP/1.1 RFC 7231 , Section 6.2.1
101	SWITCHING_PROTOCOLS	HTTP/1.1 RFC 7231, Section 6.2.2
102	PROCESSING	WebDAV RFC 2518, Section 10.1
200	OK	HTTP/1.1 RFC 7231, Section 6.3.1
201	CREATED	HTTP/1.1 RFC 7231, Section 6.3.2
202	ACCEPTED	HTTP/1.1 RFC 7231, Section 6.3.3
203	NON_AUTHORITATIVE_INFORMATION	HTTP/1.1 RFC 7231, Section 6.3.4
204	NO_CONTENT	HTTP/1.1 RFC 7231, Section 6.3.5
205	RESET_CONTENT	HTTP/1.1 RFC 7231, Section 6.3.6
206	PARTIAL_CONTENT	HTTP/1.1 RFC 7233, Section 4.1
207	MULTI_STATUS	WebDAV RFC 4918, Section 11.1
208	ALREADY_REPORTED	WebDAV Binding Extensions RFC 5842, Section
226	IM_USED	Delta Encoding in HTTP RFC 3229, Section 10.4
300	MULTIPLE_CHOICES: 有多种资源可选择	HTTP/1.1 RFC 7231, Section 6.4.1
301	MOVED_PERMANENTLY: 永久移动	HTTP/1.1 RFC 7231, Section 6.4.2
302	FOUND: 临时移动	HTTP/1.1 RFC 7231, Section 6.4.3
303	SEE_OTHER: 已经移动	HTTP/1.1 RFC 7231, Section 6.4.4
304	NOT_MODIFIED: 没有修改	HTTP/1.1 RFC 7232, Section 4.1
305	USE_PROXY: 使用代理	HTTP/1.1 RFC 7231, Section 6.4.5
307	TEMPORARY_REDIRECT: 临时重定向	HTTP/1.1 RFC 7231 , Section 6.4.7
308	PERMANENT_REDIRECT: 永久重定向	Permanent Redirect RFC 7238 , Section 3 (Experi
400	BAD_REQUEST: 错误请求	HTTP/1.1 RFC 7231, Section 6.5.1
401	UNAUTHORIZED: 未授权	HTTP/1.1 Authentication RFC 7235, Section 3.1
402	PAYMENT_REQUIRED: 保留，将来使用	HTTP/1.1 RFC 7231, Section 6.5.2
403	FORBIDDEN: 禁止	HTTP/1.1 RFC 7231, Section 6.5.3
404	NOT_FOUND: 没有找到	HTTP/1.1 RFC 7231, Section 6.5.4
405	METHOD_NOT_ALLOWED: 该请求方法不允许	HTTP/1.1 RFC 7231, Section 6.5.5
406	NOT_ACCEPTABLE: 不可接受	HTTP/1.1 RFC 7231, Section 6.5.6
407	PROXY_AUTHENTICATION_REQUIRED: 要求使用代理验证身份	HTTP/1.1 Authentication RFC 7235, Section 3.1
408	REQUEST_TIMEOUT: 请求超时	HTTP/1.1 RFC 7231 , Section 6.5.7
409	CONFLICT: 冲突	HTTP/1.1 RFC 7231, Section 6.5.8
410	GONE: 已经不存在了	HTTP/1.1 RFC 7231, Section 6.5.9
411	LENGTH_REQUIRED: 长度要求	HTTP/1.1 RFC 7231, Section 6.5.10
412	PRECONDITION_FAILED: 前提条件错误	HTTP/1.1 RFC 7232, Section 4.2
413	REQUEST_ENTITY_TOO_LARGE: 请求体太大了	HTTP/1.1 RFC 7231, Section 6.5.11
414	REQUEST_URI_TOO_LONG: 请求 URI 太长了	HTTP/1.1 RFC 7231, Section 6.5.12
415	UNSUPPORTED_MEDIA_TYPE: 不支持的媒体格式	HTTP/1.1 RFC 7231, Section 6.5.13

表 1 – 续上页

状态码	映射名	详情
416	REQUEST_RANGE_NOT_SATISFIABLE	HTTP/1.1 Range Requests RFC 7233, Section 4.4
417	EXPECTATION_FAILED: 期望失败	HTTP/1.1 RFC 7231, Section 6.5.14
422	UNPROCESSABLE_ENTITY: 可加工实体	WebDAV RFC 4918, Section 11.2
423	LOCKED: 锁着	WebDAV RFC 4918, Section 11.3
424	FAILED_DEPENDENCY: 失败的依赖	WebDAV RFC 4918, Section 11.4
426	UPGRADE_REQUIRED: 升级需要	HTTP/1.1 RFC 7231, Section 6.5.15
428	PRECONDITION_REQUIRED: 先决条件要求	Additional HTTP Status Codes RFC 6585
429	TOO_MANY_REQUESTS: 太多的请求	Additional HTTP Status Codes RFC 6585
431	REQUEST_HEADER_FIELDS_TOO_LARGE: 请求头太大	Additional HTTP Status Codes RFC 6585
500	INTERNAL_SERVER_ERROR: 内部服务错误	HTTP/1.1 RFC 7231, Section 6.6.1
501	NOT_IMPLEMENTED: 不可执行	HTTP/1.1 RFC 7231, Section 6.6.2
502	BAD_GATEWAY: 无效网关	HTTP/1.1 RFC 7231, Section 6.6.3
503	SERVICE_UNAVAILABLE: 服务不可用	HTTP/1.1 RFC 7231, Section 6.6.4
504	GATEWAY_TIMEOUT: 网关超时	HTTP/1.1 RFC 7231, Section 6.6.5
505	HTTP_VERSION_NOT_SUPPORTED: HTTP 版本不支持	HTTP/1.1 RFC 7231, Section 6.6.6
506	VARIANT_ALSO_NEGOTIATES: 服务器存在内部配置错误	透明内容协商在: HTTP RFC 2295, Section 8.1
507	INSUFFICIENT_STORAGE: 存储不足	WebDAV RFC 4918, Section 11.5
508	LOOP_DETECTED: 循环检测	WebDAV Binding Extensions RFC 5842, Section
510	NOT_EXTENDED: 不扩展	WebDAV Binding Extensions RFC 5842, Section
511	NETWORK_AUTHENTICATION_REQUIRED: 要求网络身份验证	Additional HTTP Status Codes RFC 6585, Section

为了保持向后兼容性，枚举值也以常量形式出现在`http.client` 模块中，。枚举名等于常量名 (例如 `http.HTTPStatus.OK` 也可以是 `http.client.OK`)。

21.12 http.client —HTTP 协议客户端

源码: [Lib/http/client.py](#)

这个模块定义了实现 HTTP 和 HTTPS 协议客户端的类。它通常不直接使用—模块`urllib.request` 用它来处理使用 HTTP 和 HTTPS 的 URL。

参见:

[Requests 包](#) 推荐用于更高层级的 HTTP 客户端接口。

注解: HTTPS 支持仅在 Python 启动了 SSL 支持进行编译时（通过`ssl` 模块）可用。

该模块支持以下类:

class `http.client.HTTPConnection` (*host*, *port=None*[, *timeout*], *source_address=None*)
An `HTTPConnection` instance represents one transaction with an HTTP server. It should be instantiated passing it a host and optional port number. If no port number is passed, the port is extracted from the host string if it has the form `host:port`, else the default HTTP port (80) is used. If the optional *timeout* parameter is given, blocking operations (like connection attempts) will timeout after that many seconds (if it is not given, the global default timeout setting is used). The optional *source_address* parameter may be a tuple of a (host, port) to use as the source address the HTTP connection is made from.

举个例子，以下调用都是创建连接到同一主机和端口的服务器的实例：

```
>>> h1 = http.client.HTTPConnection('www.python.org')
>>> h2 = http.client.HTTPConnection('www.python.org:80')
>>> h3 = http.client.HTTPConnection('www.python.org', 80)
>>> h4 = http.client.HTTPConnection('www.python.org', 80, timeout=10)
```

在 3.2 版更改: 添加了 *source_address*。

在 3.4 版更改: 删除了 *strict* 参数, 不再支持 HTTP 0.9 风格的“简单响应”。

```
class http.client.HTTPSConnection(host, port=None, key_file=None, cert_file=None[,
                                   timeout], source_address=None, *, context=None,
                                   check_hostname=None)
```

HTTPConnection 的子类, 使用 SSL 与安全服务器进行通信。默认端口为 443。如果指定了 *context*, 它必须为一个描述 SSL 各选项的 *ssl.SSLContext* 实例。

请参阅 *Security considerations* 了解有关最佳实践的更多信息。

在 3.2 版更改: 添加了 *source_address*, *context* 和 *check_hostname*。

在 3.2 版更改: 这个类目前会在可能的情况下 (即如果 *ssl.HAS_SNI* 为真值) 支持 HTTPS 虚拟主机。

在 3.4 版更改: 删除了 *strict* 参数, 不再支持 HTTP 0.9 风格的“简单响应”。

在 3.4.3 版更改: 目前这个类在默认情况下会执行所有必要的证书和主机检查。要回复到先前的非验证行为, 可以将 *ssl._create_unverified_context()* 传递给 *context* 参数。

3.6 版后已移除: *key_file* 和 *cert_file* 已弃用并转而推荐 *context*。请改用 *ssl.SSLContext.load_cert_chain()* 或让 *ssl.create_default_context()* 为你选择系统所信任的 CA 证书。

check_hostname 参数也已弃用; 应当改用 *context* 的 *ssl.SSLContext.check_hostname* 属性。

```
class http.client.HTTPResponse(sock, debuglevel=0, method=None, url=None)
```

在成功连接后返回类的实例, 而不是由用户直接实例化。

在 3.4 版更改: 删除了 *strict* 参数, 不再支持 HTTP 0.9 风格的“简单响应”。

下列异常可以适当地被引发:

```
exception http.client.HTTPException
```

此模块中其他异常的基类。它是 *Exception* 的一个子类。

```
exception http.client.NotConnected
```

HTTPException 的一个子类。

```
exception http.client.InvalidURL
```

HTTPException 的一个子类, 如果给出了一个非数字或为空值的端口就会被引发。

```
exception http.client.UnknownProtocol
```

HTTPException 的一个子类。

```
exception http.client.UnknownTransferEncoding
```

HTTPException 的一个子类。

```
exception http.client.UnimplementedFileMode
```

HTTPException 的一个子类。

```
exception http.client.IncompleteRead
```

HTTPException 的一个子类。

```
exception http.client.ImproperConnectionState
```

HTTPException 的一个子类。

```
exception http.client.CannotSendRequest
```

ImproperConnectionState 的一个子类。

exception `http.client.CannotSendHeader`

`ImproperConnectionState` 的一个子类。

exception `http.client.ResponseNotReady`

`ImproperConnectionState` 的一个子类。

exception `http.client.BadStatusLine`

`HTTPException` 的一个子类。如果服务器反馈了一个我们不理解的 HTTP 状态码就会被引发。

exception `http.client.LineTooLong`

`HTTPException` 的一个子类。如果在 HTTP 协议中从服务器接收到过长的行就会被引发。

exception `http.client.RemoteDisconnected`

`ConnectionResetError` 和 `BadStatusLine` 的一个子类。当尝试读取响应时的结果是未从连接读取到数据时由 `HTTPConnection.getresponse()` 引发，表明远端已关闭连接。

3.5 新版功能: 在此之前引发的异常为 `BadStatusLine('')`。

此模块中定义的常量为：

`http.client.HTTP_PORT`

HTTP 协议默认的端口号 (总是 80)。

`http.client.HTTPS_PORT`

HTTPS 协议默认的端口号 (总是 443)。

`http.client.responses`

This dictionary maps the HTTP 1.1 status codes to the W3C names.

Example: `http.client.responses[http.client.NOT_FOUND]` is 'Not Found'.

See [HTTP 状态码](#) for a list of HTTP status codes that are available in this module as constants.

21.12.1 HTTPConnection Objects

`HTTPConnection` instances have the following methods:

`HTTPConnection.request` (*method*, *url*, *body=None*, *headers={}*, *, *encode_chunked=False*)

This will send a request to the server using the HTTP request method *method* and the selector *url*.

If *body* is specified, the specified data is sent after the headers are finished. It may be a *str*, a *bytes-like object*, an open *file object*, or an iterable of *bytes*. If *body* is a string, it is encoded as ISO-8859-1, the default for HTTP. If it is a bytes-like object, the bytes are sent as is. If it is a *file object*, the contents of the file is sent; this file object should support at least the `read()` method. If the file object is an instance of `io.TextIOBase`, the data returned by the `read()` method will be encoded as ISO-8859-1, otherwise the data returned by `read()` is sent as is. If *body* is an iterable, the elements of the iterable are sent as is until the iterable is exhausted.

The *headers* argument should be a mapping of extra HTTP headers to send with the request.

If *headers* contains neither Content-Length nor Transfer-Encoding, but there is a request body, one of those header fields will be added automatically. If *body* is `None`, the Content-Length header is set to 0 for methods that expect a body (PUT, POST, and PATCH). If *body* is a string or a bytes-like object that is not also a *file*, the Content-Length header is set to its length. Any other type of *body* (files and iterables in general) will be chunk-encoded, and the Transfer-Encoding header will automatically be set instead of Content-Length.

The *encode_chunked* argument is only relevant if Transfer-Encoding is specified in *headers*. If *encode_chunked* is `False`, the `HTTPConnection` object assumes that all encoding is handled by the calling code. If it is `True`, the body will be chunk-encoded.

注解: Chunked transfer encoding has been added to the HTTP protocol version 1.1. Unless the HTTP server is known to handle HTTP 1.1, the caller must either specify the Content-Length, or must pass a `str` or bytes-like object that is not also a file as the body representation.

3.2 新版功能: `body` can now be an iterable.

在 3.6 版更改: If neither Content-Length nor Transfer-Encoding are set in `headers`, file and iterable `body` objects are now chunk-encoded. The `encode_chunked` argument was added. No attempt is made to determine the Content-Length for file objects.

`HTTPConnection.getresponse()`

Should be called after a request is sent to get the response from the server. Returns an `HTTPResponse` instance.

注解: Note that you must have read the whole response before you can send a new request to the server.

在 3.5 版更改: If a `ConnectionError` or subclass is raised, the `HTTPConnection` object will be ready to reconnect when a new request is sent.

`HTTPConnection.set_debuglevel(level)`

Set the debugging level. The default debug level is 0, meaning no debugging output is printed. Any value greater than 0 will cause all currently defined debug output to be printed to stdout. The `debuglevel` is passed to any new `HTTPResponse` objects that are created.

3.1 新版功能.

`HTTPConnection.set_tunnel(host, port=None, headers=None)`

Set the host and the port for HTTP Connect Tunnelling. This allows running the connection through a proxy server.

The host and port arguments specify the endpoint of the tunneled connection (i.e. the address included in the CONNECT request, *not* the address of the proxy server).

The headers argument should be a mapping of extra HTTP headers to send with the CONNECT request.

For example, to tunnel through a HTTPS proxy server running locally on port 8080, we would pass the address of the proxy to the `HTTPSConnection` constructor, and the address of the host that we eventually want to reach to the `set_tunnel()` method:

```
>>> import http.client
>>> conn = http.client.HTTPSConnection("localhost", 8080)
>>> conn.set_tunnel("www.python.org")
>>> conn.request("HEAD", "/index.html")
```

3.2 新版功能.

`HTTPConnection.connect()`

Connect to the server specified when the object was created. By default, this is called automatically when making a request if the client does not already have a connection.

`HTTPConnection.close()`

Close the connection to the server.

As an alternative to using the `request()` method described above, you can also send your request step by step, by using the four functions below.

`HTTPConnection.putrequest(method, url, skip_host=False, skip_accept_encoding=False)`

This should be the first call after the connection to the server has been made. It sends a line to the server consisting of the `method` string, the `url` string, and the HTTP version (HTTP/1.1). To disable automatic sending of `Host`:

or `Accept-Encoding`: headers (for example to accept additional content encodings), specify `skip_host` or `skip_accept_encoding` with non-False values.

`HTTPConnection.putheader(header, argument[, ...])`

Send an [RFC 822](#)-style header to the server. It sends a line to the server consisting of the header, a colon and a space, and the first argument. If more arguments are given, continuation lines are sent, each consisting of a tab and an argument.

`HTTPConnection.endheaders(message_body=None, *, encode_chunked=False)`

Send a blank line to the server, signalling the end of the headers. The optional `message_body` argument can be used to pass a message body associated with the request.

If `encode_chunked` is `True`, the result of each iteration of `message_body` will be chunk-encoded as specified in [RFC 7230](#), Section 3.3.1. How the data is encoded is dependent on the type of `message_body`. If `message_body` implements the buffer interface the encoding will result in a single chunk. If `message_body` is a `collections.Iterable`, each iteration of `message_body` will result in a chunk. If `message_body` is a *file object*, each call to `.read()` will result in a chunk. The method automatically signals the end of the chunk-encoded data immediately after `message_body`.

注解: Due to the chunked encoding specification, empty chunks yielded by an iterator body will be ignored by the chunk-encoder. This is to avoid premature termination of the read of the request by the target server due to malformed encoding.

3.6 新版功能: Chunked encoding support. The `encode_chunked` parameter was added.

`HTTPConnection.send(data)`

Send data to the server. This should be used directly only after the `endheaders()` method has been called and before `getresponse()` is called.

21.12.2 HTTPResponse Objects

An `HTTPResponse` instance wraps the HTTP response from the server. It provides access to the request headers and the entity body. The response is an iterable object and can be used in a with statement.

在 3.5 版更改: The `io.BufferedIOBase` interface is now implemented and all of its reader operations are supported.

`HTTPResponse.read([amt])`

Reads and returns the response body, or up to the next `amt` bytes.

`HTTPResponse.readinto(b)`

Reads up to the next `len(b)` bytes of the response body into the buffer `b`. Returns the number of bytes read.

3.3 新版功能.

`HTTPResponse.getheader(name, default=None)`

Return the value of the header `name`, or `default` if there is no header matching `name`. If there is more than one header with the name `name`, return all of the values joined by `'`, `'`. If `'default'` is any iterable other than a single string, its elements are similarly returned joined by commas.

`HTTPResponse.getheaders()`

Return a list of (header, value) tuples.

`HTTPResponse.fileno()`

Return the `fileno` of the underlying socket.

HTTPResponse.msg

A `http.client.HTTPMessage` instance containing the response headers. `http.client.HTTPMessage` is a subclass of `email.message.Message`.

HTTPResponse.version

HTTP protocol version used by server. 10 for HTTP/1.0, 11 for HTTP/1.1.

HTTPResponse.status

Status code returned by server.

HTTPResponse.reason

Reason phrase returned by server.

HTTPResponse.debuglevel

A debugging hook. If `debuglevel` is greater than zero, messages will be printed to stdout as the response is read and parsed.

HTTPResponse.closed

Is True if the stream is closed.

21.12.3 示例

Here is an example session that uses the GET method:

```
>>> import http.client
>>> conn = http.client.HTTPSConnection("www.python.org")
>>> conn.request("GET", "/")
>>> r1 = conn.getresponse()
>>> print(r1.status, r1.reason)
200 OK
>>> data1 = r1.read() # This will return entire content.
>>> # The following example demonstrates reading data in chunks.
>>> conn.request("GET", "/")
>>> r1 = conn.getresponse()
>>> while not r1.closed:
...     print(r1.read(200)) # 200 bytes
b'<!doctype html>\n<!--[if"...
...
>>> # Example of an invalid request
>>> conn.request("GET", "/parrot.spam")
>>> r2 = conn.getresponse()
>>> print(r2.status, r2.reason)
404 Not Found
>>> data2 = r2.read()
>>> conn.close()
```

Here is an example session that uses the HEAD method. Note that the HEAD method never returns any data.

```
>>> import http.client
>>> conn = http.client.HTTPSConnection("www.python.org")
>>> conn.request("HEAD", "/")
>>> res = conn.getresponse()
>>> print(res.status, res.reason)
200 OK
>>> data = res.read()
>>> print(len(data))
0
```

(下页继续)

(续上页)

```
>>> data == b''
True
```

Here is an example session that shows how to POST requests:

```
>>> import http.client, urllib.parse
>>> params = urllib.parse.urlencode({'@number': 12524, '@type': 'issue', '@action':
↳ 'show'})
>>> headers = {"Content-type": "application/x-www-form-urlencoded",
...           "Accept": "text/plain"}
>>> conn = http.client.HTTPConnection("bugs.python.org")
>>> conn.request("POST", "", params, headers)
>>> response = conn.getresponse()
>>> print(response.status, response.reason)
302 Found
>>> data = response.read()
>>> data
b'Redirecting to <a href="http://bugs.python.org/issue12524">http://bugs.python.org/
↳ issue12524</a>'
>>> conn.close()
```

Client side HTTP PUT requests are very similar to POST requests. The difference lies only the server side where HTTP server will allow resources to be created via PUT request. It should be noted that custom HTTP methods +are also handled in `urllib.request.Request` by sending the appropriate +method attribute. Here is an example session that shows how to do PUT request using `http.client`:

```
>>> # This creates an HTTP message
>>> # with the content of BODY as the enclosed representation
>>> # for the resource http://localhost:8080/file
...
>>> import http.client
>>> BODY = """filecontents"""
>>> conn = http.client.HTTPConnection("localhost", 8080)
>>> conn.request("PUT", "/file", BODY)
>>> response = conn.getresponse()
>>> print(response.status, response.reason)
200, OK
```

21.12.4 HTTPMessage Objects

An `http.client.HTTPMessage` instance holds the headers from an HTTP response. It is implemented using the `email.message.Message` class.

21.13 ftplib —FTP protocol client

Source code: `Lib/ftplib.py`

This module defines the class `FTP` and a few related items. The `FTP` class implements the client side of the FTP protocol. You can use this to write Python programs that perform a variety of automated FTP jobs, such as mirroring other FTP servers. It is also used by the module `urllib.request` to handle URLs that use FTP. For more information on FTP (File Transfer Protocol), see Internet [RFC 959](#).

Here's a sample session using the `ftplib` module:

```
>>> from ftplib import FTP
>>> ftp = FTP('ftp.debian.org')      # connect to host, default port
>>> ftp.login()                      # user anonymous, passwd anonymous@
'230 Login successful.'
>>> ftp.cwd('debian')               # change into "debian" directory
>>> ftp.retrlines('LIST')           # list directory contents
-rw-rw-r-- 1 1176      1176      1063 Jun 15 10:18 README
...
drwxr-sr-x 5 1176      1176      4096 Dec 19 2000 pool
drwxr-sr-x 4 1176      1176      4096 Nov 17 2008 project
drwxr-xr-x 3 1176      1176      4096 Oct 10 2012 tools
'226 Directory send OK.'
>>> ftp.retrbinary('RETR README', open('README', 'wb').write)
'226 Transfer complete.'
>>> ftp.quit()
```

这个模块定义了以下内容:

class `ftplib.FTP` (*host=""*, *user=""*, *passwd=""*, *acct=""*, *timeout=None*, *source_address=None*)

Return a new instance of the `FTP` class. When *host* is given, the method call `connect(host)` is made. When *user* is given, additionally the method call `login(user, passwd, acct)` is made (where *passwd* and *acct* default to the empty string when not given). The optional *timeout* parameter specifies a timeout in seconds for blocking operations like the connection attempt (if is not specified, the global default timeout setting will be used). *source_address* is a 2-tuple (*host*, *port*) for the socket to bind to as its source address before connecting.

The `FTP` class supports the `with` statement, e.g.:

```
>>> from ftplib import FTP
>>> with FTP("ftp1.at.proftpd.org") as ftp:
...     ftp.login()
...     ftp.dir()
... # doctest: +SKIP
'230 Anonymous login ok, restrictions apply.'
dr-xr-xr-x 9 ftp      ftp      154 May 6 10:43 .
dr-xr-xr-x 9 ftp      ftp      154 May 6 10:43 ..
dr-xr-xr-x 5 ftp      ftp      4096 May 6 10:43 CentOS
dr-xr-xr-x 3 ftp      ftp      18 Jul 10 2008 Fedora
>>>
```

在 3.2 版更改: 支持了 `with` 语句。

在 3.3 版更改: *source_address* parameter was added.

class `ftplib.FTP_TLS` (*host=""*, *user=""*, *passwd=""*, *acct=""*, *keyfile=None*, *certfile=None*, *context=None*, *timeout=None*, *source_address=None*)

A `FTP` subclass which adds TLS support to `FTP` as described in [RFC 4217](#). Connect as usual to port 21 implicitly securing the FTP control connection before authenticating. Securing the data connection requires the user to explicitly ask for it by calling the `prot_p()` method. *context* is a `ssl.SSLContext` object which allows bundling SSL configuration options, certificates and private keys into a single (potentially long-lived) structure. Please read [Security considerations](#) for best practices.

keyfile and *certfile* are a legacy alternative to *context* – they can point to PEM-formatted private key and certificate chain files (respectively) for the SSL connection.

3.2 新版功能.

在 3.3 版更改: *source_address* parameter was added.

在 3.4 版更改: The class now supports hostname check with `ssl.SSLContext.check_hostname` and *Server Name Indication* (see `ssl.HAS_SNI`).

3.6 版后已移除: `keyfile` and `certfile` are deprecated in favor of `context`. Please use `ssl.SSLContext.load_cert_chain()` instead, or let `ssl.create_default_context()` select the system's trusted CA certificates for you.

Here's a sample session using the `FTP_TLS` class:

```
>>> ftps = FTP_TLS('ftp.pureftpd.org')
>>> ftps.login()
'230 Anonymous user logged in'
>>> ftps.prot_p()
'200 Data protection level set to "private"'
>>> ftps.nlst()
['6jack', 'OpenBSD', 'antilink', 'blogbench', 'bsdcam', 'clockspeed', 'djbdns-jedi
↪', 'docs', 'eaccelerator-jedi', 'favicon.ico', 'francotone', 'fugu', 'ignore',
↪ 'libpuzzle', 'metalog', 'minidentd', 'misc', 'mysql-udf-global-user-variables',
↪ 'php-jenkins-hash', 'php-skein-hash', 'php-webdav', 'phpaudit', 'phpbench',
↪ 'pincaster', 'ping', 'posto', 'pub', 'public', 'public_keys', 'pure-ftp',
↪ 'qscan', 'qtc', 'sharedance', 'skycache', 'sound', 'tmp', 'ucarp']
```

exception `ftplib.error_reply`

Exception raised when an unexpected reply is received from the server.

exception `ftplib.error_temp`

Exception raised when an error code signifying a temporary error (response codes in the range 400–499) is received.

exception `ftplib.error_perm`

Exception raised when an error code signifying a permanent error (response codes in the range 500–599) is received.

exception `ftplib.error_proto`

Exception raised when a reply is received from the server that does not fit the response specifications of the File Transfer Protocol, i.e. begin with a digit in the range 1–5.

`ftplib.all_errors`

The set of all exceptions (as a tuple) that methods of `FTP` instances may raise as a result of problems with the FTP connection (as opposed to programming errors made by the caller). This set includes the four exceptions listed above as well as `OSError`.

参见:

Module `netrc` Parser for the `.netrc` file format. The file `.netrc` is typically used by FTP clients to load user authentication information before prompting the user.

21.13.1 FTP Objects

Several methods are available in two flavors: one for handling text files and another for binary files. These are named for the command which is used followed by `lines` for the text version or `binary` for the binary version.

`FTP` instances have the following methods:

`FTP.set_debuglevel` (*level*)

Set the instance's debugging level. This controls the amount of debugging output printed. The default, 0, produces no debugging output. A value of 1 produces a moderate amount of debugging output, generally a single line per request. A value of 2 or higher produces the maximum amount of debugging output, logging each line sent and received on the control connection.

`FTP.connect` (*host*=, *port*=0, *timeout*=None, *source_address*=None)

Connect to the given host and port. The default port number is 21, as specified by the FTP protocol specification.

It is rarely needed to specify a different port number. This function should be called only once for each instance; it should not be called at all if a host was given when the instance was created. All other methods can only be used after a connection has been made. The optional *timeout* parameter specifies a timeout in seconds for the connection attempt. If no *timeout* is passed, the global default timeout setting will be used. *source_address* is a 2-tuple (host, port) for the socket to bind to as its source address before connecting.

在 3.3 版更改: *source_address* parameter was added.

FTP.**getwelcome** ()

Return the welcome message sent by the server in reply to the initial connection. (This message sometimes contains disclaimers or help information that may be relevant to the user.)

FTP.**login** (user='anonymous', passwd="", acct="")

Log in as the given *user*. The *passwd* and *acct* parameters are optional and default to the empty string. If no *user* is specified, it defaults to 'anonymous'. If *user* is 'anonymous', the default *passwd* is 'anonymous@'. This function should be called only once for each instance, after a connection has been established; it should not be called at all if a host and user were given when the instance was created. Most FTP commands are only allowed after the client has logged in. The *acct* parameter supplies “accounting information”; few systems implement this.

FTP.**abort** ()

Abort a file transfer that is in progress. Using this does not always work, but it's worth a try.

FTP.**sendcmd** (cmd)

Send a simple command string to the server and return the response string.

FTP.**voidcmd** (cmd)

Send a simple command string to the server and handle the response. Return nothing if a response code corresponding to success (codes in the range 200–299) is received. Raise *error_reply* otherwise.

FTP.**retrbinary** (cmd, callback, blocksize=8192, rest=None)

Retrieve a file in binary transfer mode. *cmd* should be an appropriate RETR command: 'RETR filename'. The *callback* function is called for each block of data received, with a single bytes argument giving the data block. The optional *blocksize* argument specifies the maximum chunk size to read on the low-level socket object created to do the actual transfer (which will also be the largest size of the data blocks passed to *callback*). A reasonable default is chosen. *rest* means the same thing as in the *transfercmd*() method.

FTP.**retrlines** (cmd, callback=None)

Retrieve a file or directory listing in ASCII transfer mode. *cmd* should be an appropriate RETR command (see *retrbinary*()) or a command such as LIST or NLST (usually just the string 'LIST'). LIST retrieves a list of files and information about those files. NLST retrieves a list of file names. The *callback* function is called for each line with a string argument containing the line with the trailing CRLF stripped. The default *callback* prints the line to *sys.stdout*.

FTP.**set_pasv** (val)

Enable “passive” mode if *val* is true, otherwise disable passive mode. Passive mode is on by default.

FTP.**storbinary** (cmd, fp, blocksize=8192, callback=None, rest=None)

Store a file in binary transfer mode. *cmd* should be an appropriate STOR command: "STOR filename". *fp* is a *file object* (opened in binary mode) which is read until EOF using its *read*() method in blocks of size *blocksize* to provide the data to be stored. The *blocksize* argument defaults to 8192. *callback* is an optional single parameter callable that is called on each block of data after it is sent. *rest* means the same thing as in the *transfercmd*() method.

在 3.2 版更改: *rest* parameter added.

FTP.**storlines** (cmd, fp, callback=None)

Store a file in ASCII transfer mode. *cmd* should be an appropriate STOR command (see *storbinary*()). Lines are read until EOF from the *file object* *fp* (opened in binary mode) using its *readline*() method to provide the data to be stored. *callback* is an optional single parameter callable that is called on each line after it is sent.

FTP.**transfercmd** (*cmd*, *rest=None*)

Initiate a transfer over the data connection. If the transfer is active, send an EPRT or PORT command and the transfer command specified by *cmd*, and accept the connection. If the server is passive, send an EPSV or PASV command, connect to it, and start the transfer command. Either way, return the socket for the connection.

If optional *rest* is given, a REST command is sent to the server, passing *rest* as an argument. *rest* is usually a byte offset into the requested file, telling the server to restart sending the file's bytes at the requested offset, skipping over the initial bytes. Note however that RFC 959 requires only that *rest* be a string containing characters in the printable range from ASCII code 33 to ASCII code 126. The `transfercmd()` method, therefore, converts *rest* to a string, but no check is performed on the string's contents. If the server does not recognize the REST command, an `error_reply` exception will be raised. If this happens, simply call `transfercmd()` without a *rest* argument.

FTP.**nttransfercmd** (*cmd*, *rest=None*)

Like `transfercmd()`, but returns a tuple of the data connection and the expected size of the data. If the expected size could not be computed, `None` will be returned as the expected size. *cmd* and *rest* means the same thing as in `transfercmd()`.

FTP.**mlsd** (*path=""*, *facts=[]*)

List a directory in a standardized format by using MLSD command (RFC 3659). If *path* is omitted the current directory is assumed. *facts* is a list of strings representing the type of information desired (e.g. ["type", "size", "perm"]). Return a generator object yielding a tuple of two elements for every file found in *path*. First element is the file name, the second one is a dictionary containing facts about the file name. Content of this dictionary might be limited by the *facts* argument but server is not guaranteed to return all requested facts.

3.3 新版功能.

FTP.**nlst** (*argument* [, ...])

Return a list of file names as returned by the NLST command. The optional *argument* is a directory to list (default is the current server directory). Multiple arguments can be used to pass non-standard options to the NLST command.

注解: If your server supports the command, `mlsd()` offers a better API.

FTP.**dir** (*argument* [, ...])

Produce a directory listing as returned by the LIST command, printing it to standard output. The optional *argument* is a directory to list (default is the current server directory). Multiple arguments can be used to pass non-standard options to the LIST command. If the last argument is a function, it is used as a *callback* function as for `retrlines()`; the default prints to `sys.stdout`. This method returns `None`.

注解: If your server supports the command, `mlsd()` offers a better API.

FTP.**rename** (*fromname*, *toname*)

Rename file *fromname* on the server to *toname*.

FTP.**delete** (*filename*)

Remove the file named *filename* from the server. If successful, returns the text of the response, otherwise raises `error_perm` on permission errors or `error_reply` on other errors.

FTP.**cwd** (*pathname*)

Set the current directory on the server.

FTP.**mkd** (*pathname*)

Create a new directory on the server.

FTP.**pwd** ()

Return the pathname of the current directory on the server.

`FTP.rmd (dirname)`

Remove the directory named *dirname* on the server.

`FTP.size (filename)`

Request the size of the file named *filename* on the server. On success, the size of the file is returned as an integer, otherwise `None` is returned. Note that the `SIZE` command is not standardized, but is supported by many common server implementations.

`FTP.quit ()`

Send a `QUIT` command to the server and close the connection. This is the “polite” way to close a connection, but it may raise an exception if the server responds with an error to the `QUIT` command. This implies a call to the `close ()` method which renders the `FTP` instance useless for subsequent calls (see below).

`FTP.close ()`

Close the connection unilaterally. This should not be applied to an already closed connection such as after a successful call to `quit ()`. After this call the `FTP` instance should not be used any more (after a call to `close ()` or `quit ()` you cannot reopen the connection by issuing another `login ()` method).

21.13.2 FTP_TLS Objects

`FTP_TLS` class inherits from `FTP`, defining these additional objects:

`FTP_TLS.ssl_version`

The SSL version to use (defaults to `ssl.PROTOCOL_SSLv23`).

`FTP_TLS.auth ()`

Set up a secure control connection by using TLS or SSL, depending on what is specified in the `ssl_version` attribute.

在 3.4 版更改: The method now supports hostname check with `ssl.SSLContext.check_hostname` and *Server Name Indication* (see `ssl.HAS_SNI`).

`FTP_TLS.ccc ()`

Revert control channel back to plaintext. This can be useful to take advantage of firewalls that know how to handle NAT with non-secure FTP without opening fixed ports.

3.3 新版功能.

`FTP_TLS.prot_p ()`

Set up secure data connection.

`FTP_TLS.prot_c ()`

Set up clear text data connection.

21.14 poplib —POP3 protocol client

Source code: `Lib/poplib.py`

This module defines a class, `POP3`, which encapsulates a connection to a POP3 server and implements the protocol as defined in [RFC 1939](#). The `POP3` class supports both the minimal and optional command sets from [RFC 1939](#). The `POP3` class also supports the `STLS` command introduced in [RFC 2595](#) to enable encrypted communication on an already established connection.

Additionally, this module provides a class `POP3_SSL`, which provides support for connecting to POP3 servers that use SSL as an underlying protocol layer.

Note that POP3, though widely supported, is obsolescent. The implementation quality of POP3 servers varies widely, and too many are quite poor. If your mailserver supports IMAP, you would be better off using the `imaplib.IMAP4` class, as IMAP servers tend to be better implemented.

The `poplib` module provides two classes:

class `poplib.POP3` (*host*, *port*=`POP3_PORT`[, *timeout*])

This class implements the actual POP3 protocol. The connection is created when the instance is initialized. If *port* is omitted, the standard POP3 port (110) is used. The optional *timeout* parameter specifies a timeout in seconds for the connection attempt (if not specified, the global default timeout setting will be used).

class `poplib.POP3_SSL` (*host*, *port*=`POP3_SSL_PORT`, *keyfile*=`None`, *certfile*=`None`, *timeout*=`None`, *context*=`None`)

This is a subclass of `POP3` that connects to the server over an SSL encrypted socket. If *port* is not specified, 995, the standard POP3-over-SSL port is used. *timeout* works as in the `POP3` constructor. *context* is an optional `ssl.SSLContext` object which allows bundling SSL configuration options, certificates and private keys into a single (potentially long-lived) structure. Please read *Security considerations* for best practices.

keyfile and *certfile* are a legacy alternative to *context* - they can point to PEM-formatted private key and certificate chain files, respectively, for the SSL connection.

在 3.2 版更改: *context* parameter added.

在 3.4 版更改: The class now supports hostname check with `ssl.SSLContext.check_hostname` and *Server Name Indication* (see `ssl.HAS_SNI`).

3.6 版后已移除: *keyfile* and *certfile* are deprecated in favor of *context*. Please use `ssl.SSLContext.load_cert_chain()` instead, or let `ssl.create_default_context()` select the system's trusted CA certificates for you.

One exception is defined as an attribute of the `poplib` module:

exception `poplib.error_proto`

Exception raised on any errors from this module (errors from `socket` module are not caught). The reason for the exception is passed to the constructor as a string.

参见:

Module `imaplib` The standard Python IMAP module.

Frequently Asked Questions About Fetchmail The FAQ for the `fetchmail` POP/IMAP client collects information on POP3 server variations and RFC noncompliance that may be useful if you need to write an application based on the POP protocol.

21.14.1 POP3 Objects

All POP3 commands are represented by methods of the same name, in lower-case; most return the response text sent by the server.

An `POP3` instance has the following methods:

`POP3.set_debuglevel` (*level*)

Set the instance's debugging level. This controls the amount of debugging output printed. The default, 0, produces no debugging output. A value of 1 produces a moderate amount of debugging output, generally a single line per request. A value of 2 or higher produces the maximum amount of debugging output, logging each line sent and received on the control connection.

`POP3.getwelcome` ()

Returns the greeting string sent by the POP3 server.

POP3.**capa**()

Query the server's capabilities as specified in [RFC 2449](#). Returns a dictionary in the form {'name': ['param'...] }.

3.4 新版功能.

POP3.**user**(username)

Send user command, response should indicate that a password is required.

POP3.**pass_**(password)

Send password, response includes message count and mailbox size. Note: the mailbox on the server is locked until `quit()` is called.

POP3.**apop**(user, secret)

Use the more secure APOP authentication to log into the POP3 server.

POP3.**rpop**(user)

Use RPOP authentication (similar to UNIX r-commands) to log into POP3 server.

POP3.**stat**()

Get mailbox status. The result is a tuple of 2 integers: (message count, mailbox size).

POP3.**list**([which])

Request message list, result is in the form (response, ['mesg_num octets', ...], octets). If *which* is set, it is the message to list.

POP3.**retr**(which)

Retrieve whole message number *which*, and set its seen flag. Result is in form (response, ['line', ...], octets).

POP3.**dele**(which)

Flag message number *which* for deletion. On most servers deletions are not actually performed until QUIT (the major exception is Eudora QPOP, which deliberately violates the RFCs by doing pending deletes on any disconnect).

POP3.**rset**()

Remove any deletion marks for the mailbox.

POP3.**noop**()

Do nothing. Might be used as a keep-alive.

POP3.**quit**()

Signoff: commit changes, unlock mailbox, drop connection.

POP3.**top**(which, howmuch)

Retrieves the message header plus *howmuch* lines of the message after the header of message number *which*. Result is in form (response, ['line', ...], octets).

The POP3 TOP command this method uses, unlike the RETR command, doesn't set the message's seen flag; unfortunately, TOP is poorly specified in the RFCs and is frequently broken in off-brand servers. Test this method by hand against the POP3 servers you will use before trusting it.

POP3.**uidl**(which=None)

Return message digest (unique id) list. If *which* is specified, result contains the unique id for that message in the form 'response mesgnum uid', otherwise result is list (response, ['mesgnum uid', ...], octets).

POP3.**utf8**()

Try to switch to UTF-8 mode. Returns the server response if successful, raises `error_proto` if not. Specified in [RFC 6856](#).

3.5 新版功能.

`POP3.stls` (*context=None*)

Start a TLS session on the active connection as specified in [RFC 2595](#). This is only allowed before user authentication

context parameter is a `ssl.SSLContext` object which allows bundling SSL configuration options, certificates and private keys into a single (potentially long-lived) structure. Please read [Security considerations](#) for best practices.

This method supports hostname checking via `ssl.SSLContext.check_hostname` and *Server Name Indication* (see `ssl.HAS_SNI`).

3.4 新版功能.

Instances of `POP3_SSL` have no additional methods. The interface of this subclass is identical to its parent.

21.14.2 POP3 Example

Here is a minimal example (without error checking) that opens a mailbox and retrieves and prints all messages:

```
import getpass, poplib

M = poplib.POP3('localhost')
M.user(getpass.getuser())
M.pass_(getpass.getpass())
numMessages = len(M.list()[1])
for i in range(numMessages):
    for j in M.retr(i+1)[1]:
        print(j)
```

At the end of the module, there is a test section that contains a more extensive example of usage.

21.15 imaplib —IMAP4 protocol client

Source code: `Lib/imaplib.py`

This module defines three classes, `IMAP4`, `IMAP4_SSL` and `IMAP4_stream`, which encapsulate a connection to an IMAP4 server and implement a large subset of the IMAP4rev1 client protocol as defined in [RFC 2060](#). It is backward compatible with IMAP4 ([RFC 1730](#)) servers, but note that the `STATUS` command is not supported in IMAP4.

Three classes are provided by the `imaplib` module, `IMAP4` is the base class:

class `imaplib.IMAP4` (*host=*”, *port=IMAP4_PORT*)

This class implements the actual IMAP4 protocol. The connection is created and protocol version (IMAP4 or IMAP4rev1) is determined when the instance is initialized. If *host* is not specified, `''` (the local host) is used. If *port* is omitted, the standard IMAP4 port (143) is used.

The `IMAP4` class supports the `with` statement. When used like this, the IMAP4 `LOGOUT` command is issued automatically when the `with` statement exits. E.g.:

```
>>> from imaplib import IMAP4
>>> with IMAP4("domain.org") as M:
...     M.noop()
...
('OK', [b'Nothing Accomplished. d25if65hy903weo.87'])
```

在 3.5 版更改: 支持了 `with` 语句。

Three exceptions are defined as attributes of the `IMAP4` class:

exception `IMAP4.error`

Exception raised on any errors. The reason for the exception is passed to the constructor as a string.

exception `IMAP4.abort`

IMAP4 server errors cause this exception to be raised. This is a sub-class of `IMAP4.error`. Note that closing the instance and instantiating a new one will usually allow recovery from this exception.

exception `IMAP4.readonly`

This exception is raised when a writable mailbox has its status changed by the server. This is a sub-class of `IMAP4.error`. Some other client now has write permission, and the mailbox will need to be re-opened to re-obtain write permission.

There's also a subclass for secure connections:

class `imaplib.IMAP4_SSL` (*host*=", *port*=`IMAP4_SSL_PORT`, *keyfile*=None, *certfile*=None, *ssl_context*=None)

This is a subclass derived from `IMAP4` that connects over an SSL encrypted socket (to use this class you need a socket module that was compiled with SSL support). If *host* is not specified, `' '` (the local host) is used. If *port* is omitted, the standard IMAP4-over-SSL port (993) is used. *ssl_context* is a `ssl.SSLContext` object which allows bundling SSL configuration options, certificates and private keys into a single (potentially long-lived) structure. Please read *Security considerations* for best practices.

keyfile and *certfile* are a legacy alternative to *ssl_context* - they can point to PEM-formatted private key and certificate chain files for the SSL connection. Note that the *keyfile/certfile* parameters are mutually exclusive with *ssl_context*, a `ValueError` is raised if *keyfile/certfile* is provided along with *ssl_context*.

在 3.3 版更改: *ssl_context* parameter added.

在 3.4 版更改: The class now supports hostname check with `ssl.SSLContext.check_hostname` and *Server Name Indication* (see `ssl.HAS_SNI`).

3.6 版后已移除: *keyfile* and *certfile* are deprecated in favor of *ssl_context*. Please use `ssl.SSLContext.load_cert_chain()` instead, or let `ssl.create_default_context()` select the system's trusted CA certificates for you.

The second subclass allows for connections created by a child process:

class `imaplib.IMAP4_stream` (*command*)

This is a subclass derived from `IMAP4` that connects to the `stdin/stdout` file descriptors created by passing *command* to `subprocess.Popen()`.

The following utility functions are defined:

`imaplib.Internaldate2tuple` (*datestr*)

Parse an IMAP4 INTERNALDATE string and return corresponding local time. The return value is a `time.struct_time` tuple or None if the string has wrong format.

`imaplib.Int2AP` (*num*)

Converts an integer into a string representation using characters from the set [A .. P].

`imaplib.ParseFlags` (*flagstr*)

Converts an IMAP4 FLAGS response to a tuple of individual flags.

`imaplib.Time2Internaldate` (*date_time*)

Convert *date_time* to an IMAP4 INTERNALDATE representation. The return value is a string in the form: "DD-Mmm-YYYY HH:MM:SS +HHMM" (including double-quotes). The *date_time* argument can be a number (int or float) representing seconds since epoch (as returned by `time.time()`), a 9-tuple representing local time an instance of `time.struct_time` (as returned by `time.localtime()`), an aware instance of `datetime.datetime`, or a double-quoted string. In the last case, it is assumed to already be in the correct format.

Note that IMAP4 message numbers change as the mailbox changes; in particular, after an `EXPUNGE` command performs deletions the remaining messages are renumbered. So it is highly advisable to use UIDs instead, with the `UID` command.

At the end of the module, there is a test section that contains a more extensive example of usage.

参见:

Documents describing the protocol, and sources and binaries for servers implementing it, can all be found at the University of Washington's *IMAP Information Center* (<https://www.washington.edu/imap/>).

21.15.1 IMAP4 Objects

All IMAP4rev1 commands are represented by methods of the same name, either upper-case or lower-case.

All arguments to commands are converted to strings, except for `AUTHENTICATE`, and the last argument to `APPEND` which is passed as an IMAP4 literal. If necessary (the string contains IMAP4 protocol-sensitive characters and isn't enclosed with either parentheses or double quotes) each string is quoted. However, the *password* argument to the `LOGIN` command is always quoted. If you want to avoid having an argument string quoted (eg: the *flags* argument to `STORE`) then enclose the string in parentheses (eg: `r'(\Deleted)'`).

Each command returns a tuple: `(type, [data, ...])` where *type* is usually `'OK'` or `'NO'`, and *data* is either the text from the command response, or mandated results from the command. Each *data* is either a string, or a tuple. If a tuple, then the first part is the header of the response, and the second part contains the data (ie: 'literal' value).

The *message_set* options to commands below is a string specifying one or more messages to be acted upon. It may be a simple message number (`'1'`), a range of message numbers (`'2:4'`), or a group of non-contiguous ranges separated by commas (`'1:3,6:9'`). A range can contain an asterisk to indicate an infinite upper bound (`'3:*`').

An *IMAP4* instance has the following methods:

`IMAP4.append(mailbox, flags, date_time, message)`

Append *message* to named mailbox.

`IMAP4.authenticate(mechanism, authobject)`

Authenticate command —requires response processing.

mechanism specifies which authentication mechanism is to be used - it should appear in the instance variable *capabilities* in the form `AUTH=mechanism`.

authobject must be a callable object:

```
data = authobject(response)
```

It will be called to process server continuation responses; the *response* argument it is passed will be `bytes`. It should return `bytes` *data* that will be base64 encoded and sent to the server. It should return `None` if the client abort response `*` should be sent instead.

在 3.5 版更改: string usernames and passwords are now encoded to `utf-8` instead of being limited to ASCII.

`IMAP4.check()`

Checkpoint mailbox on server.

`IMAP4.close()`

Close currently selected mailbox. Deleted messages are removed from writable mailbox. This is the recommended command before `LOGOUT`.

`IMAP4.copy(message_set, new_mailbox)`

Copy *message_set* messages onto end of *new_mailbox*.

`IMAP4.create(mailbox)`

Create new mailbox named *mailbox*.

IMAP4.delete (*mailbox*)
Delete old mailbox named *mailbox*.

IMAP4.deleteacl (*mailbox*, *who*)
Delete the ACLs (remove any rights) set for *who* on mailbox.

IMAP4.enable (*capability*)
Enable *capability* (see [RFC 5161](#)). Most capabilities do not need to be enabled. Currently only the UTF8=ACCEPT capability is supported (see [RFC 6855](#)).

3.5 新版功能: The *enable()* method itself, and [RFC 6855](#) support.

IMAP4.expunge ()
Permanently remove deleted items from selected mailbox. Generates an EXPUNGE response for each deleted message. Returned data contains a list of EXPUNGE message numbers in order received.

IMAP4.fetch (*message_set*, *message_parts*)
Fetch (parts of) messages. *message_parts* should be a string of message part names enclosed within parentheses, eg: " (UID BODY [TEXT]) ". Returned data are tuples of message part envelope and data.

IMAP4.getacl (*mailbox*)
Get the ACLs for *mailbox*. The method is non-standard, but is supported by the Cyrus server.

IMAP4.getannotation (*mailbox*, *entry*, *attribute*)
Retrieve the specified ANNOTATIONS for *mailbox*. The method is non-standard, but is supported by the Cyrus server.

IMAP4.getquota (*root*)
Get the quota *root*'s resource usage and limits. This method is part of the IMAP4 QUOTA extension defined in rfc2087.

IMAP4.getquotaroot (*mailbox*)
Get the list of quota roots for the named *mailbox*. This method is part of the IMAP4 QUOTA extension defined in rfc2087.

IMAP4.list ([*directory*[, *pattern*]])
List mailbox names in *directory* matching *pattern*. *directory* defaults to the top-level mail folder, and *pattern* defaults to match anything. Returned data contains a list of LIST responses.

IMAP4.login (*user*, *password*)
Identify the client using a plaintext password. The *password* will be quoted.

IMAP4.login_cram_md5 (*user*, *password*)
Force use of CRAM-MD5 authentication when identifying the client to protect the password. Will only work if the server CAPABILITY response includes the phrase AUTH=CRAM-MD5.

IMAP4.logout ()
Shutdown connection to server. Returns server BYE response.

IMAP4.lsub (*directory*="*", *pattern*='*')
List subscribed mailbox names in *directory* matching *pattern*. *directory* defaults to the top level directory and *pattern* defaults to match any mailbox. Returned data are tuples of message part envelope and data.

IMAP4.myrights (*mailbox*)
Show my ACLs for a mailbox (i.e. the rights that I have on mailbox).

IMAP4.namespace ()
Returns IMAP namespaces as defined in [RFC 2342](#).

IMAP4.noop ()
Send NOOP to server.

`IMAP4.open(host, port)`

Opens socket to *port* at *host*. This method is implicitly called by the `IMAP4` constructor. The connection objects established by this method will be used in the `IMAP4.read()`, `IMAP4.readline()`, `IMAP4.send()`, and `IMAP4.shutdown()` methods. You may override this method.

`IMAP4.partial(message_num, message_part, start, length)`

Fetch truncated part of a message. Returned data is a tuple of message part envelope and data.

`IMAP4.proxyauth(user)`

Assume authentication as *user*. Allows an authorised administrator to proxy into any user's mailbox.

`IMAP4.read(size)`

Reads *size* bytes from the remote server. You may override this method.

`IMAP4.readline()`

Reads one line from the remote server. You may override this method.

`IMAP4.recent()`

Prompt server for an update. Returned data is `None` if no new messages, else value of RECENT response.

`IMAP4.rename(oldmailbox, newmailbox)`

Rename mailbox named *oldmailbox* to *newmailbox*.

`IMAP4.response(code)`

Return data for response *code* if received, or `None`. Returns the given code, instead of the usual type.

`IMAP4.search(charset, criterion[, ...])`

Search mailbox for matching messages. *charset* may be `None`, in which case no CHARSET will be specified in the request to the server. The IMAP protocol requires that at least one criterion be specified; an exception will be raised when the server returns an error. *charset* must be `None` if the UTF8=ACCEPT capability was enabled using the `enable()` command.

示例:

```
# M is a connected IMAP4 instance...
typ, msgnums = M.search(None, 'FROM', '"LDJ"')

# or:
typ, msgnums = M.search(None, '(FROM "LDJ")')
```

`IMAP4.select(mailbox='INBOX', readonly=False)`

Select a mailbox. Returned data is the count of messages in *mailbox* (EXISTS response). The default *mailbox* is 'INBOX'. If the *readonly* flag is set, modifications to the mailbox are not allowed.

`IMAP4.send(data)`

Sends data to the remote server. You may override this method.

`IMAP4.setacl(mailbox, who, what)`

Set an ACL for *mailbox*. The method is non-standard, but is supported by the Cyrus server.

`IMAP4.setannotation(mailbox, entry, attribute[, ...])`

Set ANNOTATIONS for *mailbox*. The method is non-standard, but is supported by the Cyrus server.

`IMAP4.setquota(root, limits)`

Set the quota *root*'s resource *limits*. This method is part of the IMAP4 QUOTA extension defined in rfc2087.

`IMAP4.shutdown()`

Close connection established in `open`. This method is implicitly called by `IMAP4.logout()`. You may override this method.

`IMAP4.socket()`

Returns socket instance used to connect to server.

IMAP4.**sort** (*sort_criteria*, *charset*, *search_criterion*[, ...])

The `sort` command is a variant of `search` with sorting semantics for the results. Returned data contains a space separated list of matching message numbers.

Sort has two arguments before the *search_criterion* argument(s); a parenthesized list of *sort_criteria*, and the searching *charset*. Note that unlike `search`, the searching *charset* argument is mandatory. There is also a `uid sort` command which corresponds to `sort` the way that `uid search` corresponds to `search`. The `sort` command first searches the mailbox for messages that match the given searching criteria using the *charset* argument for the interpretation of strings in the searching criteria. It then returns the numbers of matching messages.

This is an IMAP4rev1 extension command.

IMAP4.**starttls** (*ssl_context*=None)

Send a STARTTLS command. The *ssl_context* argument is optional and should be a `ssl.SSLContext` object. This will enable encryption on the IMAP connection. Please read [Security considerations](#) for best practices.

3.2 新版功能.

在 3.4 版更改: The method now supports hostname check with `ssl.SSLContext.check_hostname` and *Server Name Indication* (see `ssl.HAS_SNI`).

IMAP4.**status** (*mailbox*, *names*)

Request named status conditions for *mailbox*.

IMAP4.**store** (*message_set*, *command*, *flag_list*)

Alters flag dispositions for messages in mailbox. *command* is specified by section 6.4.6 of [RFC 2060](#) as being one of “FLAGS”, “+FLAGS”, or “-FLAGS”, optionally with a suffix of “.SILENT”.

For example, to set the delete flag on all messages:

```
typ, data = M.search(None, 'ALL')
for num in data[0].split():
    M.store(num, '+FLAGS', '\\Deleted')
M.expunge()
```

注解: Creating flags containing ‘`]`’ (for example: “[test]”) violates [RFC 3501](#) (the IMAP protocol). However, `imaplib` has historically allowed creation of such tags, and popular IMAP servers, such as Gmail, accept and produce such flags. There are non-Python programs which also create such tags. Although it is an RFC violation and IMAP clients and servers are supposed to be strict, `imaplib` nonetheless continues to allow such tags to be created for backward compatibility reasons, and as of python 3.6, handles them if they are sent from the server, since this improves real-world compatibility.

IMAP4.**subscribe** (*mailbox*)

Subscribe to new mailbox.

IMAP4.**thread** (*threading_algorithm*, *charset*, *search_criterion*[, ...])

The `thread` command is a variant of `search` with threading semantics for the results. Returned data contains a space separated list of thread members.

Thread members consist of zero or more messages numbers, delimited by spaces, indicating successive parent and child.

Thread has two arguments before the *search_criterion* argument(s); a *threading_algorithm*, and the searching *charset*. Note that unlike `search`, the searching *charset* argument is mandatory. There is also a `uid thread` command which corresponds to `thread` the way that `uid search` corresponds to `search`. The `thread` command first searches the mailbox for messages that match the given searching criteria using the *charset* argument for the interpretation of strings in the searching criteria. It then returns the matching messages threaded according to the specified threading algorithm.

This is an IMAP4rev1 extension command.

`IMAP4.uid(command, arg[, ...])`

Execute command args with messages identified by UID, rather than message number. Returns response appropriate to command. At least one argument must be supplied; if none are provided, the server will return an error and an exception will be raised.

`IMAP4.unsubscribe(mailbox)`

Unsubscribe from old mailbox.

`IMAP4.xatom(name[, ...])`

Allow simple extension commands notified by server in CAPABILITY response.

The following attributes are defined on instances of `IMAP4`:

`IMAP4.PROTOCOL_VERSION`

The most recent supported protocol in the CAPABILITY response from the server.

`IMAP4.debug`

Integer value to control debugging output. The initialize value is taken from the module variable `Debug`. Values greater than three trace each command.

`IMAP4.utf8_enabled`

Boolean value that is normally `False`, but is set to `True` if an `enable()` command is successfully issued for the UTF8=ACCEPT capability.

3.5 新版功能.

21.15.2 IMAP4 Example

Here is a minimal example (without error checking) that opens a mailbox and retrieves and prints all messages:

```
import getpass, imaplib

M = imaplib.IMAP4()
M.login(getpass.getuser(), getpass.getpass())
M.select()
typ, data = M.search(None, 'ALL')
for num in data[0].split():
    typ, data = M.fetch(num, '(RFC822)')
    print('Message %s\n%s\n' % (num, data[0][1]))
M.close()
M.logout()
```

21.16 nntplib —NNTP protocol client

Source code: [Lib/nntplib.py](#)

This module defines the class `NNTP` which implements the client side of the Network News Transfer Protocol. It can be used to implement a news reader or poster, or automated news processors. It is compatible with [RFC 3977](#) as well as the older [RFC 977](#) and [RFC 2980](#).

Here are two small examples of how it can be used. To list some statistics about a newsgroup and print the subjects of the last 10 articles:

```
>>> s = nntplib.NNTP('news.gmane.org')
>>> resp, count, first, last, name = s.group('gmane.comp.python.committers')
>>> print('Group', name, 'has', count, 'articles, range', first, 'to', last)
Group gmane.comp.python.committers has 1096 articles, range 1 to 1096
>>> resp, overviews = s.over((last - 9, last))
>>> for id, over in overviews:
...     print(id, nntplib.decode_header(over['subject']))
...
1087 Re: Commit privileges for Łukasz Langa
1088 Re: 3.2 alpha 2 freeze
1089 Re: 3.2 alpha 2 freeze
1090 Re: Commit privileges for Łukasz Langa
1091 Re: Commit privileges for Łukasz Langa
1092 Updated ssh key
1093 Re: Updated ssh key
1094 Re: Updated ssh key
1095 Hello fellow committers!
1096 Re: Hello fellow committers!
>>> s.quit()
'205 Bye!'
```

To post an article from a binary file (this assumes that the article has valid headers, and that you have right to post on the particular newsgroup):

```
>>> s = nntplib.NNTP('news.gmane.org')
>>> f = open('article.txt', 'rb')
>>> s.post(f)
'240 Article posted successfully.'
>>> s.quit()
'205 Bye!'
```

The module itself defines the following classes:

class nntplib.NNTP(*host*, *port*=119, *user*=None, *password*=None, *readermode*=None, *usenetr*=False[, *timeout*])

Return a new *NNTP* object, representing a connection to the NNTP server running on host *host*, listening at port *port*. An optional *timeout* can be specified for the socket connection. If the optional *user* and *password* are provided, or if suitable credentials are present in `/.netrc` and the optional flag *usenetr* is true, the AUTHINFO USER and AUTHINFO PASS commands are used to identify and authenticate the user to the server. If the optional flag *readermode* is true, then a mode reader command is sent before authentication is performed. Reader mode is sometimes necessary if you are connecting to an NNTP server on the local machine and intend to call reader-specific commands, such as `group`. If you get unexpected *NNTPPermanentErrors*, you might need to set *readermode*. The *NNTP* class supports the `with` statement to unconditionally consume *OSErrors* exceptions and to close the NNTP connection when done, e.g.:

```
>>> from nntplib import NNTP
>>> with NNTP('news.gmane.org') as n:
...     n.group('gmane.comp.python.committers')
... # doctest: +SKIP
('211 1755 1 1755 gmane.comp.python.committers', 1755, 1, 1755, 'gmane.comp.
python.committers')
>>>
```

在 3.2 版更改: *usenetr* is now False by default.

在 3.3 版更改: 支持了 `with` 语句。

class `nntplib.NNTP_SSL` (*host*, *port*=563, *user*=None, *password*=None, *ssl_context*=None, *reader-mode*=None, *usenetr*=False[, *timeout*])

Return a new `NNTP_SSL` object, representing an encrypted connection to the NNTP server running on host *host*, listening at port *port*. `NNTP_SSL` objects have the same methods as `NNTP` objects. If *port* is omitted, port 563 (NNTPS) is used. *ssl_context* is also optional, and is a `SSLContext` object. Please read *Security considerations* for best practices. All other parameters behave the same as for `NNTP`.

Note that SSL-on-563 is discouraged per [RFC 4642](#), in favor of STARTTLS as described below. However, some servers only support the former.

3.2 新版功能.

在 3.4 版更改: The class now supports hostname check with `ssl.SSLContext.check_hostname` and *Server Name Indication* (see `ssl.HAS_SNI`).

exception `nntplib.NNTPError`

Derived from the standard exception `Exception`, this is the base class for all exceptions raised by the `nntplib` module. Instances of this class have the following attribute:

response

The response of the server if available, as a `str` object.

exception `nntplib.NNTPReplyError`

Exception raised when an unexpected reply is received from the server.

exception `nntplib.NNTPTemporaryError`

Exception raised when a response code in the range 400–499 is received.

exception `nntplib.NNTPPermanentError`

Exception raised when a response code in the range 500–599 is received.

exception `nntplib.NNTPProtocolError`

Exception raised when a reply is received from the server that does not begin with a digit in the range 1–5.

exception `nntplib.NNTPDataError`

Exception raised when there is some error in the response data.

21.16.1 NNTP Objects

When connected, `NNTP` and `NNTP_SSL` objects support the following methods and attributes.

Attributes

`NNTP.nttp_version`

An integer representing the version of the NNTP protocol supported by the server. In practice, this should be 2 for servers advertising [RFC 3977](#) compliance and 1 for others.

3.2 新版功能.

`NNTP.nttp_implementation`

A string describing the software name and version of the NNTP server, or `None` if not advertised by the server.

3.2 新版功能.

方法

The *response* that is returned as the first item in the return tuple of almost all methods is the server's response: a string beginning with a three-digit code. If the server's response indicates an error, the method raises one of the above exceptions.

Many of the following methods take an optional keyword-only argument *file*. When the *file* argument is supplied, it must be either a *file object* opened for binary writing, or the name of an on-disk file to be written to. The method will then write any data returned by the server (except for the response line and the terminating dot) to the file; any list of lines, tuples or objects that the method normally returns will be empty.

在 3.2 版更改: Many of the following methods have been reworked and fixed, which makes them incompatible with their 3.1 counterparts.

NNTP.quit()

Send a QUIT command and close the connection. Once this method has been called, no other methods of the NNTP object should be called.

NNTP.getwelcome()

Return the welcome message sent by the server in reply to the initial connection. (This message sometimes contains disclaimers or help information that may be relevant to the user.)

NNTP.getcapabilities()

Return the [RFC 3977](#) capabilities advertised by the server, as a *dict* instance mapping capability names to (possibly empty) lists of values. On legacy servers which don't understand the CAPABILITIES command, an empty dictionary is returned instead.

```
>>> s = NNTP('news.gmane.org')
>>> 'POST' in s.getcapabilities()
True
```

3.2 新版功能.

NNTP.login(*user=None, password=None, usenetrc=True*)

Send AUTHINFO commands with the user name and password. If *user* and *password* are *None* and *usenetrc* is true, credentials from `~/.netrc` will be used if possible.

Unless intentionally delayed, login is normally performed during the *NNTP* object initialization and separately calling this function is unnecessary. To force authentication to be delayed, you must not set *user* or *password* when creating the object, and must set *usenetrc* to False.

3.2 新版功能.

NNTP.starttls(*context=None*)

Send a STARTTLS command. This will enable encryption on the NNTP connection. The *context* argument is optional and should be a *ssl.SSLContext* object. Please read [Security considerations](#) for best practices.

Note that this may not be done after authentication information has been transmitted, and authentication occurs by default if possible during a *NNTP* object initialization. See *NNTP.login()* for information on suppressing this behavior.

3.2 新版功能.

在 3.4 版更改: The method now supports hostname check with *ssl.SSLContext.check_hostname* and *Server Name Indication* (see *ssl.HAS_SNI*).

NNTP.newgroups(*date, *, file=None*)

Send a NEWGROUPS command. The *date* argument should be a *datetime.date* or *datetime.datetime* object. Return a pair (*response, groups*) where *groups* is a list representing the groups that are new since the given *date*. If *file* is supplied, though, then *groups* will be empty.

```
>>> from datetime import date, timedelta
>>> resp, groups = s.newgroups(date.today() - timedelta(days=3))
>>> len(groups) # doctest: +SKIP
85
>>> groups[0] # doctest: +SKIP
GroupInfo(group='gmane.network.tor.devel', last='4', first='1', flag='m')
```

NNTP.newnews (*group*, *date*, *, *file*=None)

Send a NEWNEWS command. Here, *group* is a group name or '*', and *date* has the same meaning as for *newgroups()*. Return a pair (response, articles) where *articles* is a list of message ids.

This command is frequently disabled by NNTP server administrators.

NNTP.list (*group_pattern*=None, *, *file*=None)

Send a LIST or LIST ACTIVE command. Return a pair (response, list) where *list* is a list of tuples representing all the groups available from this NNTP server, optionally matching the pattern string *group_pattern*. Each tuple has the form (group, last, first, flag), where *group* is a group name, *last* and *first* are the last and first article numbers, and *flag* usually takes one of these values:

- y: Local postings and articles from peers are allowed.
- m: The group is moderated and all postings must be approved.
- n: No local postings are allowed, only articles from peers.
- j: Articles from peers are filed in the junk group instead.
- x: No local postings, and articles from peers are ignored.
- =foo.bar: Articles are filed in the foo.bar group instead.

If *flag* has another value, then the status of the newsgroup should be considered unknown.

This command can return very large results, especially if *group_pattern* is not specified. It is best to cache the results offline unless you really need to refresh them.

在 3.2 版更改: *group_pattern* was added.

NNTP.descriptions (*grouppattern*)

Send a LIST NEWSGROUPS command, where *grouppattern* is a wildmat string as specified in [RFC 3977](#) (it's essentially the same as DOS or UNIX shell wildcard strings). Return a pair (response, descriptions), where *descriptions* is a dictionary mapping group names to textual descriptions.

```
>>> resp, descs = s.descriptions('gmane.comp.python.*')
>>> len(descs) # doctest: +SKIP
295
>>> descs.popitem() # doctest: +SKIP
('gmane.comp.python.bio.general', 'BioPython discussion list (Moderated)')
```

NNTP.description (*group*)

Get a description for a single group *group*. If more than one group matches (if 'group' is a real wildmat string), return the first match. If no group matches, return an empty string.

This elides the response code from the server. If the response code is needed, use *descriptions()*.

NNTP.group (*name*)

Send a GROUP command, where *name* is the group name. The group is selected as the current group, if it exists. Return a tuple (response, count, first, last, name) where *count* is the (estimated) number of articles in the group, *first* is the first article number in the group, *last* is the last article number in the group, and *name* is the group name.

NNTP.**over** (*message_spec*, *, *file=None*)

Send an OVER command, or an XOVER command on legacy servers. *message_spec* can be either a string representing a message id, or a (*first*, *last*) tuple of numbers indicating a range of articles in the current group, or a (*first*, *None*) tuple indicating a range of articles starting from *first* to the last article in the current group, or *None* to select the current article in the current group.

Return a pair (*response*, *overviews*). *overviews* is a list of (*article_number*, *overview*) tuples, one for each article selected by *message_spec*. Each *overview* is a dictionary with the same number of items, but this number depends on the server. These items are either message headers (the key is then the lower-cased header name) or metadata items (the key is then the metadata name prepended with ": "). The following items are guaranteed to be present by the NNTP specification:

- the subject, from, date, message-id and references headers
- the :bytes metadata: the number of bytes in the entire raw article (including headers and body)
- the :lines metadata: the number of lines in the article body

The value of each item is either a string, or *None* if not present.

It is advisable to use the `decode_header()` function on header values when they may contain non-ASCII characters:

```
>>> _, _, first, last, _ = s.group('gmane.comp.python.devel')
>>> resp, overviews = s.over((last, last))
>>> art_num, over = overviews[0]
>>> art_num
117216
>>> list(over.keys())
['xref', 'from', ':lines', ':bytes', 'references', 'date', 'message-id', 'subject
↵']
>>> over['from']
'=?UTF-8?B?Ik1hcnRpbIB2LiBMw7Z3aXMi?= <martin@v.loewis.de>'
>>> nntplib.decode_header(over['from'])
'"Martin v. Löwis" <martin@v.loewis.de>'
```

3.2 新版功能.

NNTP.**help** (*, *file=None*)

Send a HELP command. Return a pair (*response*, *list*) where *list* is a list of help strings.

NNTP.**stat** (*message_spec=None*)

Send a STAT command, where *message_spec* is either a message id (enclosed in '<' and '>') or an article number in the current group. If *message_spec* is omitted or *None*, the current article in the current group is considered. Return a triple (*response*, *number*, *id*) where *number* is the article number and *id* is the message id.

```
>>> _, _, first, last, _ = s.group('gmane.comp.python.devel')
>>> resp, number, message_id = s.stat(first)
>>> number, message_id
(9099, '<20030112190404.GE29873@epoch.metaslash.com>')
```

NNTP.**next** ()

Send a NEXT command. Return as for `stat()`.

NNTP.**last** ()

Send a LAST command. Return as for `stat()`.

NNTP.**article** (*message_spec=None*, *, *file=None*)

Send an ARTICLE command, where *message_spec* has the same meaning as for `stat()`. Return a tuple (*response*, *info*) where *info* is a *namedtuple* with three attributes *number*, *message_id* and *lines* (in that order). *number* is the article number in the group (or 0 if the information is not available), *message_id* the

message id as a string, and *lines* a list of lines (without terminating newlines) comprising the raw message including headers and body.

```
>>> resp, info = s.article('<20030112190404.GE29873@epoch.metaslash.com>')
>>> info.number
0
>>> info.message_id
'<20030112190404.GE29873@epoch.metaslash.com>'
>>> len(info.lines)
65
>>> info.lines[0]
b'Path: main.gmane.org!not-for-mail'
>>> info.lines[1]
b'From: Neal Norwitz <neal@metaslash.com>'
>>> info.lines[-3:]
[b'There is a patch for 2.3 as well as 2.2.', b'', b'Neal']
```

NNTP **.head** (*message_spec=None*, *, *file=None*)

Same as [article\(\)](#), but sends a HEAD command. The *lines* returned (or written to *file*) will only contain the message headers, not the body.

NNTP **.body** (*message_spec=None*, *, *file=None*)

Same as [article\(\)](#), but sends a BODY command. The *lines* returned (or written to *file*) will only contain the message body, not the headers.

NNTP **.post** (*data*)

Post an article using the POST command. The *data* argument is either a *file object* opened for binary reading, or any iterable of bytes objects (representing raw lines of the article to be posted). It should represent a well-formed news article, including the required headers. The [post\(\)](#) method automatically escapes lines beginning with . and appends the termination line.

If the method succeeds, the server's response is returned. If the server refuses posting, a [NNTPReplyError](#) is raised.

NNTP **.ihave** (*message_id*, *data*)

Send an IHAVE command. *message_id* is the id of the message to send to the server (enclosed in '<' and '>'). The *data* parameter and the return value are the same as for [post\(\)](#).

NNTP **.date** ()

Return a pair (*response*, *date*). *date* is a *datetime* object containing the current date and time of the server.

NNTP **.slave** ()

Send a SLAVE command. Return the server's response.

NNTP **.set_debuglevel** (*level*)

Set the instance's debugging level. This controls the amount of debugging output printed. The default, 0, produces no debugging output. A value of 1 produces a moderate amount of debugging output, generally a single line per request or response. A value of 2 or higher produces the maximum amount of debugging output, logging each line sent and received on the connection (including message text).

The following are optional NNTP extensions defined in [RFC 2980](#). Some of them have been superseded by newer commands in [RFC 3977](#).

NNTP **.xhdr** (*hdr*, *str*, *, *file=None*)

Send an XHDR command. The *hdr* argument is a header keyword, e.g. 'subject'. The *str* argument should have the form 'first-last' where *first* and *last* are the first and last article numbers to search. Return a pair (*response*, *list*), where *list* is a list of pairs (*id*, *text*), where *id* is an article number (as a string) and *text* is the text of the requested header for that article. If the *file* parameter is supplied, then the output of the XHDR command is stored in a file. If *file* is a string, then the method will open a file with that name, write to it then close

it. If *file* is a *file object*, then it will start calling `write()` on it to store the lines of the command output. If *file* is supplied, then the returned *list* is an empty list.

NNTP **.xover** (*start*, *end*, *, *file=None*)

Send an XOVER command. *start* and *end* are article numbers delimiting the range of articles to select. The return value is the same of for `over()`. It is recommended to use `over()` instead, since it will automatically use the newer OVER command if available.

NNTP **.xpath** (*id*)

Return a pair (*resp*, *path*), where *path* is the directory path to the article with message ID *id*. Most of the time, this extension is not enabled by NNTP server administrators.

3.3 版后已移除: The XPATH extension is not actively used.

21.16.2 Utility functions

The module also defines the following utility function:

`nntplib.decode_header` (*header_str*)

Decode a header value, un-escaping any escaped non-ASCII characters. *header_str* must be a *str* object. The unescaped value is returned. Using this function is recommended to display some headers in a human readable form:

```
>>> decode_header("Some subject")
'Some subject'
>>> decode_header("=?ISO-8859-15?Q?D=E9buter_en_Python?=")
'Débuter en Python'
>>> decode_header("Re: =?UTF-8?B?cHJvYmzDqG1lIGRlIG1hdHJpY2U=?=")
'Re: problème de matrice'
```

21.17 smtplib —SMTP 协议客户端

Source code: [Lib/smtplib.py](#)

The `smtplib` module defines an SMTP client session object that can be used to send mail to any Internet machine with an SMTP or ESMTP listener daemon. For details of SMTP and ESMTP operation, consult [RFC 821](#) (Simple Mail Transfer Protocol) and [RFC 1869](#) (SMTP Service Extensions).

class `smtplib.SMTP` (*host=""*, *port=0*, *local_hostname=None*[, *timeout*], *source_address=None*)

An *SMTP* instance encapsulates an SMTP connection. It has methods that support a full repertoire of SMTP and ESMTP operations. If the optional *host* and *port* parameters are given, the SMTP `connect()` method is called with those parameters during initialization. If specified, *local_hostname* is used as the FQDN of the local host in the HELO/EHLO command. Otherwise, the local hostname is found using `socket.getfqdn()`. If the `connect()` call returns anything other than a success code, an *SMTPConnectError* is raised. The optional *timeout* parameter specifies a timeout in seconds for blocking operations like the connection attempt (if not specified, the global default timeout setting will be used). If the timeout expires, `socket.timeout` is raised. The optional *source_address* parameter allows binding to some specific source address in a machine with multiple network interfaces, and/or to some specific source TCP port. It takes a 2-tuple (host, port), for the socket to bind to as its source address before connecting. If omitted (or if *host* or *port* are '' and/or 0 respectively) the OS default behavior will be used.

For normal use, you should only require the initialization/connect, `sendmail()`, and `SMTP.quit()` methods. An example is included below.

The `SMTP` class supports the `with` statement. When used like this, the SMTP `QUIT` command is issued automatically when the `with` statement exits. E.g.:

```
>>> from smtplib import SMTP
>>> with SMTP("domain.org") as smtp:
...     smtp.noop()
...
(250, b'Ok')
>>>
```

在 3.3 版更改: 支持了 `with` 语句。

在 3.3 版更改: `source_address` argument was added.

3.5 新版功能: The SMTPUTF8 extension ([RFC 6531](#)) is now supported.

class `smtplib.SMTP_SSL` (`host=""`, `port=0`, `local_hostname=None`, `keyfile=None`, `certfile=None`[, `timeout`], `context=None`, `source_address=None`)

An `SMTP_SSL` instance behaves exactly the same as instances of `SMTP`. `SMTP_SSL` should be used for situations where SSL is required from the beginning of the connection and using `starttls()` is not appropriate. If `host` is not specified, the local host is used. If `port` is zero, the standard SMTP-over-SSL port (465) is used. The optional arguments `local_hostname`, `timeout` and `source_address` have the same meaning as they do in the `SMTP` class. `context`, also optional, can contain a `SSLContext` and allows configuring various aspects of the secure connection. Please read [Security considerations](#) for best practices.

`keyfile` and `certfile` are a legacy alternative to `context`, and can point to a PEM formatted private key and certificate chain file for the SSL connection.

在 3.3 版更改: `context` was added.

在 3.3 版更改: `source_address` argument was added.

在 3.4 版更改: The class now supports hostname check with `ssl.SSLContext.check_hostname` and *Server Name Indication* (see `ssl.HAS_SNI`).

3.6 版后已移除: `keyfile` and `certfile` are deprecated in favor of `context`. Please use `ssl.SSLContext.load_cert_chain()` instead, or let `ssl.create_default_context()` select the system's trusted CA certificates for you.

class `smtplib.LMTP` (`host=""`, `port=LMTP_PORT`, `local_hostname=None`, `source_address=None`)

The LMTP protocol, which is very similar to ESMTP, is heavily based on the standard SMTP client. It's common to use Unix sockets for LMTP, so our `connect()` method must support that as well as a regular host:port server. The optional arguments `local_hostname` and `source_address` have the same meaning as they do in the `SMTP` class. To specify a Unix socket, you must use an absolute path for `host`, starting with a `'/'`.

Authentication is supported, using the regular SMTP mechanism. When using a Unix socket, LMTP generally don't support or require any authentication, but your mileage might vary.

A nice selection of exceptions is defined as well:

exception `smtplib.SMTPException`

Subclass of `OSError` that is the base exception class for all the other exceptions provided by this module.

在 3.4 版更改: `SMTPException` became subclass of `OSError`

exception `smtplib.SMTPServerDisconnected`

This exception is raised when the server unexpectedly disconnects, or when an attempt is made to use the `SMTP` instance before connecting it to a server.

exception `smtplib.SMTPResponseException`

Base class for all exceptions that include an SMTP error code. These exceptions are generated in some instances

when the SMTP server returns an error code. The error code is stored in the `smtp_code` attribute of the error, and the `smtp_error` attribute is set to the error message.

exception `smtpplib.SMTPSenderRefused`

Sender address refused. In addition to the attributes set by on all *SMTPResponseException* exceptions, this sets 'sender' to the string that the SMTP server refused.

exception `smtpplib.SMTPRecipientsRefused`

All recipient addresses refused. The errors for each recipient are accessible through the attribute `recipients`, which is a dictionary of exactly the same sort as *SMTP.sendmail()* returns.

exception `smtpplib.SMTPDataError`

The SMTP server refused to accept the message data.

exception `smtpplib.SMTPConnectError`

Error occurred during establishment of a connection with the server.

exception `smtpplib.SMTPHeloError`

The server refused our HELO message.

exception `smtpplib.SMTPNotSupportedError`

The command or option attempted is not supported by the server.

3.5 新版功能.

exception `smtpplib.SMTPAuthenticationError`

SMTP authentication went wrong. Most probably the server didn't accept the username/password combination provided.

参见:

RFC 821 - Simple Mail Transfer Protocol Protocol definition for SMTP. This document covers the model, operating procedure, and protocol details for SMTP.

RFC 1869 - SMTP Service Extensions Definition of the ESMTP extensions for SMTP. This describes a framework for extending SMTP with new commands, supporting dynamic discovery of the commands provided by the server, and defines a few additional commands.

21.17.1 SMTP Objects

An *SMTP* instance has the following methods:

`SMTP.set_debuglevel (level)`

Set the debug output level. A value of 1 or `True` for *level* results in debug messages for connection and for all messages sent to and received from the server. A value of 2 for *level* results in these messages being timestamped.

在 3.5 版更改: Added debuglevel 2.

`SMTP.doccmd (cmd, args=)`

Send a command *cmd* to the server. The optional argument *args* is simply concatenated to the command, separated by a space.

This returns a 2-tuple composed of a numeric response code and the actual response line (multiline responses are joined into one long line.)

In normal operation it should not be necessary to call this method explicitly. It is used to implement other methods and may be useful for testing private extensions.

If the connection to the server is lost while waiting for the reply, *SMTPServerDisconnected* will be raised.

SMTP.**connect** (*host='localhost', port=0*)

Connect to a host on a given port. The defaults are to connect to the local host at the standard SMTP port (25). If the hostname ends with a colon (':') followed by a number, that suffix will be stripped off and the number interpreted as the port number to use. This method is automatically invoked by the constructor if a host is specified during instantiation. Returns a 2-tuple of the response code and message sent by the server in its connection response.

SMTP.**hello** (*name=""*)

Identify yourself to the SMTP server using HELO. The hostname argument defaults to the fully qualified domain name of the local host. The message returned by the server is stored as the `hello_resp` attribute of the object.

In normal operation it should not be necessary to call this method explicitly. It will be implicitly called by the `sendmail()` when necessary.

SMTP.**ehlo** (*name=""*)

Identify yourself to an ESMTP server using EHLO. The hostname argument defaults to the fully qualified domain name of the local host. Examine the response for ESMTP option and store them for use by `has_extn()`. Also sets several informational attributes: the message returned by the server is stored as the `ehlo_resp` attribute, `does_esmtp` is set to true or false depending on whether the server supports ESMTP, and `esmtp_features` will be a dictionary containing the names of the SMTP service extensions this server supports, and their parameters (if any).

Unless you wish to use `has_extn()` before sending mail, it should not be necessary to call this method explicitly. It will be implicitly called by `sendmail()` when necessary.

SMTP.**ehlo_or_helo_if_needed** ()

This method calls `ehlo()` and/or `hello()` if there has been no previous EHLO or HELO command this session. It tries ESMTP EHLO first.

SMTPHelloError The server didn't reply properly to the HELO greeting.

SMTP.**has_extn** (*name*)

Return `True` if *name* is in the set of SMTP service extensions returned by the server, `False` otherwise. Case is ignored.

SMTP.**verify** (*address*)

Check the validity of an address on this server using SMTP VRFY. Returns a tuple consisting of code 250 and a full **RFC 822** address (including human name) if the user address is valid. Otherwise returns an SMTP error code of 400 or greater and an error string.

注解: Many sites disable SMTP VRFY in order to foil spammers.

SMTP.**login** (*user, password, *, initial_response_ok=True*)

Log in on an SMTP server that requires authentication. The arguments are the username and the password to authenticate with. If there has been no previous EHLO or HELO command this session, this method tries ESMTP EHLO first. This method will return normally if the authentication was successful, or may raise the following exceptions:

SMTPHelloError The server didn't reply properly to the HELO greeting.

SMTPAuthenticationError The server didn't accept the username/password combination.

SMTPNotSupportedError The AUTH command is not supported by the server.

SMTPException No suitable authentication method was found.

Each of the authentication methods supported by `smtpplib` are tried in turn if they are advertised as supported by the server. See `auth()` for a list of supported authentication methods. `initial_response_ok` is passed through to `auth()`.

Optional keyword argument *initial_response_ok* specifies whether, for authentication methods that support it, an “initial response” as specified in [RFC 4954](#) can be sent along with the AUTH command, rather than requiring a challenge/response.

在 3.5 版更改: *SMTPNotSupportedError* may be raised, and the *initial_response_ok* parameter was added.

SMTP.**auth** (*mechanism*, *authobject*, *, *initial_response_ok*=True)

Issue an SMTP AUTH command for the specified authentication *mechanism*, and handle the challenge response via *authobject*.

mechanism specifies which authentication mechanism is to be used as argument to the AUTH command; the valid values are those listed in the *auth* element of *esmtplib.features*.

authobject must be a callable object taking an optional single argument:

```
data = authobject(challenge=None)
```

If optional keyword argument *initial_response_ok* is true, *authobject* () will be called first with no argument. It can return the [RFC 4954](#) “initial response” ASCII str which will be encoded and sent with the AUTH command as below. If the *authobject* () does not support an initial response (e.g. because it requires a challenge), it should return None when called with *challenge*=None. If *initial_response_ok* is false, then *authobject* () will not be called first with None.

If the initial response check returns None, or if *initial_response_ok* is false, *authobject* () will be called to process the server’s challenge response; the *challenge* argument it is passed will be a bytes. It should return ASCII str *data* that will be base64 encoded and sent to the server.

The SMTP class provides *authobjects* for the CRAM-MD5, PLAIN, and LOGIN mechanisms; they are named SMTP.*auth_cram_md5*, SMTP.*auth_plain*, and SMTP.*auth_login* respectively. They all require that the *user* and *password* properties of the SMTP instance are set to appropriate values.

User code does not normally need to call *auth* directly, but can instead call the *login* () method, which will try each of the above mechanisms in turn, in the order listed. *auth* is exposed to facilitate the implementation of authentication methods not (or not yet) supported directly by *smtplib*.

3.5 新版功能.

SMTP.**starttls** (*keyfile*=None, *certfile*=None, *context*=None)

Put the SMTP connection in TLS (Transport Layer Security) mode. All SMTP commands that follow will be encrypted. You should then call *ehlo* () again.

If *keyfile* and *certfile* are provided, they are used to create an *ssl.SSLContext*.

Optional *context* parameter is an *ssl.SSLContext* object; This is an alternative to using a keyfile and a certfile and if specified both *keyfile* and *certfile* should be None.

If there has been no previous EHLO or HELO command this session, this method tries ESMTP EHLO first.

3.6 版后已移除: *keyfile* and *certfile* are deprecated in favor of *context*. Please use *ssl.SSLContext.load_cert_chain* () instead, or let *ssl.create_default_context* () select the system’s trusted CA certificates for you.

SMTPHeloError The server didn’t reply properly to the HELO greeting.

SMTPNotSupportedError The server does not support the STARTTLS extension.

RuntimeError SSL/TLS support is not available to your Python interpreter.

在 3.3 版更改: *context* was added.

在 3.4 版更改: The method now supports hostname check with *SSLContext.check_hostname* and *Server Name Indicator* (see [HAS_SNI](#)).

在 3.5 版更改: The error raised for lack of STARTTLS support is now the `SMTPNotSupportedError` sub-class instead of the base `SMTPException`.

`SMTP.sendmail` (*from_addr*, *to_addrs*, *msg*, *mail_options*=(), *rcpt_options*=())

Send mail. The required arguments are an [RFC 822](#) from-address string, a list of [RFC 822](#) to-address strings (a bare string will be treated as a list with 1 address), and a message string. The caller may pass a list of ESMTP options (such as `8bitmime`) to be used in `MAIL FROM` commands as *mail_options*. ESMTP options (such as DSN commands) that should be used with all `RCPT` commands can be passed as *rcpt_options*. (If you need to use different ESMTP options to different recipients you have to use the low-level methods such as `mail()`, `rcpt()` and `data()` to send the message.)

注解: The *from_addr* and *to_addrs* parameters are used to construct the message envelope used by the transport agents. `sendmail` does not modify the message headers in any way.

msg may be a string containing characters in the ASCII range, or a byte string. A string is encoded to bytes using the `ascii` codec, and lone `\r` and `\n` characters are converted to `\r\n` characters. A byte string is not modified.

If there has been no previous `EHLO` or `HELO` command this session, this method tries ESMTP `EHLO` first. If the server does ESMTP, message size and each of the specified options will be passed to it (if the option is in the feature set the server advertises). If `EHLO` fails, `HELO` will be tried and ESMTP options suppressed.

This method will return normally if the mail is accepted for at least one recipient. Otherwise it will raise an exception. That is, if this method does not raise an exception, then someone should get your mail. If this method does not raise an exception, it returns a dictionary, with one entry for each recipient that was refused. Each entry contains a tuple of the SMTP error code and the accompanying error message sent by the server.

If `SMTPUTF8` is included in *mail_options*, and the server supports it, *from_addr* and *to_addrs* may contain non-ASCII characters.

This method may raise the following exceptions:

`SMTPRecipientsRefused` All recipients were refused. Nobody got the mail. The `recipients` attribute of the exception object is a dictionary with information about the refused recipients (like the one returned when at least one recipient was accepted).

`SMTPHeloError` The server didn't reply properly to the `HELO` greeting.

`SMTPSenderRefused` The server didn't accept the *from_addr*.

`SMTPDataError` The server replied with an unexpected error code (other than a refusal of a recipient).

`SMTPNotSupportedError` `SMTPUTF8` was given in the *mail_options* but is not supported by the server.

Unless otherwise noted, the connection will be open even after an exception is raised.

在 3.2 版更改: *msg* may be a byte string.

在 3.5 版更改: `SMTPUTF8` support added, and `SMTPNotSupportedError` may be raised if `SMTPUTF8` is specified but the server does not support it.

`SMTP.send_message` (*msg*, *from_addr*=None, *to_addrs*=None, *mail_options*=(), *rcpt_options*=())

This is a convenience method for calling `sendmail()` with the message represented by an `email.message.Message` object. The arguments have the same meaning as for `sendmail()`, except that *msg* is a `Message` object.

If *from_addr* is None or *to_addrs* is None, `send_message` fills those arguments with addresses extracted from the headers of *msg* as specified in [RFC 5322](#): *from_addr* is set to the `Sender` field if it is present, and otherwise to the `From` field. *to_addrs* combines the values (if any) of the `To`, `Cc`, and `Bcc` fields from *msg*. If exactly one set of `Resent-*` headers appear in the message, the regular headers are ignored and the `Resent-*` headers are

used instead. If the message contains more than one set of *Resent-** headers, a *ValueError* is raised, since there is no way to unambiguously detect the most recent set of *Resent-* headers.

`send_message` serializes *msg* using *BytesGenerator* with `\r\n` as the *linesep*, and calls `sendmail()` to transmit the resulting message. Regardless of the values of *from_addr* and *to_addrs*, `send_message` does not transmit any *Bcc* or *Resent-Bcc* headers that may appear in *msg*. If any of the addresses in *from_addr* and *to_addrs* contain non-ASCII characters and the server does not advertise SMTPUTF8 support, an *SMTPNotSupported* error is raised. Otherwise the Message is serialized with a clone of its *policy* with the *utf8* attribute set to *True*, and SMTPUTF8 and BODY=8BITMIME are added to *mail_options*.

3.2 新版功能.

3.5 新版功能: Support for internationalized addresses (SMTPUTF8).

`SMTP.quit()`

Terminate the SMTP session and close the connection. Return the result of the SMTP QUIT command.

Low-level methods corresponding to the standard SMTP/ESMTP commands HELP, RSET, NOOP, MAIL, RCPT, and DATA are also supported. Normally these do not need to be called directly, so they are not documented here. For details, consult the module code.

21.17.2 SMTP Example

This example prompts the user for addresses needed in the message envelope (‘To’ and ‘From’ addresses), and the message to be delivered. Note that the headers to be included with the message must be included in the message as entered; this example doesn’t do any processing of the [RFC 822](#) headers. In particular, the ‘To’ and ‘From’ addresses must be included in the message headers explicitly.

```
import smtplib

def prompt(prompt):
    return input(prompt).strip()

fromaddr = prompt("From: ")
toaddrs = prompt("To: ").split()
print("Enter message, end with ^D (Unix) or ^Z (Windows):")

# Add the From: and To: headers at the start!
msg = ("From: %s\r\nTo: %s\r\n\r\n"
       % (fromaddr, ", ".join(toaddrs)))
while True:
    try:
        line = input()
    except EOFError:
        break
    if not line:
        break
    msg = msg + line

print("Message length is", len(msg))

server = smtplib.SMTP('localhost')
server.set_debuglevel(1)
server.sendmail(fromaddr, toaddrs, msg)
server.quit()
```

注解: In general, you will want to use the *email* package’s features to construct an email message, which you can

then send via `send_message()`; see *email*: 示例.

21.18 smtpd —SMTP Server

Source code: [Lib/smtpd.py](#)

This module offers several classes to implement SMTP (email) servers.

参见:

The `aiosmtpd` package is a recommended replacement for this module. It is based on `asyncio` and provides a more straightforward API. `smtpd` should be considered deprecated.

Several server implementations are present; one is a generic do-nothing implementation, which can be overridden, while the other two offer specific mail-sending strategies.

Additionally the SMTPChannel may be extended to implement very specific interaction behaviour with SMTP clients.

The code supports [RFC 5321](#), plus the [RFC 1870](#) SIZE and [RFC 6531](#) SMTPUTF8 extensions.

21.18.1 SMTPServer Objects

class `smtpd.SMTPServer`(*localaddr*, *remoteaddr*, *data_size_limit*=33554432, *map*=None, *enable_SMTPUTF8*=False, *decode_data*=False)

Create a new `SMTPServer` object, which binds to local address *localaddr*. It will treat *remoteaddr* as an upstream SMTP relay. Both *localaddr* and *remoteaddr* should be a (*host*, *port*) tuple. The object inherits from `asyncore.dispatcher`, and so will insert itself into `asyncore`'s event loop on instantiation.

data_size_limit specifies the maximum number of bytes that will be accepted in a DATA command. A value of None or 0 means no limit.

map is the socket map to use for connections (an initially empty dictionary is a suitable value). If not specified the `asyncore` global socket map is used.

enable_SMTPUTF8 determines whether the SMTPUTF8 extension (as defined in [RFC 6531](#)) should be enabled. The default is False. When True, SMTPUTF8 is accepted as a parameter to the MAIL command and when present is passed to `process_message()` in the `kwargs['mail_options']` list. *decode_data* and *enable_SMTPUTF8* cannot be set to True at the same time.

decode_data specifies whether the data portion of the SMTP transaction should be decoded using UTF-8. When *decode_data* is False (the default), the server advertises the 8BITMIME extension ([RFC 6152](#)), accepts the BODY=8BITMIME parameter to the MAIL command, and when present passes it to `process_message()` in the `kwargs['mail_options']` list. *decode_data* and *enable_SMTPUTF8* cannot be set to True at the same time.

process_message(*peer*, *mailfrom*, *rcpttos*, *data*, ***kwargs*)

Raise a `NotImplementedError` exception. Override this in subclasses to do something useful with this message. Whatever was passed in the constructor as *remoteaddr* will be available as the `_remoteaddr` attribute. *peer* is the remote host's address, *mailfrom* is the envelope originator, *rcpttos* are the envelope recipients and *data* is a string containing the contents of the e-mail (which should be in [RFC 5321](#) format).

If the *decode_data* constructor keyword is set to True, the *data* argument will be a unicode string. If it is set to False, it will be a bytes object.

kwargs is a dictionary containing additional information. It is empty if *decode_data*=True was given as an init argument, otherwise it contains the following keys:

mail_options: a list of all received parameters to the MAIL command (the elements are uppercase strings; example: ['BODY=8BITMIME', 'SMTPUTF8']).

rcpt_options: same as *mail_options* but for the RCPT command. Currently no RCPT TO options are supported, so for now this will always be an empty list.

Implementations of *process_message* should use the ***kwargs* signature to accept arbitrary keyword arguments, since future feature enhancements may add keys to the *kwargs* dictionary.

Return *None* to request a normal 250 Ok response; otherwise return the desired response string in **RFC 5321** format.

channel_class

Override this in subclasses to use a custom *SMTPChannel* for managing SMTP clients.

3.4 新版功能: The *map* constructor argument.

在 3.5 版更改: *localaddr* and *remoteaddr* may now contain IPv6 addresses.

3.5 新版功能: The *decode_data* and *enable_SMTPUTF8* constructor parameters, and the *kwargs* parameter to *process_message()* when *decode_data* is *False*.

在 3.6 版更改: *decode_data* is now *False* by default.

21.18.2 DebuggingServer Objects

class *smtpd.DebuggingServer* (*localaddr*, *remoteaddr*)

Create a new debugging server. Arguments are as per *SMTPServer*. Messages will be discarded, and printed on stdout.

21.18.3 PureProxy Objects

class *smtpd.PureProxy* (*localaddr*, *remoteaddr*)

Create a new pure proxy server. Arguments are as per *SMTPServer*. Everything will be relayed to *remoteaddr*. Note that running this has a good chance to make you into an open relay, so please be careful.

21.18.4 MailmanProxy Objects

class *smtpd.MailmanProxy* (*localaddr*, *remoteaddr*)

Create a new pure proxy server. Arguments are as per *SMTPServer*. Everything will be relayed to *remoteaddr*, unless local mailman configurations knows about an address, in which case it will be handled via mailman. Note that running this has a good chance to make you into an open relay, so please be careful.

21.18.5 SMTPChannel Objects

class *smtpd.SMTPChannel* (*server*, *conn*, *addr*, *data_size_limit*=33554432, *map*=*None*, *enable_SMTPUTF8*=*False*, *decode_data*=*False*)

Create a new *SMTPChannel* object which manages the communication between the server and a single SMTP client.

conn and *addr* are as per the instance variables described below.

data_size_limit specifies the maximum number of bytes that will be accepted in a DATA command. A value of *None* or 0 means no limit.

`enable_SMTPUTF8` determines whether the SMTPUTF8 extension (as defined in [RFC 6531](#)) should be enabled. The default is `False`. `decode_data` and `enable_SMTPUTF8` cannot be set to `True` at the same time.

A dictionary can be specified in `map` to avoid using a global socket map.

`decode_data` specifies whether the data portion of the SMTP transaction should be decoded using UTF-8. The default is `False`. `decode_data` and `enable_SMTPUTF8` cannot be set to `True` at the same time.

To use a custom SMTPChannel implementation you need to override the `SMTPServer.channel_class` of your `SMTPServer`.

在 3.5 版更改: The `decode_data` and `enable_SMTPUTF8` parameters were added.

在 3.6 版更改: `decode_data` is now `False` by default.

The `SMTPChannel` has the following instance variables:

smtp_server

Holds the `SMTPServer` that spawned this channel.

conn

Holds the socket object connecting to the client.

addr

Holds the address of the client, the second value returned by `socket.accept`

received_lines

Holds a list of the line strings (decoded using UTF-8) received from the client. The lines have their `"\r\n"` line ending translated to `"\n"`.

smtp_state

Holds the current state of the channel. This will be either `COMMAND` initially and then `DATA` after the client sends a “DATA” line.

seen_greeting

Holds a string containing the greeting sent by the client in its “HELO” .

mailfrom

Holds a string containing the address identified in the “MAIL FROM:” line from the client.

rcpttos

Holds a list of strings containing the addresses identified in the “RCPT TO:” lines from the client.

received_data

Holds a string containing all of the data sent by the client during the `DATA` state, up to but not including the terminating `"\r\n.\r\n"`.

fqdn

Holds the fully-qualified domain name of the server as returned by `socket.getfqdn()`.

peer

Holds the name of the client peer as returned by `conn.getpeername()` where `conn` is `conn`.

The `SMTPChannel` operates by invoking methods named `smtp_<command>` upon reception of a command line from the client. Built into the base `SMTPChannel` class are methods for handling the following commands (and responding to them appropriately):

Com- mand	Action taken
HELO	Accepts the greeting from the client and stores it in <i>seen_greeting</i> . Sets server to base command mode.
EHLO	Accepts the greeting from the client and stores it in <i>seen_greeting</i> . Sets server to extended command mode.
NOOP	Takes no action.
QUIT	Closes the connection cleanly.
MAIL	Accepts the “MAIL FROM:” syntax and stores the supplied address as <i>mailfrom</i> . In extended command mode, accepts the RFC 1870 SIZE attribute and responds appropriately based on the value of <i>data_size_limit</i> .
RCPT	Accepts the “RCPT TO:” syntax and stores the supplied addresses in the <i>rcpttos</i> list.
RSET	Resets the <i>mailfrom</i> , <i>rcpttos</i> , and <i>received_data</i> , but not the greeting.
DATA	Sets the internal state to DATA and stores remaining lines from the client in <i>received_data</i> until the terminator “\r\n.\r\n” is received.
HELP	Returns minimal information on command syntax
VERFY	Returns code 252 (the server doesn’t know if the address is valid)
EXPN	Reports that the command is not implemented.

21.19 telnetlib —Telnet client

Source code: [Lib/telnetlib.py](#)

The *telnetlib* module provides a *Telnet* class that implements the Telnet protocol. See **RFC 854** for details about the protocol. In addition, it provides symbolic constants for the protocol characters (see below), and for the telnet options. The symbolic names of the telnet options follow the definitions in *arpa/telnet.h*, with the leading *TELOPT_* removed. For symbolic names of options which are traditionally not included in *arpa/telnet.h*, see the module source itself.

The symbolic constants for the telnet commands are: IAC, DONT, DO, WONT, WILL, SE (Subnegotiation End), NOP (No Operation), DM (Data Mark), BRK (Break), IP (Interrupt process), AO (Abort output), AYT (Are You There), EC (Erase Character), EL (Erase Line), GA (Go Ahead), SB (Subnegotiation Begin).

class *telnetlib.Telnet* (*host=None, port=0[, timeout]*)

Telnet represents a connection to a Telnet server. The instance is initially not connected by default; the *open()* method must be used to establish a connection. Alternatively, the host name and optional port number can be passed to the constructor too, in which case the connection to the server will be established before the constructor returns. The optional *timeout* parameter specifies a timeout in seconds for blocking operations like the connection attempt (if not specified, the global default timeout setting will be used).

Do not reopen an already connected instance.

This class has many *read_*()* methods. Note that some of them raise *EOFError* when the end of the connection is read, because they can return an empty string for other reasons. See the individual descriptions below.

A *Telnet* object is a context manager and can be used in a *with* statement. When the *with* block ends, the *close()* method is called:

```
>>> from telnetlib import Telnet
>>> with Telnet('localhost', 23) as tn:
...     tn.interact()
... 
```

在 3.6 版更改: Context manager support added

参见:

RFC 854 - Telnet Protocol Specification Definition of the Telnet protocol.

21.19.1 Telnet Objects

Telnet instances have the following methods:

`Telnet.read_until(expected, timeout=None)`

Read until a given byte string, *expected*, is encountered or until *timeout* seconds have passed.

When no match is found, return whatever is available instead, possibly empty bytes. Raise *EOFError* if the connection is closed and no cooked data is available.

`Telnet.read_all()`

Read all data until EOF as bytes; block until connection closed.

`Telnet.read_some()`

Read at least one byte of cooked data unless EOF is hit. Return `b''` if EOF is hit. Block if no data is immediately available.

`Telnet.read_very_eager()`

Read everything that can be without blocking in I/O (eager).

Raise *EOFError* if connection closed and no cooked data available. Return `b''` if no cooked data available otherwise. Do not block unless in the midst of an IAC sequence.

`Telnet.read_eager()`

Read readily available data.

Raise *EOFError* if connection closed and no cooked data available. Return `b''` if no cooked data available otherwise. Do not block unless in the midst of an IAC sequence.

`Telnet.read_lazy()`

Process and return data already in the queues (lazy).

Raise *EOFError* if connection closed and no data available. Return `b''` if no cooked data available otherwise. Do not block unless in the midst of an IAC sequence.

`Telnet.read_very_lazy()`

Return any data available in the cooked queue (very lazy).

Raise *EOFError* if connection closed and no data available. Return `b''` if no cooked data available otherwise. This method never blocks.

`Telnet.read_sb_data()`

Return the data collected between a SB/SE pair (suboption begin/end). The callback should access these data when it was invoked with a SE command. This method never blocks.

`Telnet.open(host, port=0[, timeout])`

Connect to a host. The optional second argument is the port number, which defaults to the standard Telnet port (23). The optional *timeout* parameter specifies a timeout in seconds for blocking operations like the connection attempt (if not specified, the global default timeout setting will be used).

Do not try to reopen an already connected instance.

`Telnet.msg(msg, *args)`

Print a debug message when the debug level is `> 0`. If extra arguments are present, they are substituted in the message using the standard string formatting operator.

`Telnet.set_debuglevel(debuglevel)`

Set the debug level. The higher the value of *debuglevel*, the more debug output you get (on `sys.stdout`).

`Telnet.close()`

关闭连接对象。

`Telnet.get_socket()`

Return the socket object used internally.

`Telnet.fileno()`

Return the file descriptor of the socket object used internally.

`Telnet.write(buffer)`

Write a byte string to the socket, doubling any IAC characters. This can block if the connection is blocked. May raise *OSError* if the connection is closed.

在 3.3 版更改: This method used to raise *socket.error*, which is now an alias of *OSError*.

`Telnet.interact()`

Interaction function, emulates a very dumb Telnet client.

`Telnet.mt_interact()`

Multithreaded version of *interact()*.

`Telnet.expect(list, timeout=None)`

Read until one from a list of a regular expressions matches.

The first argument is a list of regular expressions, either compiled (*regex objects*) or uncompiled (byte strings). The optional second argument is a timeout, in seconds; the default is to block indefinitely.

Return a tuple of three items: the index in the list of the first regular expression that matches; the match object returned; and the bytes read up till and including the match.

If end of file is found and no bytes were read, raise *EOFError*. Otherwise, when nothing matches, return `(-1, None, data)` where *data* is the bytes received so far (may be empty bytes if a timeout happened).

If a regular expression ends with a greedy match (such as `.*`) or if more than one expression can match the same input, the results are non-deterministic, and may depend on the I/O timing.

`Telnet.set_option_negotiation_callback(callback)`

Each time a telnet option is read on the input flow, this *callback* (if set) is called with the following parameters: `callback(telnet socket, command (DO/DONT/WILL/WONT), option)`. No other action is done afterwards by `telnetlib`.

21.19.2 Telnet Example

A simple example illustrating typical use:

```
import getpass
import telnetlib

HOST = "localhost"
user = input("Enter your remote account: ")
password = getpass.getpass()

tn = telnetlib.Telnet(HOST)

tn.read_until(b"login: ")
tn.write(user.encode('ascii') + b"\n")
if password:
```

(下页继续)

(续上页)

```

tn.read_until(b"Password: ")
tn.write(password.encode('ascii') + b"\n")

tn.write(b"ls\n")
tn.write(b"exit\n")

print(tn.read_all().decode('ascii'))

```

21.20 uuid —UUID objects according to RFC 4122

Source code: [Lib/uuid.py](#)

This module provides immutable *UUID* objects (the *UUID* class) and the functions *uuid1()*, *uuid3()*, *uuid4()*, *uuid5()* for generating version 1, 3, 4, and 5 UUIDs as specified in [RFC 4122](#).

If all you want is a unique ID, you should probably call *uuid1()* or *uuid4()*. Note that *uuid1()* may compromise privacy since it creates a UUID containing the computer's network address. *uuid4()* creates a random UUID.

class *uuid.UUID* (*hex=None, bytes=None, bytes_le=None, fields=None, int=None, version=None*)

Create a UUID from either a string of 32 hexadecimal digits, a string of 16 bytes in big-endian order as the *bytes* argument, a string of 16 bytes in little-endian order as the *bytes_le* argument, a tuple of six integers (32-bit *time_low*, 16-bit *time_mid*, 16-bit *time_hi_version*, 8-bit *clock_seq_hi_variant*, 8-bit *clock_seq_low*, 48-bit *node*) as the *fields* argument, or a single 128-bit integer as the *int* argument. When a string of hex digits is given, curly braces, hyphens, and a URN prefix are all optional. For example, these expressions all yield the same UUID:

```

UUID('{12345678-1234-5678-1234-567812345678}')
UUID('12345678123456781234567812345678')
UUID('urn:uuid:12345678-1234-5678-1234-567812345678')
UUID(bytes=b'\x12\x34\x56\x78'*4)
UUID(bytes_le=b'\x78\x56\x34\x12\x34\x12\x78\x56' +
           b'\x12\x34\x56\x78\x12\x34\x56\x78')
UUID(fields=(0x12345678, 0x1234, 0x5678, 0x12, 0x34, 0x567812345678))
UUID(int=0x12345678123456781234567812345678)

```

Exactly one of *hex*, *bytes*, *bytes_le*, *fields*, or *int* must be given. The *version* argument is optional; if given, the resulting UUID will have its variant and version number set according to [RFC 4122](#), overriding bits in the given *hex*, *bytes*, *bytes_le*, *fields*, or *int*.

Comparison of UUID objects are made by way of comparing their *UUID.int* attributes. Comparison with a non-UUID object raises a *TypeError*.

str(uuid) returns a string in the form 12345678-1234-5678-1234-567812345678 where the 32 hexadecimal digits represent the UUID.

UUID instances have these read-only attributes:

UUID.bytes

The UUID as a 16-byte string (containing the six integer fields in big-endian byte order).

UUID.bytes_le

The UUID as a 16-byte string (with *time_low*, *time_mid*, and *time_hi_version* in little-endian byte order).

UUID.fields

A tuple of the six integer fields of the UUID, which are also available as six individual attributes and two derived attributes:

Field	意义
<code>time_low</code>	the first 32 bits of the UUID
<code>time_mid</code>	the next 16 bits of the UUID
<code>time_hi_version</code>	the next 16 bits of the UUID
<code>clock_seq_hi_variant</code>	the next 8 bits of the UUID
<code>clock_seq_low</code>	the next 8 bits of the UUID
<code>node</code>	the last 48 bits of the UUID
<code>time</code>	the 60-bit timestamp
<code>clock_seq</code>	the 14-bit sequence number

`UUID.hex`

The UUID as a 32-character hexadecimal string.

`UUID.int`

The UUID as a 128-bit integer.

`UUID.urn`

The UUID as a URN as specified in [RFC 4122](#).

`UUID.variant`

The UUID variant, which determines the internal layout of the UUID. This will be one of the constants `RESERVED_NCS`, `RFC_4122`, `RESERVED_MICROSOFT`, or `RESERVED_FUTURE`.

`UUID.version`

The UUID version number (1 through 5, meaningful only when the variant is `RFC_4122`).

The `uuid` module defines the following functions:

`uuid.getnode()`

Get the hardware address as a 48-bit positive integer. The first time this runs, it may launch a separate program, which could be quite slow. If all attempts to obtain the hardware address fail, we choose a random 48-bit number with its eighth bit set to 1 as recommended in [RFC 4122](#). “Hardware address” means the MAC address of a network interface, and on a machine with multiple network interfaces the MAC address of any one of them may be returned.

`uuid.uuid1 (node=None, clock_seq=None)`

Generate a UUID from a host ID, sequence number, and the current time. If `node` is not given, `getnode()` is used to obtain the hardware address. If `clock_seq` is given, it is used as the sequence number; otherwise a random 14-bit sequence number is chosen.

`uuid.uuid3 (namespace, name)`

Generate a UUID based on the MD5 hash of a namespace identifier (which is a UUID) and a name (which is a string).

`uuid.uuid4()`

Generate a random UUID.

`uuid.uuid5 (namespace, name)`

Generate a UUID based on the SHA-1 hash of a namespace identifier (which is a UUID) and a name (which is a string).

The `uuid` module defines the following namespace identifiers for use with `uuid3()` or `uuid5()`.

`uuid.NAMESPACE_DNS`

When this namespace is specified, the `name` string is a fully-qualified domain name.

`uuid.NAMESPACE_URL`

When this namespace is specified, the `name` string is a URL.

`uuid.NAMESPACE_OID`

When this namespace is specified, the *name* string is an ISO OID.

`uuid.NAMESPACE_X500`

When this namespace is specified, the *name* string is an X.500 DN in DER or a text output format.

The `uuid` module defines the following constants for the possible values of the `variant` attribute:

`uuid.RESERVED_NCS`

Reserved for NCS compatibility.

`uuid.RFC_4122`

Specifies the UUID layout given in [RFC 4122](#).

`uuid.RESERVED_MICROSOFT`

Reserved for Microsoft compatibility.

`uuid.RESERVED_FUTURE`

Reserved for future definition.

参见:

RFC 4122 - A Universally Unique Identifier (UUID) URN Namespace This specification defines a Uniform Resource Name namespace for UUIDs, the internal format of UUIDs, and methods of generating UUIDs.

21.20.1 示例

Here are some examples of typical usage of the `uuid` module:

```
>>> import uuid

>>> # make a UUID based on the host ID and current time
>>> uuid.uuid1()
UUID('a8098c1a-f86e-11da-bd1a-00112444be1e')

>>> # make a UUID using an MD5 hash of a namespace UUID and a name
>>> uuid.uuid3(uuid.NAMESPACE_DNS, 'python.org')
UUID('6fa459ea-ee8a-3ca4-894e-db77e160355e')

>>> # make a random UUID
>>> uuid.uuid4()
UUID('16fd2706-8baf-433b-82eb-8c7fada847da')

>>> # make a UUID using a SHA-1 hash of a namespace UUID and a name
>>> uuid.uuid5(uuid.NAMESPACE_DNS, 'python.org')
UUID('886313e1-3b8a-5372-9b90-0c9aee199e5d')

>>> # make a UUID from a string of hex digits (braces and hyphens ignored)
>>> x = uuid.UUID('{00010203-0405-0607-0809-0a0b0c0d0e0f}')

>>> # convert a UUID to a string of hex digits in standard form
>>> str(x)
'00010203-0405-0607-0809-0a0b0c0d0e0f'

>>> # get the raw 16 bytes of the UUID
>>> x.bytes
b'\x00\x01\x02\x03\x04\x05\x06\x07\x08\t\n\x0b\x0c\r\x0e\x0f'

>>> # make a UUID from a 16-byte string
```

(下页继续)

(续上页)

```
>>> uuid.UUID(bytes=x.bytes)
UUID('00010203-0405-0607-0809-0a0b0c0d0e0f')
```

21.21 socketserver — A framework for network servers

Source code: [Lib/socketserver.py](#)

The *socketserver* module simplifies the task of writing network servers.

There are four basic concrete server classes:

```
class socketserver.TCPServer(server_address, RequestHandlerClass, bind_and_activate=True)
    This uses the Internet TCP protocol, which provides for continuous streams of data between the client and
    server. If bind_and_activate is true, the constructor automatically attempts to invoke server_bind() and
    server_activate(). The other parameters are passed to the BaseServer base class.

class socketserver.UDPServer(server_address, RequestHandlerClass, bind_and_activate=True)
    This uses datagrams, which are discrete packets of information that may arrive out of order or be lost while in
    transit. The parameters are the same as for TCPServer.
```

```
class socketserver.UnixStreamServer(server_address, RequestHandlerClass,
                                   bind_and_activate=True)
class socketserver.UnixDatagramServer(server_address, RequestHandlerClass,
                                       bind_and_activate=True)
```

These more infrequently used classes are similar to the TCP and UDP classes, but use Unix domain sockets; they're not available on non-Unix platforms. The parameters are the same as for *TCPServer*.

These four classes process requests *synchronously*; each request must be completed before the next request can be started. This isn't suitable if each request takes a long time to complete, because it requires a lot of computation, or because it returns a lot of data which the client is slow to process. The solution is to create a separate process or thread to handle each request; the *ForkingMixIn* and *ThreadingMixIn* mix-in classes can be used to support asynchronous behaviour.

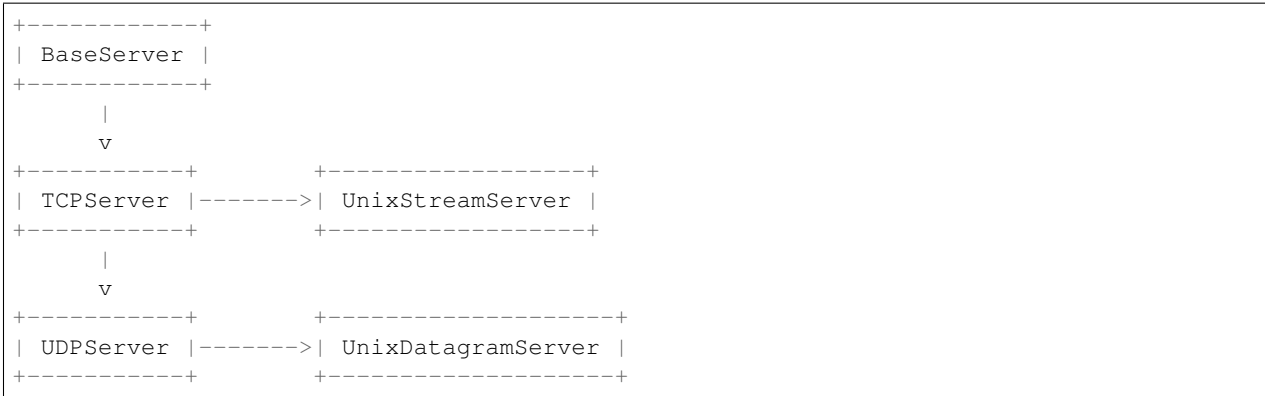
Creating a server requires several steps. First, you must create a request handler class by subclassing the *BaseRequestHandler* class and overriding its *handle()* method; this method will process incoming requests. Second, you must instantiate one of the server classes, passing it the server's address and the request handler class. It is recommended to use the server in a *with* statement. Then call the *handle_request()* or *serve_forever()* method of the server object to process one or many requests. Finally, call *server_close()* to close the socket (unless you used a *with* statement).

When inheriting from *ThreadingMixIn* for threaded connection behavior, you should explicitly declare how you want your threads to behave on an abrupt shutdown. The *ThreadingMixIn* class defines an attribute *daemon_threads*, which indicates whether or not the server should wait for thread termination. You should set the flag explicitly if you would like threads to behave autonomously; the default is *False*, meaning that Python will not exit until all threads created by *ThreadingMixIn* have exited.

Server classes have the same external methods and attributes, no matter what network protocol they use.

21.21.1 Server Creation Notes

There are five classes in an inheritance diagram, four of which represent synchronous servers of four types:



Note that *UnixDatagramServer* derives from *UDPServer*, not from *UnixStreamServer* —the only difference between an IP and a Unix stream server is the address family, which is simply repeated in both Unix server classes.

```
class socketserver.ForkingMixIn
```

```
class socketserver.ThreadingMixIn
```

Forking and threading versions of each type of server can be created using these mix-in classes. For instance, *ThreadingUDPServer* is created as follows:

```
class ThreadingUDPServer(ThreadingMixIn, UDPServer):
    pass
```

The mix-in class comes first, since it overrides a method defined in *UDPServer*. Setting the various attributes also changes the behavior of the underlying server mechanism.

ForkingMixIn and the Forking classes mentioned below are only available on POSIX platforms that support *fork()*.

```
class socketserver.ForkingTCPServer
```

```
class socketserver.ForkingUDPServer
```

```
class socketserver.ThreadingTCPServer
```

```
class socketserver.ThreadingUDPServer
```

These classes are pre-defined using the mix-in classes.

To implement a service, you must derive a class from *BaseRequestHandler* and redefine its *handle()* method. You can then run various versions of the service by combining one of the server classes with your request handler class. The request handler class must be different for datagram or stream services. This can be hidden by using the handler subclasses *StreamRequestHandler* or *DatagramRequestHandler*.

Of course, you still have to use your head! For instance, it makes no sense to use a forking server if the service contains state in memory that can be modified by different requests, since the modifications in the child process would never reach the initial state kept in the parent process and passed to each child. In this case, you can use a threading server, but you will probably have to use locks to protect the integrity of the shared data.

On the other hand, if you are building an HTTP server where all data is stored externally (for instance, in the file system), a synchronous class will essentially render the service “deaf” while one request is being handled – which may be for a very long time if a client is slow to receive all the data it has requested. Here a threading or forking server is appropriate.

In some cases, it may be appropriate to process part of a request synchronously, but to finish processing in a forked child depending on the request data. This can be implemented by using a synchronous server and doing an explicit fork in the request handler class *handle()* method.

Another approach to handling multiple simultaneous requests in an environment that supports neither threads nor *fork()* (or where these are too expensive or inappropriate for the service) is to maintain an explicit table of partially finished requests and to use *selectors* to decide which request to work on next (or whether to handle a new incoming request).

This is particularly important for stream services where each client can potentially be connected for a long time (if threads or subprocesses cannot be used). See [asyncore](#) for another way to manage this.

21.21.2 Server Objects

class `socketserver.BaseServer` (*server_address*, *RequestHandlerClass*)

This is the superclass of all Server objects in the module. It defines the interface, given below, but does not implement most of the methods, which is done in subclasses. The two parameters are stored in the respective *server_address* and *RequestHandlerClass* attributes.

fileno ()

Return an integer file descriptor for the socket on which the server is listening. This function is most commonly passed to *selectors*, to allow monitoring multiple servers in the same process.

handle_request ()

Process a single request. This function calls the following methods in order: *get_request* (), *verify_request* (), and *process_request* (). If the user-provided *handle* () method of the handler class raises an exception, the server's *handle_error* () method will be called. If no request is received within *timeout* seconds, *handle_timeout* () will be called and *handle_request* () will return.

serve_forever (*poll_interval*=0.5)

Handle requests until an explicit *shutdown* () request. Poll for shutdown every *poll_interval* seconds. Ignores the *timeout* attribute. It also calls *service_actions* (), which may be used by a subclass or mixin to provide actions specific to a given service. For example, the *ForkingMixIn* class uses *service_actions* () to clean up zombie child processes.

在 3.3 版更改: Added *service_actions* call to the *serve_forever* method.

service_actions ()

This is called in the *serve_forever* () loop. This method can be overridden by subclasses or mixin classes to perform actions specific to a given service, such as cleanup actions.

3.3 新版功能.

shutdown ()

Tell the *serve_forever* () loop to stop and wait until it does.

server_close ()

Clean up the server. May be overridden.

address_family

The family of protocols to which the server's socket belongs. Common examples are *socket.AF_INET* and *socket.AF_UNIX*.

RequestHandlerClass

The user-provided request handler class; an instance of this class is created for each request.

server_address

The address on which the server is listening. The format of addresses varies depending on the protocol family; see the documentation for the *socket* module for details. For Internet protocols, this is a tuple containing a string giving the address, and an integer port number: ('127.0.0.1', 80), for example.

socket

The socket object on which the server will listen for incoming requests.

The server classes support the following class variables:

allow_reuse_address

Whether the server will allow the reuse of an address. This defaults to *False*, and can be set in subclasses to change the policy.

request_queue_size

The size of the request queue. If it takes a long time to process a single request, any requests that arrive while the server is busy are placed into a queue, up to *request_queue_size* requests. Once the queue is full, further requests from clients will get a “Connection denied” error. The default value is usually 5, but this can be overridden by subclasses.

socket_type

The type of socket used by the server; *socket.SOCK_STREAM* and *socket.SOCK_DGRAM* are two common values.

timeout

Timeout duration, measured in seconds, or *None* if no timeout is desired. If *handle_request()* receives no incoming requests within the timeout period, the *handle_timeout()* method is called.

There are various server methods that can be overridden by subclasses of base server classes like *TCPServer*; these methods aren’t useful to external users of the server object.

finish_request (*request, client_address*)

Actually processes the request by instantiating *RequestHandlerClass* and calling its *handle()* method.

get_request ()

Must accept a request from the socket, and return a 2-tuple containing the *new* socket object to be used to communicate with the client, and the client’s address.

handle_error (*request, client_address*)

This function is called if the *handle()* method of a *RequestHandlerClass* instance raises an exception. The default action is to print the traceback to standard error and continue handling further requests.

在 3.6 版更改: Now only called for exceptions derived from the *Exception* class.

handle_timeout ()

This function is called when the *timeout* attribute has been set to a value other than *None* and the timeout period has passed with no requests being received. The default action for forking servers is to collect the status of any child processes that have exited, while in threading servers this method does nothing.

process_request (*request, client_address*)

Calls *finish_request()* to create an instance of the *RequestHandlerClass*. If desired, this function can create a new process or thread to handle the request; the *ForkingMixIn* and *ThreadingMixIn* classes do this.

server_activate ()

Called by the server’s constructor to activate the server. The default behavior for a TCP server just invokes *listen()* on the server’s socket. May be overridden.

server_bind ()

Called by the server’s constructor to bind the socket to the desired address. May be overridden.

verify_request (*request, client_address*)

Must return a Boolean value; if the value is *True*, the request will be processed, and if it’s *False*, the request will be denied. This function can be overridden to implement access controls for a server. The default implementation always returns *True*.

在 3.6 版更改: Support for the *context manager* protocol was added. Exiting the context manager is equivalent to calling *server_close()*.

21.21.3 Request Handler Objects

class socketserver.BaseRequestHandler

This is the superclass of all request handler objects. It defines the interface, given below. A concrete request handler subclass must define a new `handle()` method, and can override any of the other methods. A new instance of the subclass is created for each request.

setup()

Called before the `handle()` method to perform any initialization actions required. The default implementation does nothing.

handle()

This function must do all the work required to service a request. The default implementation does nothing. Several instance attributes are available to it; the request is available as `self.request`; the client address as `self.client_address`; and the server instance as `self.server`, in case it needs access to per-server information.

The type of `self.request` is different for datagram or stream services. For stream services, `self.request` is a socket object; for datagram services, `self.request` is a pair of string and socket.

finish()

Called after the `handle()` method to perform any clean-up actions required. The default implementation does nothing. If `setup()` raises an exception, this function will not be called.

class socketserver.StreamRequestHandler

class socketserver.DatagramRequestHandler

These `BaseRequestHandler` subclasses override the `setup()` and `finish()` methods, and provide `self.rfile` and `self.wfile` attributes. The `self.rfile` and `self.wfile` attributes can be read or written, respectively, to get the request data or return data to the client.

The `rfile` attributes of both classes support the `io.BufferedIOBase` readable interface, and `DatagramRequestHandler.wfile` supports the `io.BufferedIOBase` writable interface.

在 3.6 版更改: `StreamRequestHandler.wfile` also supports the `io.BufferedIOBase` writable interface.

21.21.4 示例

socketserver.TCPServer Example

This is the server side:

```
import socketserver

class MyTCPHandler(socketserver.BaseRequestHandler):
    """
    The request handler class for our server.

    It is instantiated once per connection to the server, and must
    override the handle() method to implement communication to the
    client.
    """

    def handle(self):
        # self.request is the TCP socket connected to the client
        self.data = self.request.recv(1024).strip()
        print("{} wrote:".format(self.client_address[0]))
```

(下页继续)

(续上页)

```

    print(self.data)
    # just send back the same data, but upper-cased
    self.request.sendall(self.data.upper())

if __name__ == "__main__":
    HOST, PORT = "localhost", 9999

    # Create the server, binding to localhost on port 9999
    with socketserver.TCPServer((HOST, PORT), MyTCPHandler) as server:
        # Activate the server; this will keep running until you
        # interrupt the program with Ctrl-C
        server.serve_forever()

```

An alternative request handler class that makes use of streams (file-like objects that simplify communication by providing the standard file interface):

```

class MyTCPHandler(socketserver.StreamRequestHandler):

    def handle(self):
        # self.rfile is a file-like object created by the handler;
        # we can now use e.g. readline() instead of raw recv() calls
        self.data = self.rfile.readline().strip()
        print("{} wrote:".format(self.client_address[0]))
        print(self.data)
        # Likewise, self.wfile is a file-like object used to write back
        # to the client
        self.wfile.write(self.data.upper())

```

The difference is that the `readline()` call in the second handler will call `recv()` multiple times until it encounters a newline character, while the single `recv()` call in the first handler will just return what has been sent from the client in one `sendall()` call.

This is the client side:

```

import socket
import sys

HOST, PORT = "localhost", 9999
data = " ".join(sys.argv[1:])

# Create a socket (SOCK_STREAM means a TCP socket)
with socket.socket(socket.AF_INET, socket.SOCK_STREAM) as sock:
    # Connect to server and send data
    sock.connect((HOST, PORT))
    sock.sendall(bytes(data + "\n", "utf-8"))

    # Receive data from the server and shut down
    received = str(sock.recv(1024), "utf-8")

print("Sent:      {}".format(data))
print("Received: {}".format(received))

```

The output of the example should look something like this:

Server:

```
$ python TCPServer.py
127.0.0.1 wrote:
b'hello world with TCP'
127.0.0.1 wrote:
b'python is nice'
```

Client:

```
$ python TCPClient.py hello world with TCP
Sent:      hello world with TCP
Received:  HELLO WORLD WITH TCP
$ python TCPClient.py python is nice
Sent:      python is nice
Received:  PYTHON IS NICE
```

socketserver.UDPServer Example

This is the server side:

```
import socketserver

class MyUDPHandler(socketserver.BaseRequestHandler):
    """
    This class works similar to the TCP handler class, except that
    self.request consists of a pair of data and client socket, and since
    there is no connection the client address must be given explicitly
    when sending data back via sendto().
    """

    def handle(self):
        data = self.request[0].strip()
        socket = self.request[1]
        print("{} wrote:".format(self.client_address[0]))
        print(data)
        socket.sendto(data.upper(), self.client_address)

if __name__ == "__main__":
    HOST, PORT = "localhost", 9999
    with socketserver.UDPServer((HOST, PORT), MyUDPHandler) as server:
        server.serve_forever()
```

This is the client side:

```
import socket
import sys

HOST, PORT = "localhost", 9999
data = " ".join(sys.argv[1:])

# SOCK_DGRAM is the socket type to use for UDP sockets
sock = socket.socket(socket.AF_INET, socket.SOCK_DGRAM)

# As you can see, there is no connect() call; UDP has no connections.
# Instead, data is directly sent to the recipient via sendto().
sock.sendto(bytes(data + "\n", "utf-8"), (HOST, PORT))
received = str(sock.recv(1024), "utf-8")
```

(下页继续)

(续上页)

```
print("Sent:      {}".format(data))
print("Received: {}".format(received))
```

The output of the example should look exactly like for the TCP server example.

Asynchronous Mixins

To build asynchronous handlers, use the *ThreadingMixIn* and *ForkingMixIn* classes.

An example for the *ThreadingMixIn* class:

```
import socket
import threading
import socketserver

class ThreadedTCPRequestHandler(socketserver.BaseRequestHandler):

    def handle(self):
        data = str(self.request.recv(1024), 'ascii')
        cur_thread = threading.current_thread()
        response = bytes("{}: {}".format(cur_thread.name, data), 'ascii')
        self.request.sendall(response)

class ThreadedTCPServer(socketserver.ThreadingMixIn, socketserver.TCPServer):
    pass

def client(ip, port, message):
    with socket.socket(socket.AF_INET, socket.SOCK_STREAM) as sock:
        sock.connect((ip, port))
        sock.sendall(bytes(message, 'ascii'))
        response = str(sock.recv(1024), 'ascii')
        print("Received: {}".format(response))

if __name__ == "__main__":
    # Port 0 means to select an arbitrary unused port
    HOST, PORT = "localhost", 0

    server = ThreadedTCPServer((HOST, PORT), ThreadedTCPRequestHandler)
    with server:
        ip, port = server.server_address

        # Start a thread with the server -- that thread will then start one
        # more thread for each request
        server_thread = threading.Thread(target=server.serve_forever)
        # Exit the server thread when the main thread terminates
        server_thread.daemon = True
        server_thread.start()
        print("Server loop running in thread:", server_thread.name)

        client(ip, port, "Hello World 1")
        client(ip, port, "Hello World 2")
        client(ip, port, "Hello World 3")

    server.shutdown()
```

The output of the example should look something like this:

```
$ python ThreadedTCPServer.py
Server loop running in thread: Thread-1
Received: Thread-2: Hello World 1
Received: Thread-3: Hello World 2
Received: Thread-4: Hello World 3
```

The *ForkingMixIn* class is used in the same way, except that the server will spawn a new process for each request. Available only on POSIX platforms that support *fork()*.

21.22 http.server —HTTP 服务器

源代码: [Lib/http/server.py](#)

这个模块定义了实现 HTTP 服务器 (Web 服务器) 的类。

警告: *http.server* is not recommended for production. It only implements only basic security checks.

HTTPServer 是 *socketserver.TCPServer* 的一个子类。它会创建和侦听 HTTP 套接字, 并将请求调度给处理程序。用于创建和运行服务器的代码看起来像这样:

```
def run(server_class=HTTPServer, handler_class=BaseHTTPRequestHandler):
    server_address = ('', 8000)
    httpd = server_class(server_address, handler_class)
    httpd.serve_forever()
```

class `http.server.HTTPServer` (*server_address*, *RequestHandlerClass*)

该类基于 *TCPServer* 类, 并将服务器地址存入名为 *server_name* 和 *server_port* 的实例变量中。服务器可被处理程序通过 *server* 实例变量访问。

The *HTTPServer* must be given a *RequestHandlerClass* on instantiation, of which this module provides three different variants:

class `http.server.BaseHTTPRequestHandler` (*request*, *client_address*, *server*)

This class is used to handle the HTTP requests that arrive at the server. By itself, it cannot respond to any actual HTTP requests; it must be subclassed to handle each request method (e.g. GET or POST). *BaseHTTPRequestHandler* provides a number of class and instance variables, and methods for use by subclasses.

The handler will parse the request and the headers, then call a method specific to the request type. The method name is constructed from the request. For example, for the request method SPAM, the `do_SPAM()` method will be called with no arguments. All of the relevant information is stored in instance variables of the handler. Subclasses should not need to override or extend the `__init__()` method.

BaseHTTPRequestHandler has the following instance variables:

client_address

Contains a tuple of the form (*host*, *port*) referring to the client's address.

server

Contains the server instance.

close_connection

Boolean that should be set before `handle_one_request()` returns, indicating if another request may be expected, or if the connection should be shut down.

requestline

Contains the string representation of the HTTP request line. The terminating CRLF is stripped. This attribute should be set by `handle_one_request()`. If no valid request line was processed, it should be set to the empty string.

command

Contains the command (request type). For example, 'GET'.

path

Contains the request path.

request_version

Contains the version string from the request. For example, 'HTTP/1.0'.

headers

Holds an instance of the class specified by the `MessageClass` class variable. This instance parses and manages the headers in the HTTP request. The `parse_headers()` function from `http.client` is used to parse the headers and it requires that the HTTP request provide a valid **RFC 2822** style header.

rfile

An `io.BufferedReader` input stream, ready to read from the start of the optional input data.

wfile

Contains the output stream for writing a response back to the client. Proper adherence to the HTTP protocol must be used when writing to this stream in order to achieve successful interoperability with HTTP clients.

在 3.6 版更改: This is an `io.BufferedReader` stream.

`BaseHTTPRequestHandler` has the following attributes:

server_version

Specifies the server software version. You may want to override this. The format is multiple whitespace-separated strings, where each string is of the form `name[/version]`. For example, 'BaseHTTP/0.2'.

sys_version

Contains the Python system version, in a form usable by the `version_string` method and the `server_version` class variable. For example, 'Python/1.4'.

error_message_format

Specifies a format string that should be used by `send_error()` method for building an error response to the client. The string is filled by default with variables from `responses` based on the status code that passed to `send_error()`.

error_content_type

Specifies the Content-Type HTTP header of error responses sent to the client. The default value is 'text/html'.

protocol_version

This specifies the HTTP protocol version used in responses. If set to 'HTTP/1.1', the server will permit HTTP persistent connections; however, your server *must* then include an accurate Content-Length header (using `send_header()`) in all of its responses to clients. For backwards compatibility, the setting defaults to 'HTTP/1.0'.

MessageClass

Specifies an `email.message.Message`-like class to parse HTTP headers. Typically, this is not overridden, and it defaults to `http.client.HTTPMessage`.

responses

This attribute contains a mapping of error code integers to two-element tuples containing a short and long message. For example, `{code: (shortmessage, longmessage)}`. The *shortmessage* is usually used as the *message* key in an error response, and *longmessage* as the *explain* key. It is used by `send_response_only()` and `send_error()` methods.

A `BaseHTTPRequestHandler` instance has the following methods:

handle()

Calls `handle_one_request()` once (or, if persistent connections are enabled, multiple times) to handle incoming HTTP requests. You should never need to override it; instead, implement appropriate `do_*` methods.

handle_one_request()

This method will parse and dispatch the request to the appropriate `do_*` method. You should never need to override it.

handle_expect_100()

When a HTTP/1.1 compliant server receives an `Expect: 100-continue` request header it responds back with a `100 Continue` followed by `200 OK` headers. This method can be overridden to raise an error if the server does not want the client to continue. For e.g. server can chose to send `417 Expectation Failed` as a response header and return `False`.

3.2 新版功能.

send_error(code, message=None, explain=None)

Sends and logs a complete error reply to the client. The numeric *code* specifies the HTTP error code, with *message* as an optional, short, human readable description of the error. The *explain* argument can be used to provide more detailed information about the error; it will be formatted using the `error_message_format` attribute and emitted, after a complete set of headers, as the response body. The `responses` attribute holds the default values for *message* and *explain* that will be used if no value is provided; for unknown codes the default value for both is the string `???`. The body will be empty if the method is `HEAD` or the response code is one of the following: `1xx`, `204 No Content`, `205 Reset Content`, `304 Not Modified`.

在 3.4 版更改: The error response includes a Content-Length header. Added the *explain* argument.

send_response(code, message=None)

Adds a response header to the headers buffer and logs the accepted request. The HTTP response line is written to the internal buffer, followed by *Server* and *Date* headers. The values for these two headers are picked up from the `version_string()` and `date_time_string()` methods, respectively. If the server does not intend to send any other headers using the `send_header()` method, then `send_response()` should be followed by an `end_headers()` call.

在 3.3 版更改: Headers are stored to an internal buffer and `end_headers()` needs to be called explicitly.

send_header(keyword, value)

Adds the HTTP header to an internal buffer which will be written to the output stream when either `end_headers()` or `flush_headers()` is invoked. *keyword* should specify the header keyword, with *value* specifying its value. Note that, after the `send_header` calls are done, `end_headers()` MUST BE called in order to complete the operation.

在 3.2 版更改: Headers are stored in an internal buffer.

send_response_only(code, message=None)

Sends the response header only, used for the purposes when `100 Continue` response is sent by the server to the client. The headers not buffered and sent directly the output stream. If the *message* is not specified, the HTTP message corresponding the response *code* is sent.

3.2 新版功能.

end_headers()

Adds a blank line (indicating the end of the HTTP headers in the response) to the headers buffer and calls `flush_headers()`.

在 3.2 版更改: The buffered headers are written to the output stream.

flush_headers()

Finally send the headers to the output stream and flush the internal headers buffer.

3.3 新版功能.

log_request() (*code*='-', *size*='-')

Logs an accepted (successful) request. *code* should specify the numeric HTTP code associated with the response. If a size of the response is available, then it should be passed as the *size* parameter.

log_error(...)

Logs an error when a request cannot be fulfilled. By default, it passes the message to `log_message()`, so it takes the same arguments (*format* and additional values).

log_message() (*format*, ...)

Logs an arbitrary message to `sys.stderr`. This is typically overridden to create custom error logging mechanisms. The *format* argument is a standard printf-style format string, where the additional arguments to `log_message()` are applied as inputs to the formatting. The client ip address and current date and time are prefixed to every message logged.

version_string()

Returns the server software's version string. This is a combination of the `server_version` and `sys_version` attributes.

date_time_string() (*timestamp*=None)

Returns the date and time given by *timestamp* (which must be None or in the format returned by `time.time()`), formatted for a message header. If *timestamp* is omitted, it uses the current date and time.

The result looks like 'Sun, 06 Nov 1994 08:49:37 GMT'.

log_date_time_string()

Returns the current date and time, formatted for logging.

address_string()

Returns the client address.

在 3.3 版更改: Previously, a name lookup was performed. To avoid name resolution delays, it now always returns the IP address.

class `http.server.SimpleHTTPRequestHandler` (*request*, *client_address*, *server*)

This class serves files from the current directory and below, directly mapping the directory structure to HTTP requests.

A lot of the work, such as parsing the request, is done by the base class `BaseHTTPRequestHandler`. This class implements the `do_GET()` and `do_HEAD()` functions.

The following are defined as class-level attributes of `SimpleHTTPRequestHandler`:

server_version

This will be "SimpleHTTP/" + `__version__`, where `__version__` is defined at the module level.

extensions_map

A dictionary mapping suffixes into MIME types. The default is signified by an empty string, and is considered to be `application/octet-stream`. The mapping is used case-insensitively, and so should contain only lower-cased keys.

The `SimpleHTTPRequestHandler` class defines the following methods:

do_HEAD()

This method serves the 'HEAD' request type: it sends the headers it would send for the equivalent GET request. See the `do_GET()` method for a more complete explanation of the possible headers.

do_GET()

The request is mapped to a local file by interpreting the request as a path relative to the current working directory.

If the request was mapped to a directory, the directory is checked for a file named `index.html` or `index.htm` (in that order). If found, the file's contents are returned; otherwise a directory listing is generated by calling the `list_directory()` method. This method uses `os.listdir()` to scan the directory, and returns a 404 error response if the `listdir()` fails.

If the request was mapped to a file, it is opened and the contents are returned. Any `OSError` exception in opening the requested file is mapped to a 404, 'File not found' error. Otherwise, the content type is guessed by calling the `guess_type()` method, which in turn uses the `extensions_map` variable.

A 'Content-type:' header with the guessed content type is output, followed by a 'Content-Length:' header with the file's size and a 'Last-Modified:' header with the file's modification time.

Then follows a blank line signifying the end of the headers, and then the contents of the file are output. If the file's MIME type starts with `text/` the file is opened in text mode; otherwise binary mode is used.

For example usage, see the implementation of the `test()` function invocation in the `http.server` module.

The `SimpleHTTPRequestHandler` class can be used in the following manner in order to create a very basic web-server serving files relative to the current directory:

```
import http.server
import socketserver

PORT = 8000

Handler = http.server.SimpleHTTPRequestHandler

with socketserver.TCPServer(("", PORT), Handler) as httpd:
    print("serving at port", PORT)
    httpd.serve_forever()
```

`http.server` can also be invoked directly using the `-m` switch of the interpreter with a `port` number argument. Similar to the previous example, this serves files relative to the current directory:

```
python -m http.server 8000
```

By default, server binds itself to all interfaces. The option `-b/--bind` specifies a specific address to which it should bind. For example, the following command causes the server to bind to localhost only:

```
python -m http.server 8000 --bind 127.0.0.1
```

3.4 新版功能: `--bind` argument was introduced.

class `http.server.CGIHTTPRequestHandler` (*request, client_address, server*)

This class is used to serve either files or output of CGI scripts from the current directory and below. Note that mapping HTTP hierarchic structure to local directory structure is exactly as in `SimpleHTTPRequestHandler`.

注解: CGI scripts run by the `CGIHTTPRequestHandler` class cannot execute redirects (HTTP code 302), because code 200 (script output follows) is sent prior to execution of the CGI script. This pre-empts the status

code.

The class will however, run the CGI script, instead of serving it as a file, if it guesses it to be a CGI script. Only directory-based CGI are used —the other common server configuration is to treat special extensions as denoting CGI scripts.

The `do_GET()` and `do_HEAD()` functions are modified to run CGI scripts and serve the output, instead of serving files, if the request leads to somewhere below the `cgi_directories` path.

The `CGIHTTPRequestHandler` defines the following data member:

`cgi_directories`

This defaults to `['/cgi-bin', '/htbin']` and describes directories to treat as containing CGI scripts.

The `CGIHTTPRequestHandler` defines the following method:

`do_POST()`

This method serves the `'POST'` request type, only allowed for CGI scripts. Error 501, “Can only POST to CGI scripts”, is output when trying to POST to a non-CGI url.

Note that CGI scripts will be run with UID of user nobody, for security reasons. Problems with the CGI script will be translated to error 403.

`CGIHTTPRequestHandler` can be enabled in the command line by passing the `--cgi` option:

```
python -m http.server --cgi 8000
```

21.23 http.cookies —HTTP state management

Source code: Lib/http/cookies.py

The `http.cookies` module defines classes for abstracting the concept of cookies, an HTTP state management mechanism. It supports both simple string-only cookies, and provides an abstraction for having any serializable data-type as cookie value.

The module formerly strictly applied the parsing rules described in the [RFC 2109](#) and [RFC 2068](#) specifications. It has since been discovered that MSIE 3.0x doesn't follow the character rules outlined in those specs and also many current day browsers and servers have relaxed parsing rules when comes to Cookie handling. As a result, the parsing rules used are a bit less strict.

The character set, `string.ascii_letters`, `string.digits` and `!#$%&'*+-.^_`|~:` denote the set of valid characters allowed by this module in Cookie name (as *key*).

在 3.3 版更改: Allowed `':'` as a valid Cookie name character.

注解: On encountering an invalid cookie, `CookieError` is raised, so if your cookie data comes from a browser you should always prepare for invalid data and catch `CookieError` on parsing.

exception `http.cookies.CookieError`

Exception failing because of [RFC 2109](#) invalidity: incorrect attributes, incorrect `Set-Cookie` header, etc.

class `http.cookies.BaseCookie([input])`

This class is a dictionary-like object whose keys are strings and whose values are *Morsel* instances. Note that upon setting a key to a value, the value is first converted to a *Morsel* containing the key and the value.

If *input* is given, it is passed to the `load()` method.

class `http.cookies.SimpleCookie([input])`

This class derives from `BaseCookie` and overrides `value_decode()` and `value_encode()` to be the identity and `str()` respectively.

参见:

Module `http.cookiejar` HTTP cookie handling for web *clients*. The `http.cookiejar` and `http.cookies` modules do not depend on each other.

RFC 2109 - HTTP State Management Mechanism This is the state management specification implemented by this module.

21.23.1 Cookie Objects

`BaseCookie.value_decode(val)`

Return a decoded value from a string representation. Return value can be any type. This method does nothing in `BaseCookie`—it exists so it can be overridden.

`BaseCookie.value_encode(val)`

Return an encoded value. *val* can be any type, but return value must be a string. This method does nothing in `BaseCookie`—it exists so it can be overridden.

In general, it should be the case that `value_encode()` and `value_decode()` are inverses on the range of `value_decode`.

`BaseCookie.output(attrs=None, header='Set-Cookie:', sep='\r\n')`

Return a string representation suitable to be sent as HTTP headers. *attrs* and *header* are sent to each `Morsel`'s `output()` method. *sep* is used to join the headers together, and is by default the combination `'\r\n'` (CRLF).

`BaseCookie.js_output(attrs=None)`

Return an embeddable JavaScript snippet, which, if run on a browser which supports JavaScript, will act the same as if the HTTP headers was sent.

The meaning for *attrs* is the same as in `output()`.

`BaseCookie.load(rawdata)`

If *rawdata* is a string, parse it as an HTTP_COOKIE and add the values found there as `Morsels`. If it is a dictionary, it is equivalent to:

```
for k, v in rawdata.items():
    cookie[k] = v
```

21.23.2 Morsel Objects

class `http.cookies.Morsel`

Abstract a key/value pair, which has some **RFC 2109** attributes.

Morsels are dictionary-like objects, whose set of keys is constant—the valid **RFC 2109** attributes, which are

- `expires`
- `path`
- `comment`
- `domain`
- `max-age`
- `secure`

- `version`
- `httponly`

The attribute `httponly` specifies that the cookie is only transferred in HTTP requests, and is not accessible through JavaScript. This is intended to mitigate some forms of cross-site scripting.

The keys are case-insensitive and their default value is `''`.

在 3.5 版更改: `__eq__()` now takes *key* and *value* into account.

`Morsel.value`

The value of the cookie.

3.5 版后已移除: assigning to `value`; use `set()` instead.

`Morsel.coded_value`

The encoded value of the cookie —this is what should be sent.

3.5 版后已移除: assigning to `coded_value`; use `set()` instead.

`Morsel.key`

The name of the cookie.

3.5 版后已移除: assigning to `key`; use `set()` instead.

`Morsel.set(key, value, coded_value)`

Set the *key*, *value* and *coded_value* attributes.

3.5 版后已移除: The undocumented *LegalChars* parameter is ignored and will be removed in a future version.

`Morsel.isReservedKey(K)`

Whether *K* is a member of the set of keys of a *Morsel*.

`Morsel.output(attrs=None, header='Set-Cookie:')`

Return a string representation of the *Morsel*, suitable to be sent as an HTTP header. By default, all the attributes are included, unless *attrs* is given, in which case it should be a list of attributes to use. *header* is by default `"Set-Cookie:"`.

`Morsel.js_output(attrs=None)`

Return an embeddable JavaScript snippet, which, if run on a browser which supports JavaScript, will act the same as if the HTTP header was sent.

The meaning for *attrs* is the same as in `output()`.

`Morsel.OutputString(attrs=None)`

Return a string representing the *Morsel*, without any surrounding HTTP or JavaScript.

The meaning for *attrs* is the same as in `output()`.

`Morsel.update(values)`

Update the values in the *Morsel* dictionary with the values in the dictionary *values*. Raise an error if any of the keys in the *values* dict is not a valid **RFC 2109** attribute.

在 3.5 版更改: an error is raised for invalid keys.

`Morsel.copy(value)`

Return a shallow copy of the *Morsel* object.

在 3.5 版更改: return a *Morsel* object instead of a dict.

`Morsel.setdefault(key, value=None)`

Raise an error if *key* is not a valid **RFC 2109** attribute, otherwise behave the same as `dict.setdefault()`.

21.23.3 示例

The following example demonstrates how to use the `http.cookies` module.

```
>>> from http import cookies
>>> C = cookies.SimpleCookie()
>>> C["fig"] = "newton"
>>> C["sugar"] = "wafer"
>>> print(C) # generate HTTP headers
Set-Cookie: fig=newton
Set-Cookie: sugar=wafer
>>> print(C.output()) # same thing
Set-Cookie: fig=newton
Set-Cookie: sugar=wafer
>>> C = cookies.SimpleCookie()
>>> C["rocky"] = "road"
>>> C["rocky"]["path"] = "/cookie"
>>> print(C.output(header="Cookie:"))
Cookie: rocky=road; Path=/cookie
>>> print(C.output(attrs=[], header="Cookie:"))
Cookie: rocky=road
>>> C = cookies.SimpleCookie()
>>> C.load("chips=ahoy; vienna=finger") # load from a string (HTTP header)
>>> print(C)
Set-Cookie: chips=ahoy
Set-Cookie: vienna=finger
>>> C = cookies.SimpleCookie()
>>> C.load('keebler="E=everybody; L=\\\"Loves\\\"; fudge=\\012;";')
>>> print(C)
Set-Cookie: keebler="E=everybody; L=\\\"Loves\\\"; fudge=\\012;";
>>> C = cookies.SimpleCookie()
>>> C["oreo"] = "doublestuff"
>>> C["oreo"]["path"] = "/"
>>> print(C)
Set-Cookie: oreo=doublestuff; Path=/
>>> C = cookies.SimpleCookie()
>>> C["twix"] = "none for you"
>>> C["twix"].value
'none for you'
>>> C = cookies.SimpleCookie()
>>> C["number"] = 7 # equivalent to C["number"] = str(7)
>>> C["string"] = "seven"
>>> C["number"].value
'7'
>>> C["string"].value
'seven'
>>> print(C)
Set-Cookie: number=7
Set-Cookie: string=seven
```

21.24 http.cookiejar —Cookie handling for HTTP clients

Source code: <Lib/http/cookiejar.py>

The `http.cookiejar` module defines classes for automatic handling of HTTP cookies. It is useful for accessing web

sites that require small pieces of data – *cookies* – to be set on the client machine by an HTTP response from a web server, and then returned to the server in later HTTP requests.

Both the regular Netscape cookie protocol and the protocol defined by [RFC 2965](#) are handled. RFC 2965 handling is switched off by default. [RFC 2109](#) cookies are parsed as Netscape cookies and subsequently treated either as Netscape or RFC 2965 cookies according to the ‘policy’ in effect. Note that the great majority of cookies on the Internet are Netscape cookies. [http.cookiejar](#) attempts to follow the de-facto Netscape cookie protocol (which differs substantially from that set out in the original Netscape specification), including taking note of the `max-age` and `port` cookie-attributes introduced with RFC 2965.

注解: The various named parameters found in *Set-Cookie* and *Set-Cookie2* headers (eg. `domain` and `expires`) are conventionally referred to as *attributes*. To distinguish them from Python attributes, the documentation for this module uses the term *cookie-attribute* instead.

The module defines the following exception:

exception `http.cookiejar.LoadError`

Instances of `FileCookieJar` raise this exception on failure to load cookies from a file. `LoadError` is a subclass of `OSError`.

在 3.3 版更改: `LoadError` was made a subclass of `OSError` instead of `IOError`.

The following classes are provided:

class `http.cookiejar.CookieJar` (*policy=None*)
policy is an object implementing the `CookiePolicy` interface.

The `CookieJar` class stores HTTP cookies. It extracts cookies from HTTP requests, and returns them in HTTP responses. `CookieJar` instances automatically expire contained cookies when necessary. Subclasses are also responsible for storing and retrieving cookies from a file or database.

class `http.cookiejar.FileCookieJar` (*filename, delayload=None, policy=None*)
policy is an object implementing the `CookiePolicy` interface. For the other arguments, see the documentation for the corresponding attributes.

A `CookieJar` which can load cookies from, and perhaps save cookies to, a file on disk. Cookies are **NOT** loaded from the named file until either the `load()` or `revert()` method is called. Subclasses of this class are documented in section *FileCookieJar subclasses and co-operation with web browsers*.

class `http.cookiejar.CookiePolicy`

This class is responsible for deciding whether each cookie should be accepted from / returned to the server.

class `http.cookiejar.DefaultCookiePolicy` (*blocked_domains=None, allowed_domains=None, netscape=True, rfc2965=False, rfc2109_as_netscape=None, hide_cookie2=False, strict_domain=False, strict_rfc2965_unverifiable=True, strict_ns_unverifiable=False, strict_ns_domain=DefaultCookiePolicy.DomainLiberal, strict_ns_set_initial_dollar=False, strict_ns_set_path=False*)

Constructor arguments should be passed as keyword arguments only. `blocked_domains` is a sequence of domain names that we never accept cookies from, nor return cookies to. `allowed_domains` if not `None`, this is a sequence of the only domains for which we accept and return cookies. For all other arguments, see the documentation for `CookiePolicy` and `DefaultCookiePolicy` objects.

`DefaultCookiePolicy` implements the standard accept / reject rules for Netscape and [RFC 2965](#) cookies. By default, [RFC 2109](#) cookies (ie. cookies received in a *Set-Cookie* header with a version cookie-attribute of 1) are treated according to the RFC 2965 rules. However, if RFC 2965 handling is turned off or

`rfc2109_as_netscape` is True, RFC 2109 cookies are ‘downgraded’ by the `CookieJar` instance to Netscape cookies, by setting the `version` attribute of the `Cookie` instance to 0. `DefaultCookiePolicy` also provides some parameters to allow some fine-tuning of policy.

class `http.cookiejar.Cookie`

This class represents Netscape, **RFC 2109** and **RFC 2965** cookies. It is not expected that users of `http.cookiejar` construct their own `Cookie` instances. Instead, if necessary, call `make_cookies()` on a `CookieJar` instance.

参见:

Module `urllib.request` URL opening with automatic cookie handling.

Module `http.cookies` HTTP cookie classes, principally useful for server-side code. The `http.cookiejar` and `http.cookies` modules do not depend on each other.

https://curl.haxx.se/rfc/cookie_spec.html The specification of the original Netscape cookie protocol. Though this is still the dominant protocol, the ‘Netscape cookie protocol’ implemented by all the major browsers (and `http.cookiejar`) only bears a passing resemblance to the one sketched out in `cookie_spec.html`.

RFC 2109 - HTTP State Management Mechanism Obsoleted by **RFC 2965**. Uses `Set-Cookie` with `version=1`.

RFC 2965 - HTTP State Management Mechanism The Netscape protocol with the bugs fixed. Uses `Set-Cookie2` in place of `Set-Cookie`. Not widely used.

<http://kristol.org/cookie/errata.html> Unfinished errata to **RFC 2965**.

RFC 2964 - Use of HTTP State Management

21.24.1 CookieJar and FileCookieJar Objects

`CookieJar` objects support the `iterator` protocol for iterating over contained `Cookie` objects.

`CookieJar` has the following methods:

`CookieJar.add_cookie_header(request)`

Add correct `Cookie` header to `request`.

If policy allows (ie. the `rfc2965` and `hide_cookie2` attributes of the `CookieJar`’s `CookiePolicy` instance are true and false respectively), the `Cookie2` header is also added when appropriate.

The `request` object (usually a `urllib.request.Request` instance) must support the methods `get_full_url()`, `get_host()`, `get_type()`, `unverifiable()`, `has_header()`, `get_header()`, `header_items()`, `add_unredirected_header()` and `origin_req_host` attribute as documented by `urllib.request`.

在 3.3 版更改: `request` object needs `origin_req_host` attribute. Dependency on a deprecated method `get_origin_req_host()` has been removed.

`CookieJar.extract_cookies(response, request)`

Extract cookies from HTTP `response` and store them in the `CookieJar`, where allowed by policy.

The `CookieJar` will look for allowable `Set-Cookie` and `Set-Cookie2` headers in the `response` argument, and store cookies as appropriate (subject to the `CookiePolicy.set_ok()` method’ s approval).

The `response` object (usually the result of a call to `urllib.request.urlopen()`, or similar) should support an `info()` method, which returns an `email.message.Message` instance.

The `request` object (usually a `urllib.request.Request` instance) must support the methods `get_full_url()`, `get_host()`, `unverifiable()`, and `origin_req_host` attribute, as documented by `urllib.request`. The request is used to set default values for cookie-attributes as well as for checking that the cookie is allowed to be set.

在 3.3 版更改: *request* object needs *origin_req_host* attribute. Dependency on a deprecated method *get_origin_req_host()* has been removed.

`CookieJar.set_policy(policy)`

Set the *CookiePolicy* instance to be used.

`CookieJar.make_cookies(response, request)`

Return sequence of *Cookie* objects extracted from *response* object.

See the documentation for *extract_cookies()* for the interfaces required of the *response* and *request* arguments.

`CookieJar.set_cookie_if_ok(cookie, request)`

Set a *Cookie* if policy says it's OK to do so.

`CookieJar.set_cookie(cookie)`

Set a *Cookie*, without checking with policy to see whether or not it should be set.

`CookieJar.clear([domain[, path[, name]]])`

Clear some cookies.

If invoked without arguments, clear all cookies. If given a single argument, only cookies belonging to that *domain* will be removed. If given two arguments, cookies belonging to the specified *domain* and URL *path* are removed. If given three arguments, then the cookie with the specified *domain*, *path* and *name* is removed.

Raises *KeyError* if no matching cookie exists.

`CookieJar.clear_session_cookies()`

Discard all session cookies.

Discards all contained cookies that have a true *discard* attribute (usually because they had either no *max-age* or *expires* cookie-attribute, or an explicit *discard* cookie-attribute). For interactive browsers, the end of a session usually corresponds to closing the browser window.

Note that the *save()* method won't save session cookies anyway, unless you ask otherwise by passing a true *ignore_discard* argument.

FileCookieJar implements the following additional methods:

`FileCookieJar.save(filename=None, ignore_discard=False, ignore_expires=False)`

Save cookies to a file.

This base class raises *NotImplementedError*. Subclasses may leave this method unimplemented.

filename is the name of file in which to save cookies. If *filename* is not specified, *self.filename* is used (whose default is the value passed to the constructor, if any); if *self.filename* is *None*, *ValueError* is raised.

ignore_discard: save even cookies set to be discarded. *ignore_expires*: save even cookies that have expired

The file is overwritten if it already exists, thus wiping all the cookies it contains. Saved cookies can be restored later using the *load()* or *revert()* methods.

`FileCookieJar.load(filename=None, ignore_discard=False, ignore_expires=False)`

Load cookies from a file.

Old cookies are kept unless overwritten by newly loaded ones.

Arguments are as for *save()*.

The named file must be in the format understood by the class, or *LoadError* will be raised. Also, *OSError* may be raised, for example if the file does not exist.

在 3.3 版更改: 过去触发的 *IOError*, 现在是 *OSError* 的别名。

`FileCookieJar.revert` (*filename=None, ignore_discard=False, ignore_expires=False*)

Clear all cookies and reload cookies from a saved file.

`revert()` can raise the same exceptions as `load()`. If there is a failure, the object's state will not be altered.

`FileCookieJar` instances have the following public attributes:

`FileCookieJar.filename`

Filename of default file in which to keep cookies. This attribute may be assigned to.

`FileCookieJar.delayload`

If true, load cookies lazily from disk. This attribute should not be assigned to. This is only a hint, since this only affects performance, not behaviour (unless the cookies on disk are changing). A `CookieJar` object may ignore it. None of the `FileCookieJar` classes included in the standard library lazily loads cookies.

21.24.2 FileCookieJar subclasses and co-operation with web browsers

The following `CookieJar` subclasses are provided for reading and writing.

class `http.cookiejar.MozillaCookieJar` (*filename, delayload=None, policy=None*)

A `FileCookieJar` that can load from and save cookies to disk in the Mozilla `cookies.txt` file format (which is also used by the Lynx and Netscape browsers).

注解: This loses information about **RFC 2965** cookies, and also about newer or non-standard cookie-attributes such as `port`.

警告: Back up your cookies before saving if you have cookies whose loss / corruption would be inconvenient (there are some subtleties which may lead to slight changes in the file over a load / save round-trip).

Also note that cookies saved while Mozilla is running will get clobbered by Mozilla.

class `http.cookiejar.LWPCookieJar` (*filename, delayload=None, policy=None*)

A `FileCookieJar` that can load from and save cookies to disk in format compatible with the libwww-perl library's Set-Cookie3 file format. This is convenient if you want to store cookies in a human-readable file.

21.24.3 CookiePolicy Objects

Objects implementing the `CookiePolicy` interface have the following methods:

`CookiePolicy.set_ok` (*cookie, request*)

Return boolean value indicating whether cookie should be accepted from server.

cookie is a `Cookie` instance. *request* is an object implementing the interface defined by the documentation for `CookieJar.extract_cookies()`.

`CookiePolicy.return_ok` (*cookie, request*)

Return boolean value indicating whether cookie should be returned to server.

cookie is a `Cookie` instance. *request* is an object implementing the interface defined by the documentation for `CookieJar.add_cookie_header()`.

`CookiePolicy.domain_return_ok` (*domain, request*)

Return false if cookies should not be returned, given cookie domain.

This method is an optimization. It removes the need for checking every cookie with a particular domain (which might involve reading many files). Returning true from `domain_return_ok()` and `path_return_ok()` leaves all the work to `return_ok()`.

If `domain_return_ok()` returns true for the cookie domain, `path_return_ok()` is called for the cookie path. Otherwise, `path_return_ok()` and `return_ok()` are never called for that cookie domain. If `path_return_ok()` returns true, `return_ok()` is called with the `Cookie` object itself for a full check. Otherwise, `return_ok()` is never called for that cookie path.

Note that `domain_return_ok()` is called for every *cookie* domain, not just for the *request* domain. For example, the function might be called with both `".example.com"` and `"www.example.com"` if the request domain is `"www.example.com"`. The same goes for `path_return_ok()`.

The *request* argument is as documented for `return_ok()`.

`CookiePolicy.path_return_ok(path, request)`

Return false if cookies should not be returned, given cookie path.

See the documentation for `domain_return_ok()`.

In addition to implementing the methods above, implementations of the `CookiePolicy` interface must also supply the following attributes, indicating which protocols should be used, and how. All of these attributes may be assigned to.

`CookiePolicy.netscape`

Implement Netscape protocol.

`CookiePolicy.rfc2965`

Implement **RFC 2965** protocol.

`CookiePolicy.hide_cookie2`

Don't add `Cookie2` header to requests (the presence of this header indicates to the server that we understand **RFC 2965** cookies).

The most useful way to define a `CookiePolicy` class is by subclassing from `DefaultCookiePolicy` and overriding some or all of the methods above. `CookiePolicy` itself may be used as a 'null policy' to allow setting and receiving any and all cookies (this is unlikely to be useful).

21.24.4 DefaultCookiePolicy Objects

Implements the standard rules for accepting and returning cookies.

Both **RFC 2965** and Netscape cookies are covered. RFC 2965 handling is switched off by default.

The easiest way to provide your own policy is to override this class and call its methods in your overridden implementations before adding your own additional checks:

```
import http.cookiejar
class MyCookiePolicy(http.cookiejar.DefaultCookiePolicy):
    def set_ok(self, cookie, request):
        if not http.cookiejar.DefaultCookiePolicy.set_ok(self, cookie, request):
            return False
        if i_dont_want_to_store_this_cookie(cookie):
            return False
        return True
```

In addition to the features required to implement the `CookiePolicy` interface, this class allows you to block and allow domains from setting and receiving cookies. There are also some strictness switches that allow you to tighten up the rather loose Netscape protocol rules a little bit (at the cost of blocking some benign cookies).

A domain blacklist and whitelist is provided (both off by default). Only domains not in the blacklist and present in the whitelist (if the whitelist is active) participate in cookie setting and returning. Use the *blocked_domains* constructor argument, and `blocked_domains()` and `set_blocked_domains()` methods (and the corresponding argument and methods for *allowed_domains*). If you set a whitelist, you can turn it off again by setting it to *None*.

Domains in block or allow lists that do not start with a dot must equal the cookie domain to be matched. For example, "example.com" matches a blacklist entry of "example.com", but "www.example.com" does not. Domains that do start with a dot are matched by more specific domains too. For example, both "www.example.com" and "www.coyote.example.com" match ".example.com" (but "example.com" itself does not). IP addresses are an exception, and must match exactly. For example, if `blocked_domains` contains "192.168.1.2" and ".168.1.2", 192.168.1.2 is blocked, but 193.168.1.2 is not.

DefaultCookiePolicy implements the following additional methods:

`DefaultCookiePolicy.blocked_domains()`

Return the sequence of blocked domains (as a tuple).

`DefaultCookiePolicy.set_blocked_domains(blocked_domains)`

Set the sequence of blocked domains.

`DefaultCookiePolicy.is_blocked(domain)`

Return whether *domain* is on the blacklist for setting or receiving cookies.

`DefaultCookiePolicy.allowed_domains()`

Return *None*, or the sequence of allowed domains (as a tuple).

`DefaultCookiePolicy.set_allowed_domains(allowed_domains)`

Set the sequence of allowed domains, or *None*.

`DefaultCookiePolicy.is_not_allowed(domain)`

Return whether *domain* is not on the whitelist for setting or receiving cookies.

DefaultCookiePolicy instances have the following attributes, which are all initialised from the constructor arguments of the same name, and which may all be assigned to.

`DefaultCookiePolicy.rfc2109_as_netscape`

If true, request that the *CookieJar* instance downgrade **RFC 2109** cookies (ie. cookies received in a *Set-Cookie* header with a version cookie-attribute of 1) to Netscape cookies by setting the version attribute of the *Cookie* instance to 0. The default value is *None*, in which case RFC 2109 cookies are downgraded if and only if **RFC 2965** handling is turned off. Therefore, RFC 2109 cookies are downgraded by default.

General strictness switches:

`DefaultCookiePolicy.strict_domain`

Don't allow sites to set two-component domains with country-code top-level domains like .co.uk, .gov.uk, .co.nz.etc. This is far from perfect and isn't guaranteed to work!

RFC 2965 protocol strictness switches:

`DefaultCookiePolicy.strict_rfc2965_unverifiable`

Follow **RFC 2965** rules on unverifiable transactions (usually, an unverifiable transaction is one resulting from a redirect or a request for an image hosted on another site). If this is false, cookies are *never* blocked on the basis of verifiability

Netscape protocol strictness switches:

`DefaultCookiePolicy.strict_ns_unverifiable`

Apply **RFC 2965** rules on unverifiable transactions even to Netscape cookies.

`DefaultCookiePolicy.strict_ns_domain`

Flags indicating how strict to be with domain-matching rules for Netscape cookies. See below for acceptable values.

`DefaultCookiePolicy.strict_ns_set_initial_dollar`

Ignore cookies in Set-Cookie: headers that have names starting with '\$'.

`DefaultCookiePolicy.strict_ns_set_path`

Don't allow setting cookies whose path doesn't path-match request URI.

`strict_ns_domain` is a collection of flags. Its value is constructed by or-ing together (for example, `DomainStrictNoDots|DomainStrictNonDomain` means both flags are set).

`DefaultCookiePolicy.DomainStrictNoDots`

When setting cookies, the 'host prefix' must not contain a dot (eg. `www.foo.bar.com` can't set a cookie for `.bar.com`, because `www.foo` contains a dot).

`DefaultCookiePolicy.DomainStrictNonDomain`

Cookies that did not explicitly specify a domain cookie-attribute can only be returned to a domain equal to the domain that set the cookie (eg. `spam.example.com` won't be returned cookies from `example.com` that had no domain cookie-attribute).

`DefaultCookiePolicy.DomainRFC2965Match`

When setting cookies, require a full [RFC 2965](#) domain-match.

The following attributes are provided for convenience, and are the most useful combinations of the above flags:

`DefaultCookiePolicy.DomainLiberal`

Equivalent to 0 (ie. all of the above Netscape domain strictness flags switched off).

`DefaultCookiePolicy.DomainStrict`

Equivalent to `DomainStrictNoDots|DomainStrictNonDomain`.

21.24.5 Cookie Objects

`Cookie` instances have Python attributes roughly corresponding to the standard cookie-attributes specified in the various cookie standards. The correspondence is not one-to-one, because there are complicated rules for assigning default values, because the `max-age` and `expires` cookie-attributes contain equivalent information, and because [RFC 2109](#) cookies may be 'downgraded' by `http.cookiejar` from version 1 to version 0 (Netscape) cookies.

Assignment to these attributes should not be necessary other than in rare circumstances in a `CookiePolicy` method. The class does not enforce internal consistency, so you should know what you're doing if you do that.

`Cookie.version`

Integer or `None`. Netscape cookies have `version` 0. [RFC 2965](#) and [RFC 2109](#) cookies have a `version` cookie-attribute of 1. However, note that `http.cookiejar` may 'downgrade' [RFC 2109](#) cookies to Netscape cookies, in which case `version` is 0.

`Cookie.name`

Cookie name (a string).

`Cookie.value`

Cookie value (a string), or `None`.

`Cookie.port`

String representing a port or a set of ports (eg. '80', or '80,8080'), or `None`.

`Cookie.path`

Cookie path (a string, eg. '/acme/rocket_launchers').

`Cookie.secure`

True if cookie should only be returned over a secure connection.

`Cookie.expires`

Integer expiry date in seconds since epoch, or `None`. See also the `is_expired()` method.

`Cookie.discard`

True if this is a session cookie.

`Cookie.comment`

String comment from the server explaining the function of this cookie, or *None*.

`Cookie.comment_url`

URL linking to a comment from the server explaining the function of this cookie, or *None*.

`Cookie.rfc2109`

True if this cookie was received as an **RFC 2109** cookie (ie. the cookie arrived in a *Set-Cookie* header, and the value of the Version cookie-attribute in that header was 1). This attribute is provided because *http.cookiejar* may ‘downgrade’ RFC 2109 cookies to Netscape cookies, in which case *version* is 0.

`Cookie.port_specified`

True if a port or set of ports was explicitly specified by the server (in the *Set-Cookie* / *Set-Cookie2* header).

`Cookie.domain_specified`

True if a domain was explicitly specified by the server.

`Cookie.domain_initial_dot`

True if the domain explicitly specified by the server began with a dot ('.').

Cookies may have additional non-standard cookie-attributes. These may be accessed using the following methods:

`Cookie.has_nonstandard_attr(name)`

Return true if cookie has the named cookie-attribute.

`Cookie.get_nonstandard_attr(name, default=None)`

If cookie has the named cookie-attribute, return its value. Otherwise, return *default*.

`Cookie.set_nonstandard_attr(name, value)`

Set the value of the named cookie-attribute.

The *Cookie* class also defines the following method:

`Cookie.is_expired(now=None)`

True if cookie has passed the time at which the server requested it should expire. If *now* is given (in seconds since the epoch), return whether the cookie has expired at the specified time.

21.24.6 示例

The first example shows the most common usage of *http.cookiejar*:

```
import http.cookiejar, urllib.request
cj = http.cookiejar.CookieJar()
opener = urllib.request.build_opener(urllib.request.HTTPCookieProcessor(cj))
r = opener.open("http://example.com/")
```

This example illustrates how to open a URL using your Netscape, Mozilla, or Lynx cookies (assumes Unix/Netscape convention for location of the cookies file):

```
import os, http.cookiejar, urllib.request
cj = http.cookiejar.MozillaCookieJar()
cj.load(os.path.join(os.path.expanduser("~"), ".netscape", "cookies.txt"))
opener = urllib.request.build_opener(urllib.request.HTTPCookieProcessor(cj))
r = opener.open("http://example.com/")
```

The next example illustrates the use of `DefaultCookiePolicy`. Turn on **RFC 2965** cookies, be more strict about domains when setting and returning Netscape cookies, and block some domains from setting cookies or having them returned:

```
import urllib.request
from http.cookiejar import CookieJar, DefaultCookiePolicy
policy = DefaultCookiePolicy(
    rfc2965=True, strict_ns_domain=Policy.DomainStrict,
    blocked_domains=["ads.net", ".ads.net"])
cj = CookieJar(policy)
opener = urllib.request.build_opener(urllib.request.HTTPCookieProcessor(cj))
r = opener.open("http://example.com/")
```

21.25 xmlrpc —XMLRPC 服务端与客户端模块

XML-RPC 是一种远程过程调用方法，它使用通过 HTTP 传递的 XML 作为载体。有了它，客户端可以在远程服务器上调用带参数的方法（服务器以 URI 命名）并获取结构化的数据。

`xmlrpc` 是一个集合了 XML-RPC 服务端与客户端实现模块的包。这些模块是：

- `xmlrpc.client`
- `xmlrpc.server`

21.26 xmlrpc.client —XML-RPC client access

Source code: `Lib/xmlrpc/client.py`

XML-RPC is a Remote Procedure Call method that uses XML passed via HTTP(S) as a transport. With it, a client can call methods with parameters on a remote server (the server is named by a URI) and get back structured data. This module supports writing XML-RPC client code; it handles all the details of translating between conformable Python objects and XML on the wire.

警告： The `xmlrpc.client` module is not secure against maliciously constructed data. If you need to parse untrusted or unauthenticated data see *XML 漏洞*.

在 3.5 版更改: For HTTPS URIs, `xmlrpc.client` now performs all the necessary certificate and hostname checks by default.

```
class xmlrpc.client.ServerProxy(uri, transport=None, encoding=None, verbose=False, allow_none=False, use_datetime=False, use_builtin_types=False, *, context=None)
```

在 3.3 版更改: The `use_builtin_types` flag was added.

A `ServerProxy` instance is an object that manages communication with a remote XML-RPC server. The required first argument is a URI (Uniform Resource Indicator), and will normally be the URL of the server. The optional second argument is a transport factory instance; by default it is an internal `SafeTransport` instance for https: URLs and an internal `HTTPTransport` instance otherwise. The optional third argument is an encoding, by default UTF-8. The optional fourth argument is a debugging flag.

The following parameters govern the use of the returned proxy instance. If `allow_none` is true, the Python constant `None` will be translated into XML; the default behaviour is for `None` to raise a `TypeError`. This

is a commonly-used extension to the XML-RPC specification, but isn't supported by all clients and servers; see <http://ontosys.com/xml-rpc/extensions.php> for a description. The `use_builtin_types` flag can be used to cause date/time values to be presented as `datetime.datetime` objects and binary data to be presented as `bytes` objects; this flag is false by default. `datetime.datetime`, `bytes` and `bytearray` objects may be passed to calls. The obsolete `use_datetime` flag is similar to `use_builtin_types` but it applies only to date/time values.

Both the HTTP and HTTPS transports support the URL syntax extension for HTTP Basic Authentication: `http://user:pass@host:port/path`. The `user:pass` portion will be base64-encoded as an HTTP 'Authorization' header, and sent to the remote server as part of the connection process when invoking an XML-RPC method. You only need to use this if the remote server requires a Basic Authentication user and password. If an HTTPS URL is provided, `context` may be `ssl.SSLContext` and configures the SSL settings of the underlying HTTPS connection.

The returned instance is a proxy object with methods that can be used to invoke corresponding RPC calls on the remote server. If the remote server supports the introspection API, the proxy can also be used to query the remote server for the methods it supports (service discovery) and fetch other server-associated metadata.

Types that are conformable (e.g. that can be marshalled through XML), include the following (and except where noted, they are unmarshalled as the same Python type):

XML-RPC type	Python 数据类型
boolean	<code>bool</code>
int, i1, i2, i4, i8 or biginteger	<code>int</code> in range from -2147483648 to 2147483647. Values get the <code><int></code> tag.
double or float	<code>float</code> . Values get the <code><double></code> tag.
string	<code>str</code>
array	<code>list</code> or <code>tuple</code> containing conformable elements. Arrays are returned as <code>lists</code> .
struct	<code>dict</code> . Keys must be strings, values may be any conformable type. Objects of user-defined classes can be passed in; only their <code>__dict__</code> attribute is transmitted.
<code>dateTime.iso8601</code>	<code>DateTime</code> or <code>datetime.datetime</code> . Returned type depends on values of <code>use_builtin_types</code> and <code>use_datetime</code> flags.
base64	<code>Binary</code> , <code>bytes</code> or <code>bytearray</code> . Returned type depends on the value of the <code>use_builtin_types</code> flag.
nil	The <code>None</code> constant. Passing is allowed only if <code>allow_none</code> is true.
bigdecimal	<code>decimal.Decimal</code> . Returned type only.

This is the full set of data types supported by XML-RPC. Method calls may also raise a special `Fault` instance, used to signal XML-RPC server errors, or `ProtocolError` used to signal an error in the HTTP/HTTPS transport layer. Both `Fault` and `ProtocolError` derive from a base class called `Error`. Note that the `xmlrpc` client module currently does not marshal instances of subclasses of built-in types.

When passing strings, characters special to XML such as `<`, `>`, and `&` will be automatically escaped. However, it's the caller's responsibility to ensure that the string is free of characters that aren't allowed in XML, such as the control characters with ASCII values between 0 and 31 (except, of course, tab, newline and carriage return); failing to do this will result in an XML-RPC request that isn't well-formed XML. If you have to pass arbitrary bytes via XML-RPC, use `bytes` or `bytearray` classes or the `Binary` wrapper class described below.

`Server` is retained as an alias for `ServerProxy` for backwards compatibility. New code should use `ServerProxy`.

在 3.5 版更改: Added the `context` argument.

在 3.6 版更改: Added support of type tags with prefixes (e.g. `ex:nil`). Added support of unmarshalling additional types used by Apache XML-RPC implementation for numerics: `i1`, `i2`, `i8`, `biginteger`, `float` and `bigdecimal`. See <http://ws.apache.org/xmlrpc/types.html> for a description.

参见:

XML-RPC HOWTO A good description of XML-RPC operation and client software in several languages. Contains pretty much everything an XML-RPC client developer needs to know.

XML-RPC Introspection Describes the XML-RPC protocol extension for introspection.

XML-RPC Specification The official specification.

Unofficial XML-RPC Errata Fredrik Lundh's "unofficial errata, intended to clarify certain details in the XML-RPC specification, as well as hint at 'best practices' to use when designing your own XML-RPC implementations."

21.26.1 ServerProxy Objects

A *ServerProxy* instance has a method corresponding to each remote procedure call accepted by the XML-RPC server. Calling the method performs an RPC, dispatched by both name and argument signature (e.g. the same method name can be overloaded with multiple argument signatures). The RPC finishes by returning a value, which may be either returned data in a conformant type or a *Fault* or *ProtocolError* object indicating an error.

Servers that support the XML introspection API support some common methods grouped under the reserved `system` attribute:

`ServerProxy.system.listMethods()`

This method returns a list of strings, one for each (non-system) method supported by the XML-RPC server.

`ServerProxy.system.methodSignature(name)`

This method takes one parameter, the name of a method implemented by the XML-RPC server. It returns an array of possible signatures for this method. A signature is an array of types. The first of these types is the return type of the method, the rest are parameters.

Because multiple signatures (ie. overloading) is permitted, this method returns a list of signatures rather than a singleton.

Signatures themselves are restricted to the top level parameters expected by a method. For instance if a method expects one array of structs as a parameter, and it returns a string, its signature is simply "string, array". If it expects three integers and returns a string, its signature is "string, int, int, int".

If no signature is defined for the method, a non-array value is returned. In Python this means that the type of the returned value will be something other than list.

`ServerProxy.system.methodHelp(name)`

This method takes one parameter, the name of a method implemented by the XML-RPC server. It returns a documentation string describing the use of that method. If no such string is available, an empty string is returned. The documentation string may contain HTML markup.

在 3.5 版更改: Instances of *ServerProxy* support the *context manager* protocol for closing the underlying transport.

A working example follows. The server code:

```
from xmlrpc.server import SimpleXMLRPCServer

def is_even(n):
    return n % 2 == 0

server = SimpleXMLRPCServer(("localhost", 8000))
print("Listening on port 8000...")
server.register_function(is_even, "is_even")
server.serve_forever()
```

The client code for the preceding server:

```
import xmlrpc.client

with xmlrpc.client.ServerProxy("http://localhost:8000/") as proxy:
    print("3 is even: %s" % str(proxy.is_even(3)))
    print("100 is even: %s" % str(proxy.is_even(100)))
```

21.26.2 DateTime 对象

class xmlrpc.client.DateTime

This class may be initialized with seconds since the epoch, a time tuple, an ISO 8601 time/date string, or a *datetime.datetime* instance. It has the following methods, supported mainly for internal use by the marshalling/unmarshalling code:

decode (*string*)

Accept a string as the instance's new time value.

encode (*out*)

Write the XML-RPC encoding of this *DateTime* item to the *out* stream object.

It also supports certain of Python's built-in operators through rich comparison and `__repr__()` methods.

A working example follows. The server code:

```
import datetime
from xmlrpc.server import SimpleXMLRPCServer
import xmlrpc.client

def today():
    today = datetime.datetime.today()
    return xmlrpc.client.DateTime(today)

server = SimpleXMLRPCServer(("localhost", 8000))
print("Listening on port 8000...")
server.register_function(today, "today")
server.serve_forever()
```

The client code for the preceding server:

```
import xmlrpc.client
import datetime

proxy = xmlrpc.client.ServerProxy("http://localhost:8000/")

today = proxy.today()
# convert the ISO8601 string to a datetime object
converted = datetime.datetime.strptime(today.value, "%Y%m%dT%H:%M:%S")
print("Today: %s" % converted.strftime("%d.%m.%Y, %H:%M"))
```

21.26.3 Binary Objects

class xmlrpc.client.Binary

This class may be initialized from bytes data (which may include NULs). The primary access to the content of a *Binary* object is provided by an attribute:

data

The binary data encapsulated by the *Binary* instance. The data is provided as a *bytes* object.

Binary objects have the following methods, supported mainly for internal use by the marshalling/unmarshalling code:

decode (*bytes*)

Accept a base64 *bytes* object and decode it as the instance's new data.

encode (*out*)

Write the XML-RPC base 64 encoding of this binary item to the *out* stream object.

The encoded data will have newlines every 76 characters as per [RFC 2045 section 6.8](#), which was the de facto standard base64 specification when the XML-RPC spec was written.

It also supports certain of Python's built-in operators through `__eq__()` and `__ne__()` methods.

Example usage of the binary objects. We're going to transfer an image over XMLRPC:

```
from xmlrpc.server import SimpleXMLRPCServer
import xmlrpc.client

def python_logo():
    with open("python_logo.jpg", "rb") as handle:
        return xmlrpc.client.Binary(handle.read())

server = SimpleXMLRPCServer(("localhost", 8000))
print("Listening on port 8000...")
server.register_function(python_logo, 'python_logo')

server.serve_forever()
```

The client gets the image and saves it to a file:

```
import xmlrpc.client

proxy = xmlrpc.client.ServerProxy("http://localhost:8000/")
with open("fetched_python_logo.jpg", "wb") as handle:
    handle.write(proxy.python_logo().data)
```

21.26.4 Fault Objects

class `xmlrpc.client.Fault`

A *Fault* object encapsulates the content of an XML-RPC fault tag. Fault objects have the following attributes:

faultCode

A string indicating the fault type.

faultString

A string containing a diagnostic message associated with the fault.

In the following example we're going to intentionally cause a *Fault* by returning a complex type object. The server code:

```
from xmlrpc.server import SimpleXMLRPCServer

# A marshalling error is going to occur because we're returning a
# complex number
def add(x, y):
    return x+y+0j
```

(下页继续)

(续上页)

```
server = SimpleXMLRPCServer(("localhost", 8000))
print("Listening on port 8000...")
server.register_function(add, 'add')

server.serve_forever()
```

The client code for the preceding server:

```
import xmlrpc.client

proxy = xmlrpc.client.ServerProxy("http://localhost:8000/")
try:
    proxy.add(2, 5)
except xmlrpc.client.Fault as err:
    print("A fault occurred")
    print("Fault code: %d" % err.faultCode)
    print("Fault string: %s" % err.faultString)
```

21.26.5 ProtocolError Objects

class `xmlrpc.client.ProtocolError`

A *ProtocolError* object describes a protocol error in the underlying transport layer (such as a 404 ‘not found’ error if the server named by the URI does not exist). It has the following attributes:

url

The URI or URL that triggered the error.

errcode

The error code.

errmsg

The error message or diagnostic string.

headers

A dict containing the headers of the HTTP/HTTPS request that triggered the error.

In the following example we’re going to intentionally cause a *ProtocolError* by providing an invalid URI:

```
import xmlrpc.client

# create a ServerProxy with a URI that doesn't respond to XMLRPC requests
proxy = xmlrpc.client.ServerProxy("http://google.com/")

try:
    proxy.some_method()
except xmlrpc.client.ProtocolError as err:
    print("A protocol error occurred")
    print("URL: %s" % err.url)
    print("HTTP/HTTPS headers: %s" % err.headers)
    print("Error code: %d" % err.errcode)
    print("Error message: %s" % err.errmsg)
```

21.26.6 MultiCall Objects

The *MultiCall* object provides a way to encapsulate multiple calls to a remote server into a single request¹.

class `xmlrpc.client.MultiCall` (*server*)

Create an object used to boxcar method calls. *server* is the eventual target of the call. Calls can be made to the result object, but they will immediately return `None`, and only store the call name and parameters in the *MultiCall* object. Calling the object itself causes all stored calls to be transmitted as a single `system.multicall` request. The result of this call is a *generator*; iterating over this generator yields the individual results.

A usage example of this class follows. The server code:

```
from xmlrpc.server import SimpleXMLRPCServer

def add(x, y):
    return x + y

def subtract(x, y):
    return x - y

def multiply(x, y):
    return x * y

def divide(x, y):
    return x // y

# A simple server with simple arithmetic functions
server = SimpleXMLRPCServer(("localhost", 8000))
print("Listening on port 8000...")
server.register_multicall_functions()
server.register_function(add, 'add')
server.register_function(subtract, 'subtract')
server.register_function(multiply, 'multiply')
server.register_function(divide, 'divide')
server.serve_forever()
```

The client code for the preceding server:

```
import xmlrpc.client

proxy = xmlrpc.client.ServerProxy("http://localhost:8000/")
multicall = xmlrpc.client.MultiCall(proxy)
multicall.add(7, 3)
multicall.subtract(7, 3)
multicall.multiply(7, 3)
multicall.divide(7, 3)
result = multicall()

print("7+3=%d, 7-3=%d, 7*3=%d, 7//3=%d" % tuple(result))
```

21.26.7 Convenience Functions

`xmlrpc.client.dumps` (*params*, *methodname=None*, *methodresponse=None*, *encoding=None*, *allow_none=False*)

Convert *params* into an XML-RPC request. or into a response if *methodresponse* is true. *params* can be either a

¹ This approach has been first presented in a discussion on xmlrpc.com.

tuple of arguments or an instance of the `Fault` exception class. If `methodresponse` is true, only a single value can be returned, meaning that `params` must be of length 1. `encoding`, if supplied, is the encoding to use in the generated XML; the default is UTF-8. Python's `None` value cannot be used in standard XML-RPC; to allow using it via an extension, provide a true value for `allow_none`.

`xmlrpc.client.loads` (*data*, *use_datetime=False*, *use_builtin_types=False*)

Convert an XML-RPC request or response into Python objects, a (`params`, `methodname`). `params` is a tuple of argument; `methodname` is a string, or `None` if no method name is present in the packet. If the XML-RPC packet represents a fault condition, this function will raise a `Fault` exception. The `use_builtin_types` flag can be used to cause date/time values to be presented as `datetime.datetime` objects and binary data to be presented as `bytes` objects; this flag is false by default.

The obsolete `use_datetime` flag is similar to `use_builtin_types` but it applies only to date/time values.

在 3.3 版更改: The `use_builtin_types` flag was added.

21.26.8 Example of Client Usage

```
# simple test program (from the XML-RPC specification)
from xmlrpc.client import ServerProxy, Error

# server = ServerProxy("http://localhost:8000") # local server
with ServerProxy("http://betty.userland.com") as proxy:

    print(proxy)

    try:
        print(proxy.examples.getStateName(41))
    except Error as v:
        print("ERROR", v)
```

To access an XML-RPC server through a HTTP proxy, you need to define a custom transport. The following example shows how:

```
import http.client
import xmlrpc.client

class ProxiedTransport(xmlrpc.client.Transport):

    def set_proxy(self, host, port=None, headers=None):
        self.proxy = host, port
        self.proxy_headers = headers

    def make_connection(self, host):
        connection = http.client.HTTPConnection(*self.proxy)
        connection.set_tunnel(host, headers=self.proxy_headers)
        self._connection = host, connection
        return connection

transport = ProxiedTransport()
transport.set_proxy('proxy-server', 8080)
server = xmlrpc.client.ServerProxy('http://betty.userland.com', transport=transport)
print(server.examples.getStateName(41))
```

21.26.9 Example of Client and Server Usage

See *SimpleXMLRPCServer Example*.

21.27 xmlrpc.server —Basic XML-RPC servers

Source code: [Lib/xmlrpc/server.py](#)

The `xmlrpc.server` module provides a basic server framework for XML-RPC servers written in Python. Servers can either be free standing, using *SimpleXMLRPCServer*, or embedded in a CGI environment, using *CGIXMLRPCRequestHandler*.

警告: The `xmlrpc.server` module is not secure against maliciously constructed data. If you need to parse untrusted or unauthenticated data see *XML 漏洞*.

```
class xmlrpc.server.SimpleXMLRPCServer (addr, requestHandler=SimpleXMLRPCRequestHandler,
                                         logRequests=True,      allow_none=False,      en-
                                         coding=None,            bind_and_activate=True,
                                         use_builtin_types=False)
```

Create a new server instance. This class provides methods for registration of functions that can be called by the XML-RPC protocol. The `requestHandler` parameter should be a factory for request handler instances; it defaults to *SimpleXMLRPCRequestHandler*. The `addr` and `requestHandler` parameters are passed to the `socketserver.TCPServer` constructor. If `logRequests` is true (the default), requests will be logged; setting this parameter to false will turn off logging. The `allow_none` and `encoding` parameters are passed on to `xmlrpc.client` and control the XML-RPC responses that will be returned from the server. The `bind_and_activate` parameter controls whether `server_bind()` and `server_activate()` are called immediately by the constructor; it defaults to true. Setting it to false allows code to manipulate the `allow_reuse_address` class variable before the address is bound. The `use_builtin_types` parameter is passed to the `loads()` function and controls which types are processed when date/times values or binary data are received; it defaults to false.

在 3.3 版更改: The `use_builtin_types` flag was added.

```
class xmlrpc.server.CGIXMLRPCRequestHandler (allow_none=False,      encoding=None,
                                              use_builtin_types=False)
```

Create a new instance to handle XML-RPC requests in a CGI environment. The `allow_none` and `encoding` parameters are passed on to `xmlrpc.client` and control the XML-RPC responses that will be returned from the server. The `use_builtin_types` parameter is passed to the `loads()` function and controls which types are processed when date/times values or binary data are received; it defaults to false.

在 3.3 版更改: The `use_builtin_types` flag was added.

```
class xmlrpc.server.SimpleXMLRPCRequestHandler
```

Create a new request handler instance. This request handler supports POST requests and modifies logging so that the `logRequests` parameter to the *SimpleXMLRPCServer* constructor parameter is honored.

21.27.1 SimpleXMLRPCServer Objects

The *SimpleXMLRPCServer* class is based on `socketserver.TCPServer` and provides a means of creating simple, stand alone XML-RPC servers.

`SimpleXMLRPCServer.register_function(function, name=None)`

Register a function that can respond to XML-RPC requests. If *name* is given, it will be the method name associated with *function*, otherwise *function.__name__* will be used. *name* can be either a normal or Unicode string, and may contain characters not legal in Python identifiers, including the period character.

`SimpleXMLRPCServer.register_instance(instance, allow_dotted_names=False)`

Register an object which is used to expose method names which have not been registered using `register_function()`. If *instance* contains a `_dispatch()` method, it is called with the requested method name and the parameters from the request. Its API is `def _dispatch(self, method, params)` (note that *params* does not represent a variable argument list). If it calls an underlying function to perform its task, that function is called as `func(*params)`, expanding the parameter list. The return value from `_dispatch()` is returned to the client as the result. If *instance* does not have a `_dispatch()` method, it is searched for an attribute matching the name of the requested method.

If the optional *allow_dotted_names* argument is true and the instance does not have a `_dispatch()` method, then if the requested method name contains periods, each component of the method name is searched for individually, with the effect that a simple hierarchical search is performed. The value found from this search is then called with the parameters from the request, and the return value is passed back to the client.

警告： Enabling the *allow_dotted_names* option allows intruders to access your module's global variables and may allow intruders to execute arbitrary code on your machine. Only use this option on a secure, closed network.

`SimpleXMLRPCServer.register_introspection_functions()`

Registers the XML-RPC introspection functions `system.listMethods`, `system.methodHelp` and `system.methodSignature`.

`SimpleXMLRPCServer.register_multicall_functions()`

Registers the XML-RPC multicall function `system.multicall`.

`SimpleXMLRPCRequestHandler.rpc_paths`

An attribute value that must be a tuple listing valid path portions of the URL for receiving XML-RPC requests. Requests posted to other paths will result in a 404 “no such page” HTTP error. If this tuple is empty, all paths will be considered valid. The default value is `('/', '/RPC2')`.

SimpleXMLRPCServer Example

Server code:

```
from xmlrpc.server import SimpleXMLRPCServer
from xmlrpc.server import SimpleXMLRPCRequestHandler

# Restrict to a particular path.
class RequestHandler(SimpleXMLRPCRequestHandler):
    rpc_paths = ('/RPC2',)

# Create server
with SimpleXMLRPCServer(("localhost", 8000),
                        requestHandler=RequestHandler) as server:
    server.register_introspection_functions()

    # Register pow() function; this will use the value of
    # pow.__name__ as the name, which is just 'pow'.
    server.register_function(pow)
```

(下页继续)

(续上页)

```

# Register a function under a different name
def adder_function(x, y):
    return x + y
server.register_function(adder_function, 'add')

# Register an instance; all the methods of the instance are
# published as XML-RPC methods (in this case, just 'mul').
class MyFuncs:
    def mul(self, x, y):
        return x * y

server.register_instance(MyFuncs())

# Run the server's main loop
server.serve_forever()

```

The following client code will call the methods made available by the preceding server:

```

import xmlrpc.client

s = xmlrpc.client.ServerProxy('http://localhost:8000')
print(s.pow(2,3)) # Returns 2**3 = 8
print(s.add(2,3)) # Returns 5
print(s.mul(5,2)) # Returns 5*2 = 10

# Print list of available methods
print(s.system.listMethods())

```

The following example included in the `Lib/xmlrpc/server.py` module shows a server allowing dotted names and registering a multicall function.

警告: Enabling the `allow_dotted_names` option allows intruders to access your module's global variables and may allow intruders to execute arbitrary code on your machine. Only use this example only within a secure, closed network.

```

import datetime

class ExampleService:
    def getData(self):
        return '42'

    class currentTime:
        @staticmethod
        def getCurrentTime():
            return datetime.datetime.now()

with SimpleXMLRPCServer(("localhost", 8000)) as server:
    server.register_function(pow)
    server.register_function(lambda x,y: x+y, 'add')
    server.register_instance(ExampleService(), allow_dotted_names=True)
    server.register_multicall_functions()
    print('Serving XML-RPC on localhost port 8000')
    try:
        server.serve_forever()
    except KeyboardInterrupt:

```

(下页继续)

(续上页)

```
print("\nKeyboard interrupt received, exiting.")
sys.exit(0)
```

This ExampleService demo can be invoked from the command line:

```
python -m xmlrpc.server
```

The client that interacts with the above server is included in *Lib/xmlrpc/client.py*:

```
server = ServerProxy("http://localhost:8000")

try:
    print(server.currentTime.getCurrentTime())
except Error as v:
    print("ERROR", v)

multi = MultiCall(server)
multi.getData()
multi.pow(2,9)
multi.add(1,2)
try:
    for response in multi():
        print(response)
except Error as v:
    print("ERROR", v)
```

This client which interacts with the demo XMLRPC server can be invoked as:

```
python -m xmlrpc.client
```

21.27.2 CGIXMLRPCRequestHandler

The *CGIXMLRPCRequestHandler* class can be used to handle XML-RPC requests sent to Python CGI scripts.

CGIXMLRPCRequestHandler.register_function (*function*, *name=None*)

Register a function that can respond to XML-RPC requests. If *name* is given, it will be the method name associated with *function*, otherwise *function.__name__* will be used. *name* can be either a normal or Unicode string, and may contain characters not legal in Python identifiers, including the period character.

CGIXMLRPCRequestHandler.register_instance (*instance*)

Register an object which is used to expose method names which have not been registered using *register_function()*. If *instance* contains a *_dispatch()* method, it is called with the requested method name and the parameters from the request; the return value is returned to the client as the result. If *instance* does not have a *_dispatch()* method, it is searched for an attribute matching the name of the requested method; if the requested method name contains periods, each component of the method name is searched for individually, with the effect that a simple hierarchical search is performed. The value found from this search is then called with the parameters from the request, and the return value is passed back to the client.

CGIXMLRPCRequestHandler.register_introspection_functions ()

Register the XML-RPC introspection functions *system.listMethods*, *system.methodHelp* and *system.methodSignature*.

CGIXMLRPCRequestHandler.register_multicall_functions ()

Register the XML-RPC multicall function *system.multicall*.

`CGIXMLRPCRequestHandler.handle_request(request_text=None)`

Handle an XML-RPC request. If *request_text* is given, it should be the POST data provided by the HTTP server, otherwise the contents of stdin will be used.

示例:

```
class MyFuncs:
    def mul(self, x, y):
        return x * y

handler = CGIXMLRPCRequestHandler()
handler.register_function(pow)
handler.register_function(lambda x,y: x+y, 'add')
handler.register_introspection_functions()
handler.register_instance(MyFuncs())
handler.handle_request()
```

21.27.3 Documenting XMLRPC server

These classes extend the above classes to serve HTML documentation in response to HTTP GET requests. Servers can either be free standing, using [`DocXMLRPCServer`](#), or embedded in a CGI environment, using [`DocCGIXMLRPCRequestHandler`](#).

class `xmlrpc.server.DocXMLRPCServer` (*addr*, *requestHandler=DocXMLRPCRequestHandler*,
logRequests=True, *allow_none=False*, *encoding=None*,
bind_and_activate=True, *use_builtin_types=True*)

Create a new server instance. All parameters have the same meaning as for [`SimpleXMLRPCServer`](#); *requestHandler* defaults to [`DocXMLRPCRequestHandler`](#).

在 3.3 版更改: The *use_builtin_types* flag was added.

class `xmlrpc.server.DocCGIXMLRPCRequestHandler`

Create a new instance to handle XML-RPC requests in a CGI environment.

class `xmlrpc.server.DocXMLRPCRequestHandler`

Create a new request handler instance. This request handler supports XML-RPC POST requests, documentation GET requests, and modifies logging so that the *logRequests* parameter to the [`DocXMLRPCServer`](#) constructor parameter is honored.

21.27.4 DocXMLRPCServer Objects

The [`DocXMLRPCServer`](#) class is derived from [`SimpleXMLRPCServer`](#) and provides a means of creating self-documenting, stand alone XML-RPC servers. HTTP POST requests are handled as XML-RPC method calls. HTTP GET requests are handled by generating pydoc-style HTML documentation. This allows a server to provide its own web-based documentation.

`DocXMLRPCServer.set_server_title(server_title)`

Set the title used in the generated HTML documentation. This title will be used inside the HTML “title” element.

`DocXMLRPCServer.set_server_name(server_name)`

Set the name used in the generated HTML documentation. This name will appear at the top of the generated documentation inside a “h1” element.

`DocXMLRPCServer.set_server_documentation(server_documentation)`

Set the description used in the generated HTML documentation. This description will appear as a paragraph, below the server name, in the documentation.

21.27.5 DocCGIXMLRPCRequestHandler

The *DocCGIXMLRPCRequestHandler* class is derived from *CGIXMLRPCRequestHandler* and provides a means of creating self-documenting, XML-RPC CGI scripts. HTTP POST requests are handled as XML-RPC method calls. HTTP GET requests are handled by generating pydoc-style HTML documentation. This allows a server to provide its own web-based documentation.

`DocCGIXMLRPCRequestHandler.set_server_title(server_title)`

Set the title used in the generated HTML documentation. This title will be used inside the HTML “title” element.

`DocCGIXMLRPCRequestHandler.set_server_name(server_name)`

Set the name used in the generated HTML documentation. This name will appear at the top of the generated documentation inside a “h1” element.

`DocCGIXMLRPCRequestHandler.set_server_documentation(server_documentation)`

Set the description used in the generated HTML documentation. This description will appear as a paragraph, below the server name, in the documentation.

21.28 ipaddress — IPv4/IPv6 manipulation library

Source code: [Lib/ipaddress.py](#)

ipaddress provides the capabilities to create, manipulate and operate on IPv4 and IPv6 addresses and networks.

The functions and classes in this module make it straightforward to handle various tasks related to IP addresses, including checking whether or not two hosts are on the same subnet, iterating over all hosts in a particular subnet, checking whether or not a string represents a valid IP address or network definition, and so on.

This is the full module API reference—for an overview and introduction, see *ipaddress-howto*.

3.3 新版功能.

21.28.1 Convenience factory functions

The *ipaddress* module provides factory functions to conveniently create IP addresses, networks and interfaces:

`ipaddress.ip_address(address)`

Return an *IPv4Address* or *IPv6Address* object depending on the IP address passed as argument. Either IPv4 or IPv6 addresses may be supplied; integers less than 2^{32} will be considered to be IPv4 by default. A *ValueError* is raised if *address* does not represent a valid IPv4 or IPv6 address.

```
>>> ipaddress.ip_address('192.168.0.1')
IPv4Address('192.168.0.1')
>>> ipaddress.ip_address('2001:db8::')
IPv6Address('2001:db8::')
```

`ipaddress.ip_network(address, strict=True)`

Return an *IPv4Network* or *IPv6Network* object depending on the IP address passed as argument. *address* is a string or integer representing the IP network. Either IPv4 or IPv6 networks may be supplied; integers less than 2^{32} will be considered to be IPv4 by default. *strict* is passed to *IPv4Network* or *IPv6Network* constructor. A *ValueError* is raised if *address* does not represent a valid IPv4 or IPv6 address, or if the network has host bits set.

```
>>> ipaddress.ip_network('192.168.0.0/28')
IPv4Network('192.168.0.0/28')
```

`ipaddress.ip_interface(address)`

Return an *IPv4Interface* or *IPv6Interface* object depending on the IP address passed as argument. *address* is a string or integer representing the IP address. Either IPv4 or IPv6 addresses may be supplied; integers less than 2^{32} will be considered to be IPv4 by default. A *ValueError* is raised if *address* does not represent a valid IPv4 or IPv6 address.

One downside of these convenience functions is that the need to handle both IPv4 and IPv6 formats means that error messages provide minimal information on the precise error, as the functions don't know whether the IPv4 or IPv6 format was intended. More detailed error reporting can be obtained by calling the appropriate version specific class constructors directly.

21.28.2 IP Addresses

Address objects

The *IPv4Address* and *IPv6Address* objects share a lot of common attributes. Some attributes that are only meaningful for IPv6 addresses are also implemented by *IPv4Address* objects, in order to make it easier to write code that handles both IP versions correctly. Address objects are *hashable*, so they can be used as keys in dictionaries.

class `ipaddress.IPv4Address(address)`

Construct an IPv4 address. An *AddressValueError* is raised if *address* is not a valid IPv4 address.

The following constitutes a valid IPv4 address:

1. A string in decimal-dot notation, consisting of four decimal integers in the inclusive range 0–255, separated by dots (e.g. `192.168.0.1`). Each integer represents an octet (byte) in the address. Leading zeroes are tolerated only for values less than 8 (as there is no ambiguity between the decimal and octal interpretations of such strings).
2. An integer that fits into 32 bits.
3. An integer packed into a *bytes* object of length 4 (most significant octet first).

```
>>> ipaddress.IPv4Address('192.168.0.1')
IPv4Address('192.168.0.1')
>>> ipaddress.IPv4Address(3232235521)
IPv4Address('192.168.0.1')
>>> ipaddress.IPv4Address(b'\xC0\xA8\x00\x01')
IPv4Address('192.168.0.1')
```

version

The appropriate version number: 4 for IPv4, 6 for IPv6.

max_prefixlen

The total number of bits in the address representation for this version: 32 for IPv4, 128 for IPv6.

The prefix defines the number of leading bits in an address that are compared to determine whether or not an address is part of a network.

compressed

exploded

The string representation in dotted decimal notation. Leading zeroes are never included in the representation.

As IPv4 does not define a shorthand notation for addresses with octets set to zero, these two attributes are always the same as `str(addr)` for IPv4 addresses. Exposing these attributes makes it easier to write display code that can handle both IPv4 and IPv6 addresses.

packed

The binary representation of this address - a *bytes* object of the appropriate length (most significant octet first). This is 4 bytes for IPv4 and 16 bytes for IPv6.

reverse_pointer

The name of the reverse DNS PTR record for the IP address, e.g.:

[illegible]

This is the name that could be used for performing a PTR lookup, not the resolved hostname itself.

3.5 新版功能.

is multicast

True if the address is reserved for multicast use. See [RFC 3171](#) (for IPv4) or [RFC 2373](#) (for IPv6).

```
is_private
```

True if the address is allocated for private networks. See [iana-ipv4-special-registry](#) (for IPv4) or [iana-ipv6-special-registry](#) (for IPv6).

is_global

True if the address is allocated for public networks. See [iana-ipv4-special-registry](#) (for IPv4) or [iana-ipv6-special-registry](#) (for IPv6).

3.4 新版功能.

is_unspecified

True if the address is unspecified. See [RFC 5735](#) (for IPv4) or [RFC 2373](#) (for IPv6).

is_reserved

True if the address is otherwise IETF reserved.

is_loopback

True if this is a loopback address. See [RFC 3330](#) (for IPv4) or [RFC 2373](#) (for IPv6).

is link local

True if the address is reserved for link-local usage. See [RFC 3927](#).

```
class ipaddress.IPv6Address(address)
```

Construct an IPv6 address. An *AddressValueError* is raised if *address* is not a valid IPv6 address.

The following constitutes a valid IPv6 address:

1. A string consisting of eight groups of four hexadecimal digits, each group representing 16 bits. The groups are separated by colons. This describes an *exploded* (longhand) notation. The string can also be *compressed* (shorthand notation) by various means. See [RFC 4291](#) for details. For example, "0000:0000:0000:0000:0000:0abc:0007:0def" can be compressed to "::abc:7:def".
2. An integer that fits into 128 bits.
3. An integer packed into a *bytes* object of length 16, big-endian.

```
>>> ipaddress.IPv6Address('2001:db8::1000')
IPv6Address('2001:db8::1000')
```

compressed

The short form of the address representation, with leading zeroes in groups omitted and the longest sequence of groups consisting entirely of zeroes collapsed to a single empty group.

This is also the value returned by `str(addr)` for IPv6 addresses.

exploded

The long form of the address representation, with all leading zeroes and groups consisting entirely of zeroes included.

For the following attributes, see the corresponding documentation of the `IPv4Address` class:

packed

reverse_pointer

version

max_prefixlen

is_multicast

is_private

is_global

is_unspecified

is_reserved

is_loopback

is_link_local

3.4 新版功能: `is_global`

is_site_local

True if the address is reserved for site-local usage. Note that the site-local address space has been deprecated by [RFC 3879](#). Use `is_private` to test if this address is in the space of unique local addresses as defined by [RFC 4193](#).

ipv4_mapped

For addresses that appear to be IPv4 mapped addresses (starting with `::FFFF/96`), this property will report the embedded IPv4 address. For any other address, this property will be `None`.

sixtofour

For addresses that appear to be 6to4 addresses (starting with `2002::/16`) as defined by [RFC 3056](#), this property will report the embedded IPv4 address. For any other address, this property will be `None`.

teredo

For addresses that appear to be Teredo addresses (starting with `2001::/32`) as defined by [RFC 4380](#), this property will report the embedded `(server, client)` IP address pair. For any other address, this property will be `None`.

Conversion to Strings and Integers

To interoperate with networking interfaces such as the `socket` module, addresses must be converted to strings or integers. This is handled using the `str()` and `int()` builtin functions:

```
>>> str(ipaddress.IPv4Address('192.168.0.1'))
'192.168.0.1'
>>> int(ipaddress.IPv4Address('192.168.0.1'))
3232235521
```

(下页继续)

(续上页)

```
>>> str(ipaddress.IPv6Address('::1'))
 '::1'
>>> int(ipaddress.IPv6Address('::1'))
1
```

运算符

Address objects support some operators. Unless stated otherwise, operators can only be applied between compatible objects (i.e. IPv4 with IPv4, IPv6 with IPv6).

Comparison operators

Address objects can be compared with the usual set of comparison operators. Some examples:

```
>>> IPv4Address('127.0.0.2') > IPv4Address('127.0.0.1')
True
>>> IPv4Address('127.0.0.2') == IPv4Address('127.0.0.1')
False
>>> IPv4Address('127.0.0.2') != IPv4Address('127.0.0.1')
True
```

Arithmetic operators

Integers can be added to or subtracted from address objects. Some examples:

```
>>> IPv4Address('127.0.0.2') + 3
IPv4Address('127.0.0.5')
>>> IPv4Address('127.0.0.2') - 3
IPv4Address('126.255.255.255')
>>> IPv4Address('255.255.255.255') + 1
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
ipaddress.AddressValueError: 4294967296 (>= 2**32) is not permitted as an IPv4 address
```

21.28.3 IP Network definitions

The *IPv4Network* and *IPv6Network* objects provide a mechanism for defining and inspecting IP network definitions. A network definition consists of a *mask* and a *network address*, and as such defines a range of IP addresses that equal the network address when masked (binary AND) with the mask. For example, a network definition with the mask 255.255.255.0 and the network address 192.168.1.0 consists of IP addresses in the inclusive range 192.168.1.0 to 192.168.1.255.

Prefix, net mask and host mask

There are several equivalent ways to specify IP network masks. A *prefix* /<nbits> is a notation that denotes how many high-order bits are set in the network mask. A *net mask* is an IP address with some number of high-order bits set. Thus the prefix /24 is equivalent to the net mask 255.255.255.0 in IPv4, or ffff:ff00:: in IPv6. In addition, a *host mask* is the logical inverse of a *net mask*, and is sometimes used (for example in Cisco access control lists) to denote a network mask. The host mask equivalent to /24 in IPv4 is 0.0.0.255.

Network objects

All attributes implemented by address objects are implemented by network objects as well. In addition, network objects implement additional attributes. All of these are common between *IPv4Network* and *IPv6Network*, so to avoid duplication they are only documented for *IPv4Network*. Network objects are *hashable*, so they can be used as keys in dictionaries.

class `ipaddress.IPv4Network` (*address*, *strict=True*)

Construct an IPv4 network definition. *address* can be one of the following:

1. A string consisting of an IP address and an optional mask, separated by a slash (/). The IP address is the network address, and the mask can be either a single number, which means it's a *prefix*, or a string representation of an IPv4 address. If it's the latter, the mask is interpreted as a *net mask* if it starts with a non-zero field, or as a *host mask* if it starts with a zero field, with the single exception of an all-zero mask which is treated as a *net mask*. If no mask is provided, it's considered to be /32.

For example, the following *address* specifications are equivalent: 192.168.1.0/24, 192.168.1.0/255.255.255.0 and 192.168.1.0/0.0.0.255.

2. An integer that fits into 32 bits. This is equivalent to a single-address network, with the network address being *address* and the mask being /32.
3. An integer packed into a *bytes* object of length 4, big-endian. The interpretation is similar to an integer *address*.
4. A two-tuple of an address description and a netmask, where the address description is either a string, a 32-bits integer, a 4-bytes packed integer, or an existing IPv4Address object; and the netmask is either an integer representing the prefix length (e.g. 24) or a string representing the prefix mask (e.g. 255.255.255.0).

An *AddressValueError* is raised if *address* is not a valid IPv4 address. A *NetmaskValueError* is raised if the mask is not valid for an IPv4 address.

If *strict* is *True* and host bits are set in the supplied address, then *ValueError* is raised. Otherwise, the host bits are masked out to determine the appropriate network address.

Unless stated otherwise, all network methods accepting other network/address objects will raise *TypeError* if the argument's IP version is incompatible to *self*.

在 3.5 版更改: Added the two-tuple form for the *address* constructor parameter.

version

max_prefixlen

Refer to the corresponding attribute documentation in *IPv4Address*.

is_multicast

is_private

is_unspecified

is_reserved

is_loopback

is_link_local

These attributes are true for the network as a whole if they are true for both the network address and the broadcast address.

network_address

The network address for the network. The network address and the prefix length together uniquely define a network.

broadcast_address

The broadcast address for the network. Packets sent to the broadcast address should be received by every host on the network.

hostmask

The host mask, as an *IPv4Address* object.

netmask

The net mask, as an *IPv4Address* object.

with_prefixlen**compressed****exploded**

A string representation of the network, with the mask in prefix notation.

`with_prefixlen` and `compressed` are always the same as `str(network)`. `exploded` uses the exploded form the network address.

with_netmask

A string representation of the network, with the mask in net mask notation.

with_hostmask

A string representation of the network, with the mask in host mask notation.

num_addresses

The total number of addresses in the network.

prefixlen

Length of the network prefix, in bits.

hosts()

Returns an iterator over the usable hosts in the network. The usable hosts are all the IP addresses that belong to the network, except the network address itself and the network broadcast address. For networks with a mask length of 31, the network address and network broadcast address are also included in the result.

```
>>> list(ip_network('192.0.2.0/29').hosts()) #doctest: +NORMALIZE_WHITESPACE
[IPv4Address('192.0.2.1'), IPv4Address('192.0.2.2'),
 IPv4Address('192.0.2.3'), IPv4Address('192.0.2.4'),
 IPv4Address('192.0.2.5'), IPv4Address('192.0.2.6')]
>>> list(ip_network('192.0.2.0/31').hosts())
[IPv4Address('192.0.2.0'), IPv4Address('192.0.2.1')]
```

overlaps(*other*)

True if this network is partly or wholly contained in *other* or *other* is wholly contained in this network.

address_exclude(*network*)

Computes the network definitions resulting from removing the given *network* from this one. Returns an iterator of network objects. Raises *ValueError* if *network* is not completely contained in this network.

```
>>> n1 = ip_network('192.0.2.0/28')
>>> n2 = ip_network('192.0.2.1/32')
>>> list(n1.address_exclude(n2)) #doctest: +NORMALIZE_WHITESPACE
[IPv4Network('192.0.2.8/29'), IPv4Network('192.0.2.4/30'),
 IPv4Network('192.0.2.2/31'), IPv4Network('192.0.2.0/32')]
```

subnets(*prefixlen_diff*=1, *new_prefix*=None)

The subnets that join to make the current network definition, depending on the argument values. *prefixlen_diff* is the amount our prefix length should be increased by. *new_prefix* is the desired new prefix of the subnets;

it must be larger than our prefix. One and only one of *prefixlen_diff* and *new_prefix* must be set. Returns an iterator of network objects.

```
>>> list(ip_network('192.0.2.0/24').subnets())
[IPv4Network('192.0.2.0/25'), IPv4Network('192.0.2.128/25')]
>>> list(ip_network('192.0.2.0/24').subnets(prefixlen_diff=2)) #doctest:␣
↪+NORMALIZE_WHITESPACE
[IPv4Network('192.0.2.0/26'), IPv4Network('192.0.2.64/26'),
 IPv4Network('192.0.2.128/26'), IPv4Network('192.0.2.192/26')]
>>> list(ip_network('192.0.2.0/24').subnets(new_prefix=26)) #doctest:␣
↪+NORMALIZE_WHITESPACE
[IPv4Network('192.0.2.0/26'), IPv4Network('192.0.2.64/26'),
 IPv4Network('192.0.2.128/26'), IPv4Network('192.0.2.192/26')]
>>> list(ip_network('192.0.2.0/24').subnets(new_prefix=23))
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
    raise ValueError('new prefix must be longer')
ValueError: new prefix must be longer
>>> list(ip_network('192.0.2.0/24').subnets(new_prefix=25))
[IPv4Network('192.0.2.0/25'), IPv4Network('192.0.2.128/25')]
```

supernet (*prefixlen_diff*=1, *new_prefix*=None)

The supernet containing this network definition, depending on the argument values. *prefixlen_diff* is the amount our prefix length should be decreased by. *new_prefix* is the desired new prefix of the supernet; it must be smaller than our prefix. One and only one of *prefixlen_diff* and *new_prefix* must be set. Returns a single network object.

```
>>> ip_network('192.0.2.0/24').supernet()
IPv4Network('192.0.2.0/23')
>>> ip_network('192.0.2.0/24').supernet(prefixlen_diff=2)
IPv4Network('192.0.0.0/22')
>>> ip_network('192.0.2.0/24').supernet(new_prefix=20)
IPv4Network('192.0.0.0/20')
```

compare_networks (*other*)

Compare this network to *other*. In this comparison only the network addresses are considered; host bits aren't. Returns either -1, 0 or 1.

```
>>> ip_network('192.0.2.1/32').compare_networks(ip_network('192.0.2.2/32'))
-1
>>> ip_network('192.0.2.1/32').compare_networks(ip_network('192.0.2.0/32'))
1
>>> ip_network('192.0.2.1/32').compare_networks(ip_network('192.0.2.1/32'))
0
```

class `ipaddress.IPv6Network` (*address*, *strict*=True)

Construct an IPv6 network definition. *address* can be one of the following:

1. A string consisting of an IP address and an optional prefix length, separated by a slash (/). The IP address is the network address, and the prefix length must be a single number, the *prefix*. If no prefix length is provided, it's considered to be /128.

Note that currently expanded netmasks are not supported. That means 2001:db00::0/24 is a valid argument while 2001:db00::0/ffff:ff00:: not.

2. An integer that fits into 128 bits. This is equivalent to a single-address network, with the network address being *address* and the mask being /128.

3. An integer packed into a *bytes* object of length 16, big-endian. The interpretation is similar to an integer *address*.
4. A two-tuple of an address description and a netmask, where the address description is either a string, a 128-bits integer, a 16-bytes packed integer, or an existing *IPv6Address* object; and the netmask is an integer representing the prefix length.

An *AddressValueError* is raised if *address* is not a valid IPv6 address. A *NetmaskValueError* is raised if the mask is not valid for an IPv6 address.

If *strict* is *True* and host bits are set in the supplied address, then *ValueError* is raised. Otherwise, the host bits are masked out to determine the appropriate network address.

在 3.5 版更改: Added the two-tuple form for the *address* constructor parameter.

version

max_prefixlen

is_multicast

is_private

is_unspecified

is_reserved

is_loopback

is_link_local

network_address

broadcast_address

hostmask

netmask

with_prefixlen

compressed

exploded

with_netmask

with_hostmask

num_addresses

prefixlen

hosts()

Returns an iterator over the usable hosts in the network. The usable hosts are all the IP addresses that belong to the network, except the Subnet-Router anycast address. For networks with a mask length of 127, the Subnet-Router anycast address is also included in the result.

overlaps(*other*)

address_exclude(*network*)

subnets(*prefixlen_diff*=1, *new_prefix*=None)

supernet(*prefixlen_diff*=1, *new_prefix*=None)

compare_networks(*other*)

Refer to the corresponding attribute documentation in *IPv4Network*.

is_site_local

This attribute is true for the network as a whole if it is true for both the network address and the broadcast address.

运算符

Network objects support some operators. Unless stated otherwise, operators can only be applied between compatible objects (i.e. IPv4 with IPv4, IPv6 with IPv6).

Logical operators

Network objects can be compared with the usual set of logical operators. Network objects are ordered first by network address, then by net mask.

Iteration

Network objects can be iterated to list all the addresses belonging to the network. For iteration, *all* hosts are returned, including unusable hosts (for usable hosts, use the `hosts()` method). An example:

```
>>> for addr in IPv4Network('192.0.2.0/28'):
...     addr
...
IPv4Address('192.0.2.0')
IPv4Address('192.0.2.1')
IPv4Address('192.0.2.2')
IPv4Address('192.0.2.3')
IPv4Address('192.0.2.4')
IPv4Address('192.0.2.5')
IPv4Address('192.0.2.6')
IPv4Address('192.0.2.7')
IPv4Address('192.0.2.8')
IPv4Address('192.0.2.9')
IPv4Address('192.0.2.10')
IPv4Address('192.0.2.11')
IPv4Address('192.0.2.12')
IPv4Address('192.0.2.13')
IPv4Address('192.0.2.14')
IPv4Address('192.0.2.15')
```

Networks as containers of addresses

Network objects can act as containers of addresses. Some examples:

```
>>> IPv4Network('192.0.2.0/28')[0]
IPv4Address('192.0.2.0')
>>> IPv4Network('192.0.2.0/28')[15]
IPv4Address('192.0.2.15')
>>> IPv4Address('192.0.2.6') in IPv4Network('192.0.2.0/28')
True
>>> IPv4Address('192.0.3.6') in IPv4Network('192.0.2.0/28')
False
```

21.28.4 Interface objects

Interface objects are *hashable*, so they can be used as keys in dictionaries.

class `ipaddress.IPv4Interface(address)`

Construct an IPv4 interface. The meaning of *address* is as in the constructor of *IPv4Network*, except that arbitrary host addresses are always accepted.

IPv4Interface is a subclass of *IPv4Address*, so it inherits all the attributes from that class. In addition, the following attributes are available:

ip

The address (*IPv4Address*) without network information.

```
>>> interface = IPv4Interface('192.0.2.5/24')
>>> interface.ip
IPv4Address('192.0.2.5')
```

network

The network (*IPv4Network*) this interface belongs to.

```
>>> interface = IPv4Interface('192.0.2.5/24')
>>> interface.network
IPv4Network('192.0.2.0/24')
```

with_prefixlen

A string representation of the interface with the mask in prefix notation.

```
>>> interface = IPv4Interface('192.0.2.5/24')
>>> interface.with_prefixlen
'192.0.2.5/24'
```

with_netmask

A string representation of the interface with the network as a net mask.

```
>>> interface = IPv4Interface('192.0.2.5/24')
>>> interface.with_netmask
'192.0.2.5/255.255.255.0'
```

with_hostmask

A string representation of the interface with the network as a host mask.

```
>>> interface = IPv4Interface('192.0.2.5/24')
>>> interface.with_hostmask
'192.0.2.5/0.0.0.255'
```

class `ipaddress.IPv6Interface(address)`

Construct an IPv6 interface. The meaning of *address* is as in the constructor of *IPv6Network*, except that arbitrary host addresses are always accepted.

IPv6Interface is a subclass of *IPv6Address*, so it inherits all the attributes from that class. In addition, the following attributes are available:

ip

network

with_prefixlen

with_netmask

with_hostmask

Refer to the corresponding attribute documentation in *IPv4Interface*.

运算符

Interface objects support some operators. Unless stated otherwise, operators can only be applied between compatible objects (i.e. IPv4 with IPv4, IPv6 with IPv6).

Logical operators

Interface objects can be compared with the usual set of logical operators.

For equality comparison (`==` and `!=`), both the IP address and network must be the same for the objects to be equal. An interface will not compare equal to any address or network object.

For ordering (`<`, `>`, etc) the rules are different. Interface and address objects with the same IP version can be compared, and the address objects will always sort before the interface objects. Two interface objects are first compared by their networks and, if those are the same, then by their IP addresses.

21.28.5 Other Module Level Functions

The module also provides the following module level functions:

ipaddress.v4_int_to_packed(*address*)

Represent an address as 4 packed bytes in network (big-endian) order. *address* is an integer representation of an IPv4 IP address. A *ValueError* is raised if the integer is negative or too large to be an IPv4 IP address.

```
>>> ipaddress.ip_address(3221225985)
IPv4Address('192.0.2.1')
>>> ipaddress.v4_int_to_packed(3221225985)
b'\xc0\x00\x02\x01'
```

ipaddress.v6_int_to_packed(*address*)

Represent an address as 16 packed bytes in network (big-endian) order. *address* is an integer representation of an IPv6 IP address. A *ValueError* is raised if the integer is negative or too large to be an IPv6 IP address.

ipaddress.summarize_address_range(*first*, *last*)

Return an iterator of the summarized network range given the first and last IP addresses. *first* is the first *IPv4Address* or *IPv6Address* in the range and *last* is the last *IPv4Address* or *IPv6Address* in the range. A *TypeError* is raised if *first* or *last* are not IP addresses or are not of the same version. A *ValueError* is raised if *last* is not greater than *first* or if *first* address version is not 4 or 6.

```
>>> [ipaddr for ipaddr in ipaddress.summarize_address_range(
...     ipaddress.IPv4Address('192.0.2.0'),
...     ipaddress.IPv4Address('192.0.2.130'))]
[IPv4Network('192.0.2.0/25'), IPv4Network('192.0.2.128/31'), IPv4Network('192.0.2.
↪130/32')]
```

ipaddress.collapse_addresses(*addresses*)

Return an iterator of the collapsed *IPv4Network* or *IPv6Network* objects. *addresses* is an iterator of *IPv4Network* or *IPv6Network* objects. A *TypeError* is raised if *addresses* contains mixed version objects.

```
>>> [ipaddr for ipaddr in
... ipaddress.collapse_addresses([ipaddress.IPv4Network('192.0.2.0/25'),
... ipaddress.IPv4Network('192.0.2.128/25')])]
[IPv4Network('192.0.2.0/24')]
```

`ipaddress.get_mixed_type_key(obj)`

Return a key suitable for sorting between networks and addresses. Address and Network objects are not sortable by default; they're fundamentally different, so the expression:

```
IPv4Address('192.0.2.0') <= IPv4Network('192.0.2.0/24')
```

doesn't make sense. There are some times however, where you may wish to have *ipaddress* sort these anyway. If you need to do this, you can use this function as the *key* argument to *sorted()*.

obj is either a network or address object.

21.28.6 Custom Exceptions

To support more specific error reporting from class constructors, the module defines the following exceptions:

exception `ipaddress.AddressValueError` (*ValueError*)

Any value error related to the address.

exception `ipaddress.NetmaskValueError` (*ValueError*)

Any value error related to the net mask.

本章描述的模块实现了主要用于多媒体应用的各种算法或接口。它们可在安装时自行决定。这是一个概述：

22.1 audioop — Manipulate raw audio data

The `audioop` module contains some useful operations on sound fragments. It operates on sound fragments consisting of signed integer samples 8, 16, 24 or 32 bits wide, stored in *bytes-like objects*. All scalar items are integers, unless specified otherwise.

在 3.4 版更改: Support for 24-bit samples was added. All functions now accept any *bytes-like object*. String input now results in an immediate error.

This module provides support for a-LAW, u-LAW and Intel/DVI ADPCM encodings.

A few of the more complicated operations only take 16-bit samples, otherwise the sample size (in bytes) is always a parameter of the operation.

The module defines the following variables and functions:

exception `audioop.error`

This exception is raised on all errors, such as unknown number of bytes per sample, etc.

`audioop.add(fragment1, fragment2, width)`

Return a fragment which is the addition of the two samples passed as parameters. *width* is the sample width in bytes, either 1, 2, 3 or 4. Both fragments should have the same length. Samples are truncated in case of overflow.

`audioop.adpcm2lin(adpcmfragment, width, state)`

Decode an Intel/DVI ADPCM coded fragment to a linear fragment. See the description of `lin2adpcm()` for details on ADPCM coding. Return a tuple (*sample*, *newstate*) where the sample has the width specified in *width*.

`audioop.alaw2lin(fragment, width)`

Convert sound fragments in a-LAW encoding to linearly encoded sound fragments. a-LAW encoding always uses 8 bits samples, so *width* refers only to the sample width of the output fragment here.

`audioop.avg(fragment, width)`

Return the average over all samples in the fragment.

`audioop.avgpp(fragment, width)`

Return the average peak-peak value over all samples in the fragment. No filtering is done, so the usefulness of this routine is questionable.

`audioop.bias(fragment, width, bias)`

Return a fragment that is the original fragment with a bias added to each sample. Samples wrap around in case of overflow.

`audioop.byteswap(fragment, width)`

“Byteswap” all samples in a fragment and returns the modified fragment. Converts big-endian samples to little-endian and vice versa.

3.4 新版功能.

`audioop.cross(fragment, width)`

Return the number of zero crossings in the fragment passed as an argument.

`audioop.findfactor(fragment, reference)`

Return a factor F such that `rms(add(fragment, mul(reference, -F)))` is minimal, i.e., return the factor with which you should multiply *reference* to make it match as well as possible to *fragment*. The fragments should both contain 2-byte samples.

The time taken by this routine is proportional to `len(fragment)`.

`audioop.findfit(fragment, reference)`

Try to match *reference* as well as possible to a portion of *fragment* (which should be the longer fragment). This is (conceptually) done by taking slices out of *fragment*, using `findfactor()` to compute the best match, and minimizing the result. The fragments should both contain 2-byte samples. Return a tuple (*offset*, *factor*) where *offset* is the (integer) offset into *fragment* where the optimal match started and *factor* is the (floating-point) factor as per `findfactor()`.

`audioop.findmax(fragment, length)`

Search *fragment* for a slice of length *length* samples (not bytes!) with maximum energy, i.e., return *i* for which `rms(fragment[i*2:(i+length)*2])` is maximal. The fragments should both contain 2-byte samples.

The routine takes time proportional to `len(fragment)`.

`audioop.getsample(fragment, width, index)`

Return the value of sample *index* from the fragment.

`audioop.lin2adpcm(fragment, width, state)`

Convert samples to 4 bit Intel/DVI ADPCM encoding. ADPCM coding is an adaptive coding scheme, whereby each 4 bit number is the difference between one sample and the next, divided by a (varying) step. The Intel/DVI ADPCM algorithm has been selected for use by the IMA, so it may well become a standard.

state is a tuple containing the state of the coder. The coder returns a tuple (*adpcmfrag*, *newstate*), and the *newstate* should be passed to the next call of `lin2adpcm()`. In the initial call, `None` can be passed as the state. *adpcmfrag* is the ADPCM coded fragment packed 2 4-bit values per byte.

`audioop.lin2alaw(fragment, width)`

Convert samples in the audio fragment to a-LAW encoding and return this as a bytes object. a-LAW is an audio encoding format whereby you get a dynamic range of about 13 bits using only 8 bit samples. It is used by the Sun audio hardware, among others.

`audioop.lin2lin(fragment, width, newwidth)`

Convert samples between 1-, 2-, 3- and 4-byte formats.

注解: In some audio formats, such as .WAV files, 16, 24 and 32 bit samples are signed, but 8 bit samples are unsigned. So when converting to 8 bit wide samples for these formats, you need to also add 128 to the result:

```
new_frames = audioop.lin2lin(frames, old_width, 1)
new_frames = audioop.bias(new_frames, 1, 128)
```

The same, in reverse, has to be applied when converting from 8 to 16, 24 or 32 bit width samples.

`audioop.lin2ulaw(fragment, width)`

Convert samples in the audio fragment to u-LAW encoding and return this as a bytes object. u-LAW is an audio encoding format whereby you get a dynamic range of about 14 bits using only 8 bit samples. It is used by the Sun audio hardware, among others.

`audioop.max(fragment, width)`

Return the maximum of the *absolute value* of all samples in a fragment.

`audioop.maxpp(fragment, width)`

Return the maximum peak-peak value in the sound fragment.

`audioop.minmax(fragment, width)`

Return a tuple consisting of the minimum and maximum values of all samples in the sound fragment.

`audioop.mul(fragment, width, factor)`

Return a fragment that has all samples in the original fragment multiplied by the floating-point value *factor*. Samples are truncated in case of overflow.

`audioop.ratecv(fragment, width, nchannels, inrate, outrate, state[, weightA[, weightB]])`

Convert the frame rate of the input fragment.

state is a tuple containing the state of the converter. The converter returns a tuple (*newfragment*, *newstate*), and *newstate* should be passed to the next call of `ratecv()`. The initial call should pass *None* as the state.

The *weightA* and *weightB* arguments are parameters for a simple digital filter and default to 1 and 0 respectively.

`audioop.reverse(fragment, width)`

Reverse the samples in a fragment and returns the modified fragment.

`audioop.rms(fragment, width)`

Return the root-mean-square of the fragment, i.e. $\sqrt{\text{sum}(S_i^2)/n}$.

This is a measure of the power in an audio signal.

`audioop.tomono(fragment, width, lfactor, rfactor)`

Convert a stereo fragment to a mono fragment. The left channel is multiplied by *lfactor* and the right channel by *rfactor* before adding the two channels to give a mono signal.

`audioop.tostereo(fragment, width, lfactor, rfactor)`

Generate a stereo fragment from a mono fragment. Each pair of samples in the stereo fragment are computed from the mono sample, whereby left channel samples are multiplied by *lfactor* and right channel samples by *rfactor*.

`audioop.ulaw2lin(fragment, width)`

Convert sound fragments in u-LAW encoding to linearly encoded sound fragments. u-LAW encoding always uses 8 bits samples, so *width* refers only to the sample width of the output fragment here.

Note that operations such as `mul()` or `max()` make no distinction between mono and stereo fragments, i.e. all samples are treated equal. If this is a problem the stereo fragment should be split into two mono fragments first and recombined later. Here is an example of how to do that:

```
def mul_stereo(sample, width, lfactor, rfactor):
    lsample = audioop.tomono(sample, width, 1, 0)
    rsample = audioop.tomono(sample, width, 0, 1)
    lsample = audioop.mul(lsample, width, lfactor)
    rsample = audioop.mul(rsample, width, rfactor)
    lsample = audioop.tostereo(lsample, width, 1, 0)
    rsample = audioop.tostereo(rsample, width, 0, 1)
    return audioop.add(lsample, rsample, width)
```

If you use the ADPCM coder to build network packets and you want your protocol to be stateless (i.e. to be able to tolerate packet loss) you should not only transmit the data but also the state. Note that you should send the *initial* state (the one you passed to `lin2adpcm()`) along to the decoder, not the final state (as returned by the coder). If you want to use `struct.Struct` to store the state in binary you can code the first element (the predicted value) in 16 bits and the second (the delta index) in 8.

The ADPCM coders have never been tried against other ADPCM coders, only against themselves. It could well be that I misinterpreted the standards in which case they will not be interoperable with the respective standards.

The `find*()` routines might look a bit funny at first sight. They are primarily meant to do echo cancellation. A reasonably fast way to do this is to pick the most energetic piece of the output sample, locate that in the input sample and subtract the whole output sample from the input sample:

```
def echocancel(outputdata, inputdata):
    pos = audioop.findmax(outputdata, 800)      # one tenth second
    out_test = outputdata[pos*2:]
    in_test = inputdata[pos*2:]
    ipos, factor = audioop.findfit(in_test, out_test)
    # Optional (for better cancellation):
    # factor = audioop.findfactor(in_test[ipos*2:ipos*2+len(out_test)],
    #                             out_test)
    prefill = '\0'*(pos+ipos)*2
    postfill = '\0'*(len(inputdata)-len(prefill)-len(outputdata))
    outputdata = prefill + audioop.mul(outputdata, 2, -factor) + postfill
    return audioop.add(inputdata, outputdata, 2)
```

22.2 aifc —Read and write AIFF and AIFC files

Source code: [Lib/aifc.py](#)

This module provides support for reading and writing AIFF and AIFF-C files. AIFF is Audio Interchange File Format, a format for storing digital audio samples in a file. AIFF-C is a newer version of the format that includes the ability to compress the audio data.

Audio files have a number of parameters that describe the audio data. The sampling rate or frame rate is the number of times per second the sound is sampled. The number of channels indicate if the audio is mono, stereo, or quadro. Each frame consists of one sample per channel. The sample size is the size in bytes of each sample. Thus a frame consists of `nchannels * samplesize` bytes, and a second's worth of audio consists of `nchannels * samplesize * framerate` bytes.

For example, CD quality audio has a sample size of two bytes (16 bits), uses two channels (stereo) and has a frame rate of 44,100 frames/second. This gives a frame size of 4 bytes (2*2), and a second's worth occupies 2*2*44100 bytes (176,400 bytes).

Module `aifc` defines the following function:

`aifc.open(file, mode=None)`

Open an AIFF or AIFF-C file and return an object instance with methods that are described below. The argument *file* is either a string naming a file or a *file object*. *mode* must be 'r' or 'rb' when the file must be opened for reading, or 'w' or 'wb' when the file must be opened for writing. If omitted, *file.mode* is used if it exists, otherwise 'rb' is used. When used for writing, the file object should be seekable, unless you know ahead of time how many samples you are going to write in total and use `writeframesraw()` and `setnframes()`. The `open()` function may be used in a `with` statement. When the `with` block completes, the `close()` method is called.

在 3.4 版更改: 支持了 `with` 语句。

Objects returned by `open()` when a file is opened for reading have the following methods:

`aifc.getnchannels()`

Return the number of audio channels (1 for mono, 2 for stereo).

`aifc.getsampwidth()`

Return the size in bytes of individual samples.

`aifc.getframerate()`

Return the sampling rate (number of audio frames per second).

`aifc.getnframes()`

Return the number of audio frames in the file.

`aifc.getcomptype()`

Return a bytes array of length 4 describing the type of compression used in the audio file. For AIFF files, the returned value is `b'NONE'`.

`aifc.getcompname()`

Return a bytes array convertible to a human-readable description of the type of compression used in the audio file. For AIFF files, the returned value is `b'not compressed'`.

`aifc.getparams()`

返回一个 *namedtuple()* (`nchannels`, `sampwidth`, `framerate`, `nframes`, `comptype`, `compname`), 与 `get*()` 方法的输出相同。

`aifc.getmarkers()`

Return a list of markers in the audio file. A marker consists of a tuple of three elements. The first is the mark ID (an integer), the second is the mark position in frames from the beginning of the data (an integer), the third is the name of the mark (a string).

`aifc.getmark(id)`

Return the tuple as described in `getmarkers()` for the mark with the given *id*.

`aifc.readframes(nframes)`

Read and return the next *nframes* frames from the audio file. The returned data is a string containing for each frame the uncompressed samples of all channels.

`aifc.rewind()`

Rewind the read pointer. The next `readframes()` will start from the beginning.

`aifc.setpos(pos)`

Seek to the specified frame number.

`aifc.tell()`

Return the current frame number.

`aifc.close()`

Close the AIFF file. After calling this method, the object can no longer be used.

Objects returned by `open()` when a file is opened for writing have all the above methods, except for `readframes()` and `setpos()`. In addition the following methods exist. The `get*()` methods can only be called after the corresponding `set*()` methods have been called. Before the first `writeframes()` or `writeframesraw()`, all parameters except for the number of frames must be filled in.

`aifc.aiff()`
Create an AIFF file. The default is that an AIFF-C file is created, unless the name of the file ends in `'.aiff'` in which case the default is an AIFF file.

`aifc.aifc()`
Create an AIFF-C file. The default is that an AIFF-C file is created, unless the name of the file ends in `'.aiff'` in which case the default is an AIFF file.

`aifc.setnchannels(nchannels)`
Specify the number of channels in the audio file.

`aifc.setsampwidth(width)`
Specify the size in bytes of audio samples.

`aifc.setframerate(rate)`
Specify the sampling frequency in frames per second.

`aifc.setnframes(nframes)`
Specify the number of frames that are to be written to the audio file. If this parameter is not set, or not set correctly, the file needs to support seeking.

`aifc.setcomptype(type, name)`
Specify the compression type. If not specified, the audio data will not be compressed. In AIFF files, compression is not possible. The name parameter should be a human-readable description of the compression type as a bytes array, the type parameter should be a bytes array of length 4. Currently the following compression types are supported: `b'NONE'`, `b'ULAW'`, `b'ALAW'`, `b'G722'`.

`aifc.setparams(nchannels, sampwidth, framerate, comptype, compname)`
Set all the above parameters at once. The argument is a tuple consisting of the various parameters. This means that it is possible to use the result of a `getparams()` call as argument to `setparams()`.

`aifc.setmark(id, pos, name)`
Add a mark with the given id (larger than 0), and the given name at the given position. This method can be called at any time before `close()`.

`aifc.tell()`
Return the current write position in the output file. Useful in combination with `setmark()`.

`aifc.writeframes(data)`
Write data to the output file. This method can only be called after the audio file parameters have been set.
在 3.4 版更改: Any *bytes-like object* is now accepted.

`aifc.writeframesraw(data)`
Like `writeframes()`, except that the header of the audio file is not updated.
在 3.4 版更改: Any *bytes-like object* is now accepted.

`aifc.close()`
Close the AIFF file. The header of the file is updated to reflect the actual size of the audio data. After calling this method, the object can no longer be used.

22.3 sunau —读写 Sun AU 文件

源代码: [Lib/sunau.py](#)

`sunau` 模拟提供了一个处理 Sun AU 声音格式的便利接口。请注意此模块与 `aifc` 和 `wave` 是兼容接口的。

An audio file consists of a header followed by the data. The fields of the header are:

Field	内容
magic word	The four bytes .snd.
header size	Size of the header, including info, in bytes.
data size	Physical size of the data, in bytes.
encoding	Indicates how the audio samples are encoded.
sample rate	The sampling rate.
# of channels	The number of channels in the samples.
info	ASCII string giving a description of the audio file (padded with null bytes).

Apart from the info field, all header fields are 4 bytes in size. They are all 32-bit unsigned integers encoded in big-endian byte order.

The `sunau` module defines the following functions:

`sunau.open(file, mode)`
If `file` is a string, open the file by that name, otherwise treat it as a seekable file-like object. `mode` can be any of

- 'r' 只读模式。
- 'w' 只写模式。

Note that it does not allow read/write files.

A `mode` of 'r' returns an `AU_read` object, while a `mode` of 'w' or 'wb' returns an `AU_write` object.

`sunau.openfp(file, mode)`
同 `open()`，用于向后兼容。

The `sunau` module defines the following exception:

exception `sunau.Error`
An error raised when something is impossible because of Sun AU specs or implementation deficiency.

The `sunau` module defines the following data items:

`sunau.AUDIO_FILE_MAGIC`
An integer every valid Sun AU file begins with, stored in big-endian form. This is the string `.snd` interpreted as an integer.

`sunau.AUDIO_FILE_ENCODING_MULAW_8`
`sunau.AUDIO_FILE_ENCODING_LINEAR_8`
`sunau.AUDIO_FILE_ENCODING_LINEAR_16`
`sunau.AUDIO_FILE_ENCODING_LINEAR_24`
`sunau.AUDIO_FILE_ENCODING_LINEAR_32`
`sunau.AUDIO_FILE_ENCODING_ALAW_8`
Values of the encoding field from the AU header which are supported by this module.

`sunau.AUDIO_FILE_ENCODING_FLOAT`
`sunau.AUDIO_FILE_ENCODING_DOUBLE`
`sunau.AUDIO_FILE_ENCODING_ADPCM_G721`
`sunau.AUDIO_FILE_ENCODING_ADPCM_G722`
`sunau.AUDIO_FILE_ENCODING_ADPCM_G723_3`
`sunau.AUDIO_FILE_ENCODING_ADPCM_G723_5`
Additional known values of the encoding field from the AU header, but which are not supported by this module.

22.3.1 AU_read Objects

AU_read objects, as returned by `open()` above, have the following methods:

`AU_read.close()`

Close the stream, and make the instance unusable. (This is called automatically on deletion.)

`AU_read.getnchannels()`

Returns number of audio channels (1 for mono, 2 for stereo).

`AU_read.getsampwidth()`

返回采样字节长度。

`AU_read.getframerate()`

返回采样频率。

`AU_read.getnframes()`

返回音频总帧数。

`AU_read.getcomptype()`

Returns compression type. Supported compression types are 'ULAW', 'ALAW' and 'NONE'.

`AU_read.getcompname()`

Human-readable version of `getcomptype()`. The supported types have the respective names 'CCITT G.711 u-law', 'CCITT G.711 A-law' and 'not compressed'.

`AU_read.getparams()`

返回一个 `namedtuple()` (nchannels, sampwidth, framerate, nframes, comptype, compname), 与 `get*()` 方法的输出相同。

`AU_read.readframes(n)`

Reads and returns at most *n* frames of audio, as a `bytes` object. The data will be returned in linear format. If the original data is in u-LAW format, it will be converted.

`AU_read.rewind()`

设置当前文件指针位置。

以下两个方法都使用指针，具体实现由其底层决定。

`AU_read.setpos(pos)`

Set the file pointer to the specified position. Only values returned from `tell()` should be used for *pos*.

`AU_read.tell()`

Return current file pointer position. Note that the returned value has nothing to do with the actual position in the file.

The following two functions are defined for compatibility with the `aifc`, and don't do anything interesting.

`AU_read.getmarkers()`

返回 None。

`AU_read.getmark(id)`

引发错误异常。

22.3.2 AU_write Objects

AU_write objects, as returned by `open()` above, have the following methods:

`AU_write.setnchannels(n)`

设置声道数。

`AU_write.setsampwidth(n)`

Set the sample width (in bytes.)

在 3.4 版更改: Added support for 24-bit samples.

`AU_write.setframerate(n)`

Set the frame rate.

`AU_write.setnframes(n)`

Set the number of frames. This can be later changed, when and if more frames are written.

`AU_write.setcomptype(type, name)`

Set the compression type and description. Only 'NONE' and 'ULAW' are supported on output.

`AU_write.setparams(tuple)`

The *tuple* should be (nchannels, sampwidth, framerate, nframes, comptype, compname), with values valid for the `set*()` methods. Set all parameters.

`AU_write.tell()`

Return current position in the file, with the same disclaimer for the `AU_read.tell()` and `AU_read.setpos()` methods.

`AU_write.writeframesraw(data)`

写入音频数据但不更新 *nframes*。

在 3.4 版更改: Any *bytes-like object* is now accepted.

`AU_write.writeframes(data)`

写入音频数据并更新 *nframes*。

在 3.4 版更改: Any *bytes-like object* is now accepted.

`AU_write.close()`

Make sure *nframes* is correct, and close the file.

This method is called upon deletion.

Note that it is invalid to set any parameters after calling `writeframes()` or `writeframesraw()`.

22.4 wave —读写 WAV 格式文件

源代码: [Lib/wave.py](#)

`wave` 模块提供了一个处理 WAV 声音格式的便利接口。它不支持压缩/解压, 但是支持单声道/立体声。

`wave` 模块定义了以下函数和异常:

`wave.open(file, mode=None)`

如果 *file* 是一个字符串, 打开对应文件名的文件。否则就把它作为文件类对象来处理。*mode* 可以为以下值:

'rb' 只读模式。

'wb' 只写模式。

注意不支持同时读写 WAV 文件。

mode 设为 'rb' 时返回一个 `Wave_read` 对象, 而 *mode* 设为 'wb' 时返回一个 `Wave_write` 对象。如果省略 *mode* 并指定 *file* 来传入一个文件类对象, 则 `file.mode` 会被用作 *mode* 的默认值。

如果操作的是文件对象，当使用 `wave` 对象的 `close()` 方法时，并不会真正关闭文件对象，这需要调用者负责来关闭文件对象。

The `open()` function may be used in a `with` statement. When the `with` block completes, the `Wave_read.close()` or `Wave_write.close()` method is called.

在 3.4 版更改: 添加了对不可搜索文件的支持。

`wave.openfp(file, mode)`

同 `open()`，用于向后兼容。

exception `wave.Error`

当不符合 WAV 格式或无法操作时引发的错误。

22.4.1 Wave_read 对象

由 `open()` 返回的 `Wave_read` 对象，有以下几种方法：

`Wave_read.close()`

关闭 `wave` 打开的数据流并使对象不可用。当对象销毁时会自动调用。

`Wave_read.getnchannels()`

返回声道数量（1 为单声道，2 为立体声）

`Wave_read.getsampwidth()`

返回采样字节长度。

`Wave_read.getframerate()`

返回采样频率。

`Wave_read.getnframes()`

返回音频总帧数。

`Wave_read.getcomptype()`

返回压缩类型（只支持 'NONE' 类型）

`Wave_read.getcompname()`

`getcomptype()` 的通俗版本。使用 'not compressed' 代替 'NONE'。

`Wave_read.getparams()`

返回一个 `namedtuple()` (`nchannels`, `sampwidth`, `framerate`, `nframes`, `comptype`, `compname`)，与 `get*()` 方法的输出相同。

`Wave_read.readframes(n)`

读取并返回以 `bytes` 对象表示的最多 `n` 帧音频。

`Wave_read.rewind()`

设置当前文件指针位置。

后面两个方法是为了和 `aifc` 保持兼容，实际不做任何事情。

`Wave_read.getmarkers()`

返回 `None`。

`Wave_read.getmark(id)`

引发错误异常。

以下两个方法都使用指针，具体实现由其底层决定。

`Wave_read.setpos(pos)`

设置文件指针到指定位置。

`Wave_read.tell()`
 返回当前文件指针位置。

22.4.2 Wave_write 对象

For seekable output streams, the `wave` header will automatically be updated to reflect the number of frames actually written. For unseekable streams, the `nframes` value must be accurate when the first frame data is written. An accurate `nframes` value can be achieved either by calling `setnframes()` or `setparams()` with the number of frames that will be written before `close()` is called and then using `writeframesraw()` to write the frame data, or by calling `writeframes()` with all of the frame data to be written. In the latter case `writeframes()` will calculate the number of frames in the data and set `nframes` accordingly before writing the frame data.

由 `open()` 返回的 `Wave_write` 对象, 有以下几种方法:

在 3.4 版更改: 添加了对不可搜索文件的支持。

`Wave_write.close()`
 Make sure `nframes` is correct, and close the file if it was opened by `wave`. This method is called upon object collection. It will raise an exception if the output stream is not seekable and `nframes` does not match the number of frames actually written.

`Wave_write.setnchannels(n)`
 设置声道数。

`Wave_write.setsampwidth(n)`
 设置采样字节长度为 `n`。

`Wave_write.setframerate(n)`
 设置采样频率为 `n`。

在 3.2 版更改: A non-integral input to this method is rounded to the nearest integer.

`Wave_write.setnframes(n)`
 Set the number of frames to `n`. This will be changed later if the number of frames actually written is different (this update attempt will raise an error if the output stream is not seekable).

`Wave_write.setcomptype(type, name)`
 设置压缩格式。目前只支持 `NONE` 即无压缩格式。

`Wave_write.setparams(tuple)`
`tuple` 应该是 `(nchannels, sampwidth, framerate, nframes, comptype, compname)`, 每项的值应可用于 `set*()` 方法。设置所有形参。

`Wave_write.tell()`
 返回当前文件指针, 其指针含义和 `Wave_read.tell()` 以及 `Wave_read.setpos()` 是一致的。

`Wave_write.writeframesraw(data)`
 写入音频数据但不更新 `nframes`。

在 3.4 版更改: Any *bytes-like object* is now accepted.

`Wave_write.writeframes(data)`
 Write audio frames and make sure `nframes` is correct. It will raise an error if the output stream is not seekable and the total number of frames that have been written after `data` has been written does not match the previously set value for `nframes`.

在 3.4 版更改: Any *bytes-like object* is now accepted.

注意在调用 `writeframes()` 或 `writeframesraw()` 之后再设置任何格式参数是无效的, 而且任何这样的尝试将引发 `wave.Error`。

22.5 chunk — Read IFF chunked data

Source code: [Lib/chunk.py](#)

This module provides an interface for reading files that use EA IFF 85 chunks.¹ This format is used in at least the Audio Interchange File Format (AIFF/AIFF-C) and the Real Media File Format (RMFF). The WAVE audio file format is closely related and can also be read using this module.

一个 chunk 具有以下结构:

偏移	长度	内容
0	4	Chunk ID
4	4	Size of chunk in big-endian byte order, not including the header
8	n	Data bytes, where n is the size given in the preceding field
$8 + n$	0 或 1	Pad byte needed if n is odd and chunk alignment is used

The ID is a 4-byte string which identifies the type of chunk.

The size field (a 32-bit value, encoded using big-endian byte order) gives the size of the chunk data, not including the 8-byte header.

Usually an IFF-type file consists of one or more chunks. The proposed usage of the `Chunk` class defined here is to instantiate an instance at the start of each chunk and read from the instance until it reaches the end, after which a new instance can be instantiated. At the end of the file, creating a new instance will fail with an `EOFError` exception.

class `chunk.Chunk` (*file*, *align=True*, *bigendian=True*, *inclheader=False*)

Class which represents a chunk. The *file* argument is expected to be a file-like object. An instance of this class is specifically allowed. The only method that is needed is `read()`. If the methods `seek()` and `tell()` are present and don't raise an exception, they are also used. If these methods are present and raise an exception, they are expected to not have altered the object. If the optional argument *align* is true, chunks are assumed to be aligned on 2-byte boundaries. If *align* is false, no alignment is assumed. The default value is true. If the optional argument *bigendian* is false, the chunk size is assumed to be in little-endian order. This is needed for WAVE audio files. The default value is true. If the optional argument *inclheader* is true, the size given in the chunk header includes the size of the header. The default value is false.

`Chunk` 对象支持下列方法:

getname()

Returns the name (ID) of the chunk. This is the first 4 bytes of the chunk.

getsize()

Returns the size of the chunk.

close()

Close and skip to the end of the chunk. This does not close the underlying file.

The remaining methods will raise `OSError` if called after the `close()` method has been called. Before Python 3.3, they used to raise `IOError`, now an alias of `OSError`.

isatty()

Returns False.

seek (*pos*, *whence=0*)

Set the chunk's current position. The *whence* argument is optional and defaults to 0 (absolute file positioning); other values are 1 (seek relative to the current position) and 2 (seek relative to the file's end). There is no return value. If the underlying file does not allow seek, only forward seeks are allowed.

¹ "EA IFF 85" Standard for Interchange Format Files, Jerry Morrison, Electronic Arts, January 1985.

tell()

Return the current position into the chunk.

read(size=-1)

Read at most *size* bytes from the chunk (less if the read hits the end of the chunk before obtaining *size* bytes). If the *size* argument is negative or omitted, read all data until the end of the chunk. An empty bytes object is returned when the end of the chunk is encountered immediately.

skip()

Skip to the end of the chunk. All further calls to `read()` for the chunk will return `b''`. If you are not interested in the contents of the chunk, this method should be called so that the file points to the start of the next chunk.

22.6 colorsys —Conversions between color systems

Source code: [Lib/colors.py](#)

The `colorsys` module defines bidirectional conversions of color values between colors expressed in the RGB (Red Green Blue) color space used in computer monitors and three other coordinate systems: YIQ, HLS (Hue Lightness Saturation) and HSV (Hue Saturation Value). Coordinates in all of these color spaces are floating point values. In the YIQ space, the Y coordinate is between 0 and 1, but the I and Q coordinates can be positive or negative. In all other spaces, the coordinates are all between 0 and 1.

参见:

More information about color spaces can be found at <http://www.poynton.com/ColorFAQ.html> and <https://www.cambridgeincolour.com/tutorials/color-spaces.htm>.

The `colorsys` module defines the following functions:

`colorsys.rgb_to_yiq(r, g, b)`

Convert the color from RGB coordinates to YIQ coordinates.

`colorsys.yiq_to_rgb(y, i, q)`

Convert the color from YIQ coordinates to RGB coordinates.

`colorsys.rgb_to_hls(r, g, b)`

Convert the color from RGB coordinates to HLS coordinates.

`colorsys.hls_to_rgb(h, l, s)`

Convert the color from HLS coordinates to RGB coordinates.

`colorsys.rgb_to_hsv(r, g, b)`

Convert the color from RGB coordinates to HSV coordinates.

`colorsys.hsv_to_rgb(h, s, v)`

Convert the color from HSV coordinates to RGB coordinates.

示例:

```
>>> import colorsys
>>> colorsys.rgb_to_hsv(0.2, 0.4, 0.4)
(0.5, 0.5, 0.4)
>>> colorsys.hsv_to_rgb(0.5, 0.5, 0.4)
(0.2, 0.4, 0.4)
```

22.7 imghdr — 推测图像类型

源代码 `Lib/imghdr.py`

`imghdr` 模块推测文件或字节流中的图像的类型。

`imghdr` 模块定义了以下类型：

`imghdr.what(filename, h=None)`

测试包含在命名为 `filename` 的文件中的图像数据，并且返回描述此类图片的字符串。如果可选的 `h` 被提供，`filename` 将被忽略并且 `h` 包含将被测试的二进制流。

在 3.6 版更改：接受一个类路径对象。

接下来的图像类型是可识别的，返回值来自 `what()`：

值	图像格式
'rgb'	SGI 图像库文件
'gif'	GIF 87a 和 89a 文件
'pbm'	便携式位图文件
'pgm'	便携式灰度图文件
'ppm'	便携式像素表文件
'tiff'	TIFF 文件
'rast'	Sun 光栅文件
'xbm'	X 位图文件
'jpeg'	JFIF 或 Exif 格式的 JPEG 数据
'bmp'	BMP 文件
'png'	便携式网络图像
'webp'	WebP 文件
'exr'	OpenEXR 文件

3.5 新版功能：`exr` 和 `webp` 格式被添加。

你可以扩展此 `imghdr` 可以被追加的这个变量识别的文件格式的列表：

`imghdr.tests`

执行单个测试的函数列表。每个函数都有两个参数：字节流和类似开放文件的对象。当 `what()` 用字节流调用时，类文件对象将是 `None`。

如果测试陈工，这个测试函数应当返回一个描述图像类型的字符串，否则返回 `None`。

示例：

```
>>> import imghdr
>>> imghdr.what('bass.gif')
'gif'
```

22.8 sndhdr — 推测声音文件的类型

源代码 `Lib/sndhdr.py`

`sndhdr` 提供了企图猜测文件中的声音数据类型的功能函数。当这些函数可以推测出存储在文件中的声音数据的类型是，它们返回一个 `collections.namedtuple()`，包含了五种属性：(`filetype`, `framerate`, `nchannels`, `nframes`, `sampwidth`)。这些 *type* 的值表示数据的类型，会是以下字符串之一：'aifc', 'aiff', 'au', 'hcom', 'sndr', 'sndt', 'voc', 'wav', '8svx', 'sb', 'ub', or 'ul'。 *sampling_rate* 可能是实际值或者当未知或者难以解码时的 0。类似的， *channels* 也会返回实际值或者在无法推测或者难以解码时返回 0。 *frames* 则是实际值或 -1。元组的最后一项， *bits_per_sample* 将会为比特表示的 *sample* 大小或者 A-LAW 时为 'A'， u-LAW 时为 'U'。

`sndhdr.what(filename)`

使用 `whathdr()` 推测存储在 *filename* 文件中的声音数据的类型。如果成功，返回上述的命名元组，否则返回 `None`。

在 3.5 版更改：将结果从元组改为命名元组。

`sndhdr.whathdr(filename)`

基于文件头推测存储在文件中的声音数据类型。文件名由 *filename* 给出。这个函数在成功时返回上述命名元组，或者在失败时返回 `None`。

在 3.5 版更改：将结果从元组改为命名元组。

22.9 ossaudiodev — Access to OSS-compatible audio devices

This module allows you to access the OSS (Open Sound System) audio interface. OSS is available for a wide range of open-source and commercial Unices, and is the standard audio interface for Linux and recent versions of FreeBSD.

在 3.3 版更改：Operations in this module now raise `OSError` where `IOError` was raised.

参见：

Open Sound System Programmer's Guide the official documentation for the OSS C API

The module defines a large number of constants supplied by the OSS device driver; see `<sys/soundcard.h>` on either Linux or FreeBSD for a listing.

`ossaudiodev` defines the following variables and functions:

exception `ossaudiodev.OSSAudioError`

This exception is raised on certain errors. The argument is a string describing what went wrong.

(If `ossaudiodev` receives an error from a system call such as `open()`, `write()`, or `ioctl()`, it raises `OSError`. Errors detected directly by `ossaudiodev` result in `OSSAudioError`.)

(For backwards compatibility, the exception class is also available as `ossaudiodev.error`.)

`ossaudiodev.open(mode)`

`ossaudiodev.open(device, mode)`

Open an audio device and return an OSS audio device object. This object supports many file-like methods, such as `read()`, `write()`, and `fileno()` (although there are subtle differences between conventional Unix read/write semantics and those of OSS audio devices). It also supports a number of audio-specific methods; see below for the complete list of methods.

device is the audio device filename to use. If it is not specified, this module first looks in the environment variable `AUDIODEV` for a device to use. If not found, it falls back to `/dev/dsp`.

mode is one of 'r' for read-only (record) access, 'w' for write-only (playback) access and 'rw' for both. Since many sound cards only allow one process to have the recorder or player open at a time, it is a good idea to open the device only for the activity needed. Further, some sound cards are half-duplex: they can be opened for reading or writing, but not both at once.

Note the unusual calling syntax: the *first* argument is optional, and the second is required. This is a historical artifact for compatibility with the older `linuxaudiodev` module which `ossaudiodev` supersedes.

`ossaudiodev.openmixer([device])`

Open a mixer device and return an OSS mixer device object. *device* is the mixer device filename to use. If it is not specified, this module first looks in the environment variable `MIXERDEV` for a device to use. If not found, it falls back to `/dev/mixer`.

22.9.1 Audio Device Objects

Before you can write to or read from an audio device, you must call three methods in the correct order:

1. `setfmt()` to set the output format
2. `channels()` to set the number of channels
3. `speed()` to set the sample rate

Alternately, you can use the `setparameters()` method to set all three audio parameters at once. This is more convenient, but may not be as flexible in all cases.

The audio device objects returned by `open()` define the following methods and (read-only) attributes:

`oss_audio_device.close()`

Explicitly close the audio device. When you are done writing to or reading from an audio device, you should explicitly close it. A closed device cannot be used again.

`oss_audio_device.fileno()`

Return the file descriptor associated with the device.

`oss_audio_device.read(size)`

Read *size* bytes from the audio input and return them as a Python string. Unlike most Unix device drivers, OSS audio devices in blocking mode (the default) will block `read()` until the entire requested amount of data is available.

`oss_audio_device.write(data)`

Write a *bytes-like object* *data* to the audio device and return the number of bytes written. If the audio device is in blocking mode (the default), the entire data is always written (again, this is different from usual Unix device semantics). If the device is in non-blocking mode, some data may not be written—see `writeall()`.

在 3.5 版更改: Writable *bytes-like object* is now accepted.

`oss_audio_device.writeall(data)`

Write a *bytes-like object* *data* to the audio device: waits until the audio device is able to accept data, writes as much data as it will accept, and repeats until *data* has been completely written. If the device is in blocking mode (the default), this has the same effect as `write()`; `writeall()` is only useful in non-blocking mode. Has no return value, since the amount of data written is always equal to the amount of data supplied.

在 3.5 版更改: Writable *bytes-like object* is now accepted.

在 3.2 版更改: Audio device objects also support the context management protocol, i.e. they can be used in a `with` statement.

The following methods each map to exactly one `ioctl()` system call. The correspondence is obvious: for example, `setfmt()` corresponds to the `SNDCTL_DSP_SETFMT` `ioctl`, and `sync()` to `SNDCTL_DSP_SYNC` (this can be useful when consulting the OSS documentation). If the underlying `ioctl()` fails, they all raise `OSError`.

`oss_audio_device.nonblock()`

Put the device into non-blocking mode. Once in non-blocking mode, there is no way to return it to blocking mode.

`oss_audio_device.getfmts()`

Return a bitmask of the audio output formats supported by the soundcard. Some of the formats supported by OSS are:

文件格式	描述
AFMT_MU_LAW	a logarithmic encoding (used by Sun .au files and /dev/audio)
AFMT_A_LAW	a logarithmic encoding
AFMT_IMA_ADPCM	a 4:1 compressed format defined by the Interactive Multimedia Association
AFMT_U8	Unsigned, 8-bit audio
AFMT_S16_LE	Signed, 16-bit audio, little-endian byte order (as used by Intel processors)
AFMT_S16_BE	Signed, 16-bit audio, big-endian byte order (as used by 68k, PowerPC, Sparc)
AFMT_S8	Signed, 8 bit audio
AFMT_U16_LE	Unsigned, 16-bit little-endian audio
AFMT_U16_BE	Unsigned, 16-bit big-endian audio

Consult the OSS documentation for a full list of audio formats, and note that most devices support only a subset of these formats. Some older devices only support AFMT_U8; the most common format used today is AFMT_S16_LE.

`oss_audio_device.setfmt(format)`

Try to set the current audio format to *format*—see `getfmts()` for a list. Returns the audio format that the device was set to, which may not be the requested format. May also be used to return the current audio format—do this by passing an “audio format” of AFMT_QUERY.

`oss_audio_device.channels(nchannels)`

Set the number of output channels to *nchannels*. A value of 1 indicates monophonic sound, 2 stereophonic. Some devices may have more than 2 channels, and some high-end devices may not support mono. Returns the number of channels the device was set to.

`oss_audio_device.speed(samplerate)`

Try to set the audio sampling rate to *samplerate* samples per second. Returns the rate actually set. Most sound devices don't support arbitrary sampling rates. Common rates are:

Rate	描述
8000	default rate for /dev/audio
11025	speech recording
22050	
44100	CD quality audio (at 16 bits/sample and 2 channels)
96000	DVD quality audio (at 24 bits/sample)

`oss_audio_device.sync()`

Wait until the sound device has played every byte in its buffer. (This happens implicitly when the device is closed.) The OSS documentation recommends closing and re-opening the device rather than using `sync()`.

`oss_audio_device.reset()`

Immediately stop playing or recording and return the device to a state where it can accept commands. The OSS documentation recommends closing and re-opening the device after calling `reset()`.

`oss_audio_device.post()`

Tell the driver that there is likely to be a pause in the output, making it possible for the device to handle the pause more intelligently. You might use this after playing a spot sound effect, before waiting for user input, or before doing disk I/O.

The following convenience methods combine several ioctls, or one ioctl and some simple calculations.

`oss_audio_device.setparameters (format, nchannels, samplerate[, strict=False])`

Set the key audio sampling parameters—sample format, number of channels, and sampling rate—in one method call. *format*, *nchannels*, and *samplerate* should be as specified in the `setfmt()`, `channels()`, and `speed()` methods. If *strict* is true, `setparameters()` checks to see if each parameter was actually set to the requested value, and raises `OSSAudioError` if not. Returns a tuple (*format*, *nchannels*, *samplerate*) indicating the parameter values that were actually set by the device driver (i.e., the same as the return values of `setfmt()`, `channels()`, and `speed()`).

For example,

```
(fmt, channels, rate) = dsp.setparameters(fmt, channels, rate)
```

is equivalent to

```
fmt = dsp.setfmt(fmt)
channels = dsp.channels(channels)
rate = dsp.rate(rate)
```

`oss_audio_device.bufsize()`

Returns the size of the hardware buffer, in samples.

`oss_audio_device.obufcount()`

Returns the number of samples that are in the hardware buffer yet to be played.

`oss_audio_device.obuffree()`

Returns the number of samples that could be queued into the hardware buffer to be played without blocking.

Audio device objects also support several read-only attributes:

`oss_audio_device.closed`

Boolean indicating whether the device has been closed.

`oss_audio_device.name`

String containing the name of the device file.

`oss_audio_device.mode`

The I/O mode for the file, either "r", "rw", or "w".

22.9.2 Mixer Device Objects

The mixer object provides two file-like methods:

`oss_mixer_device.close()`

This method closes the open mixer device file. Any further attempts to use the mixer after this file is closed will raise an `OSError`.

`oss_mixer_device.fileno()`

Returns the file handle number of the open mixer device file.

在 3.2 版更改: Mixer objects also support the context management protocol.

The remaining methods are specific to audio mixing:

`oss_mixer_device.controls()`

This method returns a bitmask specifying the available mixer controls (“Control” being a specific mixable “channel”, such as `SOUND_MIXER_PCM` or `SOUND_MIXER_SYNTH`). This bitmask indicates a subset of all available mixer controls—the `SOUND_MIXER_*` constants defined at module level. To determine if, for example, the current mixer object supports a PCM mixer, use the following Python code:

```

mixer=ossaudiodev.openmixer()
if mixer.controls() & (1 << ossaudiodev.SOUND_MIXER_PCM):
    # PCM is supported
    ... code ...

```

For most purposes, the `SOUND_MIXER_VOLUME` (master volume) and `SOUND_MIXER_PCM` controls should suffice—but code that uses the mixer should be flexible when it comes to choosing mixer controls. On the Gravis Ultrasound, for example, `SOUND_MIXER_VOLUME` does not exist.

`oss_mixer_device.stereocontrols()`

Returns a bitmask indicating stereo mixer controls. If a bit is set, the corresponding control is stereo; if it is unset, the control is either monophonic or not supported by the mixer (use in combination with `controls()` to determine which).

See the code example for the `controls()` function for an example of getting data from a bitmask.

`oss_mixer_device.recontrols()`

Returns a bitmask specifying the mixer controls that may be used to record. See the code example for `controls()` for an example of reading from a bitmask.

`oss_mixer_device.get(control)`

Returns the volume of a given mixer control. The returned volume is a 2-tuple (`left_volume`, `right_volume`). Volumes are specified as numbers from 0 (silent) to 100 (full volume). If the control is monophonic, a 2-tuple is still returned, but both volumes are the same.

Raises `OSSAudioError` if an invalid control is specified, or `OSError` if an unsupported control is specified.

`oss_mixer_device.set(control, (left, right))`

Sets the volume for a given mixer control to (`left`, `right`). `left` and `right` must be ints and between 0 (silent) and 100 (full volume). On success, the new volume is returned as a 2-tuple. Note that this may not be exactly the same as the volume specified, because of the limited resolution of some soundcard's mixers.

Raises `OSSAudioError` if an invalid mixer control was specified, or if the specified volumes were out-of-range.

`oss_mixer_device.get_recsrc()`

This method returns a bitmask indicating which control(s) are currently being used as a recording source.

`oss_mixer_device.set_recsrc(bitmask)`

Call this function to specify a recording source. Returns a bitmask indicating the new recording source (or sources) if successful; raises `OSError` if an invalid source was specified. To set the current recording source to the microphone input:

```

mixer.setrecsrc (1 << ossaudiodev.SOUND_MIXER_MIC)

```


本章中介绍的模块通过提供选择要在程序信息中使用的语言的机制或通过定制输出以匹配本地约定来帮助你编写不依赖于语言和区域设置的软件。

本章中描述的模块列表是：

23.1 gettext — 多语种国际化服务

源代码： [Lib/gettext.py](#)

The `gettext` module provides internationalization (I18N) and localization (L10N) services for your Python modules and applications. It supports both the GNU `gettext` message catalog API and a higher level, class-based API that may be more appropriate for Python files. The interface described below allows you to write your module and application messages in one natural language, and provide a catalog of translated messages for running under different natural languages.

同时还给出一些本地化 Python 模块及应用程序的小技巧。

23.1.1 GNU gettext API

模块 `gettext` 定义了下列 API，这与 `gettext` API 类似。如果你使用该 API，将会对整个应用程序产生全局的影响。如果你的应用程序支持多语种，而语言选择取决于用户的区域设置，这通常正是你所想要的。而如果你正在本地化某个 Python 模块，或者你的应用程序需要在运行时切换语言，相反你或许想用基于类的 API。

`gettext.bindtextdomain(domain, localedir=None)`

Bind the *domain* to the locale directory *localedir*. More concretely, `gettext` will look for binary `.mo` files for the given domain using the path (on Unix): `localedir/language/LC_MESSAGES/domain.mo`, where *languages* is searched for in the environment variables `LANGUAGE`, `LC_ALL`, `LC_MESSAGES`, and `LANG` respectively.

如果遗漏了 *localedir* 或者设置为 *None*, 那么将返回当前 *domain* 所绑定的值¹

`gettext.bind_textdomain_codeset (domain, codeset=None)`

Bind the *domain* to *codeset*, changing the encoding of byte strings returned by the *lgettext()*, *ldgettext()*, *lnggettext()* and *ldnggettext()* functions. If *codeset* is omitted, then the current binding is returned.

`gettext.textdomain (domain=None)`

Change or query the current global domain. If *domain* is *None*, then the current global domain is returned, otherwise the global domain is set to *domain*, which is returned.

`gettext.gettext (message)`

Return the localized translation of *message*, based on the current global domain, language, and locale directory. This function is usually aliased as *_()* in the local namespace (see examples below).

`gettext.dgettext (domain, message)`

Like *gettext()*, but look the message up in the specified *domain*.

`gettext.ngettext (singular, plural, n)`

Like *gettext()*, but consider plural forms. If a translation is found, apply the plural formula to *n*, and return the resulting message (some languages have more than two plural forms). If no translation is found, return *singular* if *n* is 1; return *plural* otherwise.

The Plural formula is taken from the catalog header. It is a C or Python expression that has a free variable *n*; the expression evaluates to the index of the plural in the catalog. See the [GNU gettext documentation](#) for the precise syntax to be used in *.po* files and the formulas for a variety of languages.

`gettext.dngettext (domain, singular, plural, n)`

Like *lnggettext()*, but look the message up in the specified *domain*.

`gettext.lgettext (message)`

`gettext.ldgettext (domain, message)`

`gettext.lnggettext (singular, plural, n)`

`gettext.ldnggettext (domain, singular, plural, n)`

Equivalent to the corresponding functions without the *l* prefix (*gettext()*, *dgettext()*, *ngettext()* and *dngettext()*), but the translation is returned as a byte string encoded in the preferred system encoding if no other encoding was explicitly set with *bind_textdomain_codeset()*.

警告: These functions should be avoided in Python 3, because they return encoded bytes. It's much better to use alternatives which return Unicode strings instead, since most Python applications will want to manipulate human readable text as strings instead of bytes. Further, it's possible that you may get unexpected Unicode-related exceptions if there are encoding problems with the translated strings. It is possible that the *l*()* functions will be deprecated in future Python versions due to their inherent problems and limitations.

Note that GNU **gettext** also defines a *dcgettext()* method, but this was deemed not useful and so it is currently unimplemented.

Here's an example of typical usage for this API:

```
import gettext
gettext.bindtextdomain('myapplication', '/path/to/my/language/directory')
gettext.textdomain('myapplication')
_ = gettext.gettext
```

(下页继续)

¹ The default locale directory is system dependent; for example, on RedHat Linux it is */usr/share/locale*, but on Solaris it is */usr/lib/locale*. The *gettext* module does not try to support these system dependent defaults; instead its default is *sys.prefix/share/locale*. For this reason, it is always best to call *bindtextdomain()* with an explicit absolute path at the start of your application.

(续上页)

```
# ...
print(_('This is a translatable string.'))
```

23.1.2 Class-based API

The class-based API of the `gettext` module gives you more flexibility and greater convenience than the GNU `gettext` API. It is the recommended way of localizing your Python applications and modules. `gettext` defines a “translations” class which implements the parsing of GNU `.mo` format files, and has methods for returning strings. Instances of this “translations” class can also install themselves in the built-in namespace as the function `_()`.

`gettext.find(domain, localedir=None, languages=None, all=False)`

This function implements the standard `.mo` file search algorithm. It takes a *domain*, identical to what `textdomain()` takes. Optional *localedir* is as in `bindtextdomain()`. Optional *languages* is a list of strings, where each string is a language code.

If *localedir* is not given, then the default system locale directory is used.² If *languages* is not given, then the following environment variables are searched: `LANGUAGE`, `LC_ALL`, `LC_MESSAGES`, and `LANG`. The first one returning a non-empty value is used for the *languages* variable. The environment variables should contain a colon separated list of languages, which will be split on the colon to produce the expected list of language code strings.

`find()` then expands and normalizes the languages, and then iterates through them, searching for an existing file built of these components:

```
localedir/language/LC_MESSAGES/domain.mo
```

The first such file name that exists is returned by `find()`. If no such file is found, then `None` is returned. If *all* is given, it returns a list of all file names, in the order in which they appear in the languages list or the environment variables.

`gettext.translation(domain, localedir=None, languages=None, class_=None, fallback=False, codeset=None)`

Return a `Translations` instance based on the *domain*, *localedir*, and *languages*, which are first passed to `find()` to get a list of the associated `.mo` file paths. Instances with identical `.mo` file names are cached. The actual class instantiated is either *class_* if provided, otherwise `GNUTranslations`. The class’ s constructor must take a single *file object* argument. If provided, *codeset* will change the charset used to encode translated strings in the `gettext()` and `gettext()` methods.

If multiple files are found, later files are used as fallbacks for earlier ones. To allow setting the fallback, `copy.copy()` is used to clone each translation object from the cache; the actual instance data is still shared with the cache.

If no `.mo` file is found, this function raises `OSError` if *fallback* is false (which is the default), and returns a `NullTranslations` instance if *fallback* is true.

在 3.3 版更改: `IOError` used to be raised instead of `OSError`.

`gettext.install(domain, localedir=None, codeset=None, names=None)`

This installs the function `_()` in Python’ s builtins namespace, based on *domain*, *localedir*, and *codeset* which are passed to the function `translation()`.

For the *names* parameter, please see the description of the translation object’ s `install()` method.

As seen below, you usually mark the strings in your application that are candidates for translation, by wrapping them in a call to the `_()` function, like this:

² See the footnote for `bindtextdomain()` above.

```
print(_('This string will be translated.'))
```

For convenience, you want the `_()` function to be installed in Python’s builtins namespace, so it is easily accessible in all modules of your application.

The `NullTranslations` class

Translation classes are what actually implement the translation of original source file message strings to translated message strings. The base class used by all translation classes is `NullTranslations`; this provides the basic interface you can use to write your own specialized translation classes. Here are the methods of `NullTranslations`:

class `gettext.NullTranslations` (*fp=None*)

Takes an optional *file object* `fp`, which is ignored by the base class. Initializes “protected” instance variables `_info` and `_charset` which are set by derived classes, as well as `_fallback`, which is set through `add_fallback()`. It then calls `self._parse(fp)` if `fp` is not `None`.

_parse (*fp*)

No-op’ d in the base class, this method takes file object `fp`, and reads the data from the file, initializing its message catalog. If you have an unsupported message catalog file format, you should override this method to parse your format.

add_fallback (*fallback*)

Add *fallback* as the fallback object for the current translation object. A translation object should consult the fallback if it cannot provide a translation for a given message.

gettext (*message*)

If a fallback has been set, forward `gettext()` to the fallback. Otherwise, return *message*. Overridden in derived classes.

ngettext (*singular, plural, n*)

If a fallback has been set, forward `ngettext()` to the fallback. Otherwise, return *singular* if *n* is 1; return *plural* otherwise. Overridden in derived classes.

lgettext (*message*)

lngettext (*singular, plural, n*)

Equivalent to `gettext()` and `ngettext()`, but the translation is returned as a byte string encoded in the preferred system encoding if no encoding was explicitly set with `set_output_charset()`. Overridden in derived classes.

警告: These methods should be avoided in Python 3. See the warning for the `lgettext()` function.

info ()

Return the “protected” `_info` variable.

charset ()

Return the encoding of the message catalog file.

output_charset ()

Return the encoding used to return translated messages in `lgettext()` and `lngettext()`.

set_output_charset (*charset*)

Change the encoding used to return translated messages.

install (*names=None*)

This method installs `gettext()` into the built-in namespace, binding it to `_`.

If the *names* parameter is given, it must be a sequence containing the names of functions you want to install in the builtins namespace in addition to `_()`. Supported names are `'gettext'`, `'ngettext'`, `'lgettext'` and `'lngettext'`.

Note that this is only one way, albeit the most convenient way, to make the `_()` function available to your application. Because it affects the entire application globally, and specifically the built-in namespace, localized modules should never install `_()`. Instead, they should use this code to make `_()` available to their module:

```
import gettext
t = gettext.translation('mymodule', ...)
_ = t.gettext
```

This puts `_()` only in the module's global namespace and so only affects calls within this module.

The GNUTranslations class

The `gettext` module provides one additional class derived from `NullTranslations`: `GNUTranslations`. This class overrides `_parse()` to enable reading GNU `gettext` format `.mo` files in both big-endian and little-endian format.

`GNUTranslations` parses optional meta-data out of the translation catalog. It is convention with GNU `gettext` to include meta-data as the translation for the empty string. This meta-data is in RFC 822-style `key: value` pairs, and should contain the `Project-Id-Version` key. If the key `Content-Type` is found, then the `charset` property is used to initialize the “protected” `_charset` instance variable, defaulting to `None` if not found. If the charset encoding is specified, then all message ids and message strings read from the catalog are converted to Unicode using this encoding, else ASCII encoding is assumed.

Since message ids are read as Unicode strings too, all `*gettext()` methods will assume message ids as Unicode strings, not byte strings.

The entire set of key/value pairs are placed into a dictionary and set as the “protected” `_info` instance variable.

If the `.mo` file's magic number is invalid, the major version number is unexpected, or if other problems occur while reading the file, instantiating a `GNUTranslations` class can raise `OSError`.

class gettext.GNUTranslations

The following methods are overridden from the base class implementation:

gettext(*message*)

Look up the *message* id in the catalog and return the corresponding message string, as a Unicode string. If there is no entry in the catalog for the *message* id, and a fallback has been set, the look up is forwarded to the fallback's `gettext()` method. Otherwise, the *message* id is returned.

ngettext(*singular*, *plural*, *n*)

Do a plural-forms lookup of a message id. *singular* is used as the message id for purposes of lookup in the catalog, while *n* is used to determine which plural form to use. The returned message string is a Unicode string.

If the message id is not found in the catalog, and a fallback is specified, the request is forwarded to the fallback's `ngettext()` method. Otherwise, when *n* is 1 *singular* is returned, and *plural* is returned in all other cases.

Here is an example:

```
n = len(os.listdir('.'))
cat = GNUTranslations(somefile)
message = cat.ngettext(
    'There is %(num)d file in this directory',
    'There are %(num)d files in this directory',
    n) % {'num': n}
```

lgettext (*message*)

lngettext (*singular, plural, n*)

Equivalent to `gettext()` and `ngettext()`, but the translation is returned as a byte string encoded in the preferred system encoding if no encoding was explicitly set with `set_output_charset()`.

警告: These methods should be avoided in Python 3. See the warning for the `lgettext()` function.

Solaris message catalog support

The Solaris operating system defines its own binary `.mo` file format, but since no documentation can be found on this format, it is not supported at this time.

The Catalog constructor

GNOME uses a version of the `gettext` module by James Henstridge, but this version has a slightly different API. Its documented usage was:

```
import gettext
cat = gettext.Catalog(domain, localedir)
_ = cat.gettext
print(_('hello world'))
```

For compatibility with this older module, the function `Catalog()` is an alias for the `translation()` function described above.

One difference between this module and Henstridge's: his catalog objects supported access through a mapping API, but this appears to be unused and so is not currently supported.

23.1.3 Internationalizing your programs and modules

Internationalization (I18N) refers to the operation by which a program is made aware of multiple languages. Localization (L10N) refers to the adaptation of your program, once internationalized, to the local language and cultural habits. In order to provide multilingual messages for your Python programs, you need to take the following steps:

1. prepare your program or module by specially marking translatable strings
2. run a suite of tools over your marked files to generate raw messages catalogs
3. create language specific translations of the message catalogs
4. use the `gettext` module so that message strings are properly translated

In order to prepare your code for I18N, you need to look at all the strings in your files. Any string that needs to be translated should be marked by wrapping it in `_('...')` —that is, a call to the function `_()`. For example:

```
filename = 'mylog.txt'
message = _('writing a log message')
fp = open(filename, 'w')
fp.write(message)
fp.close()
```

In this example, the string `'writing a log message'` is marked as a candidate for translation, while the strings `'mylog.txt'` and `'w'` are not.

There are a few tools to extract the strings meant for translation. The original GNU **gettext** only supported C or C++ source code but its extended version **xgettext** scans code written in a number of languages, including Python, to find strings marked as translatable. **Babel** is a Python internationalization library that includes a `pybabel` script to extract and compile message catalogs. François Pinard’s program called **xpot** does a similar job and is available as part of his `po-utils` package.

(Python also includes pure-Python versions of these programs, called `pygettext.py` and `msgfmt.py`; some Python distributions will install them for you. `pygettext.py` is similar to **xgettext**, but only understands Python source code and cannot handle other programming languages such as C or C++. `pygettext.py` supports a command-line interface similar to **xgettext**; for details on its use, run `pygettext.py --help`. `msgfmt.py` is binary compatible with GNU **msgfmt**. With these two programs, you may not need the GNU **gettext** package to internationalize your Python applications.)

xgettext, **pygettext**, and similar tools generate `.po` files that are message catalogs. They are structured human-readable files that contain every marked string in the source code, along with a placeholder for the translated versions of these strings.

Copies of these `.po` files are then handed over to the individual human translators who write translations for every supported natural language. They send back the completed language-specific versions as a `<language-name>.po` file that’s compiled into a machine-readable `.mo` binary catalog file using the **msgfmt** program. The `.mo` files are used by the `gettext` module for the actual translation processing at run-time.

How you use the `gettext` module in your code depends on whether you are internationalizing a single module or your entire application. The next two sections will discuss each case.

Localizing your module

If you are localizing your module, you must take care not to make global changes, e.g. to the built-in namespace. You should not use the GNU **gettext** API but instead the class-based API.

Let’s say your module is called “spam” and the module’s various natural language translation `.mo` files reside in `/usr/share/locale` in GNU **gettext** format. Here’s what you would put at the top of your module:

```
import gettext
t = gettext.translation('spam', '/usr/share/locale')
_ = t.gettext
```

Localizing your application

If you are localizing your application, you can install the `_()` function globally into the built-in namespace, usually in the main driver file of your application. This will let all your application-specific files just use `_('...')` without having to explicitly install it in each file.

In the simple case then, you need only add the following bit of code to the main driver file of your application:

```
import gettext
gettext.install('myapplication')
```

If you need to set the locale directory, you can pass it into the `install()` function:

```
import gettext
gettext.install('myapplication', '/usr/share/locale')
```

Changing languages on the fly

If your program needs to support many languages at the same time, you may want to create multiple translation instances and then switch between them explicitly, like so:

```
import gettext

lang1 = gettext.translation('myapplication', languages=['en'])
lang2 = gettext.translation('myapplication', languages=['fr'])
lang3 = gettext.translation('myapplication', languages=['de'])

# start by using language1
lang1.install()

# ... time goes by, user selects language 2
lang2.install()

# ... more time goes by, user selects language 3
lang3.install()
```

Deferred translations

In most coding situations, strings are translated where they are coded. Occasionally however, you need to mark strings for translation, but defer actual translation until later. A classic example is:

```
animals = ['mollusk',
           'albatross',
           'rat',
           'penguin',
           'python', ]

# ...
for a in animals:
    print(a)
```

Here, you want to mark the strings in the `animals` list as being translatable, but you don't actually want to translate them until they are printed.

Here is one way you can handle this situation:

```
def _(message): return message

animals = [_('mollusk'),
           _('albatross'),
           _('rat'),
           _('penguin'),
           _('python'), ]

del _

# ...
for a in animals:
    print_(a)
```

This works because the dummy definition of `_()` simply returns the string unchanged. And this dummy definition will temporarily override any definition of `_()` in the built-in namespace (until the `del` command). Take care, though if you have a previous definition of `_()` in the local namespace.

Note that the second use of `_()` will not identify “a” as being translatable to the `gettext` program, because the parameter is not a string literal.

Another way to handle this is with the following example:

```
def N_(message): return message

animals = [N_('mollusk'),
           N_('albatross'),
           N_('rat'),
           N_('penguin'),
           N_('python'), ]

# ...
for a in animals:
    print(_(a))
```

In this case, you are marking translatable strings with the function `N_()`, which won't conflict with any definition of `_()`. However, you will need to teach your message extraction program to look for translatable strings marked with `N_()`. `xgettext`, `pygettext`, `pybabel extract`, and `xpot` all support this through the use of the `-k` command-line switch. The choice of `N_()` here is totally arbitrary; it could have just as easily been `MarkThisStringForTranslation()`.

23.1.4 Acknowledgements

The following people contributed code, feedback, design suggestions, previous implementations, and valuable experience to the creation of this module:

- Peter Funk
- James Henstridge
- Juan David Ibáñez Palomar
- Marc-André Lemburg
- Martin von Löwis
- François Pinard
- Barry Warsaw
- Gustavo Niemeyer

23.2 locale — 国际化服务

源代码: [Lib/locale.py](#)

The `locale` module opens access to the POSIX locale database and functionality. The POSIX locale mechanism allows programmers to deal with certain cultural issues in an application, without requiring the programmer to know all the specifics of each country where the software is executed.

The `locale` module is implemented on top of the `_locale` module, which in turn uses an ANSI C locale implementation if available.

The `locale` module defines the following exception and functions:

exception `locale.Error`

Exception raised when the locale passed to `setlocale()` is not recognized.

`locale.setlocale(category, locale=None)`

If `locale` is given and not `None`, `setlocale()` modifies the locale setting for the `category`. The available categories are listed in the data description below. `locale` may be a string, or an iterable of two strings (language code and encoding). If it's an iterable, it's converted to a locale name using the locale aliasing engine. An empty string specifies the user's default settings. If the modification of the locale fails, the exception `Error` is raised. If successful, the new locale setting is returned.

If `locale` is omitted or `None`, the current setting for `category` is returned.

`setlocale()` is not thread-safe on most systems. Applications typically start with a call of

```
import locale
locale.setlocale(locale.LC_ALL, '')
```

This sets the locale for all categories to the user's default setting (typically specified in the `LANG` environment variable). If the locale is not changed thereafter, using multithreading should not cause problems.

`locale.localeconv()`

Returns the database of the local conventions as a dictionary. This dictionary has the following strings as keys:

类别	(Windows 注册表的) 键	含义
<code>LC_NUMERIC</code>	'decimal_point'	小数点字符。
	'grouping'	Sequence of numbers specifying which relative positions the 'thousands_sep' is expected. If the sequence is terminated with <code>CHAR_MAX</code> , no further grouping is performed. If the sequence terminates with a 0, the last group size is repeatedly used.
	'thousands_sep'	组之间使用的字符。
<code>LC_MONETARY</code>	'int_curr_symbol'	国际货币符号。
	'currency_symbol'	当地货币符号。
	'p_cs_precedes/n_cs_precedes'	Whether the currency symbol precedes the value (for positive resp. negative values).
	'p_sep_by_space/n_sep_by_space'	Whether the currency symbol is separated from the value by a space (for positive resp. negative values).
	'mon_decimal_point'	用于货币金额的小数点。
	'frac_digits'	货币值的本地格式中使用的小数位数。
	'int_frac_digits'	货币价值的国际格式中使用的小数位数。
	'mon_thousands_sep'	用于货币值的组分分隔符。
	'mon_grouping'	相当于 'grouping'，用于货币价值。
	'positive_sign'	用于标注正货币价值的符号。
	'negative_sign'	用于注释负货币价值的符号。
	'p_sign_posn/n_sign_posn'	符号的位置（对于正值或负值），见下文。

可以将所有数值设置为 `CHAR_MAX`，以指示此语言环境中未指定任何值。
下面给出了 'p_sign_posn' 和 'n_sign_posn' 的可能值。

值	解释
0	被括号括起来的货币和金额。
1	该标志应位于值和货币符号之前。
2	该标志应位于值和货币符号之后。
3	标志应该紧跟在值之前。
4	标志应该紧跟值项。
CHAR_MAX	Nothing is specified in this locale.

The function sets temporarily the `LC_CTYPE` locale to the `LC_NUMERIC` locale or the `LC_MONETARY` locale if locales are different and numeric or monetary strings are non-ASCII. This temporary change affects other threads.

在 3.6.5 版更改: The function now sets temporarily the `LC_CTYPE` locale to the `LC_NUMERIC` locale in some cases.

`locale.nl_langinfo(option)`

Return some locale-specific information as a string. This function is not available on all systems, and the set of possible options might also vary across platforms. The possible argument values are numbers, for which symbolic constants are available in the `locale` module.

The `nl_langinfo()` function accepts one of the following keys. Most descriptions are taken from the corresponding description in the GNU C library.

`locale.CODESET`

Get a string with the name of the character encoding used in the selected locale.

`locale.D_T_FMT`

Get a string that can be used as a format string for `time.strftime()` to represent date and time in a locale-specific way.

`locale.D_FMT`

Get a string that can be used as a format string for `time.strftime()` to represent a date in a locale-specific way.

`locale.T_FMT`

Get a string that can be used as a format string for `time.strftime()` to represent a time in a locale-specific way.

`locale.T_FMT_AMPM`

Get a format string for `time.strftime()` to represent time in the am/pm format.

`DAY_1 ... DAY_7`

Get the name of the n-th day of the week.

注解: This follows the US convention of `DAY_1` being Sunday, not the international convention (ISO 8601) that Monday is the first day of the week.

`ABDAY_1 ... ABDAY_7`

Get the abbreviated name of the n-th day of the week.

`MON_1 ... MON_12`

Get the name of the n-th month.

`ABMON_1 ... ABMON_12`

Get the abbreviated name of the n-th month.

`locale.RADIXCHAR`

Get the radix character (decimal dot, decimal comma, etc.).

`locale.THOUSEP`

Get the separator character for thousands (groups of three digits).

`locale.YESEXPR`

Get a regular expression that can be used with the `regex` function to recognize a positive response to a yes/no question.

注解： The expression is in the syntax suitable for the `regex()` function from the C library, which might differ from the syntax used in *re*.

`locale.NOEXPR`

Get a regular expression that can be used with the `regex(3)` function to recognize a negative response to a yes/no question.

`locale.CRNCYSTR`

Get the currency symbol, preceded by “-” if the symbol should appear before the value, “+” if the symbol should appear after the value, or “.” if the symbol should replace the radix character.

`locale.ERA`

Get a string that represents the era used in the current locale.

Most locales do not define this value. An example of a locale which does define this value is the Japanese one. In Japan, the traditional representation of dates includes the name of the era corresponding to the then-emperor’s reign.

Normally it should not be necessary to use this value directly. Specifying the E modifier in their format strings causes the `time.strftime()` function to use this information. The format of the returned string is not specified, and therefore you should not assume knowledge of it on different systems.

`locale.ERA_D_T_FMT`

Get a format string for `time.strftime()` to represent date and time in a locale-specific era-based way.

`locale.ERA_D_FMT`

Get a format string for `time.strftime()` to represent a date in a locale-specific era-based way.

`locale.ERA_T_FMT`

Get a format string for `time.strftime()` to represent a time in a locale-specific era-based way.

`locale.ALT_DIGITS`

Get a representation of up to 100 values used to represent the values 0 to 99.

`locale.getdefaultlocale([envvars])`

Tries to determine the default locale settings and returns them as a tuple of the form (language code, encoding).

According to POSIX, a program which has not called `setlocale(LC_ALL, '')` runs using the portable 'C' locale. Calling `setlocale(LC_ALL, '')` lets it use the default locale as defined by the `LANG` variable. Since we do not want to interfere with the current locale setting we thus emulate the behavior in the way described above.

To maintain compatibility with other platforms, not only the `LANG` variable is tested, but a list of variables given as `envvars` parameter. The first found to be defined will be used. `envvars` defaults to the search path used in GNU gettext; it must always contain the variable name 'LANG'. The GNU gettext search path contains 'LC_ALL', 'LC_CTYPE', 'LANG' and 'LANGUAGE', in that order.

Except for the code 'C', the language code corresponds to **RFC 1766**. *language code* and *encoding* may be `None` if their values cannot be determined.

`locale.getlocale(category=LC_CTYPE)`

Returns the current setting for the given locale category as sequence containing *language code*, *encoding*. *category* may be one of the `LC_*` values except `LC_ALL`. It defaults to `LC_CTYPE`.

Except for the code 'C', the language code corresponds to [RFC 1766](#). *language code* and *encoding* may be `None` if their values cannot be determined.

`locale.getpreferredencoding(do_setlocale=True)`

Return the encoding used for text data, according to user preferences. User preferences are expressed differently on different systems, and might not be available programmatically on some systems, so this function only returns a guess.

On some systems, it is necessary to invoke `setlocale()` to obtain the user preferences, so this function is not thread-safe. If invoking `setlocale` is not necessary or desired, `do_setlocale` should be set to `False`.

`locale.normalize(localename)`

Returns a normalized locale code for the given locale name. The returned locale code is formatted for use with `setlocale()`. If normalization fails, the original name is returned unchanged.

If the given encoding is not known, the function defaults to the default encoding for the locale code just like `setlocale()`.

`locale.resetlocale(category=LC_ALL)`

Sets the locale for *category* to the default setting.

The default setting is determined by calling `getdefaultlocale()`. *category* defaults to `LC_ALL`.

`locale.strcoll(string1, string2)`

Compares two strings according to the current `LC_COLLATE` setting. As any other compare function, returns a negative, or a positive value, or 0, depending on whether *string1* collates before or after *string2* or is equal to it.

`locale.strxfrm(string)`

Transforms a string to one that can be used in locale-aware comparisons. For example, `strxfrm(s1) < strxfrm(s2)` is equivalent to `strcoll(s1, s2) < 0`. This function can be used when the same string is compared repeatedly, e.g. when collating a sequence of strings.

`locale.format(format, val, grouping=False, monetary=False)`

Formats a number *val* according to the current `LC_NUMERIC` setting. The format follows the conventions of the `%` operator. For floating point values, the decimal point is modified if appropriate. If *grouping* is true, also takes the grouping into account.

If *monetary* is true, the conversion uses monetary thousands separator and grouping strings.

Please note that this function will only work for exactly one `%char` specifier. For whole format strings, use `format_string()`.

`locale.format_string(format, val, grouping=False)`

Processes formatting specifiers as in `format % val`, but takes the current locale settings into account.

`locale.currency(val, symbol=True, grouping=False, international=False)`

Formats a number *val* according to the current `LC_MONETARY` settings.

The returned string includes the currency symbol if *symbol* is true, which is the default. If *grouping* is true (which is not the default), grouping is done with the value. If *international* is true (which is not the default), the international currency symbol is used.

Note that this function will not work with the 'C' locale, so you have to set a locale via `setlocale()` first.

`locale.str(float)`

Formats a floating point number using the same format as the built-in function `str(float)`, but takes the decimal point into account.

`locale.delocalize(string)`

Converts a string into a normalized number string, following the `LC_NUMERIC` settings.

3.5 新版功能.

`locale.atof(string)`

Converts a string to a floating point number, following the `LC_NUMERIC` settings.

`locale.atoi(string)`

Converts a string to an integer, following the `LC_NUMERIC` conventions.

`locale.LC_CTYPE`

Locale category for the character type functions. Depending on the settings of this category, the functions of module `string` dealing with case change their behaviour.

`locale.LC_COLLATE`

Locale category for sorting strings. The functions `strcoll()` and `strxfrm()` of the `locale` module are affected.

`locale.LC_TIME`

Locale category for the formatting of time. The function `time.strftime()` follows these conventions.

`locale.LC_MONETARY`

Locale category for formatting of monetary values. The available options are available from the `localeconv()` function.

`locale.LC_MESSAGES`

Locale category for message display. Python currently does not support application specific locale-aware messages. Messages displayed by the operating system, like those returned by `os.strerror()` might be affected by this category.

`locale.LC_NUMERIC`

Locale category for formatting numbers. The functions `format()`, `atoi()`, `atof()` and `str()` of the `locale` module are affected by that category. All other numeric formatting operations are not affected.

`locale.LC_ALL`

Combination of all locale settings. If this flag is used when the locale is changed, setting the locale for all categories is attempted. If that fails for any category, no category is changed at all. When the locale is retrieved using this flag, a string indicating the setting for all categories is returned. This string can be later used to restore the settings.

`locale.CHAR_MAX`

This is a symbolic constant used for different values returned by `localeconv()`.

示例:

```
>>> import locale
>>> loc = locale.getlocale() # get current locale
# use German locale; name might vary with platform
>>> locale.setlocale(locale.LC_ALL, 'de_DE')
>>> locale.strcoll('\xe4n', 'foo') # compare a string containing an umlaut
>>> locale.setlocale(locale.LC_ALL, '') # use user's preferred locale
>>> locale.setlocale(locale.LC_ALL, 'C') # use default (C) locale
>>> locale.setlocale(locale.LC_ALL, loc) # restore saved locale
```

23.2.1 Background, details, hints, tips and caveats

The C standard defines the locale as a program-wide property that may be relatively expensive to change. On top of that, some implementation are broken in such a way that frequent locale changes may cause core dumps. This makes the locale somewhat painful to use correctly.

Initially, when a program is started, the locale is the C locale, no matter what the user's preferred locale is. There is one exception: the `LC_CTYPE` category is changed at startup to set the current locale encoding to the user's preferred locale encoding. The program must explicitly say that it wants the user's preferred locale settings for other categories by calling `setlocale(LC_ALL, '')`.

It is generally a bad idea to call `setlocale()` in some library routine, since as a side effect it affects the entire program. Saving and restoring it is almost as bad: it is expensive and affects other threads that happen to run before the settings have been restored.

If, when coding a module for general use, you need a locale independent version of an operation that is affected by the locale (such as certain formats used with `time.strftime()`), you will have to find a way to do it without using the standard library routine. Even better is convincing yourself that using locale settings is okay. Only as a last resort should you document that your module is not compatible with non-C locale settings.

The only way to perform numeric operations according to the locale is to use the special functions defined by this module: `atof()`, `atoi()`, `format()`, `str()`.

There is no way to perform case conversions and character classifications according to the locale. For (Unicode) text strings these are done according to the character value only, while for byte strings, the conversions and classifications are done according to the ASCII value of the byte, and bytes whose high bit is set (i.e., non-ASCII bytes) are never converted or considered part of a character class such as letter or whitespace.

23.2.2 For extension writers and programs that embed Python

Extension modules should never call `setlocale()`, except to find out what the current locale is. But since the return value can only be used portably to restore it, that is not very useful (except perhaps to find out whether or not the locale is C).

When Python code uses the `locale` module to change the locale, this also affects the embedding application. If the embedding application doesn't want this to happen, it should remove the `_locale` extension module (which does all the work) from the table of built-in modules in the `config.c` file, and make sure that the `_locale` module is not accessible as a shared library.

23.2.3 Access to message catalogs

```
locale.gettext(msg)
```

```
locale.dgettext(domain, msg)
```

```
locale.dcgettext(domain, msg, category)
```

```
locale.textdomain(domain)
```

```
locale.bindtextdomain(domain, dir)
```

The `locale` module exposes the C library's `gettext` interface on systems that provide this interface. It consists of the functions `gettext()`, `dgettext()`, `dcgettext()`, `textdomain()`, `bindtextdomain()`, and `bind_textdomain_codeset()`. These are similar to the same functions in the `gettext` module, but use the C library's binary format for message catalogs, and the C library's search algorithms for locating message catalogs.

Python applications should normally find no need to invoke these functions, and should use `gettext` instead. A known exception to this rule are applications that link with additional C libraries which internally invoke `gettext()` or `dcgettext()`. For these applications, it may be necessary to bind the text domain, so that the libraries can properly locate their message catalogs.

本章中描述的模块是很大程度上决定程序结构的框架。目前，这里描述的模块都面向编写命令行接口。
本章描述的完整模块列表如下：

24.1 turtle — 海龟绘图

源码：[Lib/turtle.py](#)

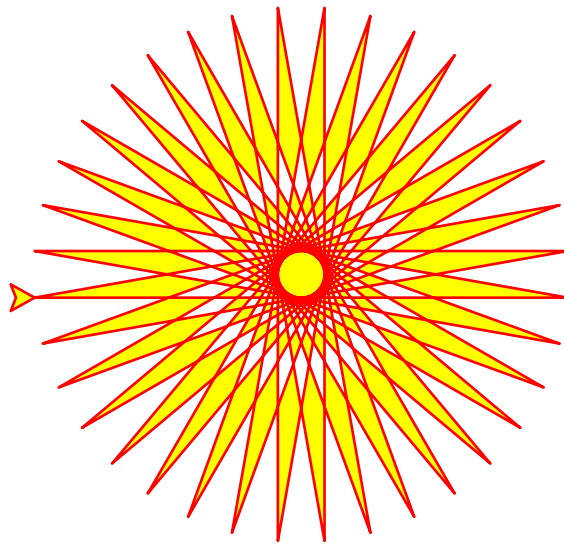
24.1.1 概述

Turtle graphics is a popular way for introducing programming to kids. It was part of the original Logo programming language developed by Wally Feurzig and Seymour Papert in 1966.

请想象绘图区有一只机器海龟，起始位置在 x-y 平面的 (0,0) 点。先执行 `import turtle`，再执行 `turtle.forward(15)`，它将（在屏幕上）朝所面对的 x 轴正方向前进 15 像素，随着它的移动画出一条线段。再执行 `turtle.right(25)`，它将原地右转 25 度。

Turtle star

使用海龟绘图可以编写重复执行简单动作的程序画出精细复杂的形状。



```
from turtle import *
color('red', 'yellow')
begin_fill()
while True:
    forward(200)
    left(170)
    if abs(pos()) < 1:
        break
end_fill()
done()
```

通过组合使用此类命令，可以轻松地绘制出精美的形状和图案。

turtle 模块是基于 Python 标准发行版 2.5 以来的同名模块重新编写并进行了功能扩展。

新模块尽量保持了原模块的特点，并且 (几乎)100% 与其兼容。这就意味着初学编程者能够以交互方式使用模块的所有命令、类和方法——运行 IDLE 时注意加 `-n` 参数。

turtle 模块提供面向对象和面向过程两种形式的海龟绘图基本组件。由于它使用 *tkinter* 实现基本图形界面，因此需要安装了 Tk 支持的 Python 版本。

面向对象的接口主要使用 “2+2” 个类：

1. *TurtleScreen* 类定义图形窗口作为绘图海龟的运动场。它的构造器需要一个 *tkinter.Canvas* 或 *ScrolledCanvas* 作为参数。应在 *turtle* 作为某个程序的一部分的时候使用。

Screen() 函数返回一个 *TurtleScreen* 子类的单例对象。此函数应在 *turtle* 作为独立绘图工具时使用。作为一个单例对象，其所属的类是不可被继承的。

TurtleScreen/Screen 的所有方法还存在对应的函数，即作为面向过程的接口组成部分。

2. *RawTurtle* (别名: *RawPen*) 类定义海龟对象在 *TurtleScreen* 上绘图。它的构造器需要一个 *Canvas*, *ScrolledCanvas* 或 *TurtleScreen* 作为参数，以指定 *RawTurtle* 对象在哪里绘图。

从 *RawTurtle* 派生出子类 *Turtle* (别名: *Pen*)，该类对象在 *Screen* 实例上绘图，如果实例不存在则会自动创建。

RawTurtle/Turtle 的所有方法也存在对应的函数，即作为面向过程的接口组成部分。

过程式接口提供与 *Screen* 和 *Turtle* 类的方法相对应的函数。函数名与对应的方法名相同。当 *Screen* 类的方法对应函数被调用时会自动创建一个 *Screen* 对象。当 *Turtle* 类的方法对应函数被调用时会自动创建一个 (匿名的) *Turtle* 对象。

如果屏幕上需要有多海龟，就必须使用面向对象的接口。

注解： 以下文档给出了函数的参数列表。对于方法来说当然还有额外的第一个参数 *self*，这里省略了。

24.1.2 可用的 Turtle 和 Screen 方法概览

Turtle 方法

海龟动作

移动和绘制

```
forward() | fd() 前进
backward() | bk() | back() 后退
right() | rt() 右转
left() | lt() 左转
goto() | setpos() | setposition() 前往/定位
setx() 设置 x 坐标
sety() 设置 y 坐标
setheading() | seth() 设置朝向
home() 返回原点
circle() 画圆
dot() 画点
stamp() 印章
clearstamp() 清除印章
clearstamps() 清除多个印章
undo() 撤消
speed() 速度
```

获取海龟的状态

```
position() | pos() 位置
towards() 目标方向
xcor() x 坐标
ycor() y 坐标
heading() 朝向
distance() 距离
```

设置与度量单位

```
degrees() 角度
radians() 弧度
```

画笔控制

绘图状态

```
pendown() | pd() | down() 画笔落下
penup() | pu() | up() 画笔抬起
```

`pensize()` | `width()` 画笔粗细
`pen()` 画笔
`isdown()` 画笔是否落下

颜色控制

`color()` 颜色
`pencolor()` 画笔颜色
`fillcolor()` 填充颜色

填充

`filling()` 是否填充
`begin_fill()` 开始填充
`end_fill()` 结束填充

更多绘图控制

`reset()` 重置
`clear()` 清空
`write()` 书写

海龟状态

可见性

`showturtle()` | `st()` 显示海龟
`hideturtle()` | `ht()` 隐藏海龟
`isvisible()` 是否可见

外观

`shape()` 形状
`resizemode()` 大小调整模式
`shapeseize()` | `turtlesize()` 形状大小
`shearfactor()` 剪切因子
`settiltangle()` 设置倾角
`tiltangle()` 倾角
`tilt()` 倾斜
`shapetransform()` 变形
`get_shapepoly()` 获取形状多边形

使用事件

`onclick()` 当鼠标点击
`onrelease()` 当鼠标释放
`ondrag()` 当鼠标拖动

特殊海龟方法

`begin_poly()` 开始记录多边形
`end_poly()` 结束记录多边形
`get_poly()` 获取多边形
`clone()` 克隆
`getturtle()` | `getpen()` 获取海龟画笔
`getscreen()` 获取屏幕
`setundobuffer()` 设置撤消缓冲区

`undobufferentries()` 撤消缓冲区条目数

TurtleScreen/Screen 方法

窗口控制

`bgcolor()` 背景颜色
`bgpic()` 背景图片
`clear()` | `clearscreen()` 清屏
`reset()` | `resetscreen()` 重置
`screensize()` 屏幕大小
`setworldcoordinates()` 设置世界坐标系

动画控制

`delay()` 延迟
`tracer()` 追踪
`update()` 更新

使用屏幕事件

`listen()` 监听
`onkey()` | `onkeyrelease()` 当键盘按下并释放
`onkeypress()` 当键盘按下
`onclick()` | `onscreenclick()` 当点击屏幕
`ontimer()` 当达到定时
`mainloop()` | `done()` 主循环

设置与特殊方法

`mode()` 模式
`colormode()` 颜色模式
`getcanvas()` 获取画布
`getshapes()` 获取形状
`register_shape()` | `addshape()` 添加形状
`turtles()` 所有海龟
`window_height()` 窗口高度
`window_width()` 窗口宽度

输入方法

`textinput()` 文本输入
`numinput()` 数字输入

Screen 专有方法

`bye()` 退出
`exitonclick()` 当点击时退出
`setup()` 设置
`title()` 标题

24.1.3 RawTurtle/Turtle 方法和对应函数

本节中的大部分示例都使用 `Turtle` 类的一个实例，命名为 `turtle`。

海龟动作

`turtle.forward(distance)`
`turtle.fd(distance)`

参数 distance – 一个数值 (整型或浮点型)

海龟前进 *distance* 指定的距离, 方向为海龟的朝向。

```
>>> turtle.position()
(0.00,0.00)
>>> turtle.forward(25)
>>> turtle.position()
(25.00,0.00)
>>> turtle.forward(-75)
>>> turtle.position()
(-50.00,0.00)
```

`turtle.back(distance)`
`turtle.bk(distance)`
`turtle.backward(distance)`

参数 distance – 一个数值

海龟后退 *distance* 指定的距离, 方向与海龟的朝向相反。不改变海龟的朝向。

```
>>> turtle.position()
(0.00,0.00)
>>> turtle.backward(30)
>>> turtle.position()
(-30.00,0.00)
```

`turtle.right(angle)`
`turtle.rt(angle)`

参数 angle – 一个数值 (整型或浮点型)

海龟右转 *angle* 个单位。(单位默认为角度, 但可通过 `degrees()` 和 `radians()` 函数改变设置。)角度的正负由海龟模式确定, 参见 `mode()`。

```
>>> turtle.heading()
22.0
>>> turtle.right(45)
>>> turtle.heading()
337.0
```

`turtle.left(angle)`
`turtle.lt(angle)`

参数 angle – 一个数值 (整型或浮点型)

海龟左转 *angle* 个单位。(单位默认为角度, 但可通过 `degrees()` 和 `radians()` 函数改变设置。)角度的正负由海龟模式确定, 参见 `mode()`。

```
>>> turtle.heading()
22.0
>>> turtle.left(45)
>>> turtle.heading()
67.0
```

`turtle.goto(x, y=None)`

```

turtle.setpos(x, y=None)
turtle.setposition(x, y=None)

```

参数

- **x** – 一个数值或数值对/向量
- **y** – 一个数值或 None

如果 *y* 为 None, *x* 应为一个表示坐标的数值对或 *Vec2D* 类对象 (例如 *pos()* 返回的对象).

海龟移动到一个绝对坐标。如果画笔已落下将会画线。不改变海龟的朝向。

```

>>> tp = turtle.pos()
>>> tp
(0.00,0.00)
>>> turtle.setpos(60,30)
>>> turtle.pos()
(60.00,30.00)
>>> turtle.setpos((20,80))
>>> turtle.pos()
(20.00,80.00)
>>> turtle.setpos(tp)
>>> turtle.pos()
(0.00,0.00)

```

```

turtle.setx(x)

```

参数 x – 一个数值 (整型或浮点型)

设置海龟的横坐标为 *x*, 纵坐标保持不变。

```

>>> turtle.position()
(0.00,240.00)
>>> turtle.setx(10)
>>> turtle.position()
(10.00,240.00)

```

```

turtle.sety(y)

```

参数 y – 一个数值 (整型或浮点型)

设置海龟的纵坐标为 *y*, 横坐标保持不变。

```

>>> turtle.position()
(0.00,40.00)
>>> turtle.sety(-10)
>>> turtle.position()
(0.00,-10.00)

```

```

turtle.setheading(to_angle)

```

```

turtle.seth(to_angle)

```

参数 to_angle – 一个数值 (整型或浮点型)

设置海龟的朝向为 *to_angle*。以下是以角度表示的几个常用方向：

标准模式	logo 模式
0 - 东	0 - 北
90 - 北	90 - 东
180 - 西	180 - 南
270 - 南	270 - 西

```
>>> turtle.setheading(90)
>>> turtle.heading()
90.0
```

`turtle.home()`

海龟移至初始坐标 (0,0)，并设置朝向为初始方向 (由海龟模式确定，参见 `mode()`)。

```
>>> turtle.heading()
90.0
>>> turtle.position()
(0.00, -10.00)
>>> turtle.home()
>>> turtle.position()
(0.00, 0.00)
>>> turtle.heading()
0.0
```

`turtle.circle(radius, extent=None, steps=None)`

参数

- **radius** – 一个数值
- **extent** – 一个数值 (或 None)
- **steps** – 一个整型数 (或 None)

绘制一个 *radius* 指定半径的圆。圆心在海龟左边 *radius* 个单位；*extent* 为一个夹角，用来决定绘制圆的一部分。如未指定 *extent* 则绘制整个圆。如果 *extent* 不是完整圆周，则以当前画笔位置为一个端点绘制圆弧。如果 *radius* 为正值则朝逆时针方向绘制圆弧，否则朝顺时针方向。最终海龟的朝向会依据 *extent* 的值而改变。

圆实际是以其内切正多边形来近似表示的，其边的数量由 *steps* 指定。如果未指定边数则会自动确定。此方法也可用来绘制正多边形。

```
>>> turtle.home()
>>> turtle.position()
(0.00, 0.00)
>>> turtle.heading()
0.0
>>> turtle.circle(50)
>>> turtle.position()
(-0.00, 0.00)
>>> turtle.heading()
0.0
>>> turtle.circle(120, 180) # draw a semicircle
>>> turtle.position()
(0.00, 240.00)
>>> turtle.heading()
180.0
```

`turtle.dot(size=None, *color)`

参数

- **size** – 一个整型数 ≥ 1 (如果指定)
- **color** – 一个颜色字符串或颜色数值元组

绘制一个直径为 *size*，颜色为 *color* 的圆点。如果 *size* 未指定，则直径取 `pensize+4` 和 `2*pensize` 中的较大值。

```
>>> turtle.home()
>>> turtle.dot()
>>> turtle.fd(50); turtle.dot(20, "blue"); turtle.fd(50)
>>> turtle.position()
(100.00,-0.00)
>>> turtle.heading()
0.0
```

turtle.stamp()

在海龟当前位置印制一个海龟形状。返回该印章的 `stamp_id`，印章可以通过调用 `clearstamp(stamp_id)` 来删除。

```
>>> turtle.color("blue")
>>> turtle.stamp()
11
>>> turtle.fd(50)
```

turtle.clearstamp(stampid)

参数 stampid—一个整型数，必须是之前 `stamp()` 调用的返回值

删除 `stampid` 指定的印章。

```
>>> turtle.position()
(150.00,-0.00)
>>> turtle.color("blue")
>>> astamp = turtle.stamp()
>>> turtle.fd(50)
>>> turtle.position()
(200.00,-0.00)
>>> turtle.clearstamp(astamp)
>>> turtle.position()
(200.00,-0.00)
```

turtle.clearstamps(n=None)

参数 n—一个整型数(或 None)

删除全部或前/后 `n` 个海龟印章。如果 `n` 为 None 则删除全部印章，如果 `n > 0` 则删除前 `n` 个印章，否则如果 `n < 0` 则删除后 `n` 个印章。

```
>>> for i in range(8):
...     turtle.stamp(); turtle.fd(30)
13
14
15
16
17
18
19
20
>>> turtle.clearstamps(2)
>>> turtle.clearstamps(-2)
>>> turtle.clearstamps()
```

turtle.undo()

撤消(或连续撤消)最近的一个(或多个)海龟动作。可撤消的次数由撤消缓冲区的大小决定。

```
>>> for i in range(4):
...     turtle.fd(50); turtle.lt(80)
...
>>> for i in range(8):
...     turtle.undo()
```

`turtle.speed(speed=None)`

参数 speed – 一个 0..10 范围内的整型数或速度字符串 (见下)

设置海龟移动的速度为 0..10 表示的整型数值。如未指定参数则返回当前速度。

如果输入数值大于 10 或小于 0.5 则速度设为 0。速度字符串与速度值的对应关系如下:

- “fastest” : 0 最快
- “fast” : 10 快
- “normal” : 6 正常
- “slow” : 3 慢
- “slowest” : 1 最慢

速度值从 1 到 10, 画线和海龟转向的动画效果逐级加快。

注意: `speed = 0` 表示 没有动画效果。`forward/back` 将使海龟向前/向后跳跃, 同样的 `left/right` 将使海龟立即改变朝向。

```
>>> turtle.speed()
3
>>> turtle.speed('normal')
>>> turtle.speed()
6
>>> turtle.speed(9)
>>> turtle.speed()
9
```

获取海龟的状态

`turtle.position()`

`turtle.pos()`

返回海龟当前的坐标 (x,y) (为 *Vec2D* 矢量类对象)。

```
>>> turtle.pos()
(440.00, -0.00)
```

`turtle.towards(x, y=None)`

参数

- **x** – 一个数值或数值对/矢量, 或一个海龟实例
- **y** – 一个数值——如果 *x* 是一个数值, 否则为 `None`

从海龟位置到由 (x,y), 矢量或另一海龟对应位置的连线的夹角。此数值依赖于海龟初始朝向 - 由 “standard” / “world” 或 “logo” 模式设置所决定)。

```
>>> turtle.goto(10, 10)
>>> turtle.towards(0, 0)
225.0
```

`turtle.xcor()`
返回海龟的 x 坐标。

```
>>> turtle.home()
>>> turtle.left(50)
>>> turtle.forward(100)
>>> turtle.pos()
(64.28, 76.60)
>>> print(round(turtle.xcor(), 5))
64.27876
```

`turtle.ycor()`
返回海龟的 y 坐标。

```
>>> turtle.home()
>>> turtle.left(60)
>>> turtle.forward(100)
>>> print(turtle.pos())
(50.00, 86.60)
>>> print(round(turtle.ycor(), 5))
86.60254
```

`turtle.heading()`
返回海龟当前的朝向 (数值依赖于海龟模式参见`mode()`)。

```
>>> turtle.home()
>>> turtle.left(67)
>>> turtle.heading()
67.0
```

`turtle.distance(x, y=None)`

参数

- **x** – 一个数值或数值对/矢量，或一个海龟实例
- **y** – 一个数值——如果 *x* 是一个数值，否则为 `None`

返回从海龟位置到由 (x,y)，适量或另一海龟对应位置的单位距离。

```
>>> turtle.home()
>>> turtle.distance(30, 40)
50.0
>>> turtle.distance((30, 40))
50.0
>>> joe = Turtle()
>>> joe.forward(77)
>>> turtle.distance(joe)
77.0
```

度量单位设置

`turtle.degrees(fullcircle=360.0)`

参数 **fullcircle** – 一个数值

设置角度的度量单位，即设置一个圆周为多少“度”。默认值为 360 度。

```
>>> turtle.home()
>>> turtle.left(90)
>>> turtle.heading()
90.0

Change angle measurement unit to grad (also known as gon,
grade, or gradian and equals 1/100-th of the right angle.)
>>> turtle.degrees(400.0)
>>> turtle.heading()
100.0
>>> turtle.degrees(360)
>>> turtle.heading()
90.0
```

`turtle.radians()`

设置角度的度量单位为弧度。其值等于 `degrees(2*math.pi)`。

```
>>> turtle.home()
>>> turtle.left(90)
>>> turtle.heading()
90.0
>>> turtle.radians()
>>> turtle.heading()
1.5707963267948966
```

画笔控制

绘图状态

`turtle.pendown()`

`turtle.pd()`

`turtle.down()`

画笔落下 – 移动时将画线。

`turtle.penup()`

`turtle.pu()`

`turtle.up()`

画笔抬起 – 移动时不画线。

`turtle.pensize(width=None)`

`turtle.width(width=None)`

参数 `width` – 一个正数值

设置线条的粗细为 `width` 或返回该值。如果 `resizemode` 设为 “auto” 并且 `turtleshape` 为多边形，该多边形也以同样粗细的线条绘制。如未指定参数，则返回当前的 `pensize`。

```
>>> turtle.pensize()
1
>>> turtle.pensize(10)    # from here on lines of width 10 are drawn
```

`turtle.pen(pen=None, **pendict)`

参数

- **pen** – 一个包含部分或全部下列键的字典
- **pendict** – 一个或多个以下列键为关键字的关键字参数

返回或设置画笔的属性，以一个包含以下键值对的“画笔字典”表示：

- “shown” : True/False
- “pendown” : True/False
- “pencolor” : 颜色字符串或颜色元组
- “fillcolor” : 颜色字符串或颜色元组
- “pensize” : 正数值
- “speed” : 0..10 范围内的数值
- “resizemode” : “auto” 或 “user” 或 “noresize”
- “stretchfactor” : (正数值, 正数值)
- “outline” : 正数值
- “tilt” : 数值

此字典可作为后续调用 `pen()` 时的参数，以恢复之前的画笔状态。另外还可将这些属性作为关键词参数提交。使用此方式可以用一条语句设置画笔的多个属性。

```
>>> turtle.pen(fillcolor="black", pencolor="red", pensize=10)
>>> sorted(turtle.pen().items())
[('fillcolor', 'black'), ('outline', 1), ('pencolor', 'red'),
 ('pendown', True), ('pensize', 10), ('resizemode', 'noresize'),
 ('shearfactor', 0.0), ('shown', True), ('speed', 9),
 ('stretchfactor', (1.0, 1.0)), ('tilt', 0.0)]
>>> penstate=turtle.pen()
>>> turtle.color("yellow", "")
>>> turtle.penup()
>>> sorted(turtle.pen().items())[:3]
[('fillcolor', ''), ('outline', 1), ('pencolor', 'yellow')]
>>> turtle.pen(penstate, fillcolor="green")
>>> sorted(turtle.pen().items())[:3]
[('fillcolor', 'green'), ('outline', 1), ('pencolor', 'red')]
```

`turtle.isdown()`

如果画笔落下返回 True，如果画笔抬起返回 False。

```
>>> turtle.penup()
>>> turtle.isdown()
False
>>> turtle.pendown()
>>> turtle.isdown()
True
```

颜色控制

`turtle.pencolor(*args)`

返回或设置画笔颜色。

允许以下四种输入格式：

`pencolor()` 返回以颜色描述字符串或元组（见示例）表示的当前画笔颜色。可用作其他 `color/pencolor/fillcolor` 调用的输入。

`pencolor(colorstring)` 设置画笔颜色为 `colorstring` 指定的 Tk 颜色描述字符串，例如 "red"、"yellow" 或 "#33cc8c"。

pencolor((r, g, b)) 设置画笔颜色为以 r, g, b 元组表示的 RGB 颜色。 r, g, b 的取值范围应为 $0..colormode$, $colormode$ 的值为 1.0 或 255 (参见 `colormode()`)。

pencolor(r, g, b)

设置画笔颜色为以 r, g, b 表示的 RGB 颜色。 r, g, b 的取值范围应为 $0..colormode$ 。

如果 `turtleshape` 为多边形, 该多边形轮廓也以新设置的画笔颜色绘制。

```
>>> colormode()
1.0
>>> turtle.pencolor()
'red'
>>> turtle.pencolor("brown")
>>> turtle.pencolor()
'brown'
>>> tup = (0.2, 0.8, 0.55)
>>> turtle.pencolor(tup)
>>> turtle.pencolor()
(0.2, 0.8, 0.5490196078431373)
>>> colormode(255)
>>> turtle.pencolor()
(51.0, 204.0, 140.0)
>>> turtle.pencolor('#32c18f')
>>> turtle.pencolor()
(50.0, 193.0, 143.0)
```

turtle.fillcolor(*args)

返回或设置填充颜色。

允许以下四种输入格式:

fillcolor() 返回以颜色描述字符串或元组 (见示例) 表示的当前填充颜色。可用作其他 `color/pencolor/fillcolor` 调用的输入。

fillcolor(colorstring) 设置填充颜色为 *colorstring* 指定的 Tk 颜色描述字符串, 例如 "red"、"yellow" 或 "#33cc8c"。

fillcolor((r, g, b)) 设置填充颜色为以 r, g, b 元组表示的 RGB 颜色。 r, g, b 的取值范围应为 $0..colormode$, $colormode$ 的值为 1.0 或 255 (参见 `colormode()`)。

fillcolor(r, g, b)

设置填充颜色为 r, g, b 表示的 RGB 颜色。 r, g, b 的取值范围应为 $0..colormode$ 。

如果 `turtleshape` 为多边形, 该多边形内部也以新设置的填充颜色填充。

```
>>> turtle.fillcolor("violet")
>>> turtle.fillcolor()
'violet'
>>> col = turtle.pencolor()
>>> col
(50.0, 193.0, 143.0)
>>> turtle.fillcolor(col)
>>> turtle.fillcolor()
(50.0, 193.0, 143.0)
>>> turtle.fillcolor('#ffffff')
>>> turtle.fillcolor()
(255.0, 255.0, 255.0)
```

turtle.color(*args)

返回或设置画笔颜色和填充颜色。

允许多种输入格式。使用如下 0 至 3 个参数:

color() 返回以一对颜色描述字符串或元组表示的当前画笔颜色和填充颜色, 两者可分别由 *pencolor()* 和 *fillcolor()* 返回。

color(colorstring), color((r,g,b)), color(r,g,b) 输入格式与 *pencolor()* 相同, 同时设置填充颜色和画笔颜色为指定的值。

color(colorstring1, colorstring2), color((r1,g1,b1), (r2,g2,b2))

相当于 *pencolor(colorstring1)* 加 *fillcolor(colorstring2)*, 使用其他输入格式的方法也与之类似。

如果 *turtleshape* 为多边形, 该多边形轮廓与填充也使用新设置的颜色。

```
>>> turtle.color("red", "green")
>>> turtle.color()
('red', 'green')
>>> color("#285078", "#a0c8f0")
>>> color()
((40.0, 80.0, 120.0), (160.0, 200.0, 240.0))
```

另参见: *Screen* 方法 *colormode()*。

填充

turtle.filling()

返回填充状态 (填充为 True, 否则为 False)。

```
>>> turtle.begin_fill()
>>> if turtle.filling():
...     turtle.pensize(5)
... else:
...     turtle.pensize(3)
```

turtle.begin_fill()

在绘制要填充的形状之前调用。

turtle.end_fill()

填充上次调用 *begin_fill()* 之后绘制的形状。

```
>>> turtle.color("black", "red")
>>> turtle.begin_fill()
>>> turtle.circle(80)
>>> turtle.end_fill()
```

更多绘图控制

turtle.reset()

从屏幕中删除海龟的绘图, 海龟回到原点并设置所有变量为默认值。

```
>>> turtle.goto(0,-22)
>>> turtle.left(100)
>>> turtle.position()
(0.00,-22.00)
>>> turtle.heading()
```

(下页继续)

(续上页)

```
100.0
>>> turtle.reset()
>>> turtle.position()
(0.00,0.00)
>>> turtle.heading()
0.0
```

`turtle.clear()`

从屏幕中删除指定海龟的绘图。不移动海龟。海龟的状态和位置以及其他海龟的绘图不受影响。

`turtle.write(arg, move=False, align="left", font=("Arial", 8, "normal"))`

参数

- **arg** – 要书写到 TurtleScreen 的对象
- **move** – True/False
- **align** – 字符串 “left”, “center” 或 “right”
- **font** – 一个三元组 (fontname, fontsize, fonttype)

书写文本 - *arg* 指定的字符串 - 到当前海龟位置, *align* 指定对齐方式 (“left”, “center” 或 right), *font* 指定字体。如果 *move* 为 True, 画笔会移动到文本的右下角。默认 *move* 为 False。

```
>>> turtle.write("Home = ", True, align="center")
>>> turtle.write((0,0), True)
```

海龟状态

可见性

`turtle.hideturtle()`

`turtle.ht()`

使海龟不可见。当你绘制复杂图形时这是个好主意, 因为隐藏海龟可显著加快绘制速度。

```
>>> turtle.hideturtle()
```

`turtle.showturtle()`

`turtle.st()`

使海龟可见。

```
>>> turtle.showturtle()
```

`turtle.isvisible()`

如果海龟显示返回 True, 如果海龟隐藏返回 False。

```
>>> turtle.hideturtle()
>>> turtle.isvisible()
False
>>> turtle.showturtle()
>>> turtle.isvisible()
True
```

外观

`turtle.shape(name=None)`

参数 name – 一个有效的形状名字符串

设置海龟形状为 *name* 指定的形状名，如未指定形状名则返回当前的形状名。*name* 指定的形状名应存在于 `TurtleScreen` 的 `shape` 字典中。多边形的形状初始时有以下几种：“arrow”，“turtle”，“circle”，“square”，“triangle”，“classic”。要了解如何处理形状请参看 `Screen` 方法 `register_shape()`。

```
>>> turtle.shape()
'classic'
>>> turtle.shape("turtle")
>>> turtle.shape()
'turtle'
```

`turtle.resizemode(rmode=None)`

参数 rmode – 字符串 “auto”，“user”，“noresize” 其中之一

设置大小调整模式为以下值之一：“auto”，“user”，“noresize”。如未指定 *rmode* 则返回当前的大小调整模式。不同的大小调整模式的效果如下：

- “auto”：根据画笔粗细值调整海龟的外观。
- “user”：根据拉伸因子和轮廓宽度 (outline) 值调整海龟的外观，两者是由 `shapesize()` 设置的。
- “noresize”：不调整海龟的外观大小。

大小调整模式 (“user”) 会在 `shapesize()` 带参数调用时生效。

```
>>> turtle.resizemode()
'noresize'
>>> turtle.resizemode("auto")
>>> turtle.resizemode()
'auto'
```

`turtle.shapesize(stretch_wid=None, stretch_len=None, outline=None)`

`turtle.turtlesize(stretch_wid=None, stretch_len=None, outline=None)`

参数

- **stretch_wid** – 正数值
- **stretch_len** – 正数值
- **outline** – 正数值

返回或设置画笔的属性 x/y-拉伸因子和/或轮廓。设置大小调整模式为 “user”。当且仅当大小调整模式设为 “user” 时海龟会基于其拉伸因子调整外观: *stretch_wid* 为垂直于其朝向的宽度拉伸因子, *stretch_len* 为平等于其朝向的长度拉伸因子, 决定形状轮廓线的粗细。

```
>>> turtle.shapesize()
(1.0, 1.0, 1)
>>> turtle.resizemode("user")
>>> turtle.shapesize(5, 5, 12)
>>> turtle.shapesize()
(5, 5, 12)
>>> turtle.shapesize(outline=8)
>>> turtle.shapesize()
(5, 5, 8)
```

`turtle.shearfactor` (*shear=None*)

参数 *shear* – 数值 (可选)

设置或返回当前的剪切因子。根据 *share* 指定的剪切因子即剪切角度的切线来剪切海龟形状。不改变海龟的朝向 (移动方向)。如未指定 *shear* 参数: 返回当前的剪切因子即剪切角度的切线, 与海龟朝向平行的线条将被剪切。

```
>>> turtle.shape("circle")
>>> turtle.shapesize(5,2)
>>> turtle.shearfactor(0.5)
>>> turtle.shearfactor()
0.5
```

`turtle.tilt` (*angle*)

参数 *angle* – 一个数值

海龟形状自其当前的倾角转动 *angle* 指定的角度, 但 不改变海龟的朝向 (移动方向)。

```
>>> turtle.reset()
>>> turtle.shape("circle")
>>> turtle.shapesize(5,2)
>>> turtle.tilt(30)
>>> turtle.fd(50)
>>> turtle.tilt(30)
>>> turtle.fd(50)
```

`turtle.settiltangle` (*angle*)

参数 *angle* – 一个数值

旋转海龟形状使其指向 *angle* 指定的方向, 忽略其当前的倾角, 不改变海龟的朝向 (移动方向)。

```
>>> turtle.reset()
>>> turtle.shape("circle")
>>> turtle.shapesize(5,2)
>>> turtle.settiltangle(45)
>>> turtle.fd(50)
>>> turtle.settiltangle(-45)
>>> turtle.fd(50)
```

3.1 版后已移除。

`turtle.tiltangle` (*angle=None*)

参数 *angle* – 一个数值 (可选)

设置或返回当前的倾角。如果指定 *angle* 则旋转海龟形状使其指向 *angle* 指定的方向, 忽略其当前的倾角。不改变海龟的朝向 (移动方向)。如果未指定 *angle*: 返回当前的倾角, 即海龟形状的方向和海龟朝向 (移动方向) 之间的夹角。

```
>>> turtle.reset()
>>> turtle.shape("circle")
>>> turtle.shapesize(5,2)
>>> turtle.tilt(45)
>>> turtle.tiltangle()
45.0
```

`turtle.shapetransform` (*t11=None, t12=None, t21=None, t22=None*)

参数

- **t11** – 一个数值 (可选)
- **t12** – 一个数值 (可选)
- **t21** – 一个数值 (可选)
- **t22** – 一个数值 (可选)

设置或返回海龟形状的当前变形矩阵。

如不指定任何矩阵元素，则返回以 4 元素元组表示的变形矩阵。否则使用指定元素设置变形矩阵改变海龟形状，矩阵第一排的值 **t11**, **t12**，第二排的值 **t21**, **t22**。行列式 $t11 * t22 - t12 * t21$ 的值不能为零，否则会出错。根据指定的矩阵修改拉伸因子，剪切因子和倾角。

```
>>> turtle = Turtle()
>>> turtle.shape("square")
>>> turtle.shapesize(4,2)
>>> turtle.shearfactor(-0.5)
>>> turtle.shapetransform()
(4.0, -1.0, -0.0, 2.0)
```

turtle.get_shapepoly()

返回以坐标值对元组表示的当前形状多边形。这可以用于定义一个新形状或一个复合形状的多个组成部分。

```
>>> turtle.shape("square")
>>> turtle.shapetransform(4, -1, 0, 2)
>>> turtle.get_shapepoly()
((50, -20), (30, 20), (-50, 20), (-30, -20))
```

使用事件

turtle.onclick(*fun, btn=1, add=None*)

参数

- **fun** – 一个函数，调用时将传入两个参数表示在画布上点击的坐标。
- **btn** – 鼠标按钮编号，默认值为 1 (鼠标左键)
- **add** – True 或 False – 如为 True 则将添加一个新绑定，否则将取代先前的绑定

将 *fun* 指定的函数绑定到鼠标点击此海龟事件。如果 *fun* 值为 None，则移除现有的绑定。以下为使用匿名海龟即过程式的示例：

```
>>> def turn(x, y):
...     left(180)
...
>>> onclick(turn)    # Now clicking into the turtle will turn it.
>>> onclick(None)    # event-binding will be removed
```

turtle.onrelease(*fun, btn=1, add=None*)

参数

- **fun** – 一个函数，调用时将传入两个参数表示在画布上点击的坐标。
- **btn** – 鼠标按钮编号，默认值为 1 (鼠标左键)
- **add** – True 或 False – 如为 True 则将添加一个新绑定，否则将取代先前的绑定

将 *fun* 指定的函数绑定到在此海龟上释放鼠标按键事件。如果 *fun* 值为 None，则移除现有的绑定。

```
>>> class MyTurtle(Turtle):
...     def glow(self, x, y):
...         self.fillcolor("red")
...     def unglow(self, x, y):
...         self.fillcolor("")
...
>>> turtle = MyTurtle()
>>> turtle.onclick(turtle.glow)      # clicking on turtle turns fillcolor red,
>>> turtle.onrelease(turtle.unglow)  # releasing turns it to transparent.
```

`turtle.ondrag(fun, btn=1, add=None)`

参数

- **fun** – 一个函数，调用时将传入两个参数表示在画布上点击的坐标。
- **btn** – 鼠标按钮编号，默认值为 1 (鼠标左键)
- **add** – True 或 False – 如为 True 则将添加一个新绑定，否则将取代先前的绑定

将 *fun* 指定的函数绑定到在此海龟上移动鼠标事件。如果 *fun* 值为 None，则移除现有的绑定。

注: 在海龟上移动鼠标事件之前应先发生在此海龟上点击鼠标事件。

```
>>> turtle.ondrag(turtle.goto)
```

在此之后点击并拖动海龟可在屏幕上手绘线条 (如果画笔为落下)。

特殊海龟方法

`turtle.begin_poly()`

开始记录多边形的顶点。当前海龟位置为多边形的第一个顶点。

`turtle.end_poly()`

停止记录多边形的顶点。当前海龟位置为多边形的最后一个顶点。它将连线到第一个顶点。

`turtle.get_poly()`

返回最新记录的多边形。

```
>>> turtle.home()
>>> turtle.begin_poly()
>>> turtle.fd(100)
>>> turtle.left(20)
>>> turtle.fd(30)
>>> turtle.left(60)
>>> turtle.fd(50)
>>> turtle.end_poly()
>>> p = turtle.get_poly()
>>> register_shape("myFavouriteShape", p)
```

`turtle.clone()`

创建并返回海龟的克隆体，具有相同的位置、朝向和海龟属性。

```
>>> mick = Turtle()
>>> joe = mick.clone()
```

`turtle.getturtle()`

`turtle.getpen()`

返回海龟对象自身。唯一合理的用法: 作为一个函数来返回“匿名海龟”:

```
>>> pet = getturtle()
>>> pet.fd(50)
>>> pet
<turtle.Turtle object at 0x...>
```

`turtle.getscreen()`

返回作为海龟绘图场所的 *TurtleScreen* 类对象。该对象将可调用 *TurtleScreen* 方法。

```
>>> ts = turtle.getscreen()
>>> ts
<turtle._Screen object at 0x...>
>>> ts.bgcolor("pink")
```

`turtle.setundobuffer(size)`

参数 *size* – 一个整型数值或 *None*

设置或禁用撤消缓冲区。如果 *size* 为一个整型数则将开辟一个指定大小的空缓冲区。*size* 表示可使用 *undo()* 方法/函数撤消的海龟命令的次数上限。如果 *size* 为 *None* 则禁用撤消缓冲区。

```
>>> turtle.setundobuffer(42)
```

`turtle.undobufferentries()`

返回撤销缓冲区里的条目数。

```
>>> while undobufferentries():
...     undo()
```

复合形状

要使用由多个不同颜色多边形构成的复合海龟形状，你必须明确地使用辅助类 *Shape*，具体步骤如下：

1. 创建一个空 *Shape* 对象，类型为 “compound”。
2. 按照需要使用 *addcomponent()* 方法向此对象添加多个部件。

例如：

```
>>> s = Shape("compound")
>>> poly1 = ((0,0), (10,-5), (0,10), (-10,-5))
>>> s.addcomponent(poly1, "red", "blue")
>>> poly2 = ((0,0), (10,-5), (-10,-5))
>>> s.addcomponent(poly2, "blue", "red")
```

3. 接下来将 *Shape* 对象添加到 *Screen* 对象的形状列表并使用它：

```
>>> register_shape("myshape", s)
>>> shape("myshape")
```

注解： *Shape* 类在 *register_shape()* 方法的内部以多种方式使用。应用程序编写者 只有在使用上述的复合形状时才需要处理 *Shape* 类。

24.1.4 TurtleScreen/Screen 方法及对应函数

本节中的大部分示例都使用 TurtleScreen 类的一个实例，命名为 `screen`。

窗口控制

`turtle.bgcolor(*args)`

参数 `args` – 一个颜色字符串或三个取值范围 0..`colormode` 内的数值或一个取值范围相同的数值 3 元组

设置或返回 TurtleScreen 的背景颜色。

```
>>> screen.bgcolor("orange")
>>> screen.bgcolor()
'orange'
>>> screen.bgcolor("#800080")
>>> screen.bgcolor()
(128.0, 0.0, 128.0)
```

`turtle.bgpic(picname=None)`

参数 `picname` – 一个字符串, gif-文件名, "nopic", 或 None

设置背景图片或返回当前背景图片名称。如果 `picname` 为一个文件名，则将相应图片设为背景。如果 `picname` 为 "nopic"，则删除当前背景图片。如果 `picname` 为 None，则返回当前背景图片文件名。:

```
>>> screen.bgpic()
'nopic'
>>> screen.bgpic("landscape.gif")
>>> screen.bgpic()
"landscape.gif"
```

`turtle.clear()`

`turtle.clearscreen()`

从中删除所有海龟的全部绘图。将已清空的 TurtleScreen 重置为初始状态: 白色背景，无背景片，无事件绑定并启用追踪。

注解: 此 TurtleScreen 方法作为全局函数时只有一个名字 `clearscreen`。全局函数 `clear` 所对应的是 Turtle 方法 `clear`。

`turtle.reset()`

`turtle.resetscreen()`

重置屏幕上的所有海龟为其初始状态。

注解: 此 TurtleScreen 方法作为全局函数时只有一个名字 `resetscreen`。全局函数 `reset` 所对应的是 Turtle 方法 `reset`。

`turtle.screensize(canvwidth=None, canvheight=None, bg=None)`

参数

- **canvwidth** – 正整型数，以像素表示画布的新宽度值
- **canvheight** – 正整型数，以像素表示画面的新高度值

- **bg** – 颜色字符串或颜色元组，新的背景颜色

如未指定任何参数，则返回当前的 (canvaswidth, canvasheight)。否则改变作为海龟绘图场所的画布大小。不改变绘图窗口。要观察画布的隐藏区域，可以使用滚动条。通过此方法可以令之前绘制于画布之外的图形变为可见。

```
>>> screen.screensize()
(400, 300)
>>> screen.screensize(2000,1500)
>>> screen.screensize()
(2000, 1500)
```

也可以用来寻找意外逃走的海龟;-)

`turtle.setworldcoordinates (llx, lly, urx, ury)`

参数

- **llx** – 一个数值, 画布左下角的 x-坐标
- **lly** – 一个数值, 画布左下角的 y-坐标
- **urx** – 一个数值, 画面右上角的 x-坐标
- **ury** – 一个数值, 画布右上角的 y-坐标

设置用户自定义坐标系并在必要时切换模式为“world”。这会执行一次 `screen.reset()`。如果“world”模式已激活，则所有图形将根据新的坐标系重绘。

注意: 在用户自定义坐标系中，角度可能显得扭曲。

```
>>> screen.reset()
>>> screen.setworldcoordinates(-50,-7.5,50,7.5)
>>> for _ in range(72):
...     left(10)
...
>>> for _ in range(8):
...     left(45); fd(2)    # a regular octagon
```

动画控制

`turtle.delay (delay=None)`

参数 delay – 正整型数

设置或返回以毫秒数表示的延迟值 *delay*。(这约等于连续两次画布刷新的间隔时间。)绘图延迟越长，动画速度越慢。

可选参数:

```
>>> screen.delay()
10
>>> screen.delay(5)
>>> screen.delay()
5
```

`turtle.tracer (n=None, delay=None)`

参数

- **n** – 非负整型数

- **delay** – 非负整型数

启用/禁用海龟动画并设置刷新图形的延迟时间。如果指定 *n* 值，则只有每第 *n* 次屏幕刷新会实际执行。(可被用来加速复杂图形的绘制。) 如果调用时不带参数，则返回当前保存的 *n* 值。第二个参数设置延迟值 (参见 `delay()`)。

```
>>> screen.tracer(8, 25)
>>> dist = 2
>>> for i in range(200):
...     fd(dist)
...     rt(90)
...     dist += 2
```

`turtle.update()`

执行一次 TurtleScreen 刷新。在禁用追踪时使用。

另参见 RawTurtle/Turtle 方法 `speed()`。

使用屏幕事件

`turtle.listen(xdummy=None, ydummy=None)`

设置焦点到 TurtleScreen (以便接收按键事件)。使用两个 Dummy 参数以便能够传递 `listen()` 给 `onclick` 方法。

`turtle.onkey(fun, key)`

`turtle.onkeyrelease(fun, key)`

参数

- **fun** – 一个无参数的函数或 None
- **key** – 一个字符串: 键 (例如 “a”) 或键标 (例如 “space”)

绑定 *fun* 指定的函数到按键释放事件。如果 *fun* 值为 None，则移除事件绑定。注: 为了能够注册按键事件，TurtleScreen 必须得到焦点。(参见 `method listen()` 方法。)

```
>>> def f():
...     fd(50)
...     lt(60)
...
>>> screen.onkey(f, "Up")
>>> screen.listen()
```

`turtle.onkeypress(fun, key=None)`

参数

- **fun** – 一个无参数的函数或 None
- **key** – 一个字符串: 键 (例如 “a”) 或键标 (例如 “space”)

绑定 *fun* 指定的函数到指定键的按下事件。如未指定键则绑定到任意键的按下事件。注: 为了能够注册按键事件，必须得到焦点。(参见 `listen()` 方法。)

```
>>> def f():
...     fd(50)
...
>>> screen.onkey(f, "Up")
>>> screen.listen()
```

`turtle.onclick(fun, btn=1, add=None)`

`turtle.onscreenclick(fun, btn=1, add=None)`

参数

- **fun** – 一个函数，调用时将传入两个参数表示在画布上点击的坐标。
- **btn** – 鼠标按钮编号，默认值为 1 (鼠标左键)
- **add** – True 或 False – 如为 True 则将添加一个新绑定，否则将取代先前的绑定

绑定 *fun* 指定的函数到鼠标点击屏幕事件。如果 *fun* 值为 None，则移除现有的绑定。

以下示例使用一个 TurtleScreen 实例 *screen* 和一个 Turtle 实例 *turtle*:

```
>>> screen.onclick(turtle.goto) # Subsequently clicking into the TurtleScreen will
>>>                               # make the turtle move to the clicked point.
>>> screen.onclick(None)        # remove event binding again
```

注解：此 TurtleScreen 方法作为全局函数时只有一个名字 `onscreenclick`。全局函数 `onclick` 所对应的是 Turtle 方法 `onclick`。

`turtle.ontimer(fun, t=0)`

参数

- **fun** – 一个无参数的函数
- **t** – 一个数值 ≥ 0

安装一个计时器，在 *t* 毫秒后调用 *fun* 函数。

```
>>> running = True
>>> def f():
...     if running:
...         fd(50)
...         lt(60)
...         screen.ontimer(f, 250)
>>> f()    ### makes the turtle march around
>>> running = False
```

`turtle.mainloop()`

`turtle.done()`

开始事件循环 - 调用 Tkinter 的 `mainloop` 函数。必须作为一个海龟绘图程序的结束语句。如果一个脚本是在以 -n 模式 (无子进程) 启动的 IDLE 中运行时 不可使用 - 用于实现海龟绘图的交互功能。:

```
>>> screen.mainloop()
```

输入方法

`turtle.textinput(title, prompt)`

参数

- **title** – 字符串
- **prompt** – 字符串

弹出一个对话框窗口用来输入一个字符串。形参 *title* 为对话框窗口的标题，*prompt* 为一条文本，通常用来提示要输入什么信息。返回输入的字符串。如果对话框被取消则返回 None。:

```
>>> screen.textinput("NIM", "Name of first player:")
```

`turtle.numinput` (*title*, *prompt*, *default=None*, *minval=None*, *maxval=None*)

参数

- **title** – 字符串
- **prompt** – 字符串
- **default** – 数值 (可选)
- **minval** – 数值 (可选)
- **maxval** – 数值 (可选)

弹出一个对话框窗口用来输入一个数值。*title* 为对话框窗口的标题, *prompt* 为一条文本, 通常用来描述要输入的数值信息。*default*: 默认值, *minval*: 可输入的最小值, *maxval*: 可输入的最大值。输入数值的必须在指定的 *minval* .. *maxval* 范围之内, 否则将给出一条提示, 对话框保持打开等待修改。返回输入的数值。如果对话框被取消则返回 `None`。

```
>>> screen.numinput("Poker", "Your stakes:", 1000, minval=10, maxval=10000)
```

设置与特殊方法

`turtle.mode` (*mode=None*)

参数 *mode* – 字符串 “standard”, “logo” 或 “world” 其中之一

设置海龟模式 (“standard”, “logo” 或 “world”) 并执行重置。如未指定模式则返回当前的模式。

“standard” 模式与旧的 *turtle* 兼容。” logo” 模式与大部分 Logo 海龟绘图兼容。” world” 模式使用用户自定义的 “世界坐标系”。**注意:** 在此模式下, 如果 *x/y* 单位比率不等于 1 则角度会显得扭曲。

模式	初始海龟朝向	正数角度
“standard”	朝右 (东)	逆时针
“logo”	朝上 (北)	顺时针

```
>>> mode("logo")    # resets turtle heading to north
>>> mode()
'logo'
```

`turtle.colormode` (*cmode=None*)

参数 *cmode* – 数值 1.0 或 255 其中之一

返回颜色模式或将其设为 1.0 或 255。构成颜色三元组的 *r, g, b* 数值必须在 0..*cmode* 范围之内。

```
>>> screen.colormode(1)
>>> turtle.pencolor(240, 160, 80)
Traceback (most recent call last):
...
TurtleGraphicsError: bad color sequence: (240, 160, 80)
>>> screen.colormode()
1.0
>>> screen.colormode(255)
>>> screen.colormode()
255
>>> turtle.pencolor(240, 160, 80)
```

`turtle.getcanvas()`

返回此 `TurtleScreen` 的 `Canvas` 对象。供了解 Tkinter 的 `Canvas` 对象内部机理的人士使用。

```
>>> cv = screen.getcanvas()
>>> cv
<turtle.ScrolledCanvas object ...>
```

`turtle.getshapes()`

返回所有当前可用海龟形状列表。

```
>>> screen.getshapes()
['arrow', 'blank', 'circle', ..., 'turtle']
```

`turtle.register_shape(name, shape=None)`

`turtle.addshape(name, shape=None)`

调用此函数有三种不同方式:

- (1) *name* 为一个 gif 文件的文件名, *shape* 为 `None`: 安装相应的图像形状。:

```
>>> screen.register_shape("turtle.gif")
```

注解: 当海龟转向时图像形状不会转动, 因此无法显示海龟的朝向!

- (2) *name* 为指定的字符串, *shape* 为由坐标值对构成的元组: 安装相应的多边形形状。

```
>>> screen.register_shape("triangle", ((5,-3), (0,5), (-5,-3)))
```

- (3) *name* 为指定的字符串, 为一个 (复合) `Shape` 类对象: 安装相应的复合形状。

将一个海龟形状加入 `TurtleScreen` 的形状列表。只有这样注册过的形状才能通过执行 `shape(shapename)` 命令来使用。

`turtle.turtles()`

返回屏幕上的海龟列表。

```
>>> for turtle in screen.turtles():
...     turtle.color("red")
```

`turtle.window_height()`

返回海龟窗口的高度。:

```
>>> screen.window_height()
480
```

`turtle.window_width()`

返回海龟窗口的宽度。:

```
>>> screen.window_width()
640
```

Screen 专有方法, 而非继承自 TurtleScreen

`turtle.bye()`

关闭海龟绘图窗口。

`turtle.exitonclick()`

将 `bye()` 方法绑定到 `Screen` 上的鼠标点击事件。

如果配置字典中 “`using_IDLE`” 的值为 `False` (默认值) 则同时进入主事件循环。注: 如果启动 `IDLE` 时使用了 `-n` 开关 (无子进程), `turtle.cfg` 中此数值应设为 `True`。在此情况下 `IDLE` 本身的主事件循环同样会作用于客户脚本。

`turtle.setup (width=_CFG["width"], height=_CFG["height"], startx=_CFG["leftright"], starty=_CFG["topbottom"])`

设置主窗口的大小和位置。默认参数值保存在配置字典中, 可通过 `turtle.cfg` 文件进行修改。

参数

- **width** – 如为一个整型数值, 表示大小为多少像素, 如为一个浮点数值, 则表示屏幕的占比; 默认为屏幕的 50%
- **height** – 如为一个整型数值, 表示高度为多少像素, 如为一个浮点数值, 则表示屏幕的占比; 默认为屏幕的 75%
- **startx** – 如为正值, 表示初始位置距离屏幕左边缘多少像素, 负值表示距离右边缘, `None` 表示窗口水平居中
- **starty** – 如为正值, 表示初始位置距离屏幕上边缘多少像素, 负值表示距离下边缘, `None` 表示窗口垂直居中

```
>>> screen.setup (width=200, height=200, startx=0, starty=0)
>>>             # sets window to 200x200 pixels, in upper left of screen
>>> screen.setup(width=.75, height=0.5, startx=None, starty=None)
>>>             # sets window to 75% of screen by 50% of screen and centers
```

`turtle.title (titlestring)`

参数 **titlestring** – 一个字符串, 显示为海龟绘图窗口的标题栏文本

设置海龟窗口标题为 *titlestring* 指定的文本。

```
>>> screen.title("Welcome to the turtle zoo!")
```

24.1.5 公共类

class `turtle.RawTurtle (canvas)`

class `turtle.RawPen (canvas)`

参数 **canvas** – 一个 `tkinter.Canvas`, `ScrolledCanvas` 或 `TurtleScreen` 类对象

创建一个海龟。海龟对象具有 “`Turtle/RawTurtle` 方法” 一节所述的全部方法。

class `turtle.Turtle`

`RawTurtle` 的子类, 具有相同的接口, 但其绘图场所为默认的 `Screen` 类对象, 在首次使用时自动创建。

class `turtle.TurtleScreen (cv)`

参数 **cv** – 一个 `tkinter.Canvas` 类对象

提供面向屏幕的方法例如 `setbg()` 等。说明见上文。

class `turtle.Screen`

`TurtleScreen` 的子类, 增加了四个方法。

class `turtle.ScrolledCanvas (master)`

参数 **master** – 可容纳 `ScrolledCanvas` 的 `Tkinter` 部件, 即添加了滚动条的 `Tkinter-canvas`

由 `Screen` 类使用，使其能够自动提供一个 `ScrolledCanvas` 作为海龟的绘图场所。

`class turtle.Shape (type_, data)`

参数 `type_` – 字符串 “polygon”，“image”，“compound” 其中之一
实现形状的数据结构。`(type_, data)` 必须遵循以下定义：

<code>type_</code>	<code>data</code>
“polygon”	一个多边形元组，即由坐标值对构成的元组
“image”	一个图片 (此形式仅限内部使用!)
“compound”	None (复合形状必须使用 <code>addcomponent()</code> 方法来构建)

`addcomponent (poly, fill, outline=None)`

参数

- `poly` – 一个多边形，即由数值对构成的元组
- `fill` – 一种颜色，将用来填充 `poly` 指定的多边形
- `outline` – 一种颜色，用于多边形的轮廓 (如有指定)

示例：

```
>>> poly = ((0,0), (10,-5), (0,10), (-10,-5))
>>> s = Shape("compound")
>>> s.addcomponent(poly, "red", "blue")
>>> # ... add more components and then use register_shape()
```

参见复合形状。

`class turtle.Vec2D (x, y)`

一个二维矢量类，用来作为实现海龟绘图的辅助类。也可能在海龟绘图程序中使用。派生自元组，因此矢量也属于元组！

提供的运算 (`a, b` 为矢量, `k` 为数值)：

- `a + b` 矢量加法
- `a - b` 矢量减法
- `a * b` 内积
- `k * a` 和 `a * k` 与标量相乘
- `abs(a)` `a` 的绝对值
- `a.rotate(angle)` 旋转

24.1.6 帮助与配置

如何使用帮助

`Screen` 和 `Turtle` 类的公用方法以文档字符串提供了详细的文档。因此可以利用 Python 帮助工具获取这些在线帮助信息：

- 当使用 IDLE 时，输入函数/方法调用将弹出工具提示显示其签名和文档字符串的头几行。
- 对文法或函数调用 `help()` 将显示其文档字符串：

```
>>> help(Screen.bgcolor)
Help on method bgcolor in module turtle:

bgcolor(self, *args) unbound turtle.Screen method
    Set or return backgroundcolor of the TurtleScreen.

    Arguments (if given): a color string or three numbers
    in the range 0..colormode or a 3-tuple of such numbers.

    >>> screen.bgcolor("orange")
    >>> screen.bgcolor()
    "orange"
    >>> screen.bgcolor(0.5,0,0.5)
    >>> screen.bgcolor()
    "#800080"

>>> help(Turtle.penup)
Help on method penup in module turtle:

penup(self) unbound turtle.Turtle method
    Pull the pen up -- no drawing when moving.

    Aliases: penup | pu | up

    No argument

    >>> turtle.penup()
```

- 方法对应函数的文档字符串的形式会有一些修改:

```
>>> help(bgcolor)
Help on function bgcolor in module turtle:

bgcolor(*args)
    Set or return backgroundcolor of the TurtleScreen.

    Arguments (if given): a color string or three numbers
    in the range 0..colormode or a 3-tuple of such numbers.

    Example::

    >>> bgcolor("orange")
    >>> bgcolor()
    "orange"
    >>> bgcolor(0.5,0,0.5)
    >>> bgcolor()
    "#800080"

>>> help(penup)
Help on function penup in module turtle:

penup()
    Pull the pen up -- no drawing when moving.

    Aliases: penup | pu | up
```

(下页继续)

(续上页)

```
No argument
```

```
Example:
```

```
>>> penup()
```

这些修改版文档字符串是在导入时与方法对应函数的定义一起自动生成的。

文档字符串翻译为不同的语言

可使用工具创建一个字典，键为方法名，值为 `Screen` 和 `Turtle` 类公共方法的文档字符串。

```
turtle.write_docstringdict(filename="turtle_docstringdict")
```

参数 filename – 一个字符串，表示文件名

创建文档字符串字典并将其写入 `filename` 指定的 Python 脚本文件。此函数必须显示地调用 (海龟绘图类并不使用此函数)。文档字符串字典将被写入到 Python 脚本文件 `filename.py`。该文件可作为模板用来将文档字符串翻译为不同语言。

如果你 (或你的学生) 想使用本国语言版本的 `turtle` 在线帮助，你必须翻译文档字符串并保存结果文件，例如 `turtle_docstringdict_german.py`。

如果你在 `turtle.cfg` 文件中加入了相应的条目，此字典将在导入模块时被读取并替代原有的英文版文档字符串。

在撰写本文档时已经有了德语和意大利语版的文档字符串字典。(更多需求请联系 glingl@aon.at)

如何配置 Screen 和 Turtle

内置的默认配置是模仿旧 `turtle` 模块的外观和行为，以便尽可能地与其保持兼容。

如果你想使用不同的配置，以便更好地反映此模块的特性或是更适合你的需求，例如在课堂中使用，你可以准备一个配置文件 `turtle.cfg`，该文件将在导入模块时被读取并根据其中的设定修改模块配置。

内置的配置对应以下的 `turtle.cfg`:

```
width = 0.5
height = 0.75
leftright = None
topbottom = None
canvwidth = 400
canvheight = 300
mode = standard
colormode = 1.0
delay = 10
undobuffersize = 1000
shape = classic
pencolor = black
fillcolor = black
resizemode = noresize
visible = True
language = english
exampleturtle = turtle
examplescreen = screen
title = Python Turtle Graphics
using_IDLE = False
```

选定条目的简短说明:

- 开头的四行对应 `Screen.setup()` 方法的参数。
- 第 5 和 6 行对应 `Screen.screensize()` 方法的参数。
- *shape* 可以是任何内置形状, 即: `arrow`, `turtle` 等。更多信息可用 `help(shape)` 查看。
- 如果你想使用无填充色 (即令海龟变透明), 你必须写 `fillcolor = ""` (但 `cfg` 文件中所有非空字符串都不可加引号)。
- 如果你想令海龟反映其状态, 你必须使用 `resizemode = auto`。
- 如果你设置语言例如 `language = italian` 则文档字符串字典 `turtle_docstringdict_italian.py` 将在导入模块时被加载 (如果导入路径即 `turtle` 的目录中存在此文件)。
- *exampleturtle* 和 *examplescreen* 条目定义了相应对象在文档字符串中显示的名称。方法文档字符串转换为函数文档字符串时将从文档字符串中删去这些名称。
- *using_IDLE*: 如果你经常使用 IDLE 并启用其 `-n` 开关 (“无子进程”) 则应将此项设为 `True`, 这将阻止 `exitonclick()` 进入主事件循环。

`turtle.cfg` 文件可以保存于 `turtle` 所在目录, 当前工作目录也可以有一个同名文件。后者会重载覆盖前者的设置。

`Lib/turtledemo` 目录中也有一个 `turtle.cfg` 文件。你可以将其作为示例进行研究, 并在运行演示时查看其作用效果 (但最好不要在演示查看器中运行)。

24.1.7 turtledemo — 演示脚本集

`turtledemo` 包汇集了一组演示脚本。这些脚本可以通过以下命令打开所提供的演示查看器运行和查看:

```
python -m turtledemo
```

此外, 你也可以单独运行其中的演示脚本。例如, :

```
python -m turtledemo.bytedesign
```

`turtledemo` 包目录中的内容:

- 一个演示查看器 `__main__.py`, 可用来查看脚本的源码并即时运行。
- 多个脚本文件, 演示 `turtle` 模块的不同特性。所有示例可通过 **Examples** 菜单打开。也可以单独运行每个脚本。
- 一个 `turtle.cfg` 文件, 作为说明如何编写并使用模块配置文件的示例模板。

演示脚本清单如下:

名称	描述	相关特性
bytedesign	复杂的传统海龟绘图模式	<code>tracer()</code> , <code>delay</code> , <code>update()</code>
chaos	绘制 Verhulst 动态模型, 演示通过计算机的运算可能会生成令人惊叹的结果	世界坐标系
clock	绘制模拟时钟显示本机的当前时间	海龟作为表针, <code>ontimer</code>
colormixer	试验 r, g, b 颜色模式	<code>ondrag()</code> 当鼠标拖动
forest	绘制 3 棵广度优先树	随机化
fractalcurves	绘制 Hilbert & Koch 曲线	递归
lindenmayer	文化数学 (印度装饰艺术)	L-系统
minimal_hanoi	汉诺塔	矩形海龟作为汉诺盘 (<code>shape</code> , <code>shapeseize</code>)
nim	玩经典的“尼姆”游戏, 开始时有三堆小棒, 与电脑对战。	海龟作为小棒, 事件驱动 (鼠标, 键盘)
paint	超极简主义绘画程序	<code>onclick()</code> 当鼠标点击
peace	初级技巧	海龟: 外观与动画
penrose	非周期性地使用风筝和飞镖形状铺满平面	<code>stamp()</code> 印章
planet_and_moon	模拟引力系统	复合开关, <code>Vec2D</code> 类
round_dance	两两相对并不断旋转舞蹈的海龟	复合形状, <code>clone</code> <code>shapeseize</code> , <code>tilt</code> , <code>get_shapepoly</code> , <code>update</code>
sorting_animate	动态演示不同的排序方法	简单对齐, 随机化
tree	一棵 (图形化的) 广度优先树 (使用生成器)	<code>clone()</code> 克隆
two_canvases	简单设计	两块画布上的海龟
wikipedia	一个来自介绍海龟绘图的维基百科文章的图案	<code>clone()</code> , <code>undo()</code>
yinyang	另一个初级示例	<code>circle()</code> 画圆

祝你玩得开心!

24.1.8 Python 2.6 之后的变化

- `Turtle.tracer()`, `Turtle.window_width()` 和 `Turtle.window_height()` 方法已被去除。具有这些名称和功能的方法现在只限于 `Screen` 类的方法。但其对应的函数仍然可用。(实际上在 Python 2.6 中这些方法就已经只是从对应的 `TurtleScreen/Screen` 类的方法复制而来。)
- `Turtle.fill()` 方法已被去除。`begin_fill()` 和 `end_fill()` 的行为则有细微改变: 现在每个填充过程必须以一个 `end_fill()` 调用来结束。
- 新增了一个 `Turtle.filling()` 方法。该方法返回一个布尔值: 如果填充过程正在进行为 `True`, 否则为 `False`。此行为相当于 Python 2.6 中不带参数的 `fill()` 调用。

24.1.9 Python 3.0 之后的变化

- 新增了 `Turtle.shearfactor()`, `Turtle.shapetransform()` 和 `Turtle.get_shapepoly()` 方法。这样就可以使用所有标准线性变换来调整海龟形状。`Turtle.tiltangle()` 的功能已被加强: 现在可被用来获取或设置倾角。`Turtle.settiltangle()` 已弃用。
- 新增了 `Screen.onkeypress()` 方法作为对 `Screen.onkey()` 的补充, 实际就是将行为绑定到 `keyrelease` 事件。后者相应增加了一个别名: `Screen.onkeyrelease()`。
- 新增了 `Screen.mainloop()` 方法。这样当仅需使用 `Screen` 和 `Turtle` 对象时不需要再额外导入 `mainloop()`。
- 新增了两个方法 `Screen.textinput()` 和 `Screen.numinput()`。用来弹出对话框接受输入并分别返回字符串和数值。

- 两个新的示例脚本 `tdemo_nim.py` 和 `tdemo_round_dance.py` 被加入到 `Lib/turtledemo` 目录中。

24.2 cmd — 支持面向行的命令解释器

源代码: [Lib/cmd.py](#)

The `Cmd` class provides a simple framework for writing line-oriented command interpreters. These are often useful for test harnesses, administrative tools, and prototypes that will later be wrapped in a more sophisticated interface.

class `cmd.Cmd` (*completekey='tab', stdin=None, stdout=None*)

A `Cmd` instance or subclass instance is a line-oriented interpreter framework. There is no good reason to instantiate `Cmd` itself; rather, it's useful as a superclass of an interpreter class you define yourself in order to inherit `Cmd`'s methods and encapsulate action methods.

The optional argument *completekey* is the *readline* name of a completion key; it defaults to `Tab`. If *completekey* is not `None` and *readline* is available, command completion is done automatically.

The optional arguments *stdin* and *stdout* specify the input and output file objects that the `Cmd` instance or subclass instance will use for input and output. If not specified, they will default to `sys.stdin` and `sys.stdout`.

If you want a given *stdin* to be used, make sure to set the instance's `use_rawinput` attribute to `False`, otherwise *stdin* will be ignored.

24.2.1 Cmd Objects

`Cmd` 实例有下列方法:

`Cmd.cmdloop` (*intro=None*)

Repeatedly issue a prompt, accept input, parse an initial prefix off the received input, and dispatch to action methods, passing them the remainder of the line as argument.

The optional argument is a banner or intro string to be issued before the first prompt (this overrides the *intro* class attribute).

If the *readline* module is loaded, input will automatically inherit **bash**-like history-list editing (e.g. `Control-P` scrolls back to the last command, `Control-N` forward to the next one, `Control-F` moves the cursor to the right non-destructively, `Control-B` moves the cursor to the left non-destructively, etc.).

An end-of-file on input is passed back as the string `'EOF'`.

An interpreter instance will recognize a command name `foo` if and only if it has a method `do_foo()`. As a special case, a line beginning with the character `'?'` is dispatched to the method `do_help()`. As another special case, a line beginning with the character `'!'` is dispatched to the method `do_shell()` (if such a method is defined).

This method will return when the `postcmd()` method returns a true value. The *stop* argument to `postcmd()` is the return value from the command's corresponding `do_*()` method.

If completion is enabled, completing commands will be done automatically, and completing of commands args is done by calling `complete_foo()` with arguments *text*, *line*, *begidx*, and *endidx*. *text* is the string prefix we are attempting to match: all returned matches must begin with it. *line* is the current input line with leading whitespace removed, *begidx* and *endidx* are the beginning and ending indexes of the prefix text, which could be used to provide different completion depending upon which position the argument is in.

All subclasses of `Cmd` inherit a predefined `do_help()`. This method, called with an argument `'bar'`, invokes the corresponding method `help_bar()`, and if that is not present, prints the docstring of `do_bar()`, if available. With no argument, `do_help()` lists all available help topics (that is, all commands with corresponding `help_*()` methods or commands that have docstrings), and also lists any undocumented commands.

`Cmd.onecmd(str)`

Interpret the argument as though it had been typed in response to the prompt. This may be overridden, but should not normally need to be; see the `precmd()` and `postcmd()` methods for useful execution hooks. The return value is a flag indicating whether interpretation of commands by the interpreter should stop. If there is a `do_*()` method for the command `str`, the return value of that method is returned, otherwise the return value from the `default()` method is returned.

`Cmd.emptyline()`

Method called when an empty line is entered in response to the prompt. If this method is not overridden, it repeats the last nonempty command entered.

`Cmd.default(line)`

Method called on an input line when the command prefix is not recognized. If this method is not overridden, it prints an error message and returns.

`Cmd.completedefault(text, line, begidx, endidx)`

Method called to complete an input line when no command-specific `complete_*()` method is available. By default, it returns an empty list.

`Cmd.precmd(line)`

Hook method executed just before the command line `line` is interpreted, but after the input prompt is generated and issued. This method is a stub in `Cmd`; it exists to be overridden by subclasses. The return value is used as the command which will be executed by the `onecmd()` method; the `precmd()` implementation may re-write the command or simply return `line` unchanged.

`Cmd.postcmd(stop, line)`

Hook method executed just after a command dispatch is finished. This method is a stub in `Cmd`; it exists to be overridden by subclasses. `line` is the command line which was executed, and `stop` is a flag which indicates whether execution will be terminated after the call to `postcmd()`; this will be the return value of the `onecmd()` method. The return value of this method will be used as the new value for the internal flag which corresponds to `stop`; returning false will cause interpretation to continue.

`Cmd.preloop()`

Hook method executed once when `cmdloop()` is called. This method is a stub in `Cmd`; it exists to be overridden by subclasses.

`Cmd.postloop()`

Hook method executed once when `cmdloop()` is about to return. This method is a stub in `Cmd`; it exists to be overridden by subclasses.

Instances of `Cmd` subclasses have some public instance variables:

`Cmd.prompt`

The prompt issued to solicit input.

`Cmd.identchars`

The string of characters accepted for the command prefix.

`Cmd.lastcmd`

The last nonempty command prefix seen.

`Cmd.cmdqueue`

A list of queued input lines. The `cmdqueue` list is checked in `cmdloop()` when new input is needed; if it is nonempty, its elements will be processed in order, as if entered at the prompt.

Cmd.intro

A string to issue as an intro or banner. May be overridden by giving the `cmdloop()` method an argument.

Cmd.doc_header

The header to issue if the help output has a section for documented commands.

Cmd.misc_header

The header to issue if the help output has a section for miscellaneous help topics (that is, there are `help_*()` methods without corresponding `do_*()` methods).

Cmd.undoc_header

The header to issue if the help output has a section for undocumented commands (that is, there are `do_*()` methods without corresponding `help_*()` methods).

Cmd.ruler

The character used to draw separator lines under the help-message headers. If empty, no ruler line is drawn. It defaults to '='.

Cmd.use_rawinput

A flag, defaulting to true. If true, `cmdloop()` uses `input()` to display a prompt and read the next command; if false, `sys.stdout.write()` and `sys.stdin.readline()` are used. (This means that by importing `readline`, on systems that support it, the interpreter will automatically support **Emacs**-like line editing and command-history keystrokes.)

24.2.2 Cmd Example

The `cmd` module is mainly useful for building custom shells that let a user work with a program interactively.

This section presents a simple example of how to build a shell around a few of the commands in the `turtle` module.

Basic turtle commands such as `forward()` are added to a `Cmd` subclass with method named `do_forward()`. The argument is converted to a number and dispatched to the turtle module. The docstring is used in the help utility provided by the shell.

The example also includes a basic record and playback facility implemented with the `precmd()` method which is responsible for converting the input to lowercase and writing the commands to a file. The `do_playback()` method reads the file and adds the recorded commands to the `cmdqueue` for immediate playback:

```
import cmd, sys
from turtle import *

class TurtleShell(cmd.Cmd):
    intro = 'Welcome to the turtle shell.  Type help or ? to list commands.\n'
    prompt = '(turtle) '
    file = None

    # ----- basic turtle commands -----
    def do_forward(self, arg):
        'Move the turtle forward by the specified distance:  FORWARD 10'
        forward(*parse(arg))
    def do_right(self, arg):
        'Turn turtle right by given number of degrees:  RIGHT 20'
        right(*parse(arg))
    def do_left(self, arg):
        'Turn turtle left by given number of degrees:  LEFT 90'
        left(*parse(arg))
    def do_goto(self, arg):
        'Move turtle to an absolute position with changing orientation.  GOTO 100 200'
```

(下页继续)

(续上页)

```

        goto(*parse(arg))
    def do_home(self, arg):
        'Return turtle to the home position:  HOME'
        home()
    def do_circle(self, arg):
        'Draw circle with given radius an options extent and steps:  CIRCLE 50'
        circle(*parse(arg))
    def do_position(self, arg):
        'Print the current turtle position:  POSITION'
        print('Current position is %d %d\n' % position())
    def do_heading(self, arg):
        'Print the current turtle heading in degrees:  HEADING'
        print('Current heading is %d\n' % (heading(),))
    def do_color(self, arg):
        'Set the color:  COLOR BLUE'
        color(arg.lower())
    def do_undo(self, arg):
        'Undo (repeatedly) the last turtle action(s):  UNDO'
    def do_reset(self, arg):
        'Clear the screen and return turtle to center:  RESET'
        reset()
    def do_bye(self, arg):
        'Stop recording, close the turtle window, and exit:  BYE'
        print('Thank you for using Turtle')
        self.close()
        bye()
        return True

# ----- record and playback -----
    def do_record(self, arg):
        'Save future commands to filename:  RECORD rose.cmd'
        self.file = open(arg, 'w')
    def do_playback(self, arg):
        'Playback commands from a file:  PLAYBACK rose.cmd'
        self.close()
        with open(arg) as f:
            self.cmdqueue.extend(f.read().splitlines())
    def precmd(self, line):
        line = line.lower()
        if self.file and 'playback' not in line:
            print(line, file=self.file)
        return line
    def close(self):
        if self.file:
            self.file.close()
            self.file = None

def parse(arg):
    'Convert a series of zero or more numbers to an argument tuple'
    return tuple(map(int, arg.split()))

if __name__ == '__main__':
    TurtleShell().cmdloop()

```

Here is a sample session with the turtle shell showing the help functions, using blank lines to repeat commands, and the simple record and playback facility:

```
Welcome to the turtle shell.  Type help or ? to list commands.
```

```
(turtle) ?
```

```
Documented commands (type help <topic>):
```

```
=====
```

```
bye      color      goto      home      playback  record  right  
circle   forward  heading  left      position  reset   undo
```

```
(turtle) help forward
```

```
Move the turtle forward by the specified distance:  FORWARD 10
```

```
(turtle) record spiral.cmd
```

```
(turtle) position
```

```
Current position is 0 0
```

```
(turtle) heading
```

```
Current heading is 0
```

```
(turtle) reset
```

```
(turtle) circle 20
```

```
(turtle) right 30
```

```
(turtle) circle 40
```

```
(turtle) right 30
```

```
(turtle) circle 60
```

```
(turtle) right 30
```

```
(turtle) circle 80
```

```
(turtle) right 30
```

```
(turtle) circle 100
```

```
(turtle) right 30
```

```
(turtle) circle 120
```

```
(turtle) right 30
```

```
(turtle) circle 120
```

```
(turtle) heading
```

```
Current heading is 180
```

```
(turtle) forward 100
```

```
(turtle)
```

```
(turtle) right 90
```

```
(turtle) forward 100
```

```
(turtle)
```

```
(turtle) right 90
```

```
(turtle) forward 400
```

```
(turtle) right 90
```

```
(turtle) forward 500
```

```
(turtle) right 90
```

```
(turtle) forward 400
```

```
(turtle) right 90
```

```
(turtle) forward 300
```

```
(turtle) playback spiral.cmd
```

```
Current position is 0 0
```

```
Current heading is 0
```

```
Current heading is 180
```

```
(turtle) bye
```

```
Thank you for using Turtle
```

24.3 shlex — Simple lexical analysis

Source code: [Lib/shlex.py](#)

The `shlex` class makes it easy to write lexical analyzers for simple syntaxes resembling that of the Unix shell. This will often be useful for writing minilanguages, (for example, in run control files for Python applications) or for parsing quoted strings.

The `shlex` module defines the following functions:

`shlex.split(s, comments=False, posix=True)`

Split the string `s` using shell-like syntax. If `comments` is `False` (the default), the parsing of comments in the given string will be disabled (setting the `commenters` attribute of the `shlex` instance to the empty string). This function operates in POSIX mode by default, but uses non-POSIX mode if the `posix` argument is false.

注解: Since the `split()` function instantiates a `shlex` instance, passing `None` for `s` will read the string to split from standard input.

`shlex.quote(s)`

Return a shell-escaped version of the string `s`. The returned value is a string that can safely be used as one token in a shell command line, for cases where you cannot use a list.

This idiom would be unsafe:

```
>>> filename = 'somefile; rm -rf ~'
>>> command = 'ls -l {}'.format(filename)
>>> print(command) # executed by a shell: boom!
ls -l somefile; rm -rf ~
```

`quote()` lets you plug the security hole:

```
>>> command = 'ls -l {}'.format(quote(filename))
>>> print(command)
ls -l 'somefile; rm -rf ~'
>>> remote_command = 'ssh home {}'.format(quote(command))
>>> print(remote_command)
ssh home 'ls -l "'"somefile; rm -rf ~"'"'
```

The quoting is compatible with UNIX shells and with `split()`:

```
>>> remote_command = split(remote_command)
>>> remote_command
['ssh', 'home', "ls -l 'somefile; rm -rf ~'"]
>>> command = split(remote_command[-1])
>>> command
['ls', '-l', 'somefile; rm -rf ~']
```

3.3 新版功能.

The `shlex` module defines the following class:

class `shlex.shlex(instream=None, infile=None, posix=False, punctuation_chars=False)`

A `shlex` instance or subclass instance is a lexical analyzer object. The initialization argument, if present, specifies where to read characters from. It must be a file-/stream-like object with `read()` and `readline()` methods, or a string. If no argument is given, input will be taken from `sys.stdin`. The second optional argument is a filename string, which sets the initial value of the `infile` attribute. If the `instream` argument is omitted or equal

to `sys.stdin`, this second argument defaults to “`stdin`”. The *posix* argument defines the operational mode: when *posix* is not true (default), the *shlex* instance will operate in compatibility mode. When operating in POSIX mode, *shlex* will try to be as close as possible to the POSIX shell parsing rules. The *punctuation_chars* argument provides a way to make the behaviour even closer to how real shells parse. This can take a number of values: the default value, `False`, preserves the behaviour seen under Python 3.5 and earlier. If set to `True`, then parsing of the characters `() ; <> | &` is changed: any run of these characters (considered punctuation characters) is returned as a single token. If set to a non-empty string of characters, those characters will be used as the punctuation characters. Any characters in the *wordchars* attribute that appear in *punctuation_chars* will be removed from *wordchars*. See *Improved Compatibility with Shells* for more information.

在 3.6 版更改: The *punctuation_chars* parameter was added.

参见:

Module *configparser* Parser for configuration files similar to the Windows `.ini` files.

24.3.1 shlex Objects

A *shlex* instance has the following methods:

`shlex.get_token()`

Return a token. If tokens have been stacked using *push_token()*, pop a token off the stack. Otherwise, read one from the input stream. If reading encounters an immediate end-of-file, *eof* is returned (the empty string `('')` in non-POSIX mode, and `None` in POSIX mode).

`shlex.push_token(str)`

Push the argument onto the token stack.

`shlex.read_token()`

Read a raw token. Ignore the pushback stack, and do not interpret source requests. (This is not ordinarily a useful entry point, and is documented here only for the sake of completeness.)

`shlex.sourcehook(filename)`

When *shlex* detects a source request (see *source* below) this method is given the following token as argument, and expected to return a tuple consisting of a filename and an open file-like object.

Normally, this method first strips any quotes off the argument. If the result is an absolute pathname, or there was no previous source request in effect, or the previous source was a stream (such as `sys.stdin`), the result is left alone. Otherwise, if the result is a relative pathname, the directory part of the name of the file immediately before it on the source inclusion stack is prepended (this behavior is like the way the C preprocessor handles `#include "file.h"`).

The result of the manipulations is treated as a filename, and returned as the first component of the tuple, with *open()* called on it to yield the second component. (Note: this is the reverse of the order of arguments in instance initialization!)

This hook is exposed so that you can use it to implement directory search paths, addition of file extensions, and other namespace hacks. There is no corresponding ‘close’ hook, but a *shlex* instance will call the *close()* method of the sourced input stream when it returns EOF.

For more explicit control of source stacking, use the *push_source()* and *pop_source()* methods.

`shlex.push_source(newstream, newfile=None)`

Push an input source stream onto the input stack. If the filename argument is specified it will later be available for use in error messages. This is the same method used internally by the *sourcehook()* method.

`shlex.pop_source()`

Pop the last-pushed input source from the input stack. This is the same method used internally when the lexer reaches EOF on a stacked input stream.

`shlex.error_leader` (*infile=None, lineno=None*)

This method generates an error message leader in the format of a Unix C compiler error label; the format is `"%s", line %d: '`, where the `%s` is replaced with the name of the current source file and the `%d` with the current input line number (the optional arguments can be used to override these).

This convenience is provided to encourage `shlex` users to generate error messages in the standard, parseable format understood by Emacs and other Unix tools.

Instances of `shlex` subclasses have some public instance variables which either control lexical analysis or can be used for debugging:

`shlex.commenters`

The string of characters that are recognized as comment beginners. All characters from the comment beginner to end of line are ignored. Includes just `'#'` by default.

`shlex.wordchars`

The string of characters that will accumulate into multi-character tokens. By default, includes all ASCII alphanumeric characters and underscore. In POSIX mode, the accented characters in the Latin-1 set are also included. If `punctuation_chars` is not empty, the characters `~-./*?=<`, which can appear in filename specifications and command line parameters, will also be included in this attribute, and any characters which appear in `punctuation_chars` will be removed from `wordchars` if they are present there.

`shlex.whitespace`

Characters that will be considered whitespace and skipped. Whitespace bounds tokens. By default, includes space, tab, linefeed and carriage-return.

`shlex.escape`

Characters that will be considered as escape. This will be only used in POSIX mode, and includes just `'\'` by default.

`shlex.quotes`

Characters that will be considered string quotes. The token accumulates until the same quote is encountered again (thus, different quote types protect each other as in the shell.) By default, includes ASCII single and double quotes.

`shlex.escapedquotes`

Characters in `quotes` that will interpret escape characters defined in `escape`. This is only used in POSIX mode, and includes just `'\"'` by default.

`shlex.whitespace_split`

If `True`, tokens will only be split in whitespaces. This is useful, for example, for parsing command lines with `shlex`, getting tokens in a similar way to shell arguments. If this attribute is `True`, `punctuation_chars` will have no effect, and splitting will happen only on whitespaces. When using `punctuation_chars`, which is intended to provide parsing closer to that implemented by shells, it is advisable to leave `whitespace_split` as `False` (the default value).

`shlex.infile`

The name of the current input file, as initially set at class instantiation time or stacked by later source requests. It may be useful to examine this when constructing error messages.

`shlex.instream`

The input stream from which this `shlex` instance is reading characters.

`shlex.source`

This attribute is `None` by default. If you assign a string to it, that string will be recognized as a lexical-level inclusion request similar to the `source` keyword in various shells. That is, the immediately following token will be opened as a filename and input will be taken from that stream until EOF, at which point the `close()` method of that stream will be called and the input source will again become the original input stream. Source requests may be stacked any number of levels deep.

`shlex.debug`

If this attribute is numeric and 1 or more, a *shlex* instance will print verbose progress output on its behavior. If you need to use this, you can read the module source code to learn the details.

`shlex.lineno`

Source line number (count of newlines seen so far plus one).

`shlex.token`

The token buffer. It may be useful to examine this when catching exceptions.

`shlex.eof`

Token used to determine end of file. This will be set to the empty string (' '), in non-POSIX mode, and to `None` in POSIX mode.

`shlex.punctuation_chars`

Characters that will be considered punctuation. Runs of punctuation characters will be returned as a single token. However, note that no semantic validity checking will be performed: for example, '>>>' could be returned as a token, even though it may not be recognised as such by shells.

3.6 新版功能.

24.3.2 Parsing Rules

When operating in non-POSIX mode, *shlex* will try to obey to the following rules.

- Quote characters are not recognized within words (`"Do" "Not" "Separate"` is parsed as the single word `"Do" "Not" "Separate"`);
- Escape characters are not recognized;
- Enclosing characters in quotes preserve the literal value of all characters within the quotes;
- Closing quotes separate words (`"Do" "Separate"` is parsed as `"Do"` and `Separate`);
- If *whitespace_split* is `False`, any character not declared to be a word character, whitespace, or a quote will be returned as a single-character token. If it is `True`, *shlex* will only split words in whitespaces;
- EOF is signaled with an empty string (' ');
- It's not possible to parse empty strings, even if quoted.

When operating in POSIX mode, *shlex* will try to obey to the following parsing rules.

- Quotes are stripped out, and do not separate words (`"Do" "Not" "Separate"` is parsed as the single word `DoNotSeparate`);
- Non-quoted escape characters (e.g. '\ ') preserve the literal value of the next character that follows;
- Enclosing characters in quotes which are not part of *escapedquotes* (e.g. '" ') preserve the literal value of all characters within the quotes;
- Enclosing characters in quotes which are part of *escapedquotes* (e.g. '" "') preserves the literal value of all characters within the quotes, with the exception of the characters mentioned in *escape*. The escape characters retain its special meaning only when followed by the quote in use, or the escape character itself. Otherwise the escape character will be considered a normal character.
- EOF is signaled with a *None* value;
- Quoted empty strings (' ') are allowed.

24.3.3 Improved Compatibility with Shells

3.6 新版功能.

The `shlex` class provides compatibility with the parsing performed by common Unix shells like `bash`, `dash`, and `sh`. To take advantage of this compatibility, specify the `punctuation_chars` argument in the constructor. This defaults to `False`, which preserves pre-3.6 behaviour. However, if it is set to `True`, then parsing of the characters `() ; <> | &` is changed: any run of these characters is returned as a single token. While this is short of a full parser for shells (which would be out of scope for the standard library, given the multiplicity of shells out there), it does allow you to perform processing of command lines more easily than you could otherwise. To illustrate, you can see the difference in the following snippet:

```
>>> import shlex
>>> text = "a && b; c && d || e; f >'abc'; (def \"ghi\")"
>>> list(shlex.shlex(text))
['a', '&', '&', 'b', ';', 'c', '&', '&', 'd', '|', '|', 'e', ';', 'f', '>',
"'abc'", ';', '(', 'def', '"ghi"', ')']
>>> list(shlex.shlex(text, punctuation_chars=True))
['a', '&&', 'b', ';', 'c', '&&', 'd', '||', 'e', ';', 'f', '>', "'abc'",
';', '(', 'def', '"ghi"', ')']
```

Of course, tokens will be returned which are not valid for shells, and you'll need to implement your own error checks on the returned tokens.

Instead of passing `True` as the value for the `punctuation_chars` parameter, you can pass a string with specific characters, which will be used to determine which characters constitute punctuation. For example:

```
>>> import shlex
>>> s = shlex.shlex("a && b || c", punctuation_chars="|")
>>> list(s)
['a', '&', '&', 'b', '||', 'c']
```

注解: When `punctuation_chars` is specified, the `wordchars` attribute is augmented with the characters `~-./*?==`. That is because these characters can appear in file names (including wildcards) and command-line arguments (e.g. `--color=auto`). Hence:

```
>>> import shlex
>>> s = shlex.shlex('~ / a && b - c --color=auto || d *.py?',
...                 punctuation_chars=True)
>>> list(s)
['~/a', '&&', 'b-c', '--color=auto', '||', 'd', '*.py?']
```

For best effect, `punctuation_chars` should be set in conjunction with `posix=True`. (Note that `posix=False` is the default for `shlex`.)

Tk 图形用户界面 (GUI)

Tcl/Tk 集成到 Python 中已经有一些年头了。Python 程序员可以通过 `tkinter` 包和它的扩展, `tkinter.tix` 模块和 `tkinter.ttk` 模块, 来使用这套鲁棒的、平台无关的窗口工具集。

`tkinter` 包使用面向对象的方式对 Tcl/Tk 进行了一层薄包装。使用 `tkinter`, 你不需要写 Tcl 代码, 但可能需要参考 Tk 文档, 甚至 Tcl 文档。`tkinter` 使用 Python 类, 对 Tk 的窗体小部件 (Widgets) 进行了一系列的封装。除此之外, 内部模块 `_tkinter` 针对 Python 和 Tcl 之间的交互, 提供了一套线程安全的机制。

`tkinter` 最大的优点就一个字: 快, 再一个, 是 Python 自带的。尽管官方文档不太完整, 但有其他资源可以参考, 比如 Tk 手册, 教程等。`tkinter` 也以比较过时的外观为人所知, 但在 Tk 8.5 中, 这一点得到了极大的改观。除此之外, 如果有兴趣, 还有其他的一些 GUI 库可供使用。更多信息, 请参考[其他图形用户界面 \(GUI\) 包](#)小节。

25.1 tkinter —Tcl/Tk 的 Python 接口

源代码: `Lib/tkinter/__init__.py`

The `tkinter` package (“Tk interface”) is the standard Python interface to the Tk GUI toolkit. Both Tk and `tkinter` are available on most Unix platforms, as well as on Windows systems. (Tk itself is not part of Python; it is maintained at ActiveState.)

在命令行中运行 `python -m tkinter`, 应该会弹出一个 Tk 界面的窗口, 表明 `tkinter` 包已经正确安装, 而且告诉你 Tcl/Tk 的版本号, 通过这个版本号, 你就可以参考对应的 Tcl/Tk 文档了。

参见:

Tkinter 文档:

Python Tkinter 资源 The Python Tkinter Topic Guide 提供了在 Python 中使用 Tk 的很多信息, 同时包含了 Tk 其他信息的链接。

TKDocs 大量的教程, 部分可视化组件的介绍说明。

Tkinter reference: a GUI for Python 在线参考资料。

Tkinter docs from effbot effbot.org 提供的 tkinter 在线参考资料。

使用 Python 编程 由 Mark Lutz 所著的书籍，对 Tkinter 进行了完美的介绍。

为繁忙的 Python 开发者所准备的现代 Tkinter 由 Mark Rozerman 所著关于如何使用 Python 和 Tkinter 搭建漂亮和时髦的图形化用户界面的书籍

Python 和 Tkinter 编程 作者：John Grayson (ISBN 1-884777-81-3).

Tcl/Tk 文档:

Tk 命令 多数命令以 `tkinter` 或者 `tkinter.ttk` 类的形式存在。改变 ‘8.6’ 以匹配所安装的 Tcl/Tk 版本。

Tcl/Tk 最新手册页面 www.tcl.tk 上面最新的 Tcl/Tk 手册。

ActiveState Tcl Home Page Tk/Tcl 的多数开发工作发生在 ActiveState 。

Tcl 及 Tk 工具集 由 Tcl 发明者 John Ousterhout 所著的书籍。

‘**Tcl 和 Tk 编程实战** <<http://www.beedub.com/book/>>’_ Brent Welch 所著的百科全局式书籍。

25.1.1 Tkinter 模块

在大多数时候你只需要 `tkinter` 就足够了，但也有一些额外的模块可供使用。Tk 接口位于一个名字 `_tkinter` 的二进制模块当中。此模块包含了低层级的 Tk 接口，它不应该被应用程序员所直接使用。它通常是一个共享库（或 DLL），但在某些情况下也可能被静态链接到 Python 解释器。

除了 Tk 接口，`tkinter` 也包含了若干 Python 模块，`tkinter.constants` 是其中最重要的。导入 `tkinter` 会自动导入 `tkinter.constants`，所以，要使用 Tkinter 通常你只需要一条简单的 `import` 语句：

```
import tkinter
```

或者更常用的：

```
from tkinter import *
```

class `tkinter.Tk` (`screenName=None`, `baseName=None`, `className='Tk'`, `useTk=1`)

The `Tk` class is instantiated without arguments. This creates a toplevel widget of Tk which usually is the main window of an application. Each instance has its own associated Tcl interpreter.

`tkinter.Tcl` (`screenName=None`, `baseName=None`, `className='Tk'`, `useTk=0`)

The `Tcl()` function is a factory function which creates an object much like that created by the `Tk` class, except that it does not initialize the Tk subsystem. This is most often useful when driving the Tcl interpreter in an environment where one doesn't want to create extraneous toplevel windows, or where one cannot (such as Unix/Linux systems without an X server). An object created by the `Tcl()` object can have a Toplevel window created (and the Tk subsystem initialized) by calling its `loadtk()` method.

Other modules that provide Tk support include:

`tkinter.scrolledtext` Text widget with a vertical scroll bar built in.

`tkinter.colorchooser` Dialog to let the user choose a color.

`tkinter.commondialog` Base class for the dialogs defined in the other modules listed here.

`tkinter.filedialog` Common dialogs to allow the user to specify a file to open or save.

`tkinter.font` Utilities to help work with fonts.

`tkinter.messagebox` Access to standard Tk dialog boxes.

`tkinter.simpledialog` Basic dialogs and convenience functions.

tkinter.dnd Drag-and-drop support for *tkinter*. This is experimental and should become deprecated when it is replaced with the Tk DND.

turtle Turtle graphics in a Tk window.

25.1.2 Tkinter Life Preserver

This section is not designed to be an exhaustive tutorial on either Tk or Tkinter. Rather, it is intended as a stop gap, providing some introductory orientation on the system.

Credits:

- Tk was written by John Ousterhout while at Berkeley.
- Tkinter was written by Steen Lumholt and Guido van Rossum.
- This Life Preserver was written by Matt Conway at the University of Virginia.
- The HTML rendering, and some liberal editing, was produced from a FrameMaker version by Ken Manheimer.
- Fredrik Lundh elaborated and revised the class interface descriptions, to get them current with Tk 4.2.
- Mike Clarkson converted the documentation to LaTeX, and compiled the User Interface chapter of the reference manual.

How To Use This Section

This section is designed in two parts: the first half (roughly) covers background material, while the second half can be taken to the keyboard as a handy reference.

When trying to answer questions of the form “how do I do blah”, it is often best to find out how to do “blah” in straight Tk, and then convert this back into the corresponding *tkinter* call. Python programmers can often guess at the correct Python command by looking at the Tk documentation. This means that in order to use Tkinter, you will have to know a little bit about Tk. This document can’t fulfill that role, so the best we can do is point you to the best documentation that exists. Here are some hints:

- The authors strongly suggest getting a copy of the Tk man pages. Specifically, the man pages in the `manN` directory are most useful. The `man3` man pages describe the C interface to the Tk library and thus are not especially helpful for script writers.
- Addison-Wesley publishes a book called *Tcl and the Tk Toolkit* by John Ousterhout (ISBN 0-201-63337-X) which is a good introduction to Tcl and Tk for the novice. The book is not exhaustive, and for many details it defers to the man pages.
- `tkinter/__init__.py` is a last resort for most, but can be a good place to go when nothing else makes sense.

A Simple Hello World Program

```
import tkinter as tk

class Application(tk.Frame):
    def __init__(self, master=None):
        super().__init__(master)
        self.master = master
        self.pack()
        self.create_widgets()
```

(下页继续)

(续上页)

```

def create_widgets(self):
    self.hi_there = tk.Button(self)
    self.hi_there["text"] = "Hello World\n(click me)"
    self.hi_there["command"] = self.say_hi
    self.hi_there.pack(side="top")

    self.quit = tk.Button(self, text="QUIT", fg="red",
                           command=self.master.destroy)
    self.quit.pack(side="bottom")

def say_hi(self):
    print("hi there, everyone!")

root = tk.Tk()
app = Application(master=root)
app.mainloop()

```

25.1.3 A (Very) Quick Look at Tcl/Tk

The class hierarchy looks complicated, but in actual practice, application programmers almost always refer to the classes at the very bottom of the hierarchy.

注释:

- These classes are provided for the purposes of organizing certain functions under one namespace. They aren't meant to be instantiated independently.
- The `Tk` class is meant to be instantiated only once in an application. Application programmers need not instantiate one explicitly, the system creates one whenever any of the other classes are instantiated.
- The `Widget` class is not meant to be instantiated, it is meant only for subclassing to make “real” widgets (in C++, this is called an ‘abstract class’).

To make use of this reference material, there will be times when you will need to know how to read short passages of Tk and how to identify the various parts of a Tk command. (See section [Mapping Basic Tk into Tkinter](#) for the *tkinter* equivalents of what's below.)

Tk scripts are Tcl programs. Like all Tcl programs, Tk scripts are just lists of tokens separated by spaces. A Tk widget is just its *class*, the *options* that help configure it, and the *actions* that make it do useful things.

To make a widget in Tk, the command is always of the form:

```
classCommand newPathname options
```

classCommand denotes which kind of widget to make (a button, a label, a menu...)

newPathname is the new name for this widget. All names in Tk must be unique. To help enforce this, widgets in Tk are named with *pathnames*, just like files in a file system. The top level widget, the *root*, is called `.` (period) and children are delimited by more periods. For example, `.myApp.controlPanel.okButton` might be the name of a widget.

options configure the widget's appearance and in some cases, its behavior. The options come in the form of a list of flags and values. Flags are preceded by a `-`, like Unix shell command flags, and values are put in quotes if they are more than one word.

例如:

```

button      .fred      -fg red -text "hi there"
  ^          ^          |
  |          |          |
class      new          options
command  widget  (-opt val -opt val ...)

```

Once created, the pathname to the widget becomes a new command. This new *widget command* is the programmer's handle for getting the new widget to perform some *action*. In C, you'd express this as `someAction(fred, someOptions)`, in C++, you would express this as `fred.someAction(someOptions)`, and in Tk, you say:

```
.fred someAction someOptions
```

Note that the object name, `.fred`, starts with a dot.

As you'd expect, the legal values for *someAction* will depend on the widget's class: `.fred disable` works if `fred` is a button (`fred` gets greyed out), but does not work if `fred` is a label (disabling of labels is not supported in Tk).

The legal values of *someOptions* is action dependent. Some actions, like `disable`, require no arguments, others, like a text-entry box's `delete` command, would need arguments to specify what range of text to delete.

25.1.4 Mapping Basic Tk into Tkinter

Class commands in Tk correspond to class constructors in Tkinter.

```
button .fred          =====> fred = Button()
```

The master of an object is implicit in the new name given to it at creation time. In Tkinter, masters are specified explicitly.

```
button .panel.fred    =====> fred = Button(panel)
```

The configuration options in Tk are given in lists of hyphenated tags followed by values. In Tkinter, options are specified as keyword-arguments in the instance constructor, and keyword-args for `configure` calls or as instance indices, in dictionary style, for established instances. See section [Setting Options](#) on setting options.

```

button .fred -fg red      =====> fred = Button(panel, fg="red")
.fred configure -fg red    =====> fred["fg"] = red
                           OR ==> fred.config(fg="red")

```

In Tk, to perform an action on a widget, use the widget name as a command, and follow it with an action name, possibly with arguments (options). In Tkinter, you call methods on the class instance to invoke actions on the widget. The actions (methods) that a given widget can perform are listed in `tkinter/__init__.py`.

```
.fred invoke          =====> fred.invoke()
```

To give a widget to the packer (geometry manager), you call `pack` with optional arguments. In Tkinter, the `Pack` class holds all this functionality, and the various forms of the `pack` command are implemented as methods. All widgets in *tkinter* are subclassed from the `Packer`, and so inherit all the packing methods. See the *tkinter.tix* module documentation for additional information on the Form geometry manager.

```
pack .fred -side left     =====> fred.pack(side="left")
```

25.1.5 How Tk and Tkinter are Related

From the top down:

Your App Here (Python) A Python application makes a *tkinter* call.

tkinter (Python Package) This call (say, for example, creating a button widget), is implemented in the *tkinter* package, which is written in Python. This Python function will parse the commands and the arguments and convert them into a form that makes them look as if they had come from a Tk script instead of a Python script.

_tkinter (C) These commands and their arguments will be passed to a C function in the *_tkinter* - note the underscore - extension module.

Tk Widgets (C and Tcl) This C function is able to make calls into other C modules, including the C functions that make up the Tk library. Tk is implemented in C and some Tcl. The Tcl part of the Tk widgets is used to bind certain default behaviors to widgets, and is executed once at the point where the Python *tkinter* package is imported. (The user never sees this stage).

Tk (C) The Tk part of the Tk Widgets implement the final mapping to ...

Xlib (C) the Xlib library to draw graphics on the screen.

25.1.6 Handy Reference

Setting Options

Options control things like the color and border width of a widget. Options can be set in three ways:

At object creation time, using keyword arguments

```
fred = Button(self, fg="red", bg="blue")
```

After object creation, treating the option name like a dictionary index

```
fred["fg"] = "red"
fred["bg"] = "blue"
```

Use the `config()` method to update multiple attrs subsequent to object creation

```
fred.config(fg="red", bg="blue")
```

For a complete explanation of a given option and its behavior, see the Tk man pages for the widget in question.

Note that the man pages list “STANDARD OPTIONS” and “WIDGET SPECIFIC OPTIONS” for each widget. The former is a list of options that are common to many widgets, the latter are the options that are idiosyncratic to that particular widget. The Standard Options are documented on the *options(3)* man page.

No distinction between standard and widget-specific options is made in this document. Some options don't apply to some kinds of widgets. Whether a given widget responds to a particular option depends on the class of the widget; buttons have a `command` option, labels do not.

The options supported by a given widget are listed in that widget's man page, or can be queried at runtime by calling the `config()` method without arguments, or by calling the `keys()` method on that widget. The return value of these calls is a dictionary whose key is the name of the option as a string (for example, `'relief'`) and whose values are 5-tuples.

Some options, like `bg` are synonyms for common options with long names (`bg` is shorthand for “background”). Passing the `config()` method the name of a shorthand option will return a 2-tuple, not 5-tuple. The 2-tuple passed back will contain the name of the synonym and the “real” option (such as `('bg', 'background')`).

索引	意义	示例
0	选项名称	'relief'
1	option name for database lookup	'relief'
2	option class for database lookup	'Relief'
3	default value	'raised'
4	current value	'groove'

示例:

```
>>> print(fred.config())
{'relief': ('relief', 'relief', 'Relief', 'raised', 'groove')}
```

Of course, the dictionary printed will include all the options available and their values. This is meant only as an example.

The Packer

The packer is one of Tk's geometry-management mechanisms. Geometry managers are used to specify the relative positioning of the positioning of widgets within their container - their mutual *master*. In contrast to the more cumbersome *placer* (which is used less commonly, and we do not cover here), the packer takes qualitative relationship specification - *above*, *to the left of*, *filling*, etc - and works everything out to determine the exact placement coordinates for you.

The size of any *master* widget is determined by the size of the “slave widgets” inside. The packer is used to control where slave widgets appear inside the master into which they are packed. You can pack widgets into frames, and frames into other frames, in order to achieve the kind of layout you desire. Additionally, the arrangement is dynamically adjusted to accommodate incremental changes to the configuration, once it is packed.

Note that widgets do not appear until they have had their geometry specified with a geometry manager. It's a common early mistake to leave out the geometry specification, and then be surprised when the widget is created but nothing appears. A widget will appear only after it has had, for example, the packer's `pack()` method applied to it.

The `pack()` method can be called with keyword-option/value pairs that control where the widget is to appear within its container, and how it is to behave when the main application window is resized. Here are some examples:

```
fred.pack()                                # defaults to side = "top"
fred.pack(side="left")
fred.pack(expand=1)
```

Packer Options

For more extensive information on the packer and the options that it can take, see the man pages and page 183 of John Ousterhout's book.

anchor Anchor type. Denotes where the packer is to place each slave in its parcel.

expand Boolean, 0 or 1.

fill Legal values: 'x', 'y', 'both', 'none'.

ipadx and ipady A distance - designating internal padding on each side of the slave widget.

padx and pady A distance - designating external padding on each side of the slave widget.

side Legal values are: 'left', 'right', 'top', 'bottom'.

Coupling Widget Variables

The current-value setting of some widgets (like text entry widgets) can be connected directly to application variables by using special options. These options are `variable`, `textvariable`, `onvalue`, `offvalue`, and `value`. This connection works both ways: if the variable changes for any reason, the widget it's connected to will be updated to reflect the new value.

Unfortunately, in the current implementation of *tkinter* it is not possible to hand over an arbitrary Python variable to a widget through a `variable` or `textvariable` option. The only kinds of variables for which this works are variables that are subclassed from a class called `Variable`, defined in *tkinter*.

There are many useful subclasses of `Variable` already defined: `StringVar`, `IntVar`, `DoubleVar`, and `BooleanVar`. To read the current value of such a variable, call the `get()` method on it, and to change its value you call the `set()` method. If you follow this protocol, the widget will always track the value of the variable, with no further intervention on your part.

例如:

```
class App(Frame):
    def __init__(self, master=None):
        super().__init__(master)
        self.pack()

        self.entrythingy = Entry()
        self.entrythingy.pack()

        # here is the application variable
        self.contents = StringVar()
        # set it to some value
        self.contents.set("this is a variable")
        # tell the entry widget to watch this variable
        self.entrythingy["textvariable"] = self.contents

        # and here we get a callback when the user hits return.
        # we will have the program print out the value of the
        # application variable when the user hits return
        self.entrythingy.bind('<Key-Return>',
                               self.print_contents)

    def print_contents(self, event):
        print("hi. contents of entry is now ---->",
              self.contents.get())
```

The Window Manager

In Tk, there is a utility command, `wm`, for interacting with the window manager. Options to the `wm` command allow you to control things like titles, placement, icon bitmaps, and the like. In *tkinter*, these commands have been implemented as methods on the `Wm` class. Toplevel widgets are subclassed from the `Wm` class, and so can call the `Wm` methods directly.

To get at the toplevel window that contains a given widget, you can often just refer to the widget's master. Of course if the widget has been packed inside of a frame, the master won't represent a toplevel window. To get at the toplevel window that contains an arbitrary widget, you can call the `_root()` method. This method begins with an underscore to denote the fact that this function is part of the implementation, and not an interface to Tk functionality.

Here are some examples of typical usage:

```
import tkinter as tk

class App(tk.Frame):
    def __init__(self, master=None):
        super().__init__(master)
        self.pack()

# create the application
myapp = App()

#
# here are method calls to the window manager class
#
myapp.master.title("My Do-Nothing Application")
myapp.master.maxsize(1000, 400)

# start the program
myapp.mainloop()
```

Tk Option Data Types

anchor Legal values are points of the compass: "n", "ne", "e", "se", "s", "sw", "w", "nw", and also "center".

bitmap There are eight built-in, named bitmaps: 'error', 'gray25', 'gray50', 'hourglass', 'info', 'questhead', 'question', 'warning'. To specify an X bitmap filename, give the full path to the file, preceded with an @, as in "@usr/contrib/bitmap/gumby.bit".

boolean You can pass integers 0 or 1 or the strings "yes" or "no".

callback This is any Python function that takes no arguments. For example:

```
def print_it():
    print("hi there")
fred["command"] = print_it
```

color Colors can be given as the names of X colors in the rgb.txt file, or as strings representing RGB values in 4 bit: "#RGB", 8 bit: "#RRGGBB", 12 bit: "#RRRGGGBBB", or 16 bit: "#RRRRGGGGBBBB" ranges, where R,G,B here represent any legal hex digit. See page 160 of Ousterhout's book for details.

cursor The standard X cursor names from cursorfont.h can be used, without the XC_ prefix. For example to get a hand cursor (XC_hand2), use the string "hand2". You can also specify a bitmap and mask file of your own. See page 179 of Ousterhout's book.

distance Screen distances can be specified in either pixels or absolute distances. Pixels are given as numbers and absolute distances as strings, with the trailing character denoting units: c for centimetres, i for inches, m for millimetres, p for printer's points. For example, 3.5 inches is expressed as "3.5i".

font Tk uses a list font name format, such as {courier 10 bold}. Font sizes with positive numbers are measured in points; sizes with negative numbers are measured in pixels.

geometry This is a string of the form widthxheight, where width and height are measured in pixels for most widgets (in characters for widgets displaying text). For example: fred["geometry"] = "200x100".

justify Legal values are the strings: "left", "center", "right", and "fill".

region This is a string with four space-delimited elements, each of which is a legal distance (see above). For example: "2 3 4 5" and "3i 2i 4.5i 2i" and "3c 2c 4c 10.43c" are all legal regions.

relief Determines what the border style of a widget will be. Legal values are: "raised", "sunken", "flat", "groove", and "ridge".

scrollcommand This is almost always the `set()` method of some scrollbar widget, but can be any widget method that takes a single argument.

wrap: Must be one of: "none", "char", or "word".

Bindings and Events

The `bind` method from the widget command allows you to watch for certain events and to have a callback function trigger when that event type occurs. The form of the `bind` method is:

```
def bind(self, sequence, func, add='')
```

where:

序列 is a string that denotes the target kind of event. (See the `bind` man page and page 201 of John Ousterhout's book for details).

func is a Python function, taking one argument, to be invoked when the event occurs. An `Event` instance will be passed as the argument. (Functions deployed this way are commonly known as *callbacks*.)

add is optional, either `' '` or `'+'`. Passing an empty string denotes that this binding is to replace any other bindings that this event is associated with. Passing a `'+'` means that this function is to be added to the list of functions bound to this event type.

例如:

```
def turn_red(self, event):
    event.widget["activeforeground"] = "red"

self.button.bind("<Enter>", self.turn_red)
```

Notice how the `widget` field of the event is being accessed in the `turn_red()` callback. This field contains the widget that caught the X event. The following table lists the other event fields you can access, and how they are denoted in Tk, which can be useful when referring to the Tk man pages.

Tk	Tkinter Event Field	Tk	Tkinter Event Field
%f	focus	%A	char
%h	height	%E	send_event
%k	keycode	%K	keysym
%s	state	%N	keysym_num
%t	time	%T	类型
%w	width	%W	widget
%x	x	%X	x_root
%y	y	%Y	y_root

The index Parameter

A number of widgets require “index” parameters to be passed. These are used to point at a specific place in a `Text` widget, or to particular characters in an `Entry` widget, or to particular menu items in a `Menu` widget.

Entry widget indexes (index, view index, etc.) Entry widgets have options that refer to character positions in the text being displayed. You can use these *tkinter* functions to access these special points in text widgets:

Text widget indexes The index notation for Text widgets is very rich and is best described in the Tk man pages.

Menu indexes (`menu.invoke()`, `menu.entryconfig()`, etc.) Some options and methods for menus manipulate specific menu entries. Anytime a menu index is needed for an option or a parameter, you may pass in:

- an integer which refers to the numeric position of the entry in the widget, counted from the top, starting with 0;
- the string "active", which refers to the menu position that is currently under the cursor;
- the string "last" which refers to the last menu item;
- An integer preceded by @, as in @6, where the integer is interpreted as a y pixel coordinate in the menu's coordinate system;
- the string "none", which indicates no menu entry at all, most often used with `menu.activate()` to deactivate all entries, and finally,
- a text string that is pattern matched against the label of the menu entry, as scanned from the top of the menu to the bottom. Note that this index type is considered after all the others, which means that matches for menu items labelled last, active, or none may be interpreted as the above literals, instead.

Images

Images of different formats can be created through the corresponding subclass of `tkinter.Image`:

- `BitmapImage` for images in XBM format.
- `PhotoImage` for images in PGM, PPM, GIF and PNG formats. The latter is supported starting with Tk 8.6.

Either type of image is created through either the `file` or the `data` option (other options are available as well).

The image object can then be used wherever an `image` option is supported by some widget (e.g. labels, buttons, menus). In these cases, Tk will not keep a reference to the image. When the last Python reference to the image object is deleted, the image data is deleted as well, and Tk will display an empty box wherever the image was used.

参见:

The [Pillow](#) package adds support for formats such as BMP, JPEG, TIFF, and WebP, among others.

25.1.7 File Handlers

Tk allows you to register and unregister a callback function which will be called from the Tk mainloop when I/O is possible on a file descriptor. Only one handler may be registered per file descriptor. Example code:

```
import tkinter
widget = tkinter.Tk()
mask = tkinter.READABLE | tkinter.WRITABLE
widget.tk.createfilehandler(file, mask, callback)
...
widget.tk.deletefilehandler(file)
```

This feature is not available on Windows.

Since you don't know how many bytes are available for reading, you may not want to use the [BufferedIOBase](#) or [TextIOBase](#) `read()` or `readline()` methods, since these will insist on reading a predefined number of bytes. For sockets, the `recv()` or `recvfrom()` methods will work fine; for other files, use raw reads or `os.read(file.fileno(), maxbytecount)`.

`Widget.tk.createfilehandler` (*file*, *mask*, *func*)

Registers the file handler callback function *func*. The *file* argument may either be an object with a *fileno()* method (such as a file or socket object), or an integer file descriptor. The *mask* argument is an ORed combination of any of the three constants below. The callback is called as follows:

```
callback(file, mask)
```

`Widget.tk.deletefilehandler` (*file*)

Unregisters a file handler.

`tkinter.READABLE`

`tkinter.WRITABLE`

`tkinter.EXCEPTION`

Constants used in the *mask* arguments.

25.2 `tkinter.ttk` —Tk themed widgets

Source code: [Lib/tkinter/ttk.py](#)

The `tkinter.ttk` module provides access to the Tk themed widget set, introduced in Tk 8.5. If Python has not been compiled against Tk 8.5, this module can still be accessed if *Tile* has been installed. The former method using Tk 8.5 provides additional benefits including anti-aliased font rendering under X11 and window transparency (requiring a composition window manager on X11).

The basic idea for `tkinter.ttk` is to separate, to the extent possible, the code implementing a widget's behavior from the code implementing its appearance.

参见:

Tk Widget Styling Support A document introducing theming support for Tk

25.2.1 Using Ttk

To start using Ttk, import its module:

```
from tkinter import ttk
```

To override the basic Tk widgets, the import should follow the Tk import:

```
from tkinter import *
from tkinter.ttk import *
```

That code causes several `tkinter.ttk` widgets (Button, Checkbutton, Entry, Frame, Label, LabelFrame, Menubutton, PanedWindow, Radiobutton, Scale and Scrollbar) to automatically replace the Tk widgets.

This has the direct benefit of using the new widgets which gives a better look and feel across platforms; however, the replacement widgets are not completely compatible. The main difference is that widget options such as “fg”, “bg” and others related to widget styling are no longer present in Ttk widgets. Instead, use the `ttk.Style` class for improved styling effects.

参见:

Converting existing applications to use Tile widgets A monograph (using Tcl terminology) about differences typically encountered when moving applications to use the new widgets.

25.2.2 Ttk Widgets

Ttk comes with 17 widgets, eleven of which already existed in tkinter: `Button`, `Checkbutton`, `Entry`, `Frame`, `Label`, `LabelFrame`, `Menubutton`, `PanedWindow`, `Radiobutton`, `Scale` and `Scrollbar`. The other six are new: `Combobox`, `Notebook`, `Progressbar`, `Separator`, `Sizegrip` and `Treeview`. And all them are subclasses of `Widget`.

Using the Ttk widgets gives the application an improved look and feel. As discussed above, there are differences in how the styling is coded.

Tk code:

```
l1 = tkinter.Label(text="Test", fg="black", bg="white")
l2 = tkinter.Label(text="Test", fg="black", bg="white")
```

Ttk code:

```
style = ttk.Style()
style.configure("BW.TLabel", foreground="black", background="white")

l1 = ttk.Label(text="Test", style="BW.TLabel")
l2 = ttk.Label(text="Test", style="BW.TLabel")
```

For more information about *TtkStyling*, see the `Style` class documentation.

25.2.3 Widget

`ttk.Widget` defines standard options and methods supported by Tk themed widgets and is not supposed to be directly instantiated.

Standard Options

All the `ttk` Widgets accepts the following options:

选项	描述
类	Specifies the window class. The class is used when querying the option database for the window's other options, to determine the default bindtags for the window, and to select the widget's default layout and style. This option is read-only, and may only be specified when the window is created.
cursor	Specifies the mouse cursor to be used for the widget. If set to the empty string (the default), the cursor is inherited for the parent widget.
takefocus	Determines whether the window accepts the focus during keyboard traversal. 0, 1 or an empty string is returned. If 0 is returned, it means that the window should be skipped entirely during keyboard traversal. If 1, it means that the window should receive the input focus as long as it is viewable. And an empty string means that the traversal scripts make the decision about whether or not to focus on the window.
style	May be used to specify a custom widget style.

Scrollable Widget Options

The following options are supported by widgets that are controlled by a scrollbar.

选项	描述
xscrollcommand	Used to communicate with horizontal scrollbars. When the view in the widget's window change, the widget will generate a Tcl command based on the scrollcommand. Usually this option consists of the method <code>Scrollbar.set()</code> of some scrollbar. This will cause the scrollbar to be updated whenever the view in the window changes.
yscrollcommand	Used to communicate with vertical scrollbars. For some more information, see above.

Label Options

The following options are supported by labels, buttons and other button-like widgets.

选项	描述
text	Specifies a text string to be displayed inside the widget.
textvariable	Specifies a name whose value will be used in place of the text option resource.
underline	If set, specifies the index (0-based) of a character to underline in the text string. The underline character is used for mnemonic activation.
image	Specifies an image to display. This is a list of 1 or more elements. The first element is the default image name. The rest of the list is a sequence of statespec/value pairs as defined by <code>Style.map()</code> , specifying different images to use when the widget is in a particular state or a combination of states. All images in the list should have the same size.
compound	Specifies how to display the image relative to the text, in the case both text and images options are present. Valid values are: <ul style="list-style-type: none"> • text: display text only • image: display image only • top, bottom, left, right: display image above, below, left of, or right of the text, respectively. • none: the default. display the image if present, otherwise the text.
width	If greater than zero, specifies how much space, in character widths, to allocate for the text label, if less than zero, specifies a minimum width. If zero or unspecified, the natural width of the text label is used.

Compatibility Options

选项	描述
state	May be set to “normal” or “disabled” to control the “disabled” state bit. This is a write-only option: setting it changes the widget state, but the <code>Widget.state()</code> method does not affect this option.

Widget States

The widget state is a bitmap of independent state flags.

标志	描述
active	The mouse cursor is over the widget and pressing a mouse button will cause some action to occur
disabled	Widget is disabled under program control
focus	Widget has keyboard focus
pressed	Widget is being pressed
selected	“On” , “true” , or “current” for things like Checkbuttons and radiobuttons
background	Windows and Mac have a notion of an “active” or foreground window. The <i>background</i> state is set for widgets in a background window, and cleared for those in the foreground window
只读	Widget should not allow user modification
alternate	A widget-specific alternate display format
invalid	The widget’s value is invalid

A state specification is a sequence of state names, optionally prefixed with an exclamation point indicating that the bit is off.

ttk.Widget

Besides the methods described below, the `ttk.Widget` supports the methods `tkinter.Widget.cget()` and `tkinter.Widget.configure()`.

class `tkinter.ttk.Widget`

identify (*x*, *y*)

Returns the name of the element at position *x* *y*, or the empty string if the point does not lie within any element.

x and *y* are pixel coordinates relative to the widget.

instate (*statespec*, *callback=None*, **args*, ***kw*)

Test the widget’s state. If a callback is not specified, returns `True` if the widget state matches *statespec* and `False` otherwise. If callback is specified then it is called with *args* if widget state matches *statespec*.

state (*statespec=None*)

Modify or inquire widget state. If *statespec* is specified, sets the widget state according to it and return a new *statespec* indicating which flags were changed. If *statespec* is not specified, returns the currently-enabled state flags.

statespec will usually be a list or a tuple.

25.2.4 Combobox

The `ttk.Combobox` widget combines a text field with a pop-down list of values. This widget is a subclass of `Entry`.

Besides the methods inherited from *Widget*: `Widget.cget()`, `Widget.configure()`, `Widget.identify()`, `Widget.instate()` and `Widget.state()`, and the following inherited from `Entry`: `Entry.bbox()`, `Entry.delete()`, `Entry.icursor()`, `Entry.index()`, `Entry.insert()`, `Entry.selection()`, `Entry.xview()`, it has some other methods, described at `ttk.Combobox`.

选项

This widget accepts the following specific options:

选项	描述
exportselection	Boolean value. If set, the widget selection is linked to the Window Manager selection (which can be returned by invoking <code>Misc.selection_get</code> , for example).
justify	Specifies how the text is aligned within the widget. One of “left”, “center”, or “right” .
height	Specifies the height of the pop-down listbox, in rows.
postcommand	A script (possibly registered with <code>Misc.register</code>) that is called immediately before displaying the values. It may specify which values to display.
state	One of “normal”, “readonly”, or “disabled” . In the “readonly” state, the value may not be edited directly, and the user can only selection of the values from the dropdown list. In the “normal” state, the text field is directly editable. In the “disabled” state, no interaction is possible.
textvariable	Specifies a name whose value is linked to the widget value. Whenever the value associated with that name changes, the widget value is updated, and vice versa. See <code>tkinter.StringVar</code> .
values	Specifies the list of values to display in the drop-down listbox.
width	Specifies an integer value indicating the desired width of the entry window, in average-size characters of the widget’s font.

Virtual events

The combobox widgets generates a `<<ComboboxSelected>>` virtual event when the user selects an element from the list of values.

ttk.Combobox

class `tkinter.ttk.Combobox`

current (*newindex=None*)

If *newindex* is specified, sets the combobox value to the element position *newindex*. Otherwise, returns the index of the current value or -1 if the current value is not in the values list.

get ()

Returns the current value of the combobox.

set (*value*)

Sets the value of the combobox to *value*.

25.2.5 Notebook

Ttk Notebook widget manages a collection of windows and displays a single one at a time. Each child window is associated with a tab, which the user may select to change the currently-displayed window.

选项

This widget accepts the following specific options:

选项	描述
height	If present and greater than zero, specifies the desired height of the pane area (not including internal padding or tabs). Otherwise, the maximum height of all panes is used.
padding	Specifies the amount of extra space to add around the outside of the notebook. The padding is a list up to four length specifications left top right bottom. If fewer than four elements are specified, bottom defaults to top, right defaults to left, and top defaults to left.
width	If present and greater than zero, specified the desired width of the pane area (not including internal padding). Otherwise, the maximum width of all panes is used.

Tab Options

There are also specific options for tabs:

选项	描述
state	Either “normal” , “disabled” or “hidden” . If “disabled” , then the tab is not selectable. If “hidden” , then the tab is not shown.
sticky	Specifies how the child window is positioned within the pane area. Value is a string containing zero or more of the characters “n” , “s” , “e” or “w” . Each letter refers to a side (north, south, east or west) that the child window will stick to, as per the <code>grid()</code> geometry manager.
padding	Specifies the amount of extra space to add between the notebook and this pane. Syntax is the same as for the option padding used by this widget.
text	Specifies a text to be displayed in the tab.
image	Specifies an image to display in the tab. See the option image described in <i>Widget</i> .
compound	Specifies how to display the image relative to the text, in the case both options text and image are present. See <i>Label Options</i> for legal values.
underline	Specifies the index (0-based) of a character to underline in the text string. The underlined character is used for mnemonic activation if <code>Notebook.enable_traversal()</code> is called.

Tab Identifiers

The `tab_id` present in several methods of `ttk.Notebook` may take any of the following forms:

- An integer between zero and the number of tabs
- The name of a child window
- A positional specification of the form “@x,y” , which identifies the tab
- The literal string “current” , which identifies the currently-selected tab
- The literal string “end” , which returns the number of tabs (only valid for `Notebook.index()`)

Virtual Events

This widget generates a `<<NotebookTabChanged>>` virtual event after a new tab is selected.

ttk.Notebook

```
class tkinter.ttk.Notebook
```

add (*child*, ***kw*)

Adds a new tab to the notebook.

If window is currently managed by the notebook but hidden, it is restored to its previous position.

See [Tab Options](#) for the list of available options.

forget (*tab_id*)

Removes the tab specified by *tab_id*, unmaps and unmanages the associated window.

hide (*tab_id*)

Hides the tab specified by *tab_id*.

The tab will not be displayed, but the associated window remains managed by the notebook and its configuration remembered. Hidden tabs may be restored with the [add\(\)](#) command.

identify (*x*, *y*)

Returns the name of the tab element at position *x*, *y*, or the empty string if none.

index (*tab_id*)

Returns the numeric index of the tab specified by *tab_id*, or the total number of tabs if *tab_id* is the string “end” .

insert (*pos*, *child*, ***kw*)

Inserts a pane at the specified position.

pos is either the string “end” , an integer index, or the name of a managed child. If *child* is already managed by the notebook, moves it to the specified position.

See [Tab Options](#) for the list of available options.

select (*tab_id=None*)

Selects the specified *tab_id*.

The associated child window will be displayed, and the previously-selected window (if different) is unmapped. If *tab_id* is omitted, returns the widget name of the currently selected pane.

tab (*tab_id*, *option=None*, ***kw*)

Query or modify the options of the specific *tab_id*.

If *kw* is not given, returns a dictionary of the tab option values. If *option* is specified, returns the value of that *option*. Otherwise, sets the options to the corresponding values.

tabs ()

Returns a list of windows managed by the notebook.

enable_traversal ()

Enable keyboard traversal for a toplevel window containing this notebook.

This will extend the bindings for the toplevel window containing the notebook as follows:

- Control-Tab: selects the tab following the currently selected one.
- Shift-Control-Tab: selects the tab preceding the currently selected one.
- Alt-K: where *K* is the mnemonic (underlined) character of any tab, will select that tab.

Multiple notebooks in a single toplevel may be enabled for traversal, including nested notebooks. However, notebook traversal only works properly if all panes have the notebook they are in as master.

25.2.6 Progressbar

The `ttk.Progressbar` widget shows the status of a long-running operation. It can operate in two modes: 1) the determinate mode which shows the amount completed relative to the total amount of work to be done and 2) the indeterminate mode which provides an animated display to let the user know that work is progressing.

选项

This widget accepts the following specific options:

选项	描述
orient	One of “horizontal” or “vertical” . Specifies the orientation of the progress bar.
length	Specifies the length of the long axis of the progress bar (width if horizontal, height if vertical).
模式	One of “determinate” or “indeterminate” .
maximum	A number specifying the maximum value. Defaults to 100.
值	The current value of the progress bar. In “determinate” mode, this represents the amount of work completed. In “indeterminate” mode, it is interpreted as modulo <i>maximum</i> ; that is, the progress bar completes one “cycle” when its value increases by <i>maximum</i> .
variable	A name which is linked to the option value. If specified, the value of the progress bar is automatically set to the value of this name whenever the latter is modified.
phase	Read-only option. The widget periodically increments the value of this option whenever its value is greater than 0 and, in determinate mode, less than maximum. This option may be used by the current theme to provide additional animation effects.

ttk.Progressbar

class `tkinter.ttk.Progressbar`

- start** (*interval=None*)
Begin autoincrement mode: schedules a recurring timer event that calls `Progressbar.step()` every *interval* milliseconds. If omitted, *interval* defaults to 50 milliseconds.
- step** (*amount=None*)
Increments the progress bar’s value by *amount*.
amount defaults to 1.0 if omitted.
- stop** ()
Stop autoincrement mode: cancels any recurring timer event initiated by `Progressbar.start()` for this progress bar.

25.2.7 Separator

The `ttk.Separator` widget displays a horizontal or vertical separator bar.

It has no other methods besides the ones inherited from `ttk.Widget`.

选项

This widget accepts the following specific option:

选项	描述
orient	One of “horizontal” or “vertical” . Specifies the orientation of the separator.

25.2.8 Sizegrip

The `ttk.Sizegrip` widget (also known as a grow box) allows the user to resize the containing toplevel window by pressing and dragging the grip.

This widget has neither specific options nor specific methods, besides the ones inherited from `ttk.Widget`.

Platform-specific notes

- On MacOS X, toplevel windows automatically include a built-in size grip by default. Adding a `Sizegrip` is harmless, since the built-in grip will just mask the widget.

Bugs

- If the containing toplevel’s position was specified relative to the right or bottom of the screen (e.g. `...()`), the `Sizegrip` widget will not resize the window.
- This widget supports only “southeast” resizing.

25.2.9 Treeview

The `ttk.Treeview` widget displays a hierarchical collection of items. Each item has a textual label, an optional image, and an optional list of data values. The data values are displayed in successive columns after the tree label.

The order in which data values are displayed may be controlled by setting the widget option `displaycolumns`. The tree widget can also display column headings. Columns may be accessed by number or symbolic names listed in the widget option `columns`. See [Column Identifiers](#).

Each item is identified by a unique name. The widget will generate item IDs if they are not supplied by the caller. There is a distinguished root item, named `{ }`. The root item itself is not displayed; its children appear at the top level of the hierarchy.

Each item also has a list of tags, which can be used to associate event bindings with individual items and control the appearance of the item.

The `Treeview` widget supports horizontal and vertical scrolling, according to the options described in [Scrollable Widget Options](#) and the methods `Treeview.xview()` and `Treeview.yview()`.

选项

This widget accepts the following specific options:

选项	描述
columns	A list of column identifiers, specifying the number of columns and their names.
displaycolumns	A list of column identifiers (either symbolic or integer indices) specifying which data columns are displayed and the order in which they appear, or the string “#all” .
height	Specifies the number of rows which should be visible. Note: the requested width is determined from the sum of the column widths.
padding	Specifies the internal padding for the widget. The padding is a list of up to four length specifications.
selectmode	Controls how the built-in class bindings manage the selection. One of “extended” , “browse” or “none” . If set to “extended” (the default), multiple items may be selected. If “browse” , only a single item will be selected at a time. If “none” , the selection will not be changed. Note that the application code and tag bindings can set the selection however they wish, regardless of the value of this option.
show	A list containing zero or more of the following values, specifying which elements of the tree to display. <ul style="list-style-type: none"> • tree: display tree labels in column #0. • headings: display the heading row. The default is “tree headings” , i.e., show all elements. Note: Column #0 always refers to the tree column, even if show=” tree” is not specified.

Item Options

The following item options may be specified for items in the insert and item widget commands.

选项	描述
text	The textual label to display for the item.
image	A Tk Image, displayed to the left of the label.
values	The list of values associated with the item. Each item should have the same number of values as the widget option columns. If there are fewer values than columns, the remaining values are assumed empty. If there are more values than columns, the extra values are ignored.
open	True/False value indicating whether the item’ s children should be displayed or hidden.
tags	A list of tags associated with this item.

Tag Options

The following options may be specified on tags:

选项	描述
foreground	Specifies the text foreground color.
background	Specifies the cell or item background color.
font	Specifies the font to use when drawing text.
image	Specifies the item image, in case the item’ s image option is empty.

Column Identifiers

Column identifiers take any of the following forms:

- A symbolic name from the list of columns option.
- An integer *n*, specifying the *n*th data column.
- A string of the form *#n*, where *n* is an integer, specifying the *n*th display column.

注释:

- Item's option values may be displayed in a different order than the order in which they are stored.
- Column #0 always refers to the tree column, even if `show="tree"` is not specified.

A data column number is an index into an item's option values list; a display column number is the column number in the tree where the values are displayed. Tree labels are displayed in column #0. If option `displaycolumns` is not set, then data column *n* is displayed in column *#n+1*. Again, **column #0 always refers to the tree column.**

Virtual Events

The Treeview widget generates the following virtual events.

Event	描述
<<TreeviewSelect>>	Generated whenever the selection changes.
<<TreeviewOpen>>	Generated just before settings the focus item to <code>open=True</code> .
<<TreeviewClose>>	Generated just after setting the focus item to <code>open=False</code> .

The `Treeview.focus()` and `Treeview.selection()` methods can be used to determine the affected item or items.

ttk.Treeview

```
class tkinter.ttk.Treeview
```

bbox (*item*, *column=None*)

Returns the bounding box (relative to the treeview widget's window) of the specified *item* in the form (x, y, width, height).

If *column* is specified, returns the bounding box of that cell. If the *item* is not visible (i.e., if it is a descendant of a closed item or is scrolled offscreen), returns an empty string.

get_children (*item=None*)

Returns the list of children belonging to *item*.

If *item* is not specified, returns root children.

set_children (*item*, **newchildren*)

Replaces *item*'s child with *newchildren*.

Children present in *item* that are not present in *newchildren* are detached from the tree. No items in *newchildren* may be an ancestor of *item*. Note that not specifying *newchildren* results in detaching *item*'s children.

column (*column*, *option=None*, ***kw*)

Query or modify the options for the specified *column*.

If *kw* is not given, returns a dict of the column option values. If *option* is specified then the value for that *option* is returned. Otherwise, sets the options to the corresponding values.

The valid options/values are:

- **id** Returns the column name. This is a read-only option.
- **anchor: One of the standard Tk anchor values.** Specifies how the text in this column should be aligned with respect to the cell.
- **minwidth: width** The minimum width of the column in pixels. The treeview widget will not make the column any smaller than specified by this option when the widget is resized or the user drags a column.
- **stretch: True/False** Specifies whether the column's width should be adjusted when the widget is resized.
- **width: width** The width of the column in pixels.

To configure the tree column, call this with `column = "#0"`

delete (**items*)

Delete all specified *items* and all their descendants.

The root item may not be deleted.

detach (**items*)

Unlinks all of the specified *items* from the tree.

The items and all of their descendants are still present, and may be reinserted at another point in the tree, but will not be displayed.

The root item may not be detached.

exists (*item*)

Returns `True` if the specified *item* is present in the tree.

focus (*item=None*)

If *item* is specified, sets the focus item to *item*. Otherwise, returns the current focus item, or `''` if there is none.

heading (*column, option=None, **kw*)

Query or modify the heading options for the specified *column*.

If *kw* is not given, returns a dict of the heading option values. If *option* is specified then the value for that *option* is returned. Otherwise, sets the options to the corresponding values.

The valid options/values are:

- **text: text** The text to display in the column heading.
- **image: imageName** Specifies an image to display to the right of the column heading.
- **anchor: anchor** Specifies how the heading text should be aligned. One of the standard Tk anchor values.
- **command: callback** A callback to be invoked when the heading label is pressed.

To configure the tree column heading, call this with `column = "#0"`.

identify (*component, x, y*)

Returns a description of the specified *component* under the point given by *x* and *y*, or the empty string if no such *component* is present at that position.

identify_row (*y*)

Returns the item ID of the item at position *y*.

identify_column (*x*)

Returns the data column identifier of the cell at position *x*.

The tree column has ID #0.

identify_region (*x*, *y*)

Returns one of:

region	meaning
heading	Tree heading area.
separator	Space between two columns headings.
tree	The tree area.
cell	A data cell.

Availability: Tk 8.6.

identify_element (*x*, *y*)

Returns the element at position *x*, *y*.

Availability: Tk 8.6.

index (*item*)

Returns the integer index of *item* within its parent's list of children.

insert (*parent*, *index*, *iid=None*, ***kw*)

Creates a new item and returns the item identifier of the newly created item.

parent is the item ID of the parent item, or the empty string to create a new top-level item. *index* is an integer, or the value "end", specifying where in the list of parent's children to insert the new item. If *index* is less than or equal to zero, the new node is inserted at the beginning; if *index* is greater than or equal to the current number of children, it is inserted at the end. If *iid* is specified, it is used as the item identifier; *iid* must not already exist in the tree. Otherwise, a new unique identifier is generated.

See [Item Options](#) for the list of available points.

item (*item*, *option=None*, ***kw*)

Query or modify the options for the specified *item*.

If no options are given, a dict with options/values for the item is returned. If *option* is specified then the value for that option is returned. Otherwise, sets the options to the corresponding values as given by *kw*.

move (*item*, *parent*, *index*)

Moves *item* to position *index* in *parent*'s list of children.

It is illegal to move an item under one of its descendants. If *index* is less than or equal to zero, *item* is moved to the beginning; if greater than or equal to the number of children, it is moved to the end. If *item* was detached it is reattached.

next (*item*)

Returns the identifier of *item*'s next sibling, or `''` if *item* is the last child of its parent.

parent (*item*)

Returns the ID of the parent of *item*, or `''` if *item* is at the top level of the hierarchy.

prev (*item*)

Returns the identifier of *item*'s previous sibling, or `''` if *item* is the first child of its parent.

reattach (*item*, *parent*, *index*)

An alias for [Treeview.move\(\)](#).

see (*item*)

Ensure that *item* is visible.

Sets all of *item*'s ancestors open option to `True`, and scrolls the widget if necessary so that *item* is within the visible portion of the tree.

selection (*selop=None, items=None*)

If *selop* is not specified, returns selected items. Otherwise, it will act according to the following selection methods.

Deprecated since version 3.6, will be removed in version 3.8: Using `selection()` for changing the selection state is deprecated. Use the following selection methods instead.

selection_set (**items*)

items becomes the new selection.

在 3.6 版更改: *items* can be passed as separate arguments, not just as a single tuple.

selection_add (**items*)

Add *items* to the selection.

在 3.6 版更改: *items* can be passed as separate arguments, not just as a single tuple.

selection_remove (**items*)

Remove *items* from the selection.

在 3.6 版更改: *items* can be passed as separate arguments, not just as a single tuple.

selection_toggle (**items*)

Toggle the selection state of each item in *items*.

在 3.6 版更改: *items* can be passed as separate arguments, not just as a single tuple.

set (*item, column=None, value=None*)

With one argument, returns a dictionary of column/value pairs for the specified *item*. With two arguments, returns the current value of the specified *column*. With three arguments, sets the value of given *column* in given *item* to the specified *value*.

tag_bind (*tagname, sequence=None, callback=None*)

Bind a callback for the given event *sequence* to the tag *tagname*. When an event is delivered to an item, the callbacks for each of the item's tags option are called.

tag_configure (*tagname, option=None, **kw*)

Query or modify the options for the specified *tagname*.

If *kw* is not given, returns a dict of the option settings for *tagname*. If *option* is specified, returns the value for that *option* for the specified *tagname*. Otherwise, sets the options to the corresponding values for the given *tagname*.

tag_has (*tagname, item=None*)

If *item* is specified, returns 1 or 0 depending on whether the specified *item* has the given *tagname*. Otherwise, returns a list of all items that have the specified tag.

Availability: Tk 8.6

xview (**args*)

Query or modify horizontal position of the treeview.

yview (**args*)

Query or modify vertical position of the treeview.

25.2.10 Ttk Styling

Each widget in `ttk` is assigned a style, which specifies the set of elements making up the widget and how they are arranged, along with dynamic and default settings for element options. By default the style name is the same as the widget's class

name, but it may be overridden by the widget's style option. If you don't know the class name of a widget, use the method `Misc.winfo_class()` (`somewidget.winfo_class()`).

参见:

Tcl '2004 conference presentation This document explains how the theme engine works

class `tkinter.ttk.Style`

This class is used to manipulate the style database.

configure (*style*, *query_opt=None*, ***kw*)

Query or set the default value of the specified option(s) in *style*.

Each key in *kw* is an option and each value is a string identifying the value for that option.

For example, to change every default button to be a flat button with some padding and a different background color:

```
from tkinter import ttk
import tkinter

root = tkinter.Tk()

ttk.Style().configure("TButton", padding=6, relief="flat",
                      background="#ccc")

btn = ttk.Button(text="Sample")
btn.pack()

root.mainloop()
```

map (*style*, *query_opt=None*, ***kw*)

Query or sets dynamic values of the specified option(s) in *style*.

Each key in *kw* is an option and each value should be a list or a tuple (usually) containing statespecs grouped in tuples, lists, or some other preference. A statespec is a compound of one or more states and then a value.

An example may make it more understandable:

```
import tkinter
from tkinter import ttk

root = tkinter.Tk()

style = ttk.Style()
style.map("C.TButton",
         foreground=[('pressed', 'red'), ('active', 'blue')],
         background=[('pressed', '!disabled', 'black'), ('active', 'white')]
        )

colored_btn = ttk.Button(text="Test", style="C.TButton").pack()

root.mainloop()
```

Note that the order of the (states, value) sequences for an option does matter, if the order is changed to `[('active', 'blue'), ('pressed', 'red')]` in the foreground option, for example, the result would be a blue foreground when the widget were in active or pressed states.

lookup (*style*, *option*, *state=None*, *default=None*)

Returns the value specified for *option* in *style*.

If *state* is specified, it is expected to be a sequence of one or more states. If the *default* argument is set, it is used as a fallback value in case no specification for option is found.

To check what font a Button uses by default:

```
from tkinter import ttk

print(ttk.Style().lookup("TButton", "font"))
```

layout (*style*, *layoutspec*=None)

Define the widget layout for given *style*. If *layoutspec* is omitted, return the layout specification for given style.

layoutspec, if specified, is expected to be a list or some other sequence type (excluding strings), where each item should be a tuple and the first item is the layout name and the second item should have the format described in [Layouts](#).

To understand the format, see the following example (it is not intended to do anything useful):

```
from tkinter import ttk
import tkinter

root = tkinter.Tk()

style = ttk.Style()
style.layout("TMenubutton", [
    ("Menubutton.background", None),
    ("Menubutton.button", {"children":
        [ ("Menubutton.focus", {"children":
            [ ("Menubutton.padding", {"children":
                [ ("Menubutton.label", {"side": "left", "expand": 1})]
            })]
        })]
    })],
])

mbtn = ttk.Menubutton(text='Text')
mbtn.pack()
root.mainloop()
```

element_create (*elementname*, *etype*, **args*, ***kw*)

Create a new element in the current theme, of the given *etype* which is expected to be either “image”, “from” or “vsapi”. The latter is only available in Tk 8.6a for Windows XP and Vista and is not described here.

If “image” is used, *args* should contain the default image name followed by statespec/value pairs (this is the *imagespec*), and *kw* may have the following options:

- **border=padding** padding is a list of up to four integers, specifying the left, top, right, and bottom borders, respectively.
- **height=height** Specifies a minimum height for the element. If less than zero, the base image’s height is used as a default.
- **padding=padding** Specifies the element’s interior padding. Defaults to border’s value if not specified.
- **sticky=spec** Specifies how the image is placed within the final parcel. spec contains zero or more characters “n”, “s”, “w”, or “e”.
- **width=width** Specifies a minimum width for the element. If less than zero, the base image’s width is used as a default.

If “from” is used as the value of *etype*, `element_create()` will clone an existing element. *args* is expected to contain a *themename*, from which the element will be cloned, and optionally an element to clone from. If this element to clone from is not specified, an empty element will be used. *kw* is discarded.

element_names()

Returns the list of elements defined in the current theme.

element_options (*elementname*)

Returns the list of *elementname*’s options.

theme_create (*themename*, *parent=None*, *settings=None*)

Create a new theme.

It is an error if *themename* already exists. If *parent* is specified, the new theme will inherit styles, elements and layouts from the parent theme. If *settings* are present they are expected to have the same syntax used for `theme_settings()`.

theme_settings (*themename*, *settings*)

Temporarily sets the current theme to *themename*, apply specified *settings* and then restore the previous theme.

Each key in *settings* is a style and each value may contain the keys ‘configure’, ‘map’, ‘layout’ and ‘element create’ and they are expected to have the same format as specified by the methods `Style.configure()`, `Style.map()`, `Style.layout()` and `Style.element_create()` respectively.

As an example, let’s change the Combobox for the default theme a bit:

```
from tkinter import ttk
import tkinter

root = tkinter.Tk()

style = ttk.Style()
style.theme_settings("default", {
    "TCombobox": {
        "configure": {"padding": 5},
        "map": {
            "background": [("active", "green2"),
                           ("!disabled", "green4")],
            "fieldbackground": [("!disabled", "green3")],
            "foreground": [("focus", "OliveDrab1"),
                           ("!disabled", "OliveDrab2")]
        }
    }
})

combo = ttk.Combobox().pack()

root.mainloop()
```

theme_names()

Returns a list of all known themes.

theme_use (*themename=None*)

If *themename* is not given, returns the theme in use. Otherwise, sets the current theme to *themename*, refreshes all widgets and emits a <<ThemeChanged>> event.

Layouts

A layout can be just `None`, if it takes no options, or a dict of options specifying how to arrange the element. The layout mechanism uses a simplified version of the pack geometry manager: given an initial cavity, each element is allocated a parcel. Valid options/values are:

- **side: whichside** Specifies which side of the cavity to place the element; one of top, right, bottom or left. If omitted, the element occupies the entire cavity.
- **sticky: nswe** Specifies where the element is placed inside its allocated parcel.
- **unit: 0 or 1** If set to 1, causes the element and all of its descendants to be treated as a single element for the purposes of `Widget.identify()` et al. It's used for things like scrollbar thumbs with grips.
- **children: [sublayout...]** Specifies a list of elements to place inside the element. Each element is a tuple (or other sequence type) where the first item is the layout name, and the other is a *Layout*.

25.3 tkinter.tix —Extension widgets for Tk

Source code: `Lib/tkinter/tix.py`

3.6 版后已移除: This Tk extension is unmaintained and should not be used in new code. Use `tkinter.ttk` instead.

The `tkinter.tix` (Tk Interface Extension) module provides an additional rich set of widgets. Although the standard Tk library has many useful widgets, they are far from complete. The `tkinter.tix` library provides most of the commonly needed widgets that are missing from standard Tk: `HList`, `ComboBox`, `Control` (a.k.a. `SpinBox`) and an assortment of scrollable widgets. `tkinter.tix` also includes many more widgets that are generally useful in a wide range of applications: `NoteBook`, `FileEntry`, `PanedWindow`, etc; there are more than 40 of them.

With all these new widgets, you can introduce new interaction techniques into applications, creating more useful and more intuitive user interfaces. You can design your application by choosing the most appropriate widgets to match the special needs of your application and users.

参见:

Tix Homepage The home page for Tix. This includes links to additional documentation and downloads.

Tix Man Pages On-line version of the man pages and reference material.

Tix Programming Guide On-line version of the programmer's reference material.

Tix Development Applications Tix applications for development of Tix and Tkinter programs. Tide applications work under Tk or Tkinter, and include **TixInspect**, an inspector to remotely modify and debug Tix/Tk/Tkinter applications.

25.3.1 Using Tix

class `tkinter.tix.Tk` (`screenName=None`, `baseName=None`, `className='Tix'`)

Toplevel widget of Tix which represents mostly the main window of an application. It has an associated Tcl interpreter.

Classes in the `tkinter.tix` module subclasses the classes in the `tkinter`. The former imports the latter, so to use `tkinter.tix` with Tkinter, all you need to do is to import one module. In general, you can just import `tkinter.tix`, and replace the toplevel call to `tkinter.Tk` with `tix.Tk`:

```
from tkinter import tix
from tkinter.constants import *
root = tix.Tk()
```

To use `tkinter.tix`, you must have the Tix widgets installed, usually alongside your installation of the Tk widgets. To test your installation, try the following:

```
from tkinter import tix
root = tix.Tk()
root.tk.eval('package require Tix')
```

If this fails, you have a Tk installation problem which must be resolved before proceeding. Use the environment variable `TIK_LIBRARY` to point to the installed Tix library directory, and make sure you have the dynamic object library (`tix8183.dll` or `libtix8183.so`) in the same directory that contains your Tk dynamic object library (`tk8183.dll` or `libtk8183.so`). The directory with the dynamic object library should also have a file called `pkgIndex.tcl` (case sensitive), which contains the line:

```
package ifneeded Tix 8.1 [list load "[file join $dir tix8183.dll]" Tix]
```

25.3.2 Tix Widgets

Tix introduces over 40 widget classes to the `tkinter` repertoire.

Basic Widgets

class `tkinter.tix.Balloon`

A `Balloon` that pops up over a widget to provide help. When the user moves the cursor inside a widget to which a `Balloon` widget has been bound, a small pop-up window with a descriptive message will be shown on the screen.

class `tkinter.tix.ButtonBox`

The `ButtonBox` widget creates a box of buttons, such as is commonly used for `Ok Cancel`.

class `tkinter.tix.ComboBox`

The `ComboBox` widget is similar to the combo box control in MS Windows. The user can select a choice by either typing in the entry subwidget or selecting from the listbox subwidget.

class `tkinter.tix.Control`

The `Control` widget is also known as the `SpinBox` widget. The user can adjust the value by pressing the two arrow buttons or by entering the value directly into the entry. The new value will be checked against the user-defined upper and lower limits.

class `tkinter.tix.LabelEntry`

The `LabelEntry` widget packages an entry widget and a label into one mega widget. It can be used to simplify the creation of “entry-form” type of interface.

class `tkinter.tix.LabelFrame`

The `LabelFrame` widget packages a frame widget and a label into one mega widget. To create widgets inside a `LabelFrame` widget, one creates the new widgets relative to the `frame` subwidget and manage them inside the `frame` subwidget.

class `tkinter.tix.Meter`

The `Meter` widget can be used to show the progress of a background job which may take a long time to execute.

class `tkinter.tix.OptionMenu`

The `OptionMenu` creates a menu button of options.

class `tkinter.tix.PopupMenu`

The `PopupMenu` widget can be used as a replacement of the `tk_popup` command. The advantage of the `Tix PopupMenu` widget is it requires less application code to manipulate.

class `tkinter.tix.Select`

The `Select` widget is a container of button subwidgets. It can be used to provide radio-box or check-box style of selection options for the user.

class `tkinter.tix.StdButtonBox`

The `StdButtonBox` widget is a group of standard buttons for Motif-like dialog boxes.

File Selectors

class `tkinter.tix.DirList`

The `DirList` widget displays a list view of a directory, its previous directories and its sub-directories. The user can choose one of the directories displayed in the list or change to another directory.

class `tkinter.tix.DirTree`

The `DirTree` widget displays a tree view of a directory, its previous directories and its sub-directories. The user can choose one of the directories displayed in the list or change to another directory.

class `tkinter.tix.DirSelectDialog`

The `DirSelectDialog` widget presents the directories in the file system in a dialog window. The user can use this dialog window to navigate through the file system to select the desired directory.

class `tkinter.tix.DirSelectBox`

The `DirSelectBox` is similar to the standard Motif(TM) directory-selection box. It is generally used for the user to choose a directory. `DirSelectBox` stores the directories mostly recently selected into a `ComboBox` widget so that they can be quickly selected again.

class `tkinter.tix.ExFileSelectBox`

The `ExFileSelectBox` widget is usually embedded in a `tixExFileSelectDialog` widget. It provides a convenient method for the user to select files. The style of the `ExFileSelectBox` widget is very similar to the standard file dialog on MS Windows 3.1.

class `tkinter.tix.FileSelectBox`

The `FileSelectBox` is similar to the standard Motif(TM) file-selection box. It is generally used for the user to choose a file. `FileSelectBox` stores the files mostly recently selected into a `ComboBox` widget so that they can be quickly selected again.

class `tkinter.tix.FileEntry`

The `FileEntry` widget can be used to input a filename. The user can type in the filename manually. Alternatively, the user can press the button widget that sits next to the entry, which will bring up a file selection dialog.

Hierarchical ListBox

class `tkinter.tix.HList`

The `HList` widget can be used to display any data that have a hierarchical structure, for example, file system directory trees. The list entries are indented and connected by branch lines according to their places in the hierarchy.

class `tkinter.tix.CheckList`

The `CheckList` widget displays a list of items to be selected by the user. `CheckList` acts similarly to the Tk check-button or radiobutton widgets, except it is capable of handling many more items than checkbuttons or radiobuttons.

class `tkinter.tix.Tree`

The `Tree` widget can be used to display hierarchical data in a tree form. The user can adjust the view of the tree by opening or closing parts of the tree.

Tabular ListBox

class `tkinter.tix.TList`

The `TList` widget can be used to display data in a tabular format. The list entries of a `TList` widget are similar to the entries in the Tk listbox widget. The main differences are (1) the `TList` widget can display the list entries in a two dimensional format and (2) you can use graphical images as well as multiple colors and fonts for the list entries.

Manager Widgets

class `tkinter.tix.PanedWindow`

The `PanedWindow` widget allows the user to interactively manipulate the sizes of several panes. The panes can be arranged either vertically or horizontally. The user changes the sizes of the panes by dragging the resize handle between two panes.

class `tkinter.tix.ListNoteBook`

The `ListNoteBook` widget is very similar to the `TixNoteBook` widget: it can be used to display many windows in a limited space using a notebook metaphor. The notebook is divided into a stack of pages (windows). At one time only one of these pages can be shown. The user can navigate through these pages by choosing the name of the desired page in the `hlist` subwidget.

class `tkinter.tix.NoteBook`

The `NoteBook` widget can be used to display many windows in a limited space using a notebook metaphor. The notebook is divided into a stack of pages. At one time only one of these pages can be shown. The user can navigate through these pages by choosing the visual “tabs” at the top of the `NoteBook` widget.

Image Types

The `tkinter.tix` module adds:

- `pixmap` capabilities to all `tkinter.tix` and `tkinter` widgets to create color images from XPM files.
- `Compound` image types can be used to create images that consists of multiple horizontal lines; each line is composed of a series of items (texts, bitmaps, images or spaces) arranged from left to right. For example, a compound image can be used to display a bitmap and a text string simultaneously in a Tk `Button` widget.

Miscellaneous Widgets

class `tkinter.tix.InputOnly`

The `InputOnly` widgets are to accept inputs from the user, which can be done with the `bind` command (Unix only).

Form Geometry Manager

In addition, `tkinter.tix` augments `tkinter` by providing:

class `tkinter.tix.Form`

The `Form` geometry manager based on attachment rules for all Tk widgets.

25.3.3 Tix Commands

class `tkinter.tix.tixCommand`

The `tix commands` provide access to miscellaneous elements of `Tix`'s internal state and the `Tix` application

context. Most of the information manipulated by these methods pertains to the application as a whole, or to a screen or display, rather than to a particular window.

To view the current settings, the common usage is:

```
from tkinter import tix
root = tix.Tk()
print(root.tix_configure())
```

`tixCommand.tix_configure` (*cnf=None, **kw*)

Query or modify the configuration options of the Tix application context. If no option is specified, returns a dictionary all of the available options. If option is specified with no value, then the method returns a list describing the one named option (this list will be identical to the corresponding sublist of the value returned if no option is specified). If one or more option-value pairs are specified, then the method modifies the given option(s) to have the given value(s); in this case the method returns an empty string. Option may be any of the configuration options.

`tixCommand.tix_cget` (*option*)

Returns the current value of the configuration option given by *option*. Option may be any of the configuration options.

`tixCommand.tix_getbitmap` (*name*)

Locates a bitmap file of the name *name*.xpm or *name* in one of the bitmap directories (see the `tix_addbitmapdir()` method). By using `tix_getbitmap()`, you can avoid hard coding the pathnames of the bitmap files in your application. When successful, it returns the complete pathname of the bitmap file, prefixed with the character @. The returned value can be used to configure the `bitmap` option of the Tk and Tix widgets.

`tixCommand.tix_addbitmapdir` (*directory*)

Tix maintains a list of directories under which the `tix_getimage()` and `tix_getbitmap()` methods will search for image files. The standard bitmap directory is `$TIX_LIBRARY/bitmaps`. The `tix_addbitmapdir()` method adds *directory* into this list. By using this method, the image files of an applications can also be located using the `tix_getimage()` or `tix_getbitmap()` method.

`tixCommand.tix_filedialog` (*[dlgclass]*)

Returns the file selection dialog that may be shared among different calls from this application. This method will create a file selection dialog widget when it is called the first time. This dialog will be returned by all subsequent calls to `tix_filedialog()`. An optional *dlgclass* parameter can be passed as a string to specify what type of file selection dialog widget is desired. Possible options are `tix`, `FileSelectDialog` or `tixExFileSelectDialog`.

`tixCommand.tix_getimage` (*self, name*)

Locates an image file of the name *name*.xpm, *name*.xbm or *name*.ppm in one of the bitmap directories (see the `tix_addbitmapdir()` method above). If more than one file with the same name (but different extensions) exist, then the image type is chosen according to the depth of the X display: xbm images are chosen on monochrome displays and color images are chosen on color displays. By using `tix_getimage()`, you can avoid hard coding the pathnames of the image files in your application. When successful, this method returns the name of the newly created image, which can be used to configure the `image` option of the Tk and Tix widgets.

`tixCommand.tix_option_get` (*name*)

Gets the options maintained by the Tix scheme mechanism.

`tixCommand.tix_resetoptions` (*newScheme, newFontSet[, newScmPrio]*)

Resets the scheme and fontset of the Tix application to *newScheme* and *newFontSet*, respectively. This affects only those widgets created after this call. Therefore, it is best to call the `resetoptions` method before the creation of any widgets in a Tix application.

The optional parameter *newScmPrio* can be given to reset the priority level of the Tk options set by the Tix schemes.

Because of the way Tk handles the X option database, after Tix has been imported and initied, it is not possible to

reset the color schemes and font sets using the `tix_config()` method. Instead, the `tix_resetoptions()` method must be used.

25.4 `tkinter.scrolledtext` — 滚动文字控件

源代码: [Lib/tkinter/scrolledtext.py](#)

`tkinter.scrolledtext` 模块提供一个同名的类, 实现了一个带有垂直滚动条的文字控件。使用 `ScrolledText` 类会比直接配置一个文本控件和滚动条简单。它的构造函数与 `tkinter.Text` 类相同。

文本控件与滚动条打包在一个 `Frame` 中, `Grid` 方法和 `Pack` 方法的布局管理器从 `Frame` 对象中获得。这允许 `ScrolledText` 控件可以直接用于实现大多数正常的布局管理行为。

如果需要更具体的控制, 可以使用以下属性:

`ScrolledText.frame`
围绕文本和滚动条控件的框架。

`ScrolledText.vbar`
滚动条控件。

25.5 IDLE

源代码: [Lib/idlelib/](#)

IDLE 是 Python 所内置的开发与学习环境。

IDLE 具有以下特性:

- 编码于 100% 纯正的 Python, 使用名为 `tkinter` 的图形用户界面工具
- 跨平台: 在 Windows、Unix 和 macOS 上工作近似。
- 提供输入输出高亮和错误信息的 Python 命令行窗口 (交互解释器)
- 提供多次撤销操作、Python 语法高亮、智能缩进、函数调用提示、自动补全等功能的多窗口文本编辑器
- 在多个窗口中检索, 在编辑器中替换文本, 以及在多个文件中检索 (通过 `grep`)
- 提供持久保存的断点调试、单步调试、查看本地和全局命名空间功能的调试器
- 配置、浏览以及其它对话框

25.5.1 目录

IDLE 具有两个主要窗口类型, 分别是命令行窗口和编辑器窗口。用户可以同时打开多个编辑器窗口。对于 Windows 和 Linux 平台, 都有各自的主菜单。如下记录的每个菜单标识着与之关联的窗口类型。

导出窗口, 例如使用编辑 => 在文件中查找是编辑器窗口的的一个子类型。它们目前有着相同的主菜单, 但是默认标题和上下文菜单不同。

On macOS, there is one application menu. It dynamically changes according to the window currently selected. It has an IDLE menu, and some entries described below are moved around to conform to Apple guidelines.

文件目录（命令行和编辑器）

新建文件 创建一个文件编辑器窗口。

打开… 使用打开窗口以打开一个已存在的文件。

近期文件 打开一个近期文件列表，选取一个以打开它。

打开模块… 打开一个已存在的模块（搜索 `sys.path`）

类浏览器 于当前所编辑的文件中使用树形结构展示函数、类以及方法。在命令行中，首先打开一个模块。

路径浏览 在树状结构中展示 `sys.path` 目录、模块、函数、类和方法。

保存 如果文件已经存在，则将当前窗口保存至对应的文件。自打开或上次保存之后经过修改的文件的窗口标题栏首尾将出现星号 *。如果没有对应的文件，则使用“另存为”代替。

保存为… 使用“保存为”对话框保存当前窗口。被保存的文件将作为当前窗口新的对应文件。

另存为副本… 保存当前窗口至另一个文件，而不修改当前对应文件。

打印窗口 通过默认打印机打印当前窗口。

关闭 关闭当前窗口（如果未保存则询问）。

退出 关闭所有窗口并退出 IDLE（如果未保存则询问）

编辑目录（命令行和编辑器）

撤销操作 撤销当前窗口的最近一次操作。最高可以撤回 1000 条操作记录。

重做 重做当前窗口最近一次所撤销的操作。

剪切 复制选区至系统剪贴板，然后删除选区。

复制 复制选区至系统剪贴板。

粘贴 插入系统剪贴板的内容至当前窗口。

剪贴板功能也可用于上下文目录。

全选 选择当前窗口的全部内容。

查找… 打开一个提供多选项的查找窗口。

再次查找 重复上次搜索，如果结果存在。

查找选区 查找当前选中的字符串，如果存在

在文件中查找… 打开文件查找对话框。将结果输出至新的输出窗口。

替换… 打开查找并替换对话框。

前往行 将光标移动至请求的行编号，并使其恢复可见。

Show Completions Open a scrollable list allowing selection of keywords and attributes. See Completions in the Tips sections below.

展开文本 Expand a prefix you have typed to match a full word in the same window; repeat to get a different expansion.

显示调用贴士 After an unclosed parenthesis for a function, open a small window with function parameter hints.

Show surrounding parens Highlight the surrounding parenthesis.

Format menu (Editor window only)

Indent Region Shift selected lines right by the indent width (default 4 spaces).

Dedent Region Shift selected lines left by the indent width (default 4 spaces).

Comment Out Region Insert `##` in front of selected lines.

Uncomment Region Remove leading `#` or `##` from selected lines.

Tabify Region Turn *leading* stretches of spaces into tabs. (Note: We recommend using 4 space blocks to indent Python code.)

Untabify Region Turn *all* tabs into the correct number of spaces.

Toggle Tabs Open a dialog to switch between indenting with spaces and tabs.

New Indent Width Open a dialog to change indent width. The accepted default by the Python community is 4 spaces.

Format Paragraph Reformat the current blank-line-delimited paragraph in comment block or multiline string or selected line in a string. All lines in the paragraph will be formatted to less than N columns, where N defaults to 72.

Strip trailing whitespace Remove trailing space and other whitespace characters after the last non-whitespace character of a line by applying `str.rstrip` to each line, including lines within multiline strings.

Run menu (Editor window only)

Python Shell Open or wake up the Python Shell window.

Check Module Check the syntax of the module currently open in the Editor window. If the module has not been saved IDLE will either prompt the user to save or autosave, as selected in the General tab of the Idle Settings dialog. If there is a syntax error, the approximate location is indicated in the Editor window.

Run Module Do Check Module (above). If no error, restart the shell to clean the environment, then execute the module. Output is displayed in the Shell window. Note that output requires use of `print` or `write`. When execution is complete, the Shell retains focus and displays a prompt. At this point, one may interactively explore the result of execution. This is similar to executing a file with `python -i file` at a command line.

Shell menu (Shell window only)

View Last Restart Scroll the shell window to the last Shell restart.

Restart Shell Restart the shell to clean the environment.

Interrupt Execution Stop a running program.

Debug menu (Shell window only)

Go to File/Line Look on the current line, with the cursor, and the line above for a filename and line number. If found, open the file if not already open, and show the line. Use this to view source lines referenced in an exception traceback and lines found by Find in Files. Also available in the context menu of the Shell window and Output windows.

Debugger (toggle) When activated, code entered in the Shell or run from an Editor will run under the debugger. In the Editor, breakpoints can be set with the context menu. This feature is still incomplete and somewhat experimental.

Stack Viewer Show the stack traceback of the last exception in a tree widget, with access to locals and globals.

Auto-open Stack Viewer Toggle automatically opening the stack viewer on an unhandled exception.

Options menu (Shell and Editor)

Configure IDLE Open a configuration dialog and change preferences for the following: fonts, indentation, keybindings, text color themes, startup windows and size, additional help sources, and extensions (see below). On macOS, open the configuration dialog by selecting Preferences in the application menu. To use a new built-in color theme (IDLE Dark) with older IDLEs, save it as a new custom theme.

Non-default user settings are saved in a .idlerc directory in the user's home directory. Problems caused by bad user configuration files are solved by editing or deleting one or more of the files in .idlerc.

Code Context (toggle)(Editor Window only) Open a pane at the top of the edit window which shows the block context of the code which has scrolled above the top of the window. Clicking a line in this pane exposes that line at the top of the editor.

Window menu (Shell and Editor)

Zoom Height Toggles the window between normal size and maximum height. The initial size defaults to 40 lines by 80 chars unless changed on the General tab of the Configure IDLE dialog.

The rest of this menu lists the names of all open windows; select one to bring it to the foreground (deiconifying it if necessary).

Help menu (Shell and Editor)

About IDLE Display version, copyright, license, credits, and more.

IDLE Help Display this IDLE document, detailing the menu options, basic editing and navigation, and other tips.

Python Docs Access local Python documentation, if installed, or start a web browser and open docs.python.org showing the latest Python documentation.

Turtle Demo Run the turtledemo module with example python code and turtle drawings.

Additional help sources may be added here with the Configure IDLE dialog under the General tab. See the “Help sources” subsection below for more on Help menu choices.

Context Menus

Open a context menu by right-clicking in a window (Control-click on macOS). Context menus have the standard clipboard functions also on the Edit menu.

剪切 复制选区至系统剪贴板，然后删除选区。

复制 复制选区至系统剪贴板。

粘贴 插入系统剪贴板的内容至当前窗口。

Editor windows also have breakpoint functions. Lines with a breakpoint set are specially marked. Breakpoints only have an effect when running under the debugger. Breakpoints for a file are saved in the user's .idlerc directory.

Set Breakpoint Set a breakpoint on the current line.

Clear Breakpoint Clear the breakpoint on that line.

Shell and Output windows also have the following.

Go to file/line Same as in Debug menu.

The Shell window also has an output squeezing facility explained in the the *Python Shell window* subsection below.

Squeeze If the cursor is over an output line, squeeze all the output between the code above and the prompt below down to a 'Squeezed text' label.

25.5.2 Editing and navigation

Editor windows

IDLE may open editor windows when it starts, depending on settings and how you start IDLE. Thereafter, use the File menu. There can be only one open editor window for a given file.

The title bar contains the name of the file, the full path, and the version of Python and IDLE running the window. The status bar contains the line number ('Ln') and column number ('Col'). Line numbers start with 1; column numbers with 0.

IDLE assumes that files with a known .py* extension contain Python code and that other files do not. Run Python code with the Run menu.

Key bindings

In this section, 'C' refers to the `Control` key on Windows and Unix and the `Command` key on macOS.

- `Backspace` deletes to the left; `Del` deletes to the right
- `C-Backspace` delete word left; `C-Del` delete word to the right
- Arrow keys and `Page Up/Page Down` to move around
- `C-LeftArrow` and `C-RightArrow` moves by words
- `Home/End` go to begin/end of line
- `C-Home/C-End` go to begin/end of file
- Some useful Emacs bindings are inherited from Tcl/Tk:
 - `C-a` beginning of line
 - `C-e` end of line
 - `C-k` kill line (but doesn't put it in clipboard)
 - `C-l` center window around the insertion point
 - `C-b` go backward one character without deleting (usually you can also use the cursor key for this)
 - `C-f` go forward one character without deleting (usually you can also use the cursor key for this)
 - `C-p` go up one line (usually you can also use the cursor key for this)
 - `C-d` delete next character

Standard keybindings (like `C-c` to copy and `C-v` to paste) may work. Keybindings are selected in the Configure IDLE dialog.

Automatic indentation

After a block-opening statement, the next line is indented by 4 spaces (in the Python Shell window by one tab). After certain keywords (`break`, `return` etc.) the next line is dedented. In leading indentation, `Backspace` deletes up to 4 spaces if they are there. `Tab` inserts spaces (in the Python Shell window one tab), number depends on Indent width. Currently, tabs are restricted to four spaces due to Tcl/Tk limitations.

See also the indent/dedent region commands in the edit menu.

Completions

Completions are supplied for functions, classes, and attributes of classes, both built-in and user-defined. Completions are also provided for filenames.

The AutoCompleteWindow (ACW) will open after a predefined delay (default is two seconds) after a `'.'` or (in a string) an `os.sep` is typed. If after one of those characters (plus zero or more other characters) a tab is typed the ACW will open immediately if a possible continuation is found.

If there is only one possible completion for the characters entered, a Tab will supply that completion without opening the ACW.

`'Show Completions'` will force open a completions window, by default the `C-space` will open a completions window. In an empty string, this will contain the files in the current directory. On a blank line, it will contain the built-in and user-defined functions and classes in the current namespaces, plus any modules imported. If some characters have been entered, the ACW will attempt to be more specific.

If a string of characters is typed, the ACW selection will jump to the entry most closely matching those characters. Entering a tab will cause the longest non-ambiguous match to be entered in the Editor window or Shell. Two tabs in a row will supply the current ACW selection, as will return or a double click. Cursor keys, Page Up/Down, mouse selection, and the scroll wheel all operate on the ACW.

“Hidden” attributes can be accessed by typing the beginning of hidden name after a `'.'`, e.g. `'_'`. This allows access to modules with `__all__` set, or to class-private attributes.

Completions and the `'Expand Word'` facility can save a lot of typing!

Completions are currently limited to those in the namespaces. Names in an Editor window which are not via `__main__` and `sys.modules` will not be found. Run the module once with your imports to correct this situation. Note that IDLE itself places quite a few modules in `sys.modules`, so much can be found by default, e.g. the `re` module.

If you don't like the ACW popping up unbidden, simply make the delay longer or disable the extension.

Calltips

A calltip is shown when one types `(` after the name of an *accessible* function. A name expression may include dots and subscripts. A calltip remains until it is clicked, the cursor is moved out of the argument area, or `)` is typed. When the cursor is in the argument part of a definition, the menu or shortcut display a calltip.

A calltip consists of the function signature and the first line of the docstring. For builtins without an accessible signature, the calltip consists of all lines up the fifth line or the first blank line. These details may change.

The set of *accessible* functions depends on what modules have been imported into the user process, including those imported by Idle itself, and what definitions have been run, all since the last restart.

For example, restart the Shell and enter `itertools.count()`. A calltip appears because Idle imports `itertools` into the user process for its own use. (This could change.) Enter `turtle.write()` and nothing appears. Idle does not import `turtle`. The menu or shortcut do nothing either. Enter `import turtle` and then `turtle.write()` will work.

In an editor, import statements have no effect until one runs the file. One might want to run a file after writing the import statements at the top, or immediately run an existing file before editing.

Python Shell window

With IDLE's Shell, one enters, edits, and recalls complete statements. Most consoles and terminals only work with a single physical line at a time.

When one pastes code into Shell, it is not compiled and possibly executed until one hits `Return`. One may edit pasted code first. If one pastes more than one statement into Shell, the result will be a `SyntaxError` when multiple statements are compiled as if they were one.

The editing features described in previous subsections work when entering code interactively. IDLE's Shell window also responds to the following keys.

- `C-c` interrupts executing command
- `C-d` sends end-of-file; closes window if typed at a `>>>` prompt
- `Alt-/` (Expand word) is also useful to reduce typing

Command history

- `Alt-p` retrieves previous command matching what you have typed. On macOS use `C-p`.
- `Alt-n` retrieves next. On macOS use `C-n`.
- `Return` while on any previous command retrieves that command

Text colors

Idle defaults to black on white text, but colors text with special meanings. For the shell, these are shell output, shell error, user output, and user error. For Python code, at the shell prompt or in an editor, these are keywords, builtin class and function names, names following `class` and `def`, strings, and comments. For any text window, these are the cursor (when present), found text (when possible), and selected text.

Text coloring is done in the background, so uncolored text is occasionally visible. To change the color scheme, use the Configure IDLE dialog Highlighting tab. The marking of debugger breakpoint lines in the editor and text in popups and dialogs is not user-configurable.

25.5.3 Startup and code execution

Upon startup with the `-s` option, IDLE will execute the file referenced by the environment variables `IDLESTARTUP` or `PYTHONSTARTUP`. IDLE first checks for `IDLESTARTUP`; if `IDLESTARTUP` is present the file referenced is run. If `IDLESTARTUP` is not present, IDLE checks for `PYTHONSTARTUP`. Files referenced by these environment variables are convenient places to store functions that are used frequently from the IDLE shell, or for executing import statements to import common modules.

In addition, Tk also loads a startup file if it is present. Note that the Tk file is loaded unconditionally. This additional file is `.Idle.py` and is looked for in the user's home directory. Statements in this file will be executed in the Tk namespace, so this file is not useful for importing functions to be used from IDLE's Python shell.

Command line usage

```
idle.py [-c command] [-d] [-e] [-h] [-i] [-r file] [-s] [-t title] [-] [arg] ...

-c command    run command in the shell window
-d            enable debugger and open shell window
-e            open editor window
-h            print help message with legal combinations and exit
-i            open shell window
-r file       run file in shell window
-s            run $IDLESTARTUP or $PYTHONSTARTUP first, in shell window
-t title      set title of shell window
-            run stdin in shell (- must be last option before args)
```

If there are arguments:

- If `-`, `-c`, or `r` is used, all arguments are placed in `sys.argv[1:..]` and `sys.argv[0]` is set to `' '`, `'-c'`, or `'-r'`. No editor window is opened, even if that is the default set in the Options dialog.
- Otherwise, arguments are files opened for editing and `sys.argv` reflects the arguments passed to IDLE itself.

Startup failure

IDLE uses a socket to communicate between the IDLE GUI process and the user code execution process. A connection must be established whenever the Shell starts or restarts. (The latter is indicated by a divider line that says ‘RESTART’). If the user process fails to connect to the GUI process, it displays a Tk error box with a ‘cannot connect’ message that directs the user here. It then exits.

A common cause of failure is a user-written file with the same name as a standard library module, such as *random.py* and *tkinter.py*. When such a file is located in the same directory as a file that is about to be run, IDLE cannot import the stdlib file. The current fix is to rename the user file.

Though less common than in the past, an antivirus or firewall program may stop the connection. If the program cannot be taught to allow the connection, then it must be turned off for IDLE to work. It is safe to allow this internal connection because no data is visible on external ports. A similar problem is a network mis-configuration that blocks connections.

Python installation issues occasionally stop IDLE: multiple versions can clash, or a single installation might need admin access. If one undo the clash, or cannot or does not want to run as admin, it might be easiest to completely remove Python and start over.

A zombie pythonw.exe process could be a problem. On Windows, use Task Manager to detect and stop one. Sometimes a restart initiated by a program crash or Keyboard Interrupt (control-C) may fail to connect. Dismissing the error box or Restart Shell on the Shell menu may fix a temporary problem.

When IDLE first starts, it attempts to read user configuration files in `~/.idlerc/` (`~` is one’s home directory). If there is a problem, an error message should be displayed. Leaving aside random disk glitches, this can be prevented by never editing the files by hand, using the configuration dialog, under Options, instead Options. Once it happens, the solution may be to delete one or more of the configuration files.

If IDLE quits with no message, and it was not started from a console, try starting from a console (`python -m idlelib`) and see if a message appears.

Running user code

With rare exceptions, the result of executing Python code with IDLE is intended to be the same as executing the same code by the default method, directly with Python in a text-mode system console or terminal window. However, the different interface and operation occasionally affect visible results. For instance, `sys.modules` starts with more entries, and `threading.activeCount()` returns 2 instead of 1.

By default, IDLE runs user code in a separate OS process rather than in the user interface process that runs the shell and editor. In the execution process, it replaces `sys.stdin`, `sys.stdout`, and `sys.stderr` with objects that get input from and send output to the Shell window. The original values stored in `sys.__stdin__`, `sys.__stdout__`, and `sys.__stderr__` are not touched, but may be `None`.

When Shell has the focus, it controls the keyboard and screen. This is normally transparent, but functions that directly access the keyboard and screen will not work. These include system-specific functions that determine whether a key has been pressed and if so, which.

IDLE’s standard stream replacements are not inherited by subprocesses created in the execution process, whether directly by user code or by modules such as multiprocessing. If such subprocess use `input` from `sys.stdin` or `print` or `write` to `sys.stdout` or `sys.stderr`, IDLE should be started in a command line window. The secondary subprocess will then be attached to that window for input and output.

If `sys` is reset by user code, such as with `importlib.reload(sys)`, IDLE's changes are lost and input from the keyboard and output to the screen will not work correctly.

User output in Shell

When a program outputs text, the result is determined by the corresponding output device. When IDLE executes user code, `sys.stdout` and `sys.stderr` are connected to the display area of IDLE's Shell. Some of its features are inherited from the underlying Tk Text widget. Others are programmed additions. Where it matters, Shell is designed for development rather than production runs.

For instance, Shell never throws away output. A program that sends unlimited output to Shell will eventually fill memory, resulting in a memory error. In contrast, some system text windows only keep the last *n* lines of output. A Windows console, for instance, keeps a user-settable 1 to 9999 lines, with 300 the default.

Text widgets display a subset of Unicode, the Basic Multilingual Plane (BMP). Which characters get a proper glyph instead of a replacement box depends on the operating system and installed fonts. Newline characters cause following text to appear on a new line, but other control characters are either replaced with a box or deleted. However, `repr()`, which is used for interactive echo of expression values, replaces control characters, some BMP codepoints, and all non-BMP characters with escape codes before they are output.

Normal and error output are generally kept separate (on separate lines) from code input and each other. They each get different highlight colors.

For SyntaxError tracebacks, the normal '^' marking where the error was detected is replaced by coloring the text with an error highlight. When code run from a file causes other exceptions, one may right click on a traceback line to jump to the corresponding line in an IDLE editor. The file will be opened if necessary.

Shell has a special facility for squeezing output lines down to a 'Squeezed text' label. This is done automatically for output over *N* lines (*N* = 50 by default). *N* can be changed in the PyShell section of the General page of the Settings dialog. Output with fewer lines can be squeezed by right clicking on the output. This can be useful lines long enough to slow down scrolling.

Squeezed output is expanded in place by double-clicking the label. It can also be sent to the clipboard or a separate view window by right-clicking the label.

Developing tkinter applications

IDLE is intentionally different from standard Python in order to facilitate development of tkinter programs. Enter `import tkinter as tk; root = tk.Tk()` in standard Python and nothing appears. Enter the same in IDLE and a tk window appears. In standard Python, one must also enter `root.update()` to see the window. IDLE does the equivalent in the background, about 20 times a second, which is about every 50 milliseconds. Next enter `b = tk.Button(root, text='button'); b.pack()`. Again, nothing visibly changes in standard Python until one enters `root.update()`.

Most tkinter programs run `root.mainloop()`, which usually does not return until the tk app is destroyed. If the program is run with `python -i` or from an IDLE editor, a `>>>` shell prompt does not appear until `mainloop()` returns, at which time there is nothing left to interact with.

When running a tkinter program from an IDLE editor, one can comment out the `mainloop` call. One then gets a shell prompt immediately and can interact with the live application. One just has to remember to re-enable the `mainloop` call when running in standard Python.

Running without a subprocess

By default, IDLE executes user code in a separate subprocess via a socket, which uses the internal loopback interface. This connection is not externally visible and no data is sent to or received from the Internet. If firewall software complains

anyway, you can ignore it.

If the attempt to make the socket connection fails, Idle will notify you. Such failures are sometimes transient, but if persistent, the problem may be either a firewall blocking the connection or misconfiguration of a particular system. Until the problem is fixed, one can run Idle with the `-n` command line switch.

If IDLE is started with the `-n` command line switch it will run in a single process and will not create the subprocess which runs the RPC Python execution server. This can be useful if Python cannot create the subprocess or the RPC socket interface on your platform. However, in this mode user code is not isolated from IDLE itself. Also, the environment is not restarted when Run/Run Module (F5) is selected. If your code has been modified, you must `reload()` the affected modules and re-import any specific items (e.g. `from foo import baz`) if the changes are to take effect. For these reasons, it is preferable to run IDLE with the default subprocess if at all possible.

3.4 版后已移除.

25.5.4 Help and preferences

Help sources

Help menu entry “IDLE Help” displays a formatted html version of the IDLE chapter of the Library Reference. The result, in a read-only tkinter text window, is close to what one sees in a web browser. Navigate through the text with a mousewheel, the scrollbar, or up and down arrow keys held down. Or click the TOC (Table of Contents) button and select a section header in the opened box.

Help menu entry “Python Docs” opens the extensive sources of help, including tutorials, available at docs.python.org/x.y, where ‘x.y’ is the currently running Python version. If your system has an off-line copy of the docs (this may be an installation option), that will be opened instead.

Selected URLs can be added or removed from the help menu at any time using the General tab of the Configure IDLE dialog .

Setting preferences

The font preferences, highlighting, keys, and general preferences can be changed via Configure IDLE on the Option menu. Keys can be user defined; IDLE ships with four built-in key sets. In addition, a user can create a custom key set in the Configure IDLE dialog under the keys tab.

IDLE on macOS

Under System Preferences: Dock, one can set “Prefer tabs when opening documents” to “Always” . This setting is not compatible with the tk/tkinter GUI framework used by IDLE, and it breaks a few IDLE features.

Extensions

IDLE contains an extension facility. Preferences for extensions can be changed with the Extensions tab of the preferences dialog. See the beginning of `config-extensions.def` in the `idlelib` directory for further information. The only current default extension is `zzdummy`, an example also used for testing.

25.6 其他图形用户界面 (GUI) 包

Python 可用的主要跨平台 (Windows, Mac OS X, 类 Unix) GUI 工具:

参见:

PyGObject PyGObject provides introspection bindings for C libraries using [GObject](#). One of these libraries is the [GTK+ 3](#) widget set. GTK+ comes with many more widgets than Tkinter provides. An online [Python GTK+ 3 Tutorial](#) is available.

PyGTK PyGTK provides bindings for an older version of the library, GTK+ 2. It provides an object oriented interface that is slightly higher level than the C one. There are also bindings to [GNOME](#). An online [tutorial](#) is available.

PyQt PyQt is a **sip**-wrapped binding to the Qt toolkit. Qt is an extensive C++ GUI application development framework that is available for Unix, Windows and Mac OS X. **sip** is a tool for generating bindings for C++ libraries as Python classes, and is specifically designed for Python.

PySide PySide is a newer binding to the Qt toolkit, provided by Nokia. Compared to PyQt, its licensing scheme is friendlier to non-open source applications.

wxPython wxPython is a cross-platform GUI toolkit for Python that is built around the popular [wxWidgets](#) (formerly wxWindows) C++ toolkit. It provides a native look and feel for applications on Windows, Mac OS X, and Unix systems by using each platform's native widgets where ever possible, (GTK+ on Unix-like systems). In addition to an extensive set of widgets, wxPython provides classes for online documentation and context sensitive help, printing, HTML viewing, low-level device context drawing, drag and drop, system clipboard access, an XML-based resource format and more, including an ever growing library of user-contributed modules.

PyGTK, PyQt, and wxPython, all have a modern look and feel and more widgets than Tkinter. In addition, there are many other GUI toolkits for Python, both cross-platform, and platform-specific. See the [GUI Programming](#) page in the Python Wiki for a much more complete list, and also for links to documents where the different GUI toolkits are compared.

本章中描述的各模块可帮你编写 Python 程序。例如，`pydoc` 模块接受一个模块并根据该模块的内容来生成文档。`doctest` 和 `unittest` 这两个模块包含了用于编写单元测试的框架，并可用于自动测试所编写的代码，验证预期的输出是否产生。`2to3` 程序能够将 Python 2.x 源代码翻译成有效的 Python 3.x 源代码。

本章中描述的模块列表是：

26.1 typing — 类型标注支持

3.5 新版功能.

源码： `Lib/typing.py`

注解：`typing` 模块以**暂定状态**包含在标准库中。如果核心开发人员认为有必要，可能会添加新功能，甚至可能会在次要版本之间改变 API。

此模块支持 **PEP 484** 和 **PEP 526** 指定的类型提示。最基本的支持由 `Any`, `Union`, `Tuple`, `Callable`, `TypeVar` 和 `Generic` 类型组成。有关完整的规范，请参阅 **PEP 484**。有关类型提示的简单介绍，请参阅 **PEP 483**。

函数接受并返回一个字符串，注释像下面这样：

```
def greeting(name: str) -> str:
    return 'Hello ' + name
```

在函数 `greeting` 中，参数 `name` 预期是 `str` 类型，并且返回 `str` 类型。子类型允许作为参数。

26.1.1 类型别名

类型别名通过将类型分配给别名来定义。在这个例子中，`Vector` 和 `List[float]` 将被视为可互换的同义词：

```
from typing import List
Vector = List[float]

def scale(scalar: float, vector: Vector) -> Vector:
    return [scalar * num for num in vector]

# typechecks; a list of floats qualifies as a Vector.
new_vector = scale(2.0, [1.0, -4.2, 5.4])
```

类型别名可用于简化复杂类型签名。例如：

```
from typing import Dict, Tuple, List

ConnectionOptions = Dict[str, str]
Address = Tuple[str, int]
Server = Tuple[Address, ConnectionOptions]

def broadcast_message(message: str, servers: List[Server]) -> None:
    ...

# The static type checker will treat the previous type signature as
# being exactly equivalent to this one.
def broadcast_message(
    message: str,
    servers: List[Tuple[Tuple[str, int], Dict[str, str]]]) -> None:
    ...
```

请注意，`None` 作为类型提示是一种特殊情况，并且由 `type(None)` 取代。

26.1.2 NewType

使用 `NewType()` 辅助函数创建不同的类型：

```
from typing import NewType

UserId = NewType('UserId', int)
some_id = UserId(524313)
```

静态类型检查器会将新类型视为它是原始类型的子类。这对于帮助捕捉逻辑错误非常有用：

```
def get_user_name(user_id: UserId) -> str:
    ...

# typechecks
user_a = get_user_name(UserId(42351))

# does not typecheck; an int is not a UserId
user_b = get_user_name(-1)
```

您仍然可以对 `UserId` 类型的变量执行所有的 `int` 支持的操作，但结果将始终为 `int` 类型。这可以让你在需要 `int` 的地方传入 `UserId`，但会阻止你以无效的方式无意中创建 `UserId`：

```
# 'output' is of type 'int', not 'UserId'
output = UserId(23413) + UserId(54341)
```

请注意，这些检查仅通过静态类型检查程序强制执行。在运行时，`Derived = NewType('Derived', Base)` 将 `Derived` 一个函数，该函数立即返回您传递它的任何参数。这意味着表达式 `Derived(some_value)` 不会创建一个新的类或引入任何超出常规函数调用的开销。

更确切地说，表达式 `some_value is Derived(some_value)` 在运行时总是为真。

这也意味着无法创建 `Derived` 的子类型，因为它是运行时的标识函数，而不是实际的类型：

```
from typing import NewType

UserId = NewType('UserId', int)

# Fails at runtime and does not typecheck
class AdminUserId(UserId): pass
```

但是，可以基于 'derived' `NewType` 创建 `NewType()`

```
from typing import NewType

UserId = NewType('UserId', int)

ProUserId = NewType('ProUserId', UserId)
```

并且 `ProUserId` 的类型检查将按预期工作。

有关更多详细信息，请参阅 [PEP 484](#)。

注解：回想一下，使用类型别名声明两种类型彼此等效。`Alias = Original` 将使静态类型检查对待所有情况下 `Alias` 完全等同于 `Original`。当您想简化复杂类型签名时，这很有用。

相反，`NewType` 声明一种类型是另一种类型的子类型。`Derived = NewType('Derived', Original)` 将使静态类型检查器将 `Derived` 当作 `Original` 的子类，这意味着 `Original` 类型的值不能用于 `Derived` 类型的值需要的地方。当您想以最小的运行时间成本防止逻辑错误时，这非常有用。

3.5.2 新版功能.

26.1.3 Callable

期望特定签名的回调函数的框架可以将类型标注为 `Callable[[Arg1Type, Arg2Type], ReturnType]`。

例如：

```
from typing import Callable

def feeder(get_next_item: Callable[[], str]) -> None:
    # Body

def async_query(on_success: Callable[[int], None],
               on_error: Callable[[int, Exception], None]) -> None:
    # Body
```

通过用文字省略号替换类型提示中的参数列表: `Callable[..., ReturnType]`, 可以声明可调用的返回类型, 而无需指定调用签名。

26.1.4 泛型 (Generic)

由于无法以通用方式静态推断有关保存在容器中的对象的类型信息, 因此抽象基类已扩展为支持订阅以表示容器元素的预期类型。

```
from typing import Mapping, Sequence

def notify_by_email(employees: Sequence[Employee],
                   overrides: Mapping[str, str]) -> None: ...
```

泛型可以通过使用 `typing` 模块中名为 `TypeVar` 的新工厂进行参数化。

```
from typing import Sequence, TypeVar

T = TypeVar('T')           # Declare type variable

def first(l: Sequence[T]) -> T:    # Generic function
    return l[0]
```

26.1.5 用户定义的泛型类型

用户定义的类可以定义为泛型类。

```
from typing import TypeVar, Generic
from logging import Logger

T = TypeVar('T')

class LoggedVar(Generic[T]):
    def __init__(self, value: T, name: str, logger: Logger) -> None:
        self.name = name
        self.logger = logger
        self.value = value

    def set(self, new: T) -> None:
        self.log('Set ' + repr(self.value))
        self.value = new

    def get(self) -> T:
        self.log('Get ' + repr(self.value))
        return self.value

    def log(self, message: str) -> None:
        self.logger.info('%s: %s', self.name, message)
```

`Generic[T]` 作为基类定义了类 `LoggedVar` 采用单个类型参数 `T`。这也使得 `T` 作为类体内的一个类型有效。

`Generic` 基类使用定义了 `__getitem__()` 的元类, 以便 `LoggedVar[t]` 作为类型有效:

```
from typing import Iterable

def zero_all_vars(vars: Iterable[LoggedVar[int]]) -> None:
    for var in vars:
        var.set(0)
```

泛型类型可以有任意数量的类型变量，并且类型变量可能会受到限制：

```
from typing import TypeVar, Generic
...

T = TypeVar('T')
S = TypeVar('S', int, str)

class StrangePair(Generic[T, S]):
    ...
```

Generic 每个参数的类型变量必须是不同的。这是无效的：

```
from typing import TypeVar, Generic
...

T = TypeVar('T')

class Pair(Generic[T, T]):    # INVALID
    ...
```

您可以对 *Generic* 使用多重继承：

```
from typing import TypeVar, Generic, Sized

T = TypeVar('T')

class LinkedList(Sized, Generic[T]):
    ...
```

从泛型类继承时，某些类型变量可能是固定的：

```
from typing import TypeVar, Mapping

T = TypeVar('T')

class MyDict(Mapping[str, T]):
    ...
```

在这种情况下，`MyDict` 只有一个参数，`T`。

在不指定类型参数的情况下使用泛型类别会为每个位置假设 *Any*。在下面的例子中，`MyIterable` 不是泛型，但是隐式继承自 `Iterable[Any]`：

```
from typing import Iterable

class MyIterable(Iterable): # Same as Iterable[Any]
```

用户定义的通用类型别名也受支持。例子：

```

from typing import TypeVar, Iterable, Tuple, Union
S = TypeVar('S')
Response = Union[Iterable[S], int]

# Return type here is same as Union[Iterable[str], int]
def response(query: str) -> Response[str]:
    ...

T = TypeVar('T', int, float, complex)
Vec = Iterable[Tuple[T, T]]

def inproduct(v: Vec[T]) -> T: # Same as Iterable[Tuple[T, T]]
    return sum(x*y for x, y in v)

```

Generic 使用的元类是 `abc.ABCMeta` 的子类。泛型类可以通过包含抽象方法或属性成为 ABC，并且泛型类也可以使用 ABCs 作为基类而不存在元类冲突。不支持泛型元类。参数化泛型的结果被缓存，并且 `typing` 模块中的大部分类型都是可散列的，并且可以比较是否相等。

26.1.6 Any 类型

Any 是一种特殊的类型。静态类型检查器将所有类型视为与 *Any* 兼容，反之亦然，*Any* 也与所有类型相兼容。

这意味着可对类型为 *Any* 的值执行任何操作或方法调用，并将其赋值给任何变量：

```

from typing import Any

a = None      # type: Any
a = []        # OK
a = 2         # OK

s = ''        # type: str
s = a         # OK

def foo(item: Any) -> int:
    # Typechecks; 'item' could be any type,
    # and that type might have a 'bar' method
    item.bar()
    ...

```

需要注意的是，将 *Any* 类型的值赋值给另一个更具体的类型时，Python 不会执行类型检查。例如，当把 `a` 赋值给 `s` 时，即使 `s` 被声明为 *str* 类型，在运行时接收到的是 *int* 值，静态类型检查器也不会报错。

此外，所有返回值无类型或形参无类型的函数将隐式地默认使用 *Any* 类型：

```

def legacy_parser(text):
    ...
    return data

# A static type checker will treat the above
# as having the same signature as:
def legacy_parser(text: Any) -> Any:
    ...
    return data

```

当需要混用动态类型和静态类型的代码时，上述行为可以让 *Any* 被用作 应急出口。

`Any` 和 `object` 的行为对比。与 `Any` 相似，所有的类型都是 `object` 的子类型。然而不同于 `Any`，反之并不成立：`object` 不是其他所有类型的子类型。

这意味着当一个值的类型是 `object` 的时候，类型检查器会拒绝对它的几乎所有的操作。把它赋值给一个指定了类型的变量（或者当作返回值）是一个类型错误。比如说：

```
def hash_a(item: object) -> int:
    # Fails; an object does not have a 'magic' method.
    item.magic()
    ...

def hash_b(item: Any) -> int:
    # Typechecks
    item.magic()
    ...

# Typechecks, since ints and strs are subclasses of object
hash_a(42)
hash_a("foo")

# Typechecks, since Any is compatible with all types
hash_b(42)
hash_b("foo")
```

使用 `object` 示意一个值可以类型安全地兼容任何类型。使用 `Any` 示意一个值的类型是动态定义的。

26.1.7 类, 函数和修饰器.

这个模块定义了如下的类, 模块和修饰器.

class `typing.TypeVar`

类型变量

用法:

```
T = TypeVar('T') # Can be anything
A = TypeVar('A', str, bytes) # Must be str or bytes
```

Type variables exist primarily for the benefit of static type checkers. They serve as the parameters for generic types as well as for generic function definitions. See class `Generic` for more information on generic types. Generic functions work as follows:

```
def repeat(x: T, n: int) -> Sequence[T]:
    """Return a list containing n references to x."""
    return [x]*n

def longest(x: A, y: A) -> A:
    """Return the longest of two strings."""
    return x if len(x) >= len(y) else y
```

The latter example's signature is essentially the overloading of `(str, str) -> str` and `(bytes, bytes) -> bytes`. Also note that if the arguments are instances of some subclass of `str`, the return type is still plain `str`.

`isinstance(x, T)` 会在运行时抛出 `TypeError` 异常。一般地说, `isinstance()` 和 `issubclass()` 不应该和类型一起使用。

Type variables may be marked covariant or contravariant by passing `covariant=True` or `contravariant=True`. See [PEP 484](#) for more details. By default type variables are invariant. Alternatively, a type variable may specify an upper bound using `bound=<type>`. This means that an actual type substituted (explicitly or implicitly) for the type variable must be a subclass of the boundary type, see [PEP 484](#).

`class typing.Generic`

Abstract base class for generic types.

A generic type is typically declared by inheriting from an instantiation of this class with one or more type variables. For example, a generic mapping type might be defined as:

```
class Mapping(Generic[KT, VT]):
    def __getitem__(self, key: KT) -> VT:
        ...
        # Etc.
```

这个类之后可以被这样用:

```
X = TypeVar('X')
Y = TypeVar('Y')

def lookup_name(mapping: Mapping[X, Y], key: X, default: Y) -> Y:
    try:
        return mapping[key]
    except KeyError:
        return default
```

`class typing.Type(Generic[CT_co])`

A variable annotated with `C` may accept a value of type `C`. In contrast, a variable annotated with `Type[C]` may accept values that are classes themselves – specifically, it will accept the *class object* of `C`. For example:

```
a = 3          # Has type 'int'
b = int        # Has type 'Type[int]'
c = type(a)    # Also has type 'Type[int]'
```

Note that `Type[C]` is covariant:

```
class User: ...
class BasicUser(User): ...
class ProUser(User): ...
class TeamUser(User): ...

# Accepts User, BasicUser, ProUser, TeamUser, ...
def make_new_user(user_class: Type[User]) -> User:
    # ...
    return user_class()
```

The fact that `Type[C]` is covariant implies that all subclasses of `C` should implement the same constructor signature and class method signatures as `C`. The type checker should flag violations of this, but should also allow constructor calls in subclasses that match the constructor calls in the indicated base class. How the type checker is required to handle this particular case may change in future revisions of [PEP 484](#).

The only legal parameters for `Type` are classes, *Any*, *type variables*, and unions of any of these types. For example:

```
def new_non_team_user(user_class: Type[Union[BaseUser, ProUser]]): ...
```

`Type[Any]` is equivalent to `Type` which in turn is equivalent to `type`, which is the root of Python's metaclass hierarchy.

3.5.2 新版功能.

```
class typing.Iterable (Generic[T_co])
    collections.abc.Iterable 的泛型版本。

class typing.Iterator (Iterable[T_co])
    collections.abc.Iterator 的泛型版本。

class typing.Reversible (Iterable[T_co])
    collections.abc.Reversible 的泛型版本。

class typing.SupportsInt
    An ABC with one abstract method __int__.

class typing.SupportsFloat
    An ABC with one abstract method __float__.

class typing.SupportsComplex
    An ABC with one abstract method __complex__.

class typing.SupportsBytes
    An ABC with one abstract method __bytes__.

class typing.SupportsAbs
    An ABC with one abstract method __abs__ that is covariant in its return type.

class typing.SupportsRound
    An ABC with one abstract method __round__ that is covariant in its return type.

class typing.Container (Generic[T_co])
    collections.abc.Container 的泛型版本。

class typing.Hashable
    collections.abc.Hashable 的别名。

class typing.Sized
    collections.abc.Sized 的别名。

class typing.Collection (Sized, Iterable[T_co], Container[T_co])
    collections.abc.Collection 的泛型版本。
```

3.6 新版功能.

```
class typing.AbstractSet (Sized, Collection[T_co])
    collections.abc.Set 的泛型版本。

class typing.MutableSet (AbstractSet[T])
    collections.abc.MutableSet 的泛型版本。

class typing.Mapping (Sized, Collection[KT], Generic[VT_co])
    A generic version of collections.abc.Mapping.

class typing.MutableMapping (Mapping[KT, VT])
    collections.abc.MutableMapping 的泛型版本。

class typing.Sequence (Reversible[T_co], Collection[T_co])
    collections.abc.Sequence 的泛型版本。

class typing.MutableSequence (Sequence[T])
    collections.abc.MutableSequence 的泛型版本。

class typing.ByteString (Sequence[int])
    collections.abc.ByteString 的泛型版本。

    This type represents the types bytes, bytearray, and memoryview.
```

As a shorthand for this type, `bytes` can be used to annotate arguments of any of the types mentioned above.

class `typing.Deque` (`deque`, `MutableSequence[T]`)
`collections.deque` 的泛型版本。

3.6.1 新版功能.

class `typing.List` (`list`, `MutableSequence[T]`)
 Generic version of `list`. Useful for annotating return types. To annotate arguments it is preferred to use abstract collection types such as `Mapping`, `Sequence`, or `AbstractSet`.

这个类型可以这样用:

```
T = TypeVar('T', int, float)

def vec2(x: T, y: T) -> List[T]:
    return [x, y]

def keep_positives(vector: Sequence[T]) -> List[T]:
    return [item for item in vector if item > 0]
```

class `typing.Set` (`set`, `MutableSet[T]`)
 A generic version of `builtins.set`.

class `typing.Frozenset` (`frozenset`, `AbstractSet[T_co]`)
 A generic version of `builtins.frozenset`.

class `typing.MappingView` (`Sized`, `Iterable[T_co]`)
`collections.abc.MappingView` 的泛型版本。

class `typing.KeysView` (`MappingView[KT_co]`, `AbstractSet[KT_co]`)
`collections.abc.KeysView` 的泛型版本。

class `typing.ItemsView` (`MappingView`, `Generic[KT_co, VT_co]`)
`collections.abc.ItemsView` 的泛型版本。

class `typing.ValuesView` (`MappingView[VT_co]`)
`collections.abc.ValuesView` 的泛型版本。

class `typing.Awaitable` (`Generic[T_co]`)
`collections.abc.Awaitable` 的泛型版本。

class `typing.Coroutine` (`Awaitable[V_co]`, `Generic[T_co T_contra, V_co]`)
 A generic version of `collections.abc.Coroutine`. The variance and order of type variables correspond to those of `Generator`, for example:

```
from typing import List, Coroutine
c = None # type: Coroutine[List[str], str, int]
...
x = c.send('hi') # type: List[str]
async def bar() -> None:
    x = await c # type: int
```

class `typing.AsyncIterable` (`Generic[T_co]`)
`collections.abc.AsyncIterable` 的泛型版本。

class `typing.AsyncIterator` (`AsyncIterable[T_co]`)
`collections.abc.AsyncIterator` 的泛型版本。

class `typing.ContextManager` (`Generic[T_co]`)
`contextlib.AbstractContextManager` 的泛型版本。

3.6 新版功能.

class `typing.AsyncContextManager` (`Generic[T_co]`)

An ABC with async abstract `__aenter__()` and `__aexit__()` methods.

3.6 新版功能.

class `typing.Dict` (`dict`, `MutableMapping[KT, VT]`)

A generic version of `dict`. The usage of this type is as follows:

```
def get_position_in_index(word_list: Dict[str, int], word: str) -> int:
    return word_list[word]
```

class `typing.DefaultDict` (`collections.defaultdict`, `MutableMapping[KT, VT]`)

`collections.defaultdict` 的泛型版本。

3.5.2 新版功能.

class `typing.Counter` (`collections.Counter`, `Dict[T, int]`)

`collections.Counter` 的泛型版本。

3.6.1 新版功能.

class `typing.ChainMap` (`collections.ChainMap`, `MutableMapping[KT, VT]`)

`collections.ChainMap` 的泛型版本。

3.6.1 新版功能.

class `typing.Generator` (`Iterator[T_co]`, `Generic[T_co, T_contra, V_co]`)

A generator can be annotated by the generic type `Generator[YieldType, SendType, ReturnType]`. For example:

```
def echo_round() -> Generator[int, float, str]:
    sent = yield 0
    while sent >= 0:
        sent = yield round(sent)
    return 'Done'
```

Note that unlike many other generics in the typing module, the `SendType` of `Generator` behaves contravariantly, not covariantly or invariantly.

If your generator will only yield values, set the `SendType` and `ReturnType` to `None`:

```
def infinite_stream(start: int) -> Generator[int, None, None]:
    while True:
        yield start
        start += 1
```

Alternatively, annotate your generator as having a return type of either `Iterable[YieldType]` or `Iterator[YieldType]`:

```
def infinite_stream(start: int) -> Iterator[int]:
    while True:
        yield start
        start += 1
```

class `typing.AsyncGenerator` (`AsyncIterator[T_co]`, `Generic[T_co, T_contra]`)

An async generator can be annotated by the generic type `AsyncGenerator[YieldType, SendType]`. For example:

```

async def echo_round() -> AsyncGenerator[int, float]:
    sent = yield 0
    while sent >= 0.0:
        rounded = await round(sent)
        sent = yield rounded

```

Unlike normal generators, async generators cannot return a value, so there is no `ReturnType` type parameter. As with *Generator*, the `SendType` behaves contravariantly.

If your generator will only yield values, set the `SendType` to `None`:

```

async def infinite_stream(start: int) -> AsyncGenerator[int, None]:
    while True:
        yield start
        start = await increment(start)

```

Alternatively, annotate your generator as having a return type of either `AsyncIterable[YieldType]` or `AsyncIterator[YieldType]`:

```

async def infinite_stream(start: int) -> AsyncIterator[int]:
    while True:
        yield start
        start = await increment(start)

```

3.5.4 新版功能.

class `typing.Text`

`Text` is an alias for `str`. It is provided to supply a forward compatible path for Python 2 code: in Python 2, `Text` is an alias for `unicode`.

Use `Text` to indicate that a value must contain a unicode string in a manner that is compatible with both Python 2 and Python 3:

```

def add_unicode_checkmark(text: Text) -> Text:
    return text + u' \u2713'

```

3.5.2 新版功能.

class `typing.IO`

class `typing.TextIO`

class `typing.BinaryIO`

Generic type `IO[AnyStr]` and its subclasses `TextIO(IO[str])` and `BinaryIO(IO[bytes])` represent the types of I/O streams such as returned by `open()`.

class `typing.Pattern`

class `typing.Match`

These type aliases correspond to the return types from `re.compile()` and `re.match()`. These types (and the corresponding functions) are generic in `AnyStr` and can be made specific by writing `Pattern[str]`, `Pattern[bytes]`, `Match[str]`, or `Match[bytes]`.

class `typing.NamedTuple`

Typed version of `namedtuple`.

用法:

```

class Employee(NamedTuple):
    name: str
    id: int

```

这相当于:

```
Employee = collections.namedtuple('Employee', ['name', 'id'])
```

To give a field a default value, you can assign to it in the class body:

```
class Employee(NamedTuple):
    name: str
    id: int = 3

employee = Employee('Guido')
assert employee.id == 3
```

Fields with a default value must come after any fields without a default.

The resulting class has two extra attributes: `_field_types`, giving a dict mapping field names to types, and `_field_defaults`, a dict mapping field names to default values. (The field names are in the `_fields` attribute, which is part of the `namedtuple` API.)

`NamedTuple` subclasses can also have docstrings and methods:

```
class Employee(NamedTuple):
    """Represents an employee."""
    name: str
    id: int = 3

    def __repr__(self) -> str:
        return f'<Employee {self.name}, id={self.id}>'
```

Backward-compatible usage:

```
Employee = NamedTuple('Employee', [('name', str), ('id', int)])
```

在 3.6 版更改: Added support for **PEP 526** variable annotation syntax.

在 3.6.1 版更改: Added support for default values, methods, and docstrings.

`typing.NewType` (*typ*)

A helper function to indicate a distinct types to a typechecker, see *NewType*. At runtime it returns a function that returns its argument. Usage:

```
UserId = NewType('UserId', int)
first_user = UserId(1)
```

3.5.2 新版功能.

`typing.cast` (*typ*, *val*)

Cast a value to a type.

This returns the value unchanged. To the type checker this signals that the return value has the designated type, but at runtime we intentionally don't check anything (we want this to be as fast as possible).

`typing.get_type_hints` (*obj*[, *globals*[, *locals*]])

返回一个字典, 字典内含有函数、方法、模块或类对象的类型提示。

This is often the same as `obj.__annotations__`. In addition, forward references encoded as string literals are handled by evaluating them in `globals` and `locals` namespaces. If necessary, `Optional[t]` is added for function and method annotations if a default value equal to `None` is set. For a class `C`, return a dictionary constructed by merging all the `__annotations__` along `C.__mro__` in reverse order.

@typing.overload

The `@overload` decorator allows describing functions and methods that support multiple different combinations of argument types. A series of `@overload`-decorated definitions must be followed by exactly one non-`@overload`-decorated definition (for the same function/method). The `@overload`-decorated definitions are for the benefit of the type checker only, since they will be overwritten by the non-`@overload`-decorated definition, while the latter is used at runtime but should be ignored by a type checker. At runtime, calling a `@overload`-decorated function directly will raise `NotImplementedError`. An example of overload that gives a more precise type than can be expressed using a union or a type variable:

```
@overload
def process(response: None) -> None:
    ...
@overload
def process(response: int) -> Tuple[int, str]:
    ...
@overload
def process(response: bytes) -> str:
    ...
def process(response):
    <actual implementation>
```

See [PEP 484](#) for details and comparison with other typing semantics.

@typing.no_type_check

用于指明标注不是类型提示的装饰器。

此`decorator`装饰器生效于类或函数上。如果作用于类上的话，它会递归地作用于这个类的所定义的所有方法上（但是对于超类或子类所定义的方法不会生效）。

此方法会就地地修改函数。

@typing.no_type_check_decorator

使其它装饰器起到`no_type_check()`效果的装饰器。

This wraps the decorator with something that wraps the decorated function in `no_type_check()`.

typing.Any

特殊类型，表明类型没有任何限制。

- 每一个类型都对`Any`兼容。
- `Any`对每一个类型都兼容。

typing.NoReturn

标记一个函数没有返回值的特殊类型。比如说：

```
from typing import NoReturn

def stop() -> NoReturn:
    raise RuntimeError('no way')
```

3.6.5 新版功能.

typing.Union

联合类型；`Union[X, Y]`意味着：要不是 X，要不是 Y。

使用形如 `Union[int, str]` 的形式来定义一个联合类型。细节如下：

- 参数必须是类型，而且必须至少有一个参数。
- 联合类型的联合类型会被展开打平，比如：

```
Union[Union[int, str], float] == Union[int, str, float]
```

- 仅有一个参数的联合类型会坍塌成参数自身，比如：

```
Union[int] == int # The constructor actually returns int
```

- 多余的参数会被跳过，比如：

```
Union[int, str, int] == Union[int, str]
```

- 在比较联合类型的时候，参数顺序会被忽略，比如：

```
Union[int, str] == Union[str, int]
```

- When a class and its subclass are present, the latter is skipped, e.g.:

```
Union[int, object] == object
```

- 你不能继承或者实例化一个联合类型。
- 你不能写成 `Union[X][Y]`。
- 你可以使用 `Optional[X]` 作为 `Union[X, None]` 的缩写。

typing.Optional

Optional type.

`Optional[X]` is equivalent to `Union[X, None]`.

Note that this is not the same concept as an optional argument, which is one that has a default. An optional argument with a default does not require the `Optional` qualifier on its type annotation just because it is optional. For example:

```
def foo(arg: int = 0) -> None:
    ...
```

On the other hand, if an explicit value of `None` is allowed, the use of `Optional` is appropriate, whether the argument is optional or not. For example:

```
def foo(arg: Optional[int] = None) -> None:
    ...
```

typing.Tuple

Tuple type; `Tuple[X, Y]` is the type of a tuple of two items with the first item of type `X` and the second of type `Y`.

Example: `Tuple[T1, T2]` is a tuple of two elements corresponding to type variables `T1` and `T2`. `Tuple[int, float, str]` is a tuple of an int, a float and a string.

To specify a variable-length tuple of homogeneous type, use literal ellipsis, e.g. `Tuple[int, ...]`. A plain `Tuple` is equivalent to `Tuple[Any, ...]`, and in turn to `tuple`.

typing.Callable

Callable type; `Callable[[int], str]` is a function of `(int) -> str`.

The subscription syntax must always be used with exactly two values: the argument list and the return type. The argument list must be a list of types or an ellipsis; the return type must be a single type.

There is no syntax to indicate optional or keyword arguments; such function types are rarely used as callback types. `Callable[..., ReturnType]` (literal ellipsis) can be used to type hint a callable taking any number of

arguments and returning `ReturnType`. A plain `Callable` is equivalent to `Callable[... , Any]`, and in turn to `collections.abc.Callable`.

typing.`ClassVar`

Special type construct to mark class variables.

As introduced in [PEP 526](#), a variable annotation wrapped in `ClassVar` indicates that a given attribute is intended to be used as a class variable and should not be set on instances of that class. Usage:

```
class Starship:
    stats: ClassVar[Dict[str, int]] = {} # class variable
    damage: int = 10                    # instance variable
```

`ClassVar` accepts only types and cannot be further subscribed.

`ClassVar` is not a class itself, and should not be used with `isinstance()` or `issubclass()`. `ClassVar` does not change Python runtime behavior, but it can be used by third-party type checkers. For example, a type checker might flag the following code as an error:

```
enterprise_d = Starship(3000)
enterprise_d.stats = {} # Error, setting class variable on instance
Starship.stats = {}    # This is OK
```

3.5.3 新版功能.

typing.`AnyStr`

`AnyStr` is a type variable defined as `AnyStr = TypeVar('AnyStr', str, bytes)`.

It is meant to be used for functions that may accept any kind of string without allowing different kinds of strings to mix. For example:

```
def concat(a: AnyStr, b: AnyStr) -> AnyStr:
    return a + b

concat(u"foo", u"bar") # Ok, output has type 'unicode'
concat(b"foo", b"bar") # Ok, output has type 'bytes'
concat(u"foo", b"bar") # Error, cannot mix unicode and bytes
```

typing.`TYPE_CHECKING`

A special constant that is assumed to be `True` by 3rd party static type checkers. It is `False` at runtime. Usage:

```
if TYPE_CHECKING:
    import expensive_mod

def fun(arg: 'expensive_mod.SomeType') -> None:
    local_var: expensive_mod.AnotherType = other_fun()
```

Note that the first type annotation must be enclosed in quotes, making it a “forward reference”, to hide the `expensive_mod` reference from the interpreter runtime. Type annotations for local variables are not evaluated, so the second annotation does not need to be enclosed in quotes.

3.5.2 新版功能.

26.2 pydoc — Documentation generator and online help system

Source code: [Lib/pydoc.py](#)

The `pydoc` module automatically generates documentation from Python modules. The documentation can be presented as pages of text on the console, served to a Web browser, or saved to HTML files.

For modules, classes, functions and methods, the displayed documentation is derived from the docstring (i.e. the `__doc__` attribute) of the object, and recursively of its documentable members. If there is no docstring, `pydoc` tries to obtain a description from the block of comment lines just above the definition of the class, function or method in the source file, or at the top of the module (see `inspect.getcomments()`).

The built-in function `help()` invokes the online help system in the interactive interpreter, which uses `pydoc` to generate its documentation as text on the console. The same text documentation can also be viewed from outside the Python interpreter by running `pydoc` as a script at the operating system's command prompt. For example, running

```
pydoc sys
```

at a shell prompt will display documentation on the `sys` module, in a style similar to the manual pages shown by the Unix `man` command. The argument to `pydoc` can be the name of a function, module, or package, or a dotted reference to a class, method, or function within a module or module in a package. If the argument to `pydoc` looks like a path (that is, it contains the path separator for your operating system, such as a slash in Unix), and refers to an existing Python source file, then documentation is produced for that file.

注解: In order to find objects and their documentation, `pydoc` imports the module(s) to be documented. Therefore, any code on module level will be executed on that occasion. Use an `if __name__ == '__main__':` guard to only execute code when a file is invoked as a script and not just imported.

When printing output to the console, `pydoc` attempts to paginate the output for easier reading. If the `PAGER` environment variable is set, `pydoc` will use its value as a pagination program.

Specifying a `-w` flag before the argument will cause HTML documentation to be written out to a file in the current directory, instead of displaying text on the console.

Specifying a `-k` flag before the argument will search the synopsis lines of all available modules for the keyword given as the argument, again in a manner similar to the Unix `man` command. The synopsis line of a module is the first line of its documentation string.

You can also use `pydoc` to start an HTTP server on the local machine that will serve documentation to visiting Web browsers. `pydoc -p 1234` will start a HTTP server on port 1234, allowing you to browse the documentation at `http://localhost:1234/` in your preferred Web browser. Specifying 0 as the port number will select an arbitrary unused port.

`pydoc -b` will start the server and additionally open a web browser to a module index page. Each served page has a navigation bar at the top where you can *Get* help on an individual item, *Search* all modules with a keyword in their synopsis line, and go to the *Module index*, *Topics* and *Keywords* pages.

When `pydoc` generates documentation, it uses the current environment and path to locate modules. Thus, invoking `pydoc spam` documents precisely the version of the module you would get if you started the Python interpreter and typed `import spam`.

Module docs for core modules are assumed to reside in `https://docs.python.org/X.Y/library/` where X and Y are the major and minor version numbers of the Python interpreter. This can be overridden by setting the `PYTHONDOS` environment variable to a different URL or to a local directory containing the Library Reference Manual pages.

在 3.2 版更改: Added the `-b` option.

在 3.3 版更改: The `-g` command line option was removed.

在 3.4 版更改: `pydoc` now uses `inspect.signature()` rather than `inspect.getfullargspec()` to extract signature information from callables.

26.3 doctest — 测试交互性的 Python 示例

**** 源代码 **** Lib/doctest.py

`doctest` 模块寻找像 Python 交互式代码的文本，然后执行这些代码来确保它们的确就像展示的那样正确运行，有许多方法来使用 `doctest`：

- To check that a module's docstrings are up-to-date by verifying that all interactive examples still work as documented.
- To perform regression testing by verifying that interactive examples from a test file or a test object work as expected.
- To write tutorial documentation for a package, liberally illustrated with input-output examples. Depending on whether the examples or the expository text are emphasized, this has the flavor of “literate testing” or “executable documentation” .

下面是一个小却完整的示例模块：

```
"""
This is the "example" module.

The example module supplies one function, factorial().  For example,

>>> factorial(5)
120
"""

def factorial(n):
    """Return the factorial of n, an exact integer >= 0.

    >>> [factorial(n) for n in range(6)]
    [1, 1, 2, 6, 24, 120]
    >>> factorial(30)
    265252859812191058636308480000000
    >>> factorial(-1)
    Traceback (most recent call last):
    ...
    ValueError: n must be >= 0

    Factorials of floats are OK, but the float must be an exact integer:
    >>> factorial(30.1)
    Traceback (most recent call last):
    ...
    ValueError: n must be exact integer
    >>> factorial(30.0)
    26525285981219105863630848000000

    It must also not be ridiculously large:
    >>> factorial(1e100)
    Traceback (most recent call last):
    ...
    OverflowError: n too large
    """

    import math
    if not n >= 0:
        raise ValueError("n must be >= 0")
```

(下页继续)

(续上页)

```

if math.floor(n) != n:
    raise ValueError("n must be exact integer")
if n+1 == n: # catch a value like 1e300
    raise OverflowError("n too large")
result = 1
factor = 2
while factor <= n:
    result *= factor
    factor += 1
return result

if __name__ == "__main__":
    import doctest
    doctest.testmod()

```

如果你直接在命令行里运行 `example.py` , `doctest` 将发挥他的作用。

```

$ python example.py
$

```

There's no output! That's normal, and it means all the examples worked. Pass `-v` to the script, and `doctest` prints a detailed log of what it's trying, and prints a summary at the end:

```

$ python example.py -v
Trying:
    factorial(5)
Expecting:
    120
ok
Trying:
    [factorial(n) for n in range(6)]
Expecting:
    [1, 1, 2, 6, 24, 120]
ok

```

And so on, eventually ending with:

```

Trying:
    factorial(1e100)
Expecting:
    Traceback (most recent call last):
      ...
    OverflowError: n too large
ok
2 items passed all tests:
  1 tests in __main__
  8 tests in __main__.factorial
9 tests in 2 items.
9 passed and 0 failed.
Test passed.
$

```

That's all you need to know to start making productive use of `doctest`! Jump in. The following sections provide full details. Note that there are many examples of doctests in the standard Python test suite and libraries. Especially useful examples can be found in the standard test file `Lib/test/test_doctest.py`.

26.3.1 Simple Usage: Checking Examples in Docstrings

The simplest way to start using `doctest` (but not necessarily the way you'll continue to do it) is to end each module `M` with:

```
if __name__ == "__main__":
    import doctest
    doctest.testmod()
```

`doctest` then examines docstrings in module `M`.

Running the module as a script causes the examples in the docstrings to get executed and verified:

```
python M.py
```

This won't display anything unless an example fails, in which case the failing example(s) and the cause(s) of the failure(s) are printed to stdout, and the final line of output is `***Test Failed*** N failures.`, where `N` is the number of examples that failed.

Run it with the `-v` switch instead:

```
python M.py -v
```

and a detailed report of all examples tried is printed to standard output, along with assorted summaries at the end.

You can force verbose mode by passing `verbose=True` to `testmod()`, or prohibit it by passing `verbose=False`. In either of those cases, `sys.argv` is not examined by `testmod()` (so passing `-v` or not has no effect).

There is also a command line shortcut for running `testmod()`. You can instruct the Python interpreter to run the `doctest` module directly from the standard library and pass the module name(s) on the command line:

```
python -m doctest -v example.py
```

This will import `example.py` as a standalone module and run `testmod()` on it. Note that this may not work correctly if the file is part of a package and imports other submodules from that package.

For more information on `testmod()`, see section *Basic API*.

26.3.2 Simple Usage: Checking Examples in a Text File

Another simple application of `doctest` is testing interactive examples in a text file. This can be done with the `testfile()` function:

```
import doctest
doctest.testfile("example.txt")
```

That short script executes and verifies any interactive Python examples contained in the file `example.txt`. The file content is treated as if it were a single giant docstring; the file doesn't need to contain a Python program! For example, perhaps `example.txt` contains this:

```
The ``example`` module
=====

Using ``factorial``
-----

This is an example text file in reStructuredText format.  First import
```

(下页继续)

(续上页)

```

``factorial`` from the ``example`` module:

    >>> from example import factorial

Now use it:

    >>> factorial(6)
    120

```

Running `doctest.testfile("example.txt")` then finds the error in this documentation:

```

File "./example.txt", line 14, in example.txt
Failed example:
    factorial(6)
Expected:
    120
Got:
    720

```

As with `testmod()`, `testfile()` won't display anything unless an example fails. If an example does fail, then the failing example(s) and the cause(s) of the failure(s) are printed to stdout, using the same format as `testmod()`.

By default, `testfile()` looks for files in the calling module's directory. See section [Basic API](#) for a description of the optional arguments that can be used to tell it to look for files in other locations.

Like `testmod()`, `testfile()`'s verbosity can be set with the `-v` command-line switch or with the optional keyword argument `verbose`.

There is also a command line shortcut for running `testfile()`. You can instruct the Python interpreter to run the doctest module directly from the standard library and pass the file name(s) on the command line:

```
python -m doctest -v example.txt
```

Because the file name does not end with `.py`, `doctest` infers that it must be run with `testfile()`, not `testmod()`.

For more information on `testfile()`, see section [Basic API](#).

26.3.3 How It Works

This section examines in detail how doctest works: which docstrings it looks at, how it finds interactive examples, what execution context it uses, how it handles exceptions, and how option flags can be used to control its behavior. This is the information that you need to know to write doctest examples; for information about actually running doctest on these examples, see the following sections.

Which Docstrings Are Examined?

The module docstring, and all function, class and method docstrings are searched. Objects imported into the module are not searched.

In addition, if `M.__test__` exists and “is true”, it must be a dict, and each entry maps a (string) name to a function object, class object, or string. Function and class object docstrings found from `M.__test__` are searched, and strings are treated as if they were docstrings. In output, a key `K` in `M.__test__` appears with name

```
<name of M>.__test__.K
```

Any classes found are recursively searched similarly, to test docstrings in their contained methods and nested classes.

CPython implementation detail: Prior to version 3.4, extension modules written in C were not fully searched by doctest.

How are Docstring Examples Recognized?

In most cases a copy-and-paste of an interactive console session works fine, but doctest isn't trying to do an exact emulation of any specific Python shell.

```
>>> # comments are ignored
>>> x = 12
>>> x
12
>>> if x == 13:
...     print("yes")
... else:
...     print("no")
...     print("NO")
...     print("NO!!!")
...
no
NO
NO!!!
>>>
```

Any expected output must immediately follow the final '`>>>`' or '`...`' line containing the code, and the expected output (if any) extends to the next '`>>>`' or all-whitespace line.

The fine print:

- Expected output cannot contain an all-whitespace line, since such a line is taken to signal the end of expected output. If expected output does contain a blank line, put `<BLANKLINE>` in your doctest example each place a blank line is expected.
- All hard tab characters are expanded to spaces, using 8-column tab stops. Tabs in output generated by the tested code are not modified. Because any hard tabs in the sample output *are* expanded, this means that if the code output includes hard tabs, the only way the doctest can pass is if the `NORMALIZE_WHITESPACE` option or *directive* is in effect. Alternatively, the test can be rewritten to capture the output and compare it to an expected value as part of the test. This handling of tabs in the source was arrived at through trial and error, and has proven to be the least error prone way of handling them. It is possible to use a different algorithm for handling tabs by writing a custom `DocTestParser` class.
- Output to stdout is captured, but not output to stderr (exception tracebacks are captured via a different means).
- If you continue a line via backslashing in an interactive session, or for any other reason use a backslash, you should use a raw docstring, which will preserve your backslashes exactly as you type them:

```
>>> def f(x):
...     r'''Backslashes in a raw docstring: m\n'''
>>> print(f.__doc__)
Backslashes in a raw docstring: m\n
```

Otherwise, the backslash will be interpreted as part of the string. For example, the `\n` above would be interpreted as a newline character. Alternatively, you can double each backslash in the doctest version (and not use a raw string):

```
>>> def f(x):
...     '''Backslashes in a raw docstring: m\\n'''
```

(下页继续)

(续上页)

```
>>> print(f.__doc__)
Backslashes in a raw docstring: m\n
```

- The starting column doesn't matter:

```
>>> assert "Easy!"
>>> import math
>>> math.floor(1.9)
1
```

and as many leading whitespace characters are stripped from the expected output as appeared in the initial '>>>' line that started the example.

What's the Execution Context?

By default, each time `doctest` finds a docstring to test, it uses a *shallow copy* of `M`'s globals, so that running tests doesn't change the module's real globals, and so that one test in `M` can't leave behind crumbs that accidentally allow another test to work. This means examples can freely use any names defined at top-level in `M`, and names defined earlier in the docstring being run. Examples cannot see names defined in other docstrings.

You can force use of your own dict as the execution context by passing `globs=your_dict` to `testmod()` or `testfile()` instead.

What About Exceptions?

No problem, provided that the traceback is the only output produced by the example: just paste in the traceback.¹ Since tracebacks contain details that are likely to change rapidly (for example, exact file paths and line numbers), this is one case where `doctest` works hard to be flexible in what it accepts.

Simple example:

```
>>> [1, 2, 3].remove(42)
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
ValueError: list.remove(x): x not in list
```

That `doctest` succeeds if `ValueError` is raised, with the `list.remove(x): x not in list` detail as shown.

The expected output for an exception must start with a traceback header, which may be either of the following two lines, indented the same as the first line of the example:

```
Traceback (most recent call last):
Traceback (innermost last):
```

The traceback header is followed by an optional traceback stack, whose contents are ignored by `doctest`. The traceback stack is typically omitted, or copied verbatim from an interactive session.

The traceback stack is followed by the most interesting part: the line(s) containing the exception type and detail. This is usually the last line of a traceback, but can extend across multiple lines if the exception has a multi-line detail:

¹ Examples containing both expected output and an exception are not supported. Trying to guess where one ends and the other begins is too error-prone, and that also makes for a confusing test.

```
>>> raise ValueError('multi\n    line\ndetail')
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
ValueError: multi
    line
detail
```

The last three lines (starting with `ValueError`) are compared against the exception's type and detail, and the rest are ignored.

Best practice is to omit the traceback stack, unless it adds significant documentation value to the example. So the last example is probably better as:

```
>>> raise ValueError('multi\n    line\ndetail')
Traceback (most recent call last):
...
ValueError: multi
    line
detail
```

Note that tracebacks are treated very specially. In particular, in the rewritten example, the use of `...` is independent of doctest's `ELLIPSIS` option. The ellipsis in that example could be left out, or could just as well be three (or three hundred) commas or digits, or an indented transcript of a Monty Python skit.

Some details you should read once, but won't need to remember:

- Doctest can't guess whether your expected output came from an exception traceback or from ordinary printing. So, e.g., an example that expects `ValueError: 42 is prime` will pass whether `ValueError` is actually raised or if the example merely prints that traceback text. In practice, ordinary output rarely begins with a traceback header line, so this doesn't create real problems.
- Each line of the traceback stack (if present) must be indented further than the first line of the example, *or* start with a non-alphanumeric character. The first line following the traceback header indented the same and starting with an alphanumeric is taken to be the start of the exception detail. Of course this does the right thing for genuine tracebacks.
- When the `IGNORE_EXCEPTION_DETAIL` doctest option is specified, everything following the leftmost colon and any module information in the exception name is ignored.
- The interactive shell omits the traceback header line for some `SyntaxErrors`. But doctest uses the traceback header line to distinguish exceptions from non-exceptions. So in the rare case where you need to test a `SyntaxError` that omits the traceback header, you will need to manually add the traceback header line to your test example.
- For some `SyntaxErrors`, Python displays the character position of the syntax error, using a `^` marker:

```
>>> 1 1
    File "<stdin>", line 1
      1 1
      ^
SyntaxError: invalid syntax
```

Since the lines showing the position of the error come before the exception type and detail, they are not checked by doctest. For example, the following test would pass, even though it puts the `^` marker in the wrong location:

```
>>> 1 1
    File "<stdin>", line 1
      1 1
```

(下页继续)

(续上页)

```

      ^
SyntaxError: invalid syntax

```

Option Flags

A number of option flags control various aspects of doctest's behavior. Symbolic names for the flags are supplied as module constants, which can be bitwise ORed together and passed to various functions. The names can also be used in *doctest directives*, and may be passed to the doctest command line interface via the `-o` option.

3.4 新版功能: The `-o` command line option.

The first group of options define test semantics, controlling aspects of how doctest decides whether actual output matches an example's expected output:

`doctest.DONT_ACCEPT_TRUE_FOR_1`

By default, if an expected output block contains just 1, an actual output block containing just 1 or just True is considered to be a match, and similarly for 0 versus False. When `DONT_ACCEPT_TRUE_FOR_1` is specified, neither substitution is allowed. The default behavior caters to that Python changed the return type of many functions from integer to boolean; doctests expecting “little integer” output still work in these cases. This option will probably go away, but not for several years.

`doctest.DONT_ACCEPT_BLANKLINE`

By default, if an expected output block contains a line containing only the string `<BLANKLINE>`, then that line will match a blank line in the actual output. Because a genuinely blank line delimits the expected output, this is the only way to communicate that a blank line is expected. When `DONT_ACCEPT_BLANKLINE` is specified, this substitution is not allowed.

`doctest.NORMALIZE_WHITESPACE`

When specified, all sequences of whitespace (blanks and newlines) are treated as equal. Any sequence of whitespace within the expected output will match any sequence of whitespace within the actual output. By default, whitespace must match exactly. `NORMALIZE_WHITESPACE` is especially useful when a line of expected output is very long, and you want to wrap it across multiple lines in your source.

`doctest.ELLIPSIS`

When specified, an ellipsis marker (`. . .`) in the expected output can match any substring in the actual output. This includes substrings that span line boundaries, and empty substrings, so it's best to keep usage of this simple. Complicated uses can lead to the same kinds of “oops, it matched too much!” surprises that `. *` is prone to in regular expressions.

`doctest.IGNORE_EXCEPTION_DETAIL`

When specified, an example that expects an exception passes if an exception of the expected type is raised, even if the exception detail does not match. For example, an example expecting `ValueError: 42` will pass if the actual exception raised is `ValueError: 3*14`, but will fail, e.g., if `TypeError` is raised.

It will also ignore the module name used in Python 3 doctest reports. Hence both of these variations will work with the flag specified, regardless of whether the test is run under Python 2.7 or Python 3.2 (or later versions):

```

>>> raise CustomError('message')
Traceback (most recent call last):
CustomError: message

>>> raise CustomError('message')
Traceback (most recent call last):
my_module.CustomError: message

```

Note that `ELLIPSIS` can also be used to ignore the details of the exception message, but such a test may still fail based on whether or not the module details are printed as part of the exception name. Using `IGNORE_EXCEPTION_DETAIL` and the details from Python 2.3 is also the only clear way to write a doctest that doesn't care about the exception detail yet continues to pass under Python 2.3 or earlier (those releases do not support *doctest directives* and ignore them as irrelevant comments). For example:

```
>>> (1, 2)[3] = 'moo'
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
TypeError: object doesn't support item assignment
```

passes under Python 2.3 and later Python versions with the flag specified, even though the detail changed in Python 2.4 to say “does not” instead of “doesn't”.

在 3.2 版更改: `IGNORE_EXCEPTION_DETAIL` now also ignores any information relating to the module containing the exception under test.

`doctest.SKIP`

When specified, do not run the example at all. This can be useful in contexts where doctest examples serve as both documentation and test cases, and an example should be included for documentation purposes, but should not be checked. E.g., the example's output might be random; or the example might depend on resources which would be unavailable to the test driver.

The SKIP flag can also be used for temporarily “commenting out” examples.

`doctest.COMPARISON_FLAGS`

A bitmask or'ing together all the comparison flags above.

The second group of options controls how test failures are reported:

`doctest.REPORT_UDIFF`

When specified, failures that involve multi-line expected and actual outputs are displayed using a unified diff.

`doctest.REPORT_CDIF`

When specified, failures that involve multi-line expected and actual outputs will be displayed using a context diff.

`doctest.REPORT_NDIFF`

When specified, differences are computed by `diff.lib.Differ`, using the same algorithm as the popular `ndiff.py` utility. This is the only method that marks differences within lines as well as across lines. For example, if a line of expected output contains digit 1 where actual output contains letter l, a line is inserted with a caret marking the mismatching column positions.

`doctest.REPORT_ONLY_FIRST_FAILURE`

When specified, display the first failing example in each doctest, but suppress output for all remaining examples. This will prevent doctest from reporting correct examples that break because of earlier failures; but it might also hide incorrect examples that fail independently of the first failure. When `REPORT_ONLY_FIRST_FAILURE` is specified, the remaining examples are still run, and still count towards the total number of failures reported; only the output is suppressed.

`doctest.FAIL_FAST`

When specified, exit after the first failing example and don't attempt to run the remaining examples. Thus, the number of failures reported will be at most 1. This flag may be useful during debugging, since examples after the first failure won't even produce debugging output.

The doctest command line accepts the option `-f` as a shorthand for `-o FAIL_FAST`.

3.4 新版功能.

`doctest.REPORTING_FLAGS`

A bitmask or'ing together all the reporting flags above.

There is also a way to register new option flag names, though this isn't useful unless you intend to extend *doctest* internals via subclassing:

`doctest.register_optionflag(name)`

Create a new option flag with a given name, and return the new flag's integer value. *register_optionflag()* can be used when subclassing *OutputChecker* or *DocTestRunner* to create new options that are supported by your subclasses. *register_optionflag()* should always be called using the following idiom:

```
MY_FLAG = register_optionflag('MY_FLAG')
```

Directives

Doctest directives may be used to modify the *option flags* for an individual example. Doctest directives are special Python comments following an example's source code:

```
directive          ::=  "# "doctest:" directive_options
directive_options  ::=  directive_option ("," directive_option)*
directive_option   ::=  on_or_off directive_option_name
on_or_off          ::=  "+" \| "-"
directive_option_name ::=  "DONT_ACCEPT_BLANKLINE" \| "NORMALIZE_WHITESPACE" \| ...
```

Whitespace is not allowed between the + or - and the directive option name. The directive option name can be any of the option flag names explained above.

An example's doctest directives modify doctest's behavior for that single example. Use + to enable the named behavior, or - to disable it.

For example, this test passes:

```
>>> print(list(range(20)))
[0, 1, 2, 3, 4, 5, 6, 7, 8, 9,
10, 11, 12, 13, 14, 15, 16, 17, 18, 19]
```

Without the directive it would fail, both because the actual output doesn't have two blanks before the single-digit list elements, and because the actual output is on a single line. This test also passes, and also requires a directive to do so:

```
>>> print(list(range(20)))
[0, 1, ..., 18, 19]
```

Multiple directives can be used on a single physical line, separated by commas:

```
>>> print(list(range(20)))
[0, 1, ..., 18, 19]
```

If multiple directive comments are used for a single example, then they are combined:

```
>>> print(list(range(20)))
...
[0, 1, ..., 18, 19]
```

As the previous example shows, you can add ... lines to your example containing only directives. This can be useful when an example is too long for a directive to comfortably fit on the same line:

```
>>> print(list(range(5)) + list(range(10, 20)) + list(range(30, 40)))
...
[0, ..., 4, 10, ..., 19, 30, ..., 39]
```

Note that since all options are disabled by default, and directives apply only to the example they appear in, enabling options (via `+` in a directive) is usually the only meaningful choice. However, option flags can also be passed to functions that run doctests, establishing different defaults. In such cases, disabling an option via `-` in a directive can be useful.

警告

`doctest` is serious about requiring exact matches in expected output. If even a single character doesn't match, the test fails. This will probably surprise you a few times, as you learn exactly what Python does and doesn't guarantee about output. For example, when printing a dict, Python doesn't guarantee that the key-value pairs will be printed in any particular order, so a test like

```
>>> foo()
{"Hermione": "hippogryph", "Harry": "broomstick"}
```

is vulnerable! One workaround is to do

```
>>> foo() == {"Hermione": "hippogryph", "Harry": "broomstick"}
True
```

instead. Another is to do

```
>>> d = sorted(foo().items())
>>> d
[('Harry', 'broomstick'), ('Hermione', 'hippogryph')]
```

There are others, but you get the idea.

Another bad idea is to print things that embed an object address, like

```
>>> id(1.0) # certain to fail some of the time
7948648
>>> class C: pass
>>> C() # the default repr() for instances embeds an address
<__main__.C instance at 0x00AC18F0>
```

The *ELLIPSIS* directive gives a nice approach for the last example:

```
>>> C()
<__main__.C instance at 0x...>
```

Floating-point numbers are also subject to small output variations across platforms, because Python defers to the platform C library for float formatting, and C libraries vary widely in quality here.

```
>>> 1./7 # risky
0.14285714285714285
>>> print(1./7) # safer
0.142857142857
>>> print(round(1./7, 6)) # much safer
0.142857
```

Numbers of the form $I/2.**J$ are safe across all platforms, and I often contrive doctest examples to produce numbers of that form:

```
>>> 3./4 # utterly safe
0.75
```

Simple fractions are also easier for people to understand, and that makes for better documentation.

26.3.4 Basic API

The functions `testmod()` and `testfile()` provide a simple interface to doctest that should be sufficient for most basic uses. For a less formal introduction to these two functions, see sections *Simple Usage: Checking Examples in Docstrings* and *Simple Usage: Checking Examples in a Text File*.

```
doctest.testfile(filename, module_relative=True, name=None, package=None, globs=None, ver-
                 bose=None, report=True, optionflags=0, extraglobs=None, raise_on_error=False,
                 parser=DocTestParser(), encoding=None)
```

All arguments except `filename` are optional, and should be specified in keyword form.

Test examples in the file named `filename`. Return `(failure_count, test_count)`.

Optional argument `module_relative` specifies how the `filename` should be interpreted:

- If `module_relative` is `True` (the default), then `filename` specifies an OS-independent module-relative path. By default, this path is relative to the calling module's directory; but if the `package` argument is specified, then it is relative to that package. To ensure OS-independence, `filename` should use `/` characters to separate path segments, and may not be an absolute path (i.e., it may not begin with `/`).
- If `module_relative` is `False`, then `filename` specifies an OS-specific path. The path may be absolute or relative; relative paths are resolved with respect to the current working directory.

Optional argument `name` gives the name of the test; by default, or if `None`, `os.path.basename(filename)` is used.

Optional argument `package` is a Python package or the name of a Python package whose directory should be used as the base directory for a module-relative filename. If no package is specified, then the calling module's directory is used as the base directory for module-relative filenames. It is an error to specify `package` if `module_relative` is `False`.

Optional argument `globs` gives a dict to be used as the globals when executing examples. A new shallow copy of this dict is created for the doctest, so its examples start with a clean slate. By default, or if `None`, a new empty dict is used.

Optional argument `extraglobs` gives a dict merged into the globals used to execute examples. This works like `dict.update()`: if `globs` and `extraglobs` have a common key, the associated value in `extraglobs` appears in the combined dict. By default, or if `None`, no extra globals are used. This is an advanced feature that allows parameterization of doctests. For example, a doctest can be written for a base class, using a generic name for the class, then reused to test any number of subclasses by passing an `extraglobs` dict mapping the generic name to the subclass to be tested.

Optional argument `verbose` prints lots of stuff if true, and prints only failures if false; by default, or if `None`, it's true if and only if `'-v'` is in `sys.argv`.

Optional argument `report` prints a summary at the end when true, else prints nothing at the end. In verbose mode, the summary is detailed, else the summary is very brief (in fact, empty if all tests passed).

Optional argument `optionflags` (default value 0) takes the bitwise OR of option flags. See section *Option Flags*.

Optional argument `raise_on_error` defaults to false. If true, an exception is raised upon the first failure or unexpected exception in an example. This allows failures to be post-mortem debugged. Default behavior is to continue running examples.

Optional argument `parser` specifies a `DocTestParser` (or subclass) that should be used to extract tests from the files. It defaults to a normal parser (i.e., `DocTestParser()`).

Optional argument *encoding* specifies an encoding that should be used to convert the file to unicode.

`doctest.testmod(m=None, name=None, globs=None, verbose=None, report=True, optionflags=0, extra-globs=None, raise_on_error=False, exclude_empty=False)`

All arguments are optional, and all except for *m* should be specified in keyword form.

Test examples in docstrings in functions and classes reachable from module *m* (or module `__main__` if *m* is not supplied or is `None`), starting with `m.__doc__`.

Also test examples reachable from dict `m.__test__`, if it exists and is not `None`. `m.__test__` maps names (strings) to functions, classes and strings; function and class docstrings are searched for examples; strings are searched directly, as if they were docstrings.

Only docstrings attached to objects belonging to module *m* are searched.

Return (failure_count, test_count).

Optional argument *name* gives the name of the module; by default, or if `None`, `m.__name__` is used.

Optional argument *exclude_empty* defaults to false. If true, objects for which no doctests are found are excluded from consideration. The default is a backward compatibility hack, so that code still using `doctest.master.summarize()` in conjunction with `testmod()` continues to get output for objects with no tests. The *exclude_empty* argument to the newer `DocTestFinder` constructor defaults to true.

Optional arguments *extraglobs*, *verbose*, *report*, *optionflags*, *raise_on_error*, and *globs* are the same as for function `testfile()` above, except that *globs* defaults to `m.__dict__`.

`doctest.run_docstring_examples(f, globs, verbose=False, name="NoName", compileflags=None, optionflags=0)`

Test examples associated with object *f*; for example, *f* may be a string, a module, a function, or a class object.

A shallow copy of dictionary argument *globs* is used for the execution context.

Optional argument *name* is used in failure messages, and defaults to "NoName".

If optional argument *verbose* is true, output is generated even if there are no failures. By default, output is generated only in case of an example failure.

Optional argument *compileflags* gives the set of flags that should be used by the Python compiler when running the examples. By default, or if `None`, flags are deduced corresponding to the set of future features found in *globs*.

Optional argument *optionflags* works as for function `testfile()` above.

26.3.5 Unittest API

As your collection of doctest'ed modules grows, you'll want a way to run all their doctests systematically. `doctest` provides two functions that can be used to create `unittest` test suites from modules and text files containing doctests. To integrate with `unittest` test discovery, include a `load_tests()` function in your test module:

```
import unittest
import doctest
import my_module_with_doctests

def load_tests(loader, tests, ignore):
    tests.addTests(doctest.DocTestSuite(my_module_with_doctests))
    return tests
```

There are two main functions for creating `unittest.TestSuite` instances from text files and modules with doctests:

`doctest.DocFileSuite(*paths, module_relative=True, package=None, setUp=None, tearDown=None, globs=None, optionflags=0, parser=DocTestParser(), encoding=None)`
 Convert doctest tests from one or more text files to a `unittest.TestSuite`.

The returned `unittest.TestSuite` is to be run by the unittest framework and runs the interactive examples in each file. If an example in any file fails, then the synthesized unit test fails, and a `failureException` exception is raised showing the name of the file containing the test and a (sometimes approximate) line number.

Pass one or more paths (as strings) to text files to be examined.

Options may be provided as keyword arguments:

Optional argument `module_relative` specifies how the filenames in `paths` should be interpreted:

- If `module_relative` is `True` (the default), then each filename in `paths` specifies an OS-independent module-relative path. By default, this path is relative to the calling module's directory; but if the `package` argument is specified, then it is relative to that package. To ensure OS-independence, each filename should use `/` characters to separate path segments, and may not be an absolute path (i.e., it may not begin with `/`).
- If `module_relative` is `False`, then each filename in `paths` specifies an OS-specific path. The path may be absolute or relative; relative paths are resolved with respect to the current working directory.

Optional argument `package` is a Python package or the name of a Python package whose directory should be used as the base directory for module-relative filenames in `paths`. If no package is specified, then the calling module's directory is used as the base directory for module-relative filenames. It is an error to specify `package` if `module_relative` is `False`.

Optional argument `setUp` specifies a set-up function for the test suite. This is called before running the tests in each file. The `setUp` function will be passed a `DocTest` object. The `setUp` function can access the test globals as the `globs` attribute of the test passed.

Optional argument `tearDown` specifies a tear-down function for the test suite. This is called after running the tests in each file. The `tearDown` function will be passed a `DocTest` object. The `setUp` function can access the test globals as the `globs` attribute of the test passed.

Optional argument `globs` is a dictionary containing the initial global variables for the tests. A new copy of this dictionary is created for each test. By default, `globs` is a new empty dictionary.

Optional argument `optionflags` specifies the default doctest options for the tests, created by or-ing together individual option flags. See section [Option Flags](#). See function `set_unittest_reportflags()` below for a better way to set reporting options.

Optional argument `parser` specifies a `DocTestParser` (or subclass) that should be used to extract tests from the files. It defaults to a normal parser (i.e., `DocTestParser()`).

Optional argument `encoding` specifies an encoding that should be used to convert the file to unicode.

The global `__file__` is added to the globals provided to doctests loaded from a text file using `DocFileSuite()`.

`doctest.DocTestSuite(module=None, globs=None, extraglobs=None, test_finder=None, setUp=None, tearDown=None, checker=None)`
 Convert doctest tests for a module to a `unittest.TestSuite`.

The returned `unittest.TestSuite` is to be run by the unittest framework and runs each doctest in the module. If any of the doctests fail, then the synthesized unit test fails, and a `failureException` exception is raised showing the name of the file containing the test and a (sometimes approximate) line number.

Optional argument `module` provides the module to be tested. It can be a module object or a (possibly dotted) module name. If not specified, the module calling this function is used.

Optional argument `globs` is a dictionary containing the initial global variables for the tests. A new copy of this dictionary is created for each test. By default, `globs` is a new empty dictionary.

Optional argument *extraglobs* specifies an extra set of global variables, which is merged into *globs*. By default, no extra globals are used.

Optional argument *test_finder* is the *DocTestFinder* object (or a drop-in replacement) that is used to extract doctests from the module.

Optional arguments *setUp*, *tearDown*, and *optionflags* are the same as for function *DocFileSuite()* above.

This function uses the same search technique as *testmod()*.

在 3.5 版更改: *DocTestSuite()* returns an empty *unittest.TestSuite* if *module* contains no docstrings instead of raising *ValueError*.

Under the covers, *DocTestSuite()* creates a *unittest.TestSuite* out of *doctest.DocTestCase* instances, and *DocTestCase* is a subclass of *unittest.TestCase*. *DocTestCase* isn't documented here (it's an internal detail), but studying its code can answer questions about the exact details of *unittest* integration.

Similarly, *DocFileSuite()* creates a *unittest.TestSuite* out of *doctest.DocFileCase* instances, and *DocFileCase* is a subclass of *DocTestCase*.

So both ways of creating a *unittest.TestSuite* run instances of *DocTestCase*. This is important for a subtle reason: when you run *doctest* functions yourself, you can control the *doctest* options in use directly, by passing option flags to *doctest* functions. However, if you're writing a *unittest* framework, *unittest* ultimately controls when and how tests get run. The framework author typically wants to control *doctest* reporting options (perhaps, e.g., specified by command line options), but there's no way to pass options through *unittest* to *doctest* test runners.

For this reason, *doctest* also supports a notion of *doctest* reporting flags specific to *unittest* support, via this function:

`doctest.set_unittest_reportflags(flags)`

Set the *doctest* reporting flags to use.

Argument *flags* takes the bitwise OR of option flags. See section *Option Flags*. Only “reporting flags” can be used.

This is a module-global setting, and affects all future doctests run by module *unittest*: the *runTest()* method of *DocTestCase* looks at the option flags specified for the test case when the *DocTestCase* instance was constructed. If no reporting flags were specified (which is the typical and expected case), *doctest*'s *unittest* reporting flags are bitwise ORed into the option flags, and the option flags so augmented are passed to the *DocTestRunner* instance created to run the doctest. If any reporting flags were specified when the *DocTestCase* instance was constructed, *doctest*'s *unittest* reporting flags are ignored.

The value of the *unittest* reporting flags in effect before the function was called is returned by the function.

26.3.6 Advanced API

The basic API is a simple wrapper that's intended to make *doctest* easy to use. It is fairly flexible, and should meet most users' needs; however, if you require more fine-grained control over testing, or wish to extend *doctest*'s capabilities, then you should use the advanced API.

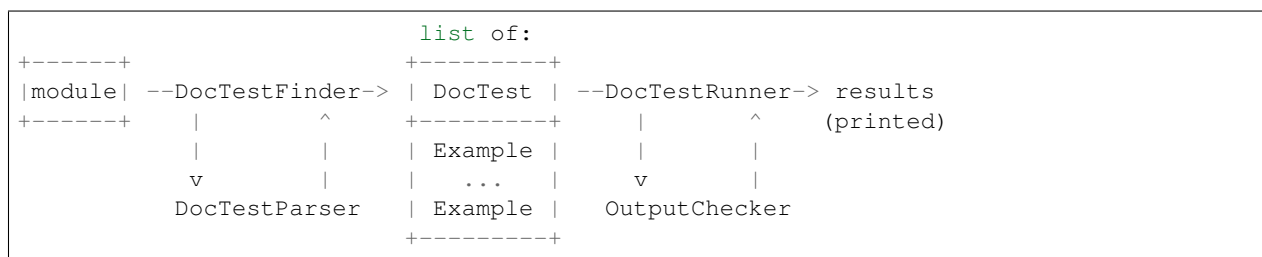
The advanced API revolves around two container classes, which are used to store the interactive examples extracted from *doctest* cases:

- *Example*: A single Python *statement*, paired with its expected output.
- *DocTest*: A collection of *Examples*, typically extracted from a single docstring or text file.

Additional processing classes are defined to find, parse, and run, and check *doctest* examples:

- *DocTestFinder*: Finds all docstrings in a given module, and uses a *DocTestParser* to create a *DocTest* from every docstring that contains interactive examples.
- *DocTestParser*: Creates a *DocTest* object from a string (such as an object's docstring).
- *DocTestRunner*: Executes the examples in a *DocTest*, and uses an *OutputChecker* to verify their output.
- *OutputChecker*: Compares the actual output from a doctest example with the expected output, and decides whether they match.

The relationships among these processing classes are summarized in the following diagram:



DocTest Objects

```
class doctest.DocTest (examples, globs, name, filename, lineno, docstring)
```

A collection of doctest examples that should be run in a single namespace. The constructor arguments are used to initialize the attributes of the same names.

DocTest defines the following attributes. They are initialized by the constructor, and should not be modified directly.

examples

A list of *Example* objects encoding the individual interactive Python examples that should be run by this test.

globs

The namespace (aka globals) that the examples should be run in. This is a dictionary mapping names to values. Any changes to the namespace made by the examples (such as binding new variables) will be reflected in `globals` after the test is run.

name

A string name identifying the *DocTest*. Typically, this is the name of the object or file that the test was extracted from.

filename

The name of the file that this *DocTest* was extracted from; or `None` if the filename is unknown, or if the *DocTest* was not extracted from a file.

lineno

The line number within *filename* where this *DocTest* begins, or *None* if the line number is unavailable. This line number is zero-based with respect to the beginning of the file.

docstring

The string that the test was extracted from, or `None` if the string is unavailable, or if the test was not extracted from a string.

Example Objects

class `doctest.Example` (*source*, *want*, *exc_msg=None*, *lineno=0*, *indent=0*, *options=None*)

A single interactive example, consisting of a Python statement and its expected output. The constructor arguments are used to initialize the attributes of the same names.

Example defines the following attributes. They are initialized by the constructor, and should not be modified directly.

source

A string containing the example's source code. This source code consists of a single Python statement, and always ends with a newline; the constructor adds a newline when necessary.

want

The expected output from running the example's source code (either from stdout, or a traceback in case of exception). *want* ends with a newline unless no output is expected, in which case it's an empty string. The constructor adds a newline when necessary.

exc_msg

The exception message generated by the example, if the example is expected to generate an exception; or *None* if it is not expected to generate an exception. This exception message is compared against the return value of `traceback.format_exception_only()`. *exc_msg* ends with a newline unless it's *None*. The constructor adds a newline if needed.

lineno

The line number within the string containing this example where the example begins. This line number is zero-based with respect to the beginning of the containing string.

indent

The example's indentation in the containing string, i.e., the number of space characters that precede the example's first prompt.

options

A dictionary mapping from option flags to *True* or *False*, which is used to override default options for this example. Any option flags not contained in this dictionary are left at their default value (as specified by the *DocTestRunner*'s *optionflags*). By default, no options are set.

DocTestFinder objects

class `doctest.DocTestFinder` (*verbose=False*, *parser=DocTestParser()*, *recurse=True*, *exclude_empty=True*)

A processing class used to extract the *DocTests* that are relevant to a given object, from its docstring and the docstrings of its contained objects. *DocTests* can be extracted from modules, classes, functions, methods, staticmethods, classmethods, and properties.

The optional argument *verbose* can be used to display the objects searched by the finder. It defaults to *False* (no output).

The optional argument *parser* specifies the *DocTestParser* object (or a drop-in replacement) that is used to extract doctests from docstrings.

If the optional argument *recurse* is false, then *DocTestFinder.find()* will only examine the given object, and not any contained objects.

If the optional argument *exclude_empty* is false, then *DocTestFinder.find()* will include tests for objects with empty docstrings.

DocTestFinder defines the following method:

find(*obj*[, *name*][, *module*][, *globs*][, *extraglobs*])

Return a list of the *DocTests* that are defined by *obj*' s docstring, or by any of its contained objects' docstrings.

The optional argument *name* specifies the object' s name; this name will be used to construct names for the returned *DocTests*. If *name* is not specified, then *obj* . *__name__* is used.

The optional parameter *module* is the module that contains the given object. If the module is not specified or is *None*, then the test finder will attempt to automatically determine the correct module. The object' s module is used:

- As a default namespace, if *globs* is not specified.
- To prevent the *DocTestFinder* from extracting *DocTests* from objects that are imported from other modules. (Contained objects with modules other than *module* are ignored.)
- To find the name of the file containing the object.
- To help find the line number of the object within its file.

If *module* is *False*, no attempt to find the module will be made. This is obscure, of use mostly in testing *doctest* itself: if *module* is *False*, or is *None* but cannot be found automatically, then all objects are considered to belong to the (non-existent) module, so all contained objects will (recursively) be searched for *doctests*.

The globals for each *DocTest* is formed by combining *globs* and *extraglobs* (bindings in *extraglobs* override bindings in *globs*). A new shallow copy of the globals dictionary is created for each *DocTest*. If *globs* is not specified, then it defaults to the module' s *__dict__*, if specified, or *{ }* otherwise. If *extraglobs* is not specified, then it defaults to *{ }*.

DocTestParser objects

class *doctest.DocTestParser*

A processing class used to extract interactive examples from a string, and use them to create a *DocTest* object.

DocTestParser defines the following methods:

get_doctest (*string*, *globs*, *name*, *filename*, *lineno*)

Extract all *doctest* examples from the given string, and collect them into a *DocTest* object.

globs, *name*, *filename*, and *lineno* are attributes for the new *DocTest* object. See the documentation for *DocTest* for more information.

get_examples (*string*, *name*=<string>')

Extract all *doctest* examples from the given string, and return them as a list of *Example* objects. Line numbers are 0-based. The optional argument *name* is a name identifying this string, and is only used for error messages.

parse (*string*, *name*=<string>')

Divide the given string into examples and intervening text, and return them as a list of alternating *Examples* and strings. Line numbers for the *Examples* are 0-based. The optional argument *name* is a name identifying this string, and is only used for error messages.

DocTestRunner objects

class *doctest.DocTestRunner* (*checker*=*None*, *verbose*=*None*, *optionflags*=0)

A processing class used to execute and verify the interactive examples in a *DocTest*.

The comparison between expected outputs and actual outputs is done by an *OutputChecker*. This comparison may be customized with a number of option flags; see section *Option Flags* for more information. If the option flags are insufficient, then the comparison may also be customized by passing a subclass of *OutputChecker* to the constructor.

The test runner's display output can be controlled in two ways. First, an output function can be passed to `TestRunner.run()`; this function will be called with strings that should be displayed. It defaults to `sys.stdout.write`. If capturing the output is not sufficient, then the display output can be also customized by subclassing *DocTestRunner*, and overriding the methods *report_start()*, *report_success()*, *report_unexpected_exception()*, and *report_failure()*.

The optional keyword argument *checker* specifies the *OutputChecker* object (or drop-in replacement) that should be used to compare the expected outputs to the actual outputs of doctest examples.

The optional keyword argument *verbose* controls the *DocTestRunner*'s verbosity. If *verbose* is `True`, then information is printed about each example, as it is run. If *verbose* is `False`, then only failures are printed. If *verbose* is unspecified, or `None`, then verbose output is used iff the command-line switch `-v` is used.

The optional keyword argument *optionflags* can be used to control how the test runner compares expected output to actual output, and how it displays failures. For more information, see section *Option Flags*.

DocTestParser defines the following methods:

report_start (*out, test, example*)

Report that the test runner is about to process the given example. This method is provided to allow subclasses of *DocTestRunner* to customize their output; it should not be called directly.

example is the example about to be processed. *test* is the test containing *example*. *out* is the output function that was passed to *DocTestRunner.run()*.

report_success (*out, test, example, got*)

Report that the given example ran successfully. This method is provided to allow subclasses of *DocTestRunner* to customize their output; it should not be called directly.

example is the example about to be processed. *got* is the actual output from the example. *test* is the test containing *example*. *out* is the output function that was passed to *DocTestRunner.run()*.

report_failure (*out, test, example, got*)

Report that the given example failed. This method is provided to allow subclasses of *DocTestRunner* to customize their output; it should not be called directly.

example is the example about to be processed. *got* is the actual output from the example. *test* is the test containing *example*. *out* is the output function that was passed to *DocTestRunner.run()*.

report_unexpected_exception (*out, test, example, exc_info*)

Report that the given example raised an unexpected exception. This method is provided to allow subclasses of *DocTestRunner* to customize their output; it should not be called directly.

example is the example about to be processed. *exc_info* is a tuple containing information about the unexpected exception (as returned by `sys.exc_info()`). *test* is the test containing *example*. *out* is the output function that was passed to *DocTestRunner.run()*.

run (*test, compileflags=None, out=None, clear_globs=True*)

Run the examples in *test* (a *DocTest* object), and display the results using the writer function *out*.

The examples are run in the namespace `test.globs`. If *clear_globs* is true (the default), then this namespace will be cleared after the test runs, to help with garbage collection. If you would like to examine the namespace after the test completes, then use *clear_globs=False*.

compileflags gives the set of flags that should be used by the Python compiler when running the examples. If not specified, then it will default to the set of future-import flags that apply to *globs*.

The output of each example is checked using the `DocTestRunner`'s output checker, and the results are formatted by the `DocTestRunner.report_*()` methods.

summarize (*verbose=None*)

Print a summary of all the test cases that have been run by this `DocTestRunner`, and return a *named tuple* `TestResults(failed, attempted)`.

The optional *verbose* argument controls how detailed the summary is. If the verbosity is not specified, then the `DocTestRunner`'s verbosity is used.

OutputChecker objects

class `doctest.OutputChecker`

A class used to check the whether the actual output from a doctest example matches the expected output. `OutputChecker` defines two methods: `check_output()`, which compares a given pair of outputs, and returns true if they match; and `output_difference()`, which returns a string describing the differences between two outputs.

`OutputChecker` defines the following methods:

check_output (*want, got, optionflags*)

Return True iff the actual output from an example (*got*) matches the expected output (*want*). These strings are always considered to match if they are identical; but depending on what option flags the test runner is using, several non-exact match types are also possible. See section *Option Flags* for more information about option flags.

output_difference (*example, got, optionflags*)

Return a string describing the differences between the expected output for a given example (*example*) and the actual output (*got*). *optionflags* is the set of option flags used to compare *want* and *got*.

26.3.7 调试

Doctest provides several mechanisms for debugging doctest examples:

- Several functions convert doctests to executable Python programs, which can be run under the Python debugger, `pdb`.
- The `DebugRunner` class is a subclass of `DocTestRunner` that raises an exception for the first failing example, containing information about that example. This information can be used to perform post-mortem debugging on the example.
- The `unittest` cases generated by `DocTestSuite()` support the `debug()` method defined by `unittest.TestCase`.
- You can add a call to `pdb.set_trace()` in a doctest example, and you'll drop into the Python debugger when that line is executed. Then you can inspect current values of variables, and so on. For example, suppose `a.py` contains just this module docstring:

```
"""
>>> def f(x):
...     g(x*2)
>>> def g(x):
...     print(x+3)
...     import pdb; pdb.set_trace()
>>> f(3)
9
"""
```

Then an interactive Python session may look like this:

```
>>> import a, doctest
>>> doctest.testmod(a)
--Return--
> <doctest a[1]>(3)g()->None
-> import pdb; pdb.set_trace()
(Pdb) list
1      def g(x):
2          print(x+3)
3  ->      import pdb; pdb.set_trace()
[EOF]
(Pdb) p x
6
(Pdb) step
--Return--
> <doctest a[0]>(2)f()->None
-> g(x*2)
(Pdb) list
1      def f(x):
2  ->      g(x*2)
[EOF]
(Pdb) p x
3
(Pdb) step
--Return--
> <doctest a[2]>(1)?()->None
-> f(3)
(Pdb) cont
(0, 3)
>>>
```

Functions that convert doctests to Python code, and possibly run the synthesized code under the debugger:

`doctest.script_from_examples(s)`

Convert text with examples to a script.

Argument *s* is a string containing doctest examples. The string is converted to a Python script, where doctest examples in *s* are converted to regular code, and everything else is converted to Python comments. The generated script is returned as a string. For example,

```
import doctest
print(doctest.script_from_examples(r"""
    Set x and y to 1 and 2.
    >>> x, y = 1, 2

    Print their sum:
    >>> print(x+y)
    3
    """))
```

displays:

```
# Set x and y to 1 and 2.
x, y = 1, 2
#
# Print their sum:
print(x+y)
```

(下页继续)

(续上页)

```
# Expected:
## 3
```

This function is used internally by other functions (see below), but can also be useful when you want to transform an interactive Python session into a Python script.

`doctest.testsource (module, name)`

Convert the doctest for an object to a script.

Argument *module* is a module object, or dotted name of a module, containing the object whose doctests are of interest. Argument *name* is the name (within the module) of the object with the doctests of interest. The result is a string, containing the object's docstring converted to a Python script, as described for `script_from_examples()` above. For example, if module `a.py` contains a top-level function `f()`, then

```
import a, doctest
print(doctest.testsource(a, "a.f"))
```

prints a script version of function `f()`'s docstring, with doctests converted to code, and the rest placed in comments.

`doctest.debug (module, name, pm=False)`

Debug the doctests for an object.

The *module* and *name* arguments are the same as for function `testsource()` above. The synthesized Python script for the named object's docstring is written to a temporary file, and then that file is run under the control of the Python debugger, `pdb`.

A shallow copy of `module.__dict__` is used for both local and global execution context.

Optional argument *pm* controls whether post-mortem debugging is used. If *pm* has a true value, the script file is run directly, and the debugger gets involved only if the script terminates via raising an unhandled exception. If it does, then post-mortem debugging is invoked, via `pdb.post_mortem()`, passing the traceback object from the unhandled exception. If *pm* is not specified, or is false, the script is run under the debugger from the start, via passing an appropriate `exec()` call to `pdb.run()`.

`doctest.debug_src (src, pm=False, globs=None)`

Debug the doctests in a string.

This is like function `debug()` above, except that a string containing doctest examples is specified directly, via the *src* argument.

Optional argument *pm* has the same meaning as in function `debug()` above.

Optional argument *globs* gives a dictionary to use as both local and global execution context. If not specified, or None, an empty dictionary is used. If specified, a shallow copy of the dictionary is used.

The `DebugRunner` class, and the special exceptions it may raise, are of most interest to testing framework authors, and will only be sketched here. See the source code, and especially `DebugRunner`'s docstring (which is a doctest!) for more details:

class `doctest.DebugRunner (checker=None, verbose=None, optionflags=0)`

A subclass of `DocTestRunner` that raises an exception as soon as a failure is encountered. If an unexpected exception occurs, an `UnexpectedException` exception is raised, containing the test, the example, and the original exception. If the output doesn't match, then a `DocTestFailure` exception is raised, containing the test, the example, and the actual output.

For information about the constructor parameters and methods, see the documentation for `DocTestRunner` in section [Advanced API](#).

There are two exceptions that may be raised by `DebugRunner` instances:

exception `doctest.DocTestFailure` (*test*, *example*, *got*)

An exception raised by `DocTestRunner` to signal that a doctest example’s actual output did not match its expected output. The constructor arguments are used to initialize the attributes of the same names.

`DocTestFailure` defines the following attributes:

`DocTestFailure.test`

The `DocTest` object that was being run when the example failed.

`DocTestFailure.example`

The `Example` that failed.

`DocTestFailure.got`

The example’s actual output.

exception `doctest.UnexpectedException` (*test*, *example*, *exc_info*)

An exception raised by `DocTestRunner` to signal that a doctest example raised an unexpected exception. The constructor arguments are used to initialize the attributes of the same names.

`UnexpectedException` defines the following attributes:

`UnexpectedException.test`

The `DocTest` object that was being run when the example failed.

`UnexpectedException.example`

The `Example` that failed.

`UnexpectedException.exc_info`

A tuple containing information about the unexpected exception, as returned by `sys.exc_info()`.

26.3.8 Soapbox

As mentioned in the introduction, `doctest` has grown to have three primary uses:

1. Checking examples in docstrings.
2. Regression testing.
3. Executable documentation / literate testing.

These uses have different requirements, and it is important to distinguish them. In particular, filling your docstrings with obscure test cases makes for bad documentation.

When writing a docstring, choose docstring examples with care. There’s an art to this that needs to be learned—it may not be natural at first. Examples should add genuine value to the documentation. A good example can often be worth many words. If done with care, the examples will be invaluable for your users, and will pay back the time it takes to collect them many times over as the years go by and things change. I’m still amazed at how often one of my `doctest` examples stops working after a “harmless” change.

Doctest also makes an excellent tool for regression testing, especially if you don’t skimp on explanatory text. By interleaving prose and examples, it becomes much easier to keep track of what’s actually being tested, and why. When a test fails, good prose can make it much easier to figure out what the problem is, and how it should be fixed. It’s true that you could write extensive comments in code-based testing, but few programmers do. Many have found that using doctest approaches instead leads to much clearer tests. Perhaps this is simply because doctest makes writing prose a little easier than writing code, while writing comments in code is a little harder. I think it goes deeper than just that: the natural attitude when writing a doctest-based test is that you want to explain the fine points of your software, and illustrate them with examples. This in turn naturally leads to test files that start with the simplest features, and logically progress to complications and edge cases. A coherent narrative is the result, instead of a collection of isolated functions that test isolated bits of functionality seemingly at random. It’s a different attitude, and produces different results, blurring the distinction between testing and explaining.

Regression testing is best confined to dedicated objects or files. There are several options for organizing tests:

- Write text files containing test cases as interactive examples, and test the files using `testfile()` or `DocFileSuite()`. This is recommended, although is easiest to do for new projects, designed from the start to use doctest.
- Define functions named `_regtest_topic` that consist of single docstrings, containing test cases for the named topics. These functions can be included in the same file as the module, or separated out into a separate test file.
- Define a `__test__` dictionary mapping from regression test topics to docstrings containing test cases.

When you have placed your tests in a module, the module can itself be the test runner. When a test fails, you can arrange for your test runner to re-run only the failing doctest while you debug the problem. Here is a minimal example of such a test runner:

```
if __name__ == '__main__':
    import doctest
    flags = doctest.REPORT_NDIFF|doctest.FAIL_FAST
    if len(sys.argv) > 1:
        name = sys.argv[1]
        if name in globals():
            obj = globals()[name]
        else:
            obj = __test__[name]
        doctest.run_docstring_examples(obj, globals(), name=name,
                                      optionflags=flags)
    else:
        fail, total = doctest.testmod(optionflags=flags)
        print("{} failures out of {} tests".format(fail, total))
```

26.4 unittest — 单元测试框架

源代码: `Lib/unittest/__init__.py`

(如果你已经对测试的概念比较熟悉了, 你可能想直接跳转到这一部分[断言方法](#)。)

`unittest` 单元测试框架是受到 JUnit 的启发, 与其他语言中的主流单元测试框架有着相似的风格。其支持测试自动化, 配置共享和关机代码测试。支持将测试样例聚合到测试集中, 并将测试与报告框架独立。

为了实现这些, `unittest` 通过面向对象的方式支持了一些重要的概念。

测试脚手架 `test fixture` 表示为了开展一项或多项测试所需要进行的准备工作, 以及所有相关的清理操作。举个例子, 这可能包含创建临时或代理的数据库、目录, 再或者启动一个服务器进程。

测试用例 一个测试用例是一个独立的测试单元。它检查输入特定的数据时的响应。`unittest` 提供一个基类: `TestCase`, 用于新建测试用例。

测试套件 `test suite` 是一系列的测试用例, 或测试套件, 或两者皆有。它用于归档需要一起执行的测试。

测试运行器 (test runner) `test runner` 是一个用于执行和输出测试结果的组件。这个运行器可能使用图形接口、文本接口, 或返回一个特定的值表示运行测试的结果。

参见:

doctest — 文档测试模块 另一个风格完全不同的测试模块。

Simple Smalltalk Testing: With Patterns Kent Beck's original paper on testing frameworks using the pattern shared by `unittest`.

Nose 和 pytest Third-party unittest frameworks with a lighter-weight syntax for writing tests. For example, `assert func(10) == 42`.

The Python Testing Tools Taxonomy An extensive list of Python testing tools including functional testing frameworks and mock object libraries.

Testing in Python Mailing List A special-interest-group for discussion of testing, and testing tools, in Python.

The script `Tools/unittestgui/unittestgui.py` in the Python source distribution is a GUI tool for test discovery and execution. This is intended largely for ease of use for those new to unit testing. For production environments it is recommended that tests be driven by a continuous integration system such as [Buildbot](#), [Jenkins](#) or [Hudson](#).

26.4.1 基本实例

`unittest` 模块提供了一系列创建和运行测试的工具。这一段落演示了这些工具的一小部分，但也足以满足大部分用户的需求。

这是一段简短的代码，来测试三种字符串方法：

```
import unittest

class TestStringMethods(unittest.TestCase):

    def test_upper(self):
        self.assertEqual('foo'.upper(), 'FOO')

    def test_isupper(self):
        self.assertTrue('FOO'.isupper())
        self.assertFalse('Foo'.isupper())

    def test_split(self):
        s = 'hello world'
        self.assertEqual(s.split(), ['hello', 'world'])
        # check that s.split fails when the separator is not a string
        with self.assertRaises(TypeError):
            s.split(2)

if __name__ == '__main__':
    unittest.main()
```

继承 `unittest.TestCase` 就创建了一个测试样例。上述三个独立的测试是三个类的方法，这些方法的命名都以 `test` 开头。这个命名约定告诉测试运行者类的哪些方法表示测试。

每个测试的关键是：调用 `assertEqual()` 来检查预期的输出；调用 `assertTrue()` 或 `assertFalse()` 来验证一个条件；调用 `assertRaises()` 来验证抛出了一个特定的异常。使用这些方法而不是 `assert` 语句是为了让测试运行者能聚合所有的测试结果并产生结果报告。

通过 `setUp()` 和 `tearDown()` 方法，可以设置测试开始前与完成后需要执行的指令。在[组织你的测试代码](#)中，对此有更为详细的描述。

最后的代码块中，演示了运行测试的一个简单的方法。`unittest.main()` 提供了一个测试脚本的命令行接口。当在命令行运行该测试脚本，上文的脚本生成如下格式的输出：

```
...
-----
Ran 3 tests in 0.000s

OK
```

在调用测试脚本时添加 `-v` 参数使 `unittest.main()` 显示更为详细的信息，生成如以下形式的输出：

```
test_isupper (__main__.TestStringMethods) ... ok
test_split (__main__.TestStringMethods) ... ok
test_upper (__main__.TestStringMethods) ... ok

-----
Ran 3 tests in 0.001s

OK
```

以上例子演示了 `unittest` 中最常用的、足够满足许多日常测试需求的特性。文档的剩余部分详述该框架的完整特性。

26.4.2 命令行界面

`unittest` 模块可以通过命令行运行模块、类和独立测试方法的测试：

```
python -m unittest test_module1 test_module2
python -m unittest test_module.TestClass
python -m unittest test_module.TestClass.test_method
```

你可以传入模块名、类或方法名或他们的任意组合。

同样的，测试模块可以通过文件路径指定：

```
python -m unittest tests/test_something.py
```

这样就可以使用 `shell` 的文件名补全指定测试模块。所指定的文件仍需要可以被作为模块导入。路径通过去除 `‘.py’`、把分隔符转换为 `‘.’` 转换为模块名。若你需要执行不能被作为模块导入的测试文件，你需要直接执行该测试文件。

在运行测试时，你可以通过添加 `-v` 参数获取更详细（更多的冗余）的信息。

```
python -m unittest -v test_module
```

当运行时不包含参数，开始探索性测试

```
python -m unittest
```

用于获取命令行选项列表：

```
python -m unittest -h
```

在 3.2 版更改：在早期版本中，只支持运行独立的测试方法，而不支持模块和类。

命令行选项

`unittest` supports these command-line options:

-b, --buffer

在测试运行时，标准输出流与标准错误流会被放入缓冲区。成功的测试的运行时输出会被丢弃；测试不通过时，测试运行中的输出会正常显示，错误会被加入到测试失败信息。

-c, --catch

当测试正在运行时，Control-C 会等待当前测试完成，并在完成后报告已执行的测试的结果。当再次按下 Control-C 时，引发平常的 `KeyboardInterrupt` 异常。

See *Signal Handling* for the functions that provide this functionality.

-f, --failfast

当出现第一个错误或者失败时，停止运行测试。

--locals

在回溯中显示局部变量。

3.2 新版功能: 添加命令行选项 `-b`, `-c` 和 `-f`。

3.5 新版功能: 命令行选项 `--locals`。

命令行亦可用于探索性测试，以运行一个项目的所有测试或其子集。

26.4.3 探索性测试

3.2 新版功能.

Unittest 支持简单的测试搜索。若需要使用探索性测试，所有的测试文件必须是 `modules` 或 `packages`（包括 *namespace packages*）并可从项目根目录导入（即它们的文件名必须是有效的 `identifiers`）。

探索性测试在 `TestLoader.discover()` 中实现，但也可以通过命令行使用。它在命令行中的基本用法如下：

```
cd project_directory
python -m unittest discover
```

注解：方便起见，`python -m unittest` 与 `python -m unittest discover` 等价。如果你需要向探索性测试传入参数，必须显式地使用 `discover` 子命令。

`discover` 有以下选项：

-v, --verbose

更详细地输出结果。

-s, --start-directory directory

开始进行搜索的目录（默认为当前目录 `.`）。

-p, --pattern pattern

用于匹配测试文件的模式（默认为 `test*.py`）。

-t, --top-level-directory directory

指定项目的最上层目录（通常为开始时所在目录）。

`-s`，`-p` 和 `-t` 选项可以按顺序作为位置参数传入。以下两条命令是等价的：

```
python -m unittest discover -s project_directory -p "*_test.py"
python -m unittest discover project_directory "*_test.py"
```

正如可以传入路径那样，传入一个包名作为起始目录也是可行的，如 `myproject.subpackage.test`。你提供的包名会被导入，它在文件系统中的位置会被作为起始目录。

警告：探索性测试通过导入测试对测试进行加载。在找到所有你指定的开始目录下的所有测试文件后，它把路径转换为包名并进行导入。如 `foo/bar/baz.py` 会被导入为 `foo.bar.baz`。

如果你有一个全局安装的包，并尝试对这个包的副本进行探索性测试，可能会从错误的地方开始导入。如果出现这种情况，测试会输出警告并退出。

如果你使用包名而不是路径作为开始目录，搜索时会假定它导入的是你想要的目录，所以你不会收到警告。

测试模块和包可以通过`load_tests protocol` 自定义测试的加载和搜索。

在 3.4 版更改: 探索性测试支持命名空间包 (*namespace packages*)。

26.4.4 组织你的测试代码

单元测试的构建单位是 *test cases*：独立的、包含执行条件与正确性检查的方案。在 `unittest` 中，测试用例表示为 `unittest.TestCase` 的实例。通过编写 `TestCase` 的子类或使用 `FunctionTestCase` 编写你自己的测试用例。

一个 `TestCase` 实例的测试代码必须是完全自含的，因此它可以独立运行，或与其它任意组合任意数量的测试用例一起运行。

`TestCase` 的最简单的子类需要实现一个测试方法（例如一个命名以 `test` 开头的方法）以执行特定的测试代码：

```
import unittest

class DefaultWidgetSizeTestCase(unittest.TestCase):
    def test_default_widget_size(self):
        widget = Widget('The widget')
        self.assertEqual(widget.size(), (50, 50))
```

可以看到，为了进行测试，我们使用了基类 `TestCase` 提供的其中一个 `assert*`() 方法。若测试不通过，将会引发一个带有说明信息的异常，并且 `unittest` 会将这个测试用例标记为测试不通过。任何其它类型的异常将会被当做错误处理。

可能同时存在多个前置操作相同的测试，我们可以把测试的前置操作从测试代码中拆解出来，并实现测试前置方法 `setUp()`。在运行测试时，测试框架会自动地为每个单独测试调用前置方法。

```
import unittest

class WidgetTestCase(unittest.TestCase):
    def setUp(self):
        self.widget = Widget('The widget')

    def test_default_widget_size(self):
        self.assertEqual(self.widget.size(), (50, 50),
                         'incorrect default size')

    def test_widget_resize(self):
        self.widget.resize(100, 150)
        self.assertEqual(self.widget.size(), (100, 150),
                         'wrong size after resize')
```

注解：多个测试运行的顺序由内置字符串排序方法对测试名进行排序的结果决定。

在测试运行时，若 `setUp()` 方法引发异常，测试框架会认为测试发生了错误，因此测试方法不会被运行。相似的，我们提供了一个 `tearDown()` 方法在测试方法运行后进行清理工作。

```
import unittest

class WidgetTestCase(unittest.TestCase):
    def setUp(self):
        self.widget = Widget('The widget')

    def tearDown(self):
        self.widget.dispose()
```

若 `setUp()` 成功运行，无论测试方法是否成功，都会运行 `tearDown()`。

这样的一个测试代码运行的环境被称为 *test fixture*。一个新的 `TestCase` 实例作为一个测试脚手架，用于运行各个独立的测试方法。在运行每个测试时，`setUp()`、`tearDown()` 和 `__init__()` 会被调用一次。

It is recommended that you use `TestCase` implementations to group tests together according to the features they test. `unittest` provides a mechanism for this: the *test suite*, represented by `unittest`'s `TestSuite` class. In most cases, calling `unittest.main()` will do the right thing and collect all the module's test cases for you and execute them.

然而，如果你需要自定义你的测试套件的话，你可以参考以下方法组织你的测试：

```
def suite():
    suite = unittest.TestSuite()
    suite.addTest(WidgetTestCase('test_default_widget_size'))
    suite.addTest(WidgetTestCase('test_widget_resize'))
    return suite

if __name__ == '__main__':
    runner = unittest.TextTestRunner()
    runner.run(suite())
```

You can place the definitions of test cases and test suites in the same modules as the code they are to test (such as `widget.py`), but there are several advantages to placing the test code in a separate module, such as `test_widget.py`:

- The test module can be run standalone from the command line.
- The test code can more easily be separated from shipped code.
- There is less temptation to change test code to fit the code it tests without a good reason.
- Test code should be modified much less frequently than the code it tests.
- Tested code can be refactored more easily.
- Tests for modules written in C must be in separate modules anyway, so why not be consistent?
- If the testing strategy changes, there is no need to change the source code.

26.4.5 复用已有的测试代码

一些用户希望直接使用 `unittest` 运行已有的测试代码，而不需要把已有的每个测试函数转化为一个 `TestCase` 的子类。

因此，`unittest` 提供 `FunctionTestCase` 类。这个 `TestCase` 的子类可用于打包已有的测试函数，并支持设置前置与后置函数。

假定有一个测试函数：

```
def testSomething():
    something = makeSomething()
    assert something.name is not None
    # ...
```

可以创建等价的测试用例如下，其中前置和后置方法是可选的。

```
testcase = unittest.FunctionTestCase(testSomething,
                                     setUp=makeSomethingDB,
                                     tearDown=deleteSomethingDB)
```

注解： Even though `FunctionTestCase` can be used to quickly convert an existing test base over to a `unittest`-based system, this approach is not recommended. Taking the time to set up proper `TestCase` subclasses will make future test refactorings infinitely easier.

In some cases, the existing tests may have been written using the `doctest` module. If so, `doctest` provides a `DocTestSuite` class that can automatically build `unittest.TestSuite` instances from the existing `doctest`-based tests.

26.4.6 跳过测试与预计的失败

3.1 新版功能.

Unittest supports skipping individual test methods and even whole classes of tests. In addition, it supports marking a test as an “expected failure,” a test that is broken and will fail, but shouldn’t be counted as a failure on a `TestResult`.

Skipping a test is simply a matter of using the `skip()` decorator or one of its conditional variants.

跳过测试的基本用法如下：

```
class MyTestCase(unittest.TestCase):

    @unittest.skip("demonstrating skipping")
    def test_nothing(self):
        self.fail("shouldn't happen")

    @unittest.skipIf(mylib.__version__ < (1, 3),
                    "not supported in this library version")
    def test_format(self):
        # Tests that work for only a certain version of the library.
        pass

    @unittest.skipUnless(sys.platform.startswith("win"), "requires Windows")
    def test_windows_support(self):
        # windows specific testing code
        pass
```

在`Python`交互式模式下运行以上测试例子时，程序输出如下：

```
test_format (__main__.MyTestCase) ... skipped 'not supported in this library version'
test_nothing (__main__.MyTestCase) ... skipped 'demonstrating skipping'
test_windows_support (__main__.MyTestCase) ... skipped 'requires Windows'

-----
Ran 3 tests in 0.005s
```

(下页继续)

(续上页)

```
OK (skipped=3)
```

跳过测试类的写法跟跳过测试方法的写法相似:

```
@unittest.skip("showing class skipping")
class MySkippedTestCase(unittest.TestCase):
    def test_not_run(self):
        pass
```

`TestCase.setUp()` 也可以跳过测试。可以用于所需资源不可用的情况下跳过接下来的测试。

使用 `expectedFailure()` 装饰器表明这个测试预计失败。:

```
class ExpectedFailureTestCase(unittest.TestCase):
    @unittest.expectedFailure
    def test_fail(self):
        self.assertEqual(1, 0, "broken")
```

It's easy to roll your own skipping decorators by making a decorator that calls `skip()` on the test when it wants it to be skipped. This decorator skips the test unless the passed object has a certain attribute:

```
def skipUnlessHasattr(obj, attr):
    if hasattr(obj, attr):
        return lambda func: func
    return unittest.skip("{!r} doesn't have {!r}".format(obj, attr))
```

The following decorators implement test skipping and expected failures:

`@unittest.skip(reason)`

跳过被此装饰器装饰的测试。*reason* 为测试被跳过的原因。

`@unittest.skipIf(condition, reason)`

当 *condition* 为真时, 跳过被装饰的测试。

`@unittest.skipUnless(condition, reason)`

跳过被装饰的测试, 除非 *condition* 为真。

`@unittest.expectedFailure`

Mark the test as an expected failure. If the test fails when run, the test is not counted as a failure.

exception `unittest.SkipTest(reason)`

引发此异常以跳过一个测试。

通常来说, 你可以使用 `TestCase.skipTest()` 或其中一个跳过测试的装饰器实现跳过测试的功能, 而不是直接引发此异常。

被跳过的测试的 `setUp()` 和 `tearDown()` 不会被运行。被跳过的类的 `setUpClass()` 和 `tearDownClass()` 不会被运行。被跳过的模块的 `setUpModule()` 和 `tearDownModule()` 不会被运行。

26.4.7 Distinguishing test iterations using subtests

3.4 新版功能.

When there are very small differences among your tests, for instance some parameters, unittest allows you to distinguish them inside the body of a test method using the `subTest()` context manager.

例如, 以下测试:

```
class NumbersTest(unittest.TestCase):

    def test_even(self):
        """
        Test that numbers between 0 and 5 are all even.
        """
        for i in range(0, 6):
            with self.subTest(i=i):
                self.assertEqual(i % 2, 0)
```

可以得到以下输出:

```
=====
FAIL: test_even (__main__.NumbersTest) (i=1)
-----
Traceback (most recent call last):
  File "subtests.py", line 32, in test_even
    self.assertEqual(i % 2, 0)
AssertionError: 1 != 0

=====
FAIL: test_even (__main__.NumbersTest) (i=3)
-----
Traceback (most recent call last):
  File "subtests.py", line 32, in test_even
    self.assertEqual(i % 2, 0)
AssertionError: 1 != 0

=====
FAIL: test_even (__main__.NumbersTest) (i=5)
-----
Traceback (most recent call last):
  File "subtests.py", line 32, in test_even
    self.assertEqual(i % 2, 0)
AssertionError: 1 != 0
```

Without using a subtest, execution would stop after the first failure, and the error would be less easy to diagnose because the value of `i` wouldn't be displayed:

```
=====
FAIL: test_even (__main__.NumbersTest)
-----
Traceback (most recent call last):
  File "subtests.py", line 32, in test_even
    self.assertEqual(i % 2, 0)
AssertionError: 1 != 0
```

26.4.8 类与函数

本节深入介绍了 `unittest` 的 API。

测试用例

```
class unittest.TestCase (methodName='runTest')
```

Instances of the `TestCase` class represent the logical test units in the `unittest` universe. This class is intended

to be used as a base class, with specific tests being implemented by concrete subclasses. This class implements the interface needed by the test runner to allow it to drive the tests, and methods that the test code can use to check for and report various kinds of failure.

Each instance of `TestCase` will run a single base method: the method named `methodName`. In most uses of `TestCase`, you will neither change the `methodName` nor reimplement the default `runTest()` method.

在 3.2 版更改: `TestCase` can be instantiated successfully without providing a `methodName`. This makes it easier to experiment with `TestCase` from the interactive interpreter.

`TestCase` instances provide three groups of methods: one group used to run the test, another used by the test implementation to check conditions and report failures, and some inquiry methods allowing information about the test itself to be gathered.

Methods in the first group (running the test) are:

setUp()

Method called to prepare the test fixture. This is called immediately before calling the test method; other than `AssertionError` or `SkipTest`, any exception raised by this method will be considered an error rather than a test failure. The default implementation does nothing.

tearDown()

Method called immediately after the test method has been called and the result recorded. This is called even if the test method raised an exception, so the implementation in subclasses may need to be particularly careful about checking internal state. Any exception, other than `AssertionError` or `SkipTest`, raised by this method will be considered an additional error rather than a test failure (thus increasing the total number of reported errors). This method will only be called if the `setUp()` succeeds, regardless of the outcome of the test method. The default implementation does nothing.

setUpClass()

A class method called before tests in an individual class are run. `setUpClass` is called with the class as the only argument and must be decorated as a `classmethod()`:

```
@classmethod
def setUpClass(cls):
    ...
```

查看 [Class and Module Fixtures](#) 获取更详细的说明。

3.2 新版功能。

tearDownClass()

A class method called after tests in an individual class have run. `tearDownClass` is called with the class as the only argument and must be decorated as a `classmethod()`:

```
@classmethod
def tearDownClass(cls):
    ...
```

查看 [Class and Module Fixtures](#) 获取更详细的说明。

3.2 新版功能。

run(result=None)

Run the test, collecting the result into the `TestResult` object passed as `result`. If `result` is omitted or `None`, a temporary result object is created (by calling the `defaultTestResult()` method) and used. The result object is returned to `run()`'s caller.

The same effect may be had by simply calling the `TestCase` instance.

在 3.3 版更改: Previous versions of `run` did not return the result. Neither did calling an instance.

skipTest (*reason*)

Calling this during a test method or `setUp()` skips the current test. See [跳过测试与预计的失败](#) for more information.

3.1 新版功能.

subTest (*msg=None, **params*)

Return a context manager which executes the enclosed code block as a subtest. *msg* and *params* are optional, arbitrary values which are displayed whenever a subtest fails, allowing you to identify them clearly.

A test case can contain any number of subtest declarations, and they can be arbitrarily nested.

查看[Distinguishing test iterations using subtests](#) 获取更详细的信息。

3.4 新版功能.

debug ()

Run the test without collecting the result. This allows exceptions raised by the test to be propagated to the caller, and can be used to support running tests under a debugger.

The `TestCase` class provides several assert methods to check for and report failures. The following table lists the most commonly used methods (see the tables below for more assert methods):

Method	Checks that	New in
<code>assertEqual(a, b)</code>	<code>a == b</code>	
<code>assertNotEqual(a, b)</code>	<code>a != b</code>	
<code>assertTrue(x)</code>	<code>bool(x) is True</code>	
<code>assertFalse(x)</code>	<code>bool(x) is False</code>	
<code>assertIs(a, b)</code>	<code>a is b</code>	3.1
<code>assertIsNot(a, b)</code>	<code>a is not b</code>	3.1
<code>assertIsNone(x)</code>	<code>x is None</code>	3.1
<code>assertIsNotNone(x)</code>	<code>x is not None</code>	3.1
<code>assertIn(a, b)</code>	<code>a in b</code>	3.1
<code>assertNotIn(a, b)</code>	<code>a not in b</code>	3.1
<code>assertIsInstance(a, b)</code>	<code>isinstance(a, b)</code>	3.2
<code>assertNotIsInstance(a, b)</code>	<code>not isinstance(a, b)</code>	3.2

All the assert methods accept a *msg* argument that, if specified, is used as the error message on failure (see also `longMessage`). Note that the *msg* keyword argument can be passed to `assertRaises()`, `assertRaisesRegex()`, `assertWarns()`, `assertWarnsRegex()` only when they are used as a context manager.

assertEqual (*first, second, msg=None*)

Test that *first* and *second* are equal. If the values do not compare equal, the test will fail.

In addition, if *first* and *second* are the exact same type and one of list, tuple, dict, set, frozenset or str or any type that a subclass registers with `addTypeEqualityFunc()` the type-specific equality function will be called in order to generate a more useful default error message (see also the [list of type-specific methods](#)).

在 3.1 版更改: Added the automatic calling of type-specific equality function.

在 3.2 版更改: `assertMultiLineEqual()` added as the default type equality function for comparing strings.

assertNotEqual (*first, second, msg=None*)

Test that *first* and *second* are not equal. If the values do compare equal, the test will fail.

assertTrue (*expr, msg=None*)

assertFalse (*expr*, *msg=None*)

Test that *expr* is true (or false).

Note that this is equivalent to `bool(expr) is True` and not to `expr is True` (use `assertIs(expr, True)` for the latter). This method should also be avoided when more specific methods are available (e.g. `assertEqual(a, b)` instead of `assertTrue(a == b)`), because they provide a better error message in case of failure.

assertIs (*first*, *second*, *msg=None*)

assertIsNot (*first*, *second*, *msg=None*)

Test that *first* and *second* evaluate (or don't evaluate) to the same object.

3.1 新版功能.

assertIsNone (*expr*, *msg=None*)

assertIsNotNone (*expr*, *msg=None*)

Test that *expr* is (or is not) `None`.

3.1 新版功能.

assertIn (*first*, *second*, *msg=None*)

assertNotIn (*first*, *second*, *msg=None*)

Test that *first* is (or is not) in *second*.

3.1 新版功能.

assertIsInstance (*obj*, *cls*, *msg=None*)

assertNotIsInstance (*obj*, *cls*, *msg=None*)

Test that *obj* is (or is not) an instance of *cls* (which can be a class or a tuple of classes, as supported by `isinstance()`). To check for the exact type, use `assertIs(type(obj), cls)`.

3.2 新版功能.

It is also possible to check the production of exceptions, warnings, and log messages using the following methods:

Method	Checks that	New in
<code>assertRaises(exc, fun, *args, **kwargs)</code>	<code>fun(*args, **kwargs)</code> raises <i>exc</i>	
<code>assertRaisesRegex(exc, r, fun, *args, **kwargs)</code>	<code>fun(*args, **kwargs)</code> raises <i>exc</i> and the message matches regex <i>r</i>	3.1
<code>assertWarns(warn, fun, *args, **kwargs)</code>	<code>fun(*args, **kwargs)</code> raises <i>warn</i>	3.2
<code>assertWarnsRegex(warn, r, fun, *args, **kwargs)</code>	<code>fun(*args, **kwargs)</code> raises <i>warn</i> and the message matches regex <i>r</i>	3.2
<code>assertLogs(logger, level)</code>	The <code>with</code> block logs on <i>logger</i> with minimum <i>level</i>	3.4

assertRaises (*exception*, *callable*, **args*, ***kwargs*)

assertRaises (*exception*, *, *msg=None*)

Test that an exception is raised when *callable* is called with any positional or keyword arguments that are also passed to `assertRaises()`. The test passes if *exception* is raised, is an error if another exception is raised, or fails if no exception is raised. To catch any of a group of exceptions, a tuple containing the exception classes may be passed as *exception*.

If only the *exception* and possibly the *msg* arguments are given, return a context manager so that the code under test can be written inline rather than as a function:

```
with self.assertRaises(SomeException):
    do_something()
```

When used as a context manager, `assertRaises()` accepts the additional keyword argument `msg`.

The context manager will store the caught exception object in its `exception` attribute. This can be useful if the intention is to perform additional checks on the exception raised:

```
with self.assertRaises(SomeException) as cm:
    do_something()

the_exception = cm.exception
self.assertEqual(the_exception.error_code, 3)
```

在 3.1 版更改: Added the ability to use `assertRaises()` as a context manager.

在 3.2 版更改: Added the `exception` attribute.

在 3.3 版更改: Added the `msg` keyword argument when used as a context manager.

assertRaisesRegex (*exception, regex, callable, *args, **kwargs*)

assertRaisesRegex (*exception, regex, *, msg=None*)

Like `assertRaises()` but also tests that `regex` matches on the string representation of the raised exception. `regex` may be a regular expression object or a string containing a regular expression suitable for use by `re.search()`. Examples:

```
self.assertRaisesRegex(ValueError, "invalid literal for.*XYZ'",
                        int, 'XYZ')
```

或者:

```
with self.assertRaisesRegex(ValueError, 'literal'):
    int('XYZ')
```

3.1 新版功能: under the name `assertRaisesRegexp`.

在 3.2 版更改: Renamed to `assertRaisesRegex()`.

在 3.3 版更改: Added the `msg` keyword argument when used as a context manager.

assertWarns (*warning, callable, *args, **kwargs*)

assertWarns (*warning, *, msg=None*)

Test that a warning is triggered when `callable` is called with any positional or keyword arguments that are also passed to `assertWarns()`. The test passes if `warning` is triggered and fails if it isn't. Any exception is an error. To catch any of a group of warnings, a tuple containing the warning classes may be passed as `warnings`.

If only the `warning` and possibly the `msg` arguments are given, return a context manager so that the code under test can be written inline rather than as a function:

```
with self.assertWarns(SomeWarning):
    do_something()
```

When used as a context manager, `assertWarns()` accepts the additional keyword argument `msg`.

The context manager will store the caught warning object in its `warning` attribute, and the source line which triggered the warnings in the `filename` and `lineno` attributes. This can be useful if the intention is to perform additional checks on the warning caught:

```
with self.assertWarns(SomeWarning) as cm:
    do_something()
```

(下页继续)

(续上页)

```
self.assertIn('myfile.py', cm.filename)
self.assertEqual(320, cm.lineno)
```

This method works regardless of the warning filters in place when it is called.

3.2 新版功能.

在 3.3 版更改: Added the *msg* keyword argument when used as a context manager.

assertWarnsRegex (*warning, regex, callable, *args, **kws*)

assertWarnsRegex (*warning, regex, *, msg=None*)

Like *assertWarns()* but also tests that *regex* matches on the message of the triggered warning. *regex* may be a regular expression object or a string containing a regular expression suitable for use by *re.search()*.

Example:

```
self.assertWarnsRegex(DeprecationWarning,
                       r'legacy_function\(\) is deprecated',
                       legacy_function, 'XYZ')
```

或者:

```
with self.assertWarnsRegex(RuntimeWarning, 'unsafe frobnicating'):
    frobnicate('/etc/passwd')
```

3.2 新版功能.

在 3.3 版更改: Added the *msg* keyword argument when used as a context manager.

assertLogs (*logger=None, level=None*)

A context manager to test that at least one message is logged on the *logger* or one of its children, with at least the given *level*.

If given, *logger* should be a *logging.Logger* object or a *str* giving the name of a logger. The default is the root logger, which will catch all messages.

If given, *level* should be either a numeric logging level or its string equivalent (for example either "ERROR" or *logging.ERROR*). The default is *logging.INFO*.

The test passes if at least one message emitted inside the *with* block matches the *logger* and *level* conditions, otherwise it fails.

The object returned by the context manager is a recording helper which keeps tracks of the matching log messages. It has two attributes:

records

A list of *logging.LogRecord* objects of the matching log messages.

output

A list of *str* objects with the formatted output of matching messages.

示例:

```
with self.assertLogs('foo', level='INFO') as cm:
    logging.getLogger('foo').info('first message')
    logging.getLogger('foo.bar').error('second message')
self.assertEqual(cm.output, ['INFO:foo:first message',
                              'ERROR:foo.bar:second message'])
```

3.4 新版功能.

There are also other methods used to perform more specific checks, such as:

Method	Checks that	New in
<code>assertAlmostEqual(a, b)</code>	<code>round(a-b, 7) == 0</code>	
<code>assertNotAlmostEqual(a, b)</code>	<code>round(a-b, 7) != 0</code>	
<code>assertGreater(a, b)</code>	<code>a > b</code>	3.1
<code>assertGreaterEqual(a, b)</code>	<code>a >= b</code>	3.1
<code>assertLess(a, b)</code>	<code>a < b</code>	3.1
<code>assertLessEqual(a, b)</code>	<code>a <= b</code>	3.1
<code>assertRegex(s, r)</code>	<code>r.search(s)</code>	3.1
<code>assertNotRegex(s, r)</code>	<code>not r.search(s)</code>	3.2
<code>assertCountEqual(a, b)</code>	<i>a</i> and <i>b</i> have the same elements in the same number, regardless of their order	3.2

assertAlmostEqual (*first*, *second*, *places*=7, *msg*=None, *delta*=None)

assertNotAlmostEqual (*first*, *second*, *places*=7, *msg*=None, *delta*=None)

Test that *first* and *second* are approximately (or not approximately) equal by computing the difference, rounding to the given number of decimal *places* (default 7), and comparing to zero. Note that these methods round the values to the given number of *decimal places* (i.e. like the `round()` function) and not *significant digits*.

If *delta* is supplied instead of *places* then the difference between *first* and *second* must be less or equal to (or greater than) *delta*.

Supplying both *delta* and *places* raises a `TypeError`.

在 3.2 版更改: `assertAlmostEqual()` automatically considers almost equal objects that compare equal. `assertNotAlmostEqual()` automatically fails if the objects compare equal. Added the *delta* keyword argument.

assertGreater (*first*, *second*, *msg*=None)

assertGreaterEqual (*first*, *second*, *msg*=None)

assertLess (*first*, *second*, *msg*=None)

assertLessEqual (*first*, *second*, *msg*=None)

Test that *first* is respectively `>`, `>=`, `<` or `<=` than *second* depending on the method name. If not, the test will fail:

```
>>> self.assertGreaterEqual(3, 4)
AssertionError: "3" unexpectedly not greater than or equal to "4"
```

3.1 新版功能.

assertRegex (*text*, *regex*, *msg*=None)

assertNotRegex (*text*, *regex*, *msg*=None)

Test that a *regex* search matches (or does not match) *text*. In case of failure, the error message will include the pattern and the *text* (or the pattern and the part of *text* that unexpectedly matched). *regex* may be a regular expression object or a string containing a regular expression suitable for use by `re.search()`.

3.1 新版功能: under the name `assertRegexpMatches`.

在 3.2 版更改: The method `assertRegexpMatches()` has been renamed to `assertRegex()`.

3.2 新版功能: `assertNotRegex()`

3.5 新版功能: The name `assertNotRegexpMatches` is a deprecated alias for `assertNotRegex()`.

assertCountEqual (*first*, *second*, *msg=None*)

Test that sequence *first* contains the same elements as *second*, regardless of their order. When they don't, an error message listing the differences between the sequences will be generated.

Duplicate elements are *not* ignored when comparing *first* and *second*. It verifies whether each element has the same count in both sequences. Equivalent to: `assertEqual(Counter(list(first)), Counter(list(second)))` but works with sequences of unhashable objects as well.

3.2 新版功能.

The `assertEqual()` method dispatches the equality check for objects of the same type to different type-specific methods. These methods are already implemented for most of the built-in types, but it's also possible to register new methods using `addTypeEqualityFunc()`:

addTypeEqualityFunc (*typeobj*, *function*)

Registers a type-specific method called by `assertEqual()` to check if two objects of exactly the same *typeobj* (not subclasses) compare equal. *function* must take two positional arguments and a third *msg=None* keyword argument just as `assertEqual()` does. It must raise `self.failureException(msg)` when inequality between the first two parameters is detected – possibly providing useful information and explaining the inequalities in details in the error message.

3.1 新版功能.

The list of type-specific methods automatically used by `assertEqual()` are summarized in the following table. Note that it's usually not necessary to invoke these methods directly.

Method	Used to compare	New in
<code>assertMultiLineEqual(a, b)</code>	strings	3.1
<code>assertSequenceEqual(a, b)</code>	sequences	3.1
<code>assertListEqual(a, b)</code>	lists	3.1
<code>assertTupleEqual(a, b)</code>	tuples	3.1
<code>assertSetEqual(a, b)</code>	sets or frozensets	3.1
<code>assertDictEqual(a, b)</code>	dicts	3.1

assertMultiLineEqual (*first*, *second*, *msg=None*)

Test that the multiline string *first* is equal to the string *second*. When not equal a diff of the two strings highlighting the differences will be included in the error message. This method is used by default when comparing strings with `assertEqual()`.

3.1 新版功能.

assertSequenceEqual (*first*, *second*, *msg=None*, *seq_type=None*)

Tests that two sequences are equal. If a *seq_type* is supplied, both *first* and *second* must be instances of *seq_type* or a failure will be raised. If the sequences are different an error message is constructed that shows the difference between the two.

This method is not called directly by `assertEqual()`, but it's used to implement `assertListEqual()` and `assertTupleEqual()`.

3.1 新版功能.

assertListEqual (*first*, *second*, *msg=None*)

assertTupleEqual (*first*, *second*, *msg=None*)

Tests that two lists or tuples are equal. If not, an error message is constructed that shows only the differences between the two. An error is also raised if either of the parameters are of the wrong type. These methods are used by default when comparing lists or tuples with `assertEqual()`.

3.1 新版功能.

assertSetEqual (*first, second, msg=None*)

Tests that two sets are equal. If not, an error message is constructed that lists the differences between the sets. This method is used by default when comparing sets or frozensets with `assertEqual()`.

Fails if either of *first* or *second* does not have a `set.difference()` method.

3.1 新版功能.

assertDictEqual (*first, second, msg=None*)

Test that two dictionaries are equal. If not, an error message is constructed that shows the differences in the dictionaries. This method will be used by default to compare dictionaries in calls to `assertEqual()`.

3.1 新版功能.

Finally the `TestCase` provides the following methods and attributes:

fail (*msg=None*)

Signals a test failure unconditionally, with *msg* or `None` for the error message.

failureException

This class attribute gives the exception raised by the test method. If a test framework needs to use a specialized exception, possibly to carry additional information, it must subclass this exception in order to “play fair” with the framework. The initial value of this attribute is `AssertionError`.

longMessage

This class attribute determines what happens when a custom failure message is passed as the *msg* argument to an `assertXXX` call that fails. `True` is the default value. In this case, the custom message is appended to the end of the standard failure message. When set to `False`, the custom message replaces the standard message.

The class setting can be overridden in individual test methods by assigning an instance attribute, `self.longMessage`, to `True` or `False` before calling the assert methods.

The class setting gets reset before each test call.

3.1 新版功能.

maxDiff

This attribute controls the maximum length of diffs output by assert methods that report diffs on failure. It defaults to 80*8 characters. Assert methods affected by this attribute are `assertSequenceEqual()` (including all the sequence comparison methods that delegate to it), `assertDictEqual()` and `assertMultiLineEqual()`.

Setting `maxDiff` to `None` means that there is no maximum length of diffs.

3.2 新版功能.

Testing frameworks can use the following methods to collect information on the test:

countTestCases ()

Return the number of tests represented by this test object. For `TestCase` instances, this will always be 1.

defaultTestResult ()

Return an instance of the test result class that should be used for this test case class (if no other result instance is provided to the `run()` method).

For `TestCase` instances, this will always be an instance of `TestResult`; subclasses of `TestCase` should override this as necessary.

id ()

Return a string identifying the specific test case. This is usually the full name of the test method, including the module and class name.

shortDescription()

Returns a description of the test, or `None` if no description has been provided. The default implementation of this method returns the first line of the test method's docstring, if available, or `None`.

在 3.1 版更改: In 3.1 this was changed to add the test name to the short description even in the presence of a docstring. This caused compatibility issues with unittest extensions and adding the test name was moved to the `TextTestResult` in Python 3.2.

addCleanup() (*function, *args, **kwargs*)

Add a function to be called after `tearDown()` to cleanup resources used during the test. Functions will be called in reverse order to the order they are added (LIFO). They are called with any arguments and keyword arguments passed into `addCleanup()` when they are added.

If `setUp()` fails, meaning that `tearDown()` is not called, then any cleanup functions added will still be called.

3.1 新版功能.

doCleanups()

This method is called unconditionally after `tearDown()`, or after `setUp()` if `setUp()` raises an exception.

It is responsible for calling all the cleanup functions added by `addCleanup()`. If you need cleanup functions to be called *prior* to `tearDown()` then you can call `doCleanups()` yourself.

`doCleanups()` pops methods off the stack of cleanup functions one at a time, so it can be called at any time.

3.1 新版功能.

class unittest.FunctionTestCase (*testFunc, setUp=None, tearDown=None, description=None*)

This class implements the portion of the `TestCase` interface which allows the test runner to drive the test, but does not provide the methods which test code can use to check and report errors. This is used to create test cases using legacy test code, allowing it to be integrated into a `unittest`-based test framework.

Deprecated aliases

For historical reasons, some of the `TestCase` methods had one or more aliases that are now deprecated. The following table lists the correct names along with their deprecated aliases:

Method Name	Deprecated alias	Deprecated alias
<code>assertEqual()</code>	<code>failUnlessEqual</code>	<code>assertEquals</code>
<code>assertNotEqual()</code>	<code>failIfEqual</code>	<code>assertNotEquals</code>
<code>assertTrue()</code>	<code>failUnless</code>	<code>assert_</code>
<code>assertFalse()</code>	<code>failIf</code>	
<code>assertRaises()</code>	<code>failUnlessRaises</code>	
<code>assertAlmostEqual()</code>	<code>failUnlessAlmostEqual</code>	<code>assertAlmostEquals</code>
<code>assertNotAlmostEqual()</code>	<code>failIfAlmostEqual</code>	<code>assertNotAlmostEquals</code>
<code>assertRegex()</code>		<code>assertRegexpMatches</code>
<code>assertNotRegex()</code>		<code>assertNotRegexpMatches</code>
<code>assertRaisesRegex()</code>		<code>assertRaisesRegexp</code>

3.1 版后已移除: the `fail*` aliases listed in the second column.

3.2 版后已移除: the `assert*` aliases listed in the third column.

3.2 版后已移除: `assertRegexpMatches` and `assertRaisesRegexp` have been renamed to `assertRegex()` and `assertRaisesRegex()`.

3.5 版后已移除: the `assertNotRegexpMatches` name in favor of `assertNotRegex()`.

Grouping tests

class `unittest.TestSuite (tests=())`

This class represents an aggregation of individual test cases and test suites. The class presents the interface needed by the test runner to allow it to be run as any other test case. Running a `TestSuite` instance is the same as iterating over the suite, running each test individually.

If `tests` is given, it must be an iterable of individual test cases or other test suites that will be used to build the suite initially. Additional methods are provided to add test cases and suites to the collection later on.

`TestSuite` objects behave much like `TestCase` objects, except they do not actually implement a test. Instead, they are used to aggregate tests into groups of tests that should be run together. Some additional methods are available to add tests to `TestSuite` instances:

addTest (`test`)

Add a `TestCase` or `TestSuite` to the suite.

addTests (`tests`)

Add all the tests from an iterable of `TestCase` and `TestSuite` instances to this test suite.

This is equivalent to iterating over `tests`, calling `addTest()` for each element.

`TestSuite` shares the following methods with `TestCase`:

run (`result`)

Run the tests associated with this suite, collecting the result into the test result object passed as `result`. Note that unlike `TestCase.run()`, `TestSuite.run()` requires the result object to be passed in.

debug ()

Run the tests associated with this suite without collecting the result. This allows exceptions raised by the test to be propagated to the caller and can be used to support running tests under a debugger.

countTestCases ()

Return the number of tests represented by this test object, including all individual tests and sub-suites.

__iter__ ()

Tests grouped by a `TestSuite` are always accessed by iteration. Subclasses can lazily provide tests by overriding `__iter__()`. Note that this method may be called several times on a single suite (for example when counting tests or comparing for equality) so the tests returned by repeated iterations before `TestSuite.run()` must be the same for each call iteration. After `TestSuite.run()`, callers should not rely on the tests returned by this method unless the caller uses a subclass that overrides `TestSuite._removeTestAtIndex()` to preserve test references.

在 3.2 版更改: In earlier versions the `TestSuite` accessed tests directly rather than through iteration, so overriding `__iter__()` wasn't sufficient for providing tests.

在 3.4 版更改: In earlier versions the `TestSuite` held references to each `TestCase` after `TestSuite.run()`. Subclasses can restore that behavior by overriding `TestSuite._removeTestAtIndex()`.

In the typical usage of a `TestSuite` object, the `run()` method is invoked by a `TestRunner` rather than by the end-user test harness.

Loading and running tests

class `unittest.TestLoader`

The `TestLoader` class is used to create test suites from classes and modules. Normally, there is no need to

create an instance of this class; the `unittest` module provides an instance that can be shared as `unittest.defaultTestLoader`. Using a subclass or instance, however, allows customization of some configurable properties.

`TestLoader` objects have the following attributes:

errors

A list of the non-fatal errors encountered while loading tests. Not reset by the loader at any point. Fatal errors are signalled by the relevant a method raising an exception to the caller. Non-fatal errors are also indicated by a synthetic test that will raise the original error when run.

3.5 新版功能.

`TestLoader` objects have the following methods:

loadTestsFromTestCase (testCaseClass)

Return a suite of all test cases contained in the `TestCase`-derived `testCaseClass`.

A test case instance is created for each method named by `getTestCaseNames()`. By default these are the method names beginning with `test`. If `getTestCaseNames()` returns no methods, but the `runTest()` method is implemented, a single test case is created for that method instead.

loadTestsFromModule (module, pattern=None)

Return a suite of all test cases contained in the given module. This method searches `module` for classes derived from `TestCase` and creates an instance of the class for each test method defined for the class.

注解: While using a hierarchy of `TestCase`-derived classes can be convenient in sharing fixtures and helper functions, defining test methods on base classes that are not intended to be instantiated directly does not play well with this method. Doing so, however, can be useful when the fixtures are different and defined in subclasses.

If a module provides a `load_tests` function it will be called to load the tests. This allows modules to customize test loading. This is the *load_tests protocol*. The `pattern` argument is passed as the third argument to `load_tests`.

在 3.2 版更改: Support for `load_tests` added.

在 3.5 版更改: The undocumented and unofficial `use_load_tests` default argument is deprecated and ignored, although it is still accepted for backward compatibility. The method also now accepts a keyword-only argument `pattern` which is passed to `load_tests` as the third argument.

loadTestsFromName (name, module=None)

Return a suite of all test cases given a string specifier.

The specifier `name` is a “dotted name” that may resolve either to a module, a test case class, a test method within a test case class, a `TestSuite` instance, or a callable object which returns a `TestCase` or `TestSuite` instance. These checks are applied in the order listed here; that is, a method on a possible test case class will be picked up as “a test method within a test case class”, rather than “a callable object”.

For example, if you have a module `SampleTests` containing a `TestCase`-derived class `SampleTestCase` with three test methods (`test_one()`, `test_two()`, and `test_three()`), the specifier `'SampleTests.SampleTestCase'` would cause this method to return a suite which will run all three test methods. Using the specifier `'SampleTests.SampleTestCase.test_two'` would cause it to return a test suite which will run only the `test_two()` test method. The specifier can refer to modules and packages which have not been imported; they will be imported as a side-effect.

The method optionally resolves `name` relative to the given `module`.

在 3.5 版更改: If an `ImportError` or `AttributeError` occurs while traversing `name` then a synthetic test that raises that error when run will be returned. These errors are included in the errors accumulated by

`self.errors.`

loadTestsFromNames (*names*, *module=None*)

Similar to `loadTestsFromName()`, but takes a sequence of names rather than a single name. The return value is a test suite which supports all the tests defined for each name.

getTestCaseNames (*testCaseClass*)

Return a sorted sequence of method names found within *testCaseClass*; this should be a subclass of `TestCase`.

discover (*start_dir*, *pattern='test*.py'*, *top_level_dir=None*)

Find all the test modules by recursing into subdirectories from the specified start directory, and return a `TestSuite` object containing them. Only test files that match *pattern* will be loaded. (Using shell style pattern matching.) Only module names that are importable (i.e. are valid Python identifiers) will be loaded.

All test modules must be importable from the top level of the project. If the start directory is not the top level directory then the top level directory must be specified separately.

If importing a module fails, for example due to a syntax error, then this will be recorded as a single error and discovery will continue. If the import failure is due to `SkipTest` being raised, it will be recorded as a skip instead of an error.

If a package (a directory containing a file named `__init__.py`) is found, the package will be checked for a `load_tests` function. If this exists then it will be called `package.load_tests(loader, tests, pattern)`. Test discovery takes care to ensure that a package is only checked for tests once during an invocation, even if the `load_tests` function itself calls `loader.discover`.

If `load_tests` exists then discovery does *not* recurse into the package, `load_tests` is responsible for loading all tests in the package.

The pattern is deliberately not stored as a loader attribute so that packages can continue discovery themselves. *top_level_dir* is stored so `load_tests` does not need to pass this argument in to `loader.discover()`.

start_dir can be a dotted module name as well as a directory.

3.2 新版功能.

在 3.4 版更改: Modules that raise `SkipTest` on import are recorded as skips, not errors. Discovery works for *namespace packages*. Paths are sorted before being imported so that execution order is the same even if the underlying file system's ordering is not dependent on file name.

在 3.5 版更改: Found packages are now checked for `load_tests` regardless of whether their path matches *pattern*, because it is impossible for a package name to match the default pattern.

The following attributes of a `TestLoader` can be configured either by subclassing or assignment on an instance:

testMethodPrefix

String giving the prefix of method names which will be interpreted as test methods. The default value is `'test'`.

This affects `getTestCaseNames()` and all the `loadTestsFrom*()` methods.

sortTestMethodsUsing

Function to be used to compare method names when sorting them in `getTestCaseNames()` and all the `loadTestsFrom*()` methods.

suiteClass

Callable object that constructs a test suite from a list of tests. No methods on the resulting object are needed. The default value is the `TestSuite` class.

This affects all the `loadTestsFrom*()` methods.

class `unittest.TestResult`

This class is used to compile information about which tests have succeeded and which have failed.

A `TestResult` object stores the results of a set of tests. The `TestCase` and `TestSuite` classes ensure that results are properly recorded; test authors do not need to worry about recording the outcome of tests.

Testing frameworks built on top of `unittest` may want access to the `TestResult` object generated by running a set of tests for reporting purposes; a `TestResult` instance is returned by the `TestRunner.run()` method for this purpose.

`TestResult` instances have the following attributes that will be of interest when inspecting the results of running a set of tests:

errors

A list containing 2-tuples of `TestCase` instances and strings holding formatted tracebacks. Each tuple represents a test which raised an unexpected exception.

failures

A list containing 2-tuples of `TestCase` instances and strings holding formatted tracebacks. Each tuple represents a test where a failure was explicitly signalled using the `TestCase.assert*()` methods.

skipped

A list containing 2-tuples of `TestCase` instances and strings holding the reason for skipping the test.

3.1 新版功能.

expectedFailures

A list containing 2-tuples of `TestCase` instances and strings holding formatted tracebacks. Each tuple represents an expected failure of the test case.

unexpectedSuccesses

A list containing `TestCase` instances that were marked as expected failures, but succeeded.

shouldStop

Set to `True` when the execution of tests should stop by `stop()`.

testsRun

The total number of tests run so far.

buffer

If set to `true`, `sys.stdout` and `sys.stderr` will be buffered in between `startTest()` and `stopTest()` being called. Collected output will only be echoed onto the real `sys.stdout` and `sys.stderr` if the test fails or errors. Any output is also attached to the failure / error message.

3.2 新版功能.

failfast

If set to `true` `stop()` will be called on the first failure or error, halting the test run.

3.2 新版功能.

tb_locals

If set to `true` then local variables will be shown in tracebacks.

3.5 新版功能.

wasSuccessful()

Return `True` if all tests run so far have passed, otherwise returns `False`.

在 3.4 版更改: Returns `False` if there were any `unexpectedSuccesses` from tests marked with the `expectedFailure()` decorator.

stop()

This method can be called to signal that the set of tests being run should be aborted by setting the

`shouldStop` attribute to `True`. `TestRunner` objects should respect this flag and return without running any additional tests.

For example, this feature is used by the `TextTestRunner` class to stop the test framework when the user signals an interrupt from the keyboard. Interactive tools which provide `TestRunner` implementations can use this in a similar manner.

The following methods of the `TestResult` class are used to maintain the internal data structures, and may be extended in subclasses to support additional reporting requirements. This is particularly useful in building tools which support interactive reporting while tests are being run.

startTest (*test*)

Called when the test case *test* is about to be run.

stopTest (*test*)

Called after the test case *test* has been executed, regardless of the outcome.

startTestRun ()

Called once before any tests are executed.

3.1 新版功能.

stopTestRun ()

Called once after all tests are executed.

3.1 新版功能.

addError (*test*, *err*)

Called when the test case *test* raises an unexpected exception. *err* is a tuple of the form returned by `sys.exc_info()`: (type, value, traceback).

The default implementation appends a tuple (*test*, *formatted_err*) to the instance's `errors` attribute, where *formatted_err* is a formatted traceback derived from *err*.

addFailure (*test*, *err*)

Called when the test case *test* signals a failure. *err* is a tuple of the form returned by `sys.exc_info()`: (type, value, traceback).

The default implementation appends a tuple (*test*, *formatted_err*) to the instance's `failures` attribute, where *formatted_err* is a formatted traceback derived from *err*.

addSuccess (*test*)

Called when the test case *test* succeeds.

The default implementation does nothing.

addSkip (*test*, *reason*)

Called when the test case *test* is skipped. *reason* is the reason the test gave for skipping.

The default implementation appends a tuple (*test*, *reason*) to the instance's `skipped` attribute.

addExpectedFailure (*test*, *err*)

Called when the test case *test* fails, but was marked with the `expectedFailure()` decorator.

The default implementation appends a tuple (*test*, *formatted_err*) to the instance's `expectedFailures` attribute, where *formatted_err* is a formatted traceback derived from *err*.

addUnexpectedSuccess (*test*)

Called when the test case *test* was marked with the `expectedFailure()` decorator, but succeeded.

The default implementation appends the test to the instance's `unexpectedSuccesses` attribute.

addSubTest (*test*, *subtest*, *outcome*)

Called when a subtest finishes. *test* is the test case corresponding to the test method. *subtest* is a custom *TestCase* instance describing the subtest.

If *outcome* is *None*, the subtest succeeded. Otherwise, it failed with an exception where *outcome* is a tuple of the form returned by *sys.exc_info()*: (type, value, traceback).

The default implementation does nothing when the outcome is a success, and records subtest failures as normal failures.

3.4 新版功能.

class unittest.**TextTestResult** (*stream*, *descriptions*, *verbosity*)

A concrete implementation of *TestResult* used by the *TextTestRunner*.

3.2 新版功能: This class was previously named *_TextTestResult*. The old name still exists as an alias but is deprecated.

unittest.defaultTestLoader

Instance of the *TestLoader* class intended to be shared. If no customization of the *TestLoader* is needed, this instance can be used instead of repeatedly creating new instances.

class unittest.**TextTestRunner** (*stream=None*, *descriptions=True*, *verbosity=1*, *failfast=False*, *buffer=False*, *resultclass=None*, *warnings=None*, **, tb_locals=False*)

A basic test runner implementation that outputs results to a stream. If *stream* is *None*, the default, *sys.stderr* is used as the output stream. This class has a few configurable parameters, but is essentially very simple. Graphical applications which run test suites should provide alternate implementations. Such implementations should accept ***kwargs* as the interface to construct runners changes when features are added to unittest.

By default this runner shows *DeprecationWarning*, *PendingDeprecationWarning*, *ResourceWarning* and *ImportWarning* even if they are *ignored by default*. Deprecation warnings caused by *deprecated unittest methods* are also special-cased and, when the warning filters are 'default' or 'always', they will appear only once per-module, in order to avoid too many warning messages. This behavior can be overridden using Python's *-Wd* or *-Wa* options (see Warning control) and leaving *warnings* to *None*.

在 3.2 版更改: Added the *warnings* argument.

在 3.2 版更改: The default stream is set to *sys.stderr* at instantiation time rather than import time.

在 3.5 版更改: Added the *tb_locals* parameter.

_makeResult ()

This method returns the instance of *TestResult* used by *run()*. It is not intended to be called directly, but can be overridden in subclasses to provide a custom *TestResult*.

_makeResult() instantiates the class or callable passed in the *TextTestRunner* constructor as the *resultclass* argument. It defaults to *TextTestResult* if no *resultclass* is provided. The result class is instantiated with the following arguments:

```
stream, descriptions, verbosity
```

run (*test*)

This method is the main public interface to the *TextTestRunner*. This method takes a *TestSuite* or *TestCase* instance. A *TestResult* is created by calling *_makeResult()* and the test(s) are run and the results printed to stdout.

unittest.main (*module='__main__'*, *defaultTest=None*, *argv=None*, *testRunner=None*, *testLoader=unittest.defaultTestLoader*, *exit=True*, *verbosity=1*, *failfast=None*, *catchbreak=None*, *buffer=None*, *warnings=None*)

A command-line program that loads a set of tests from *module* and runs them; this is primarily for making test

modules conveniently executable. The simplest use for this function is to include the following line at the end of a test script:

```
if __name__ == '__main__':
    unittest.main()
```

You can run tests with more detailed information by passing in the verbosity argument:

```
if __name__ == '__main__':
    unittest.main(verbosity=2)
```

The *defaultTest* argument is either the name of a single test or an iterable of test names to run if no test names are specified via *argv*. If not specified or *None* and no test names are provided via *argv*, all tests found in *module* are run.

The *argv* argument can be a list of options passed to the program, with the first element being the program name. If not specified or *None*, the values of *sys.argv* are used.

The *testRunner* argument can either be a test runner class or an already created instance of it. By default *main* calls *sys.exit()* with an exit code indicating success or failure of the tests run.

The *testLoader* argument has to be a *TestLoader* instance, and defaults to *defaultTestLoader*.

main supports being used from the interactive interpreter by passing in the argument *exit=False*. This displays the result on standard output without calling *sys.exit()*:

```
>>> from unittest import main
>>> main(module='test_module', exit=False)
```

The *failfast*, *catchbreak* and *buffer* parameters have the same effect as the same-name *command-line options*.

The *warnings* argument specifies the *warning filter* that should be used while running the tests. If it's not specified, it will remain *None* if a *-W* option is passed to **python** (see Warning control), otherwise it will be set to 'default'.

Calling *main* actually returns an instance of the *TestProgram* class. This stores the result of the tests run as the *result* attribute.

在 3.1 版更改: The *exit* parameter was added.

在 3.2 版更改: The *verbosity*, *failfast*, *catchbreak*, *buffer* and *warnings* parameters were added.

在 3.4 版更改: The *defaultTest* parameter was changed to also accept an iterable of test names.

load_tests Protocol

3.2 新版功能.

Modules or packages can customize how tests are loaded from them during normal test runs or test discovery by implementing a function called *load_tests*.

If a test module defines *load_tests* it will be called by *TestLoader.loadTestsFromModule()* with the following arguments:

```
load_tests(loader, standard_tests, pattern)
```

where *pattern* is passed straight through from *loadTestsFromModule*. It defaults to *None*.

It should return a *TestSuite*.

loader is the instance of *TestLoader* doing the loading. *standard_tests* are the tests that would be loaded by default from the module. It is common for test modules to only want to add or remove tests from the standard set of tests. The third argument is used when loading packages as part of test discovery.

A typical `load_tests` function that loads tests from a specific set of *TestCase* classes may look like:

```
test_cases = (TestCase1, TestCase2, TestCase3)

def load_tests(loader, tests, pattern):
    suite = TestSuite()
    for test_class in test_cases:
        tests = loader.loadTestsFromTestCase(test_class)
        suite.addTests(tests)
    return suite
```

If discovery is started in a directory containing a package, either from the command line or by calling *TestLoader.discover()*, then the package `__init__.py` will be checked for `load_tests`. If that function does not exist, discovery will recurse into the package as though it were just another directory. Otherwise, discovery of the package's tests will be left up to `load_tests` which is called with the following arguments:

```
load_tests(loader, standard_tests, pattern)
```

This should return a *TestSuite* representing all the tests from the package. (*standard_tests* will only contain tests collected from `__init__.py`.)

Because the *pattern* is passed into `load_tests` the package is free to continue (and potentially modify) test discovery. A 'do nothing' `load_tests` function for a test package would look like:

```
def load_tests(loader, standard_tests, pattern):
    # top level directory cached on loader instance
    this_dir = os.path.dirname(__file__)
    package_tests = loader.discover(start_dir=this_dir, pattern=pattern)
    standard_tests.addTests(package_tests)
    return standard_tests
```

在 3.5 版更改: Discovery no longer checks package names for matching *pattern* due to the impossibility of package names matching the default pattern.

26.4.9 Class and Module Fixtures

Class and module level fixtures are implemented in *TestSuite*. When the test suite encounters a test from a new class then `tearDownClass()` from the previous class (if there is one) is called, followed by `setUpClass()` from the new class.

Similarly if a test is from a different module from the previous test then `tearDownModule` from the previous module is run, followed by `setUpModule` from the new module.

After all the tests have run the final `tearDownClass` and `tearDownModule` are run.

Note that shared fixtures do not play well with [potential] features like test parallelization and they break test isolation. They should be used with care.

The default ordering of tests created by the unittest test loaders is to group all tests from the same modules and classes together. This will lead to `setUpClass` / `setUpModule` (etc) being called exactly once per class and module. If you randomize the order, so that tests from different modules and classes are adjacent to each other, then these shared fixture functions may be called multiple times in a single test run.

Shared fixtures are not intended to work with suites with non-standard ordering. A *BaseTestSuite* still exists for frameworks that don't want to support shared fixtures.

If there are any exceptions raised during one of the shared fixture functions the test is reported as an error. Because there is no corresponding test instance an `_ErrorHandler` object (that has the same interface as a `TestCase`) is created to represent the error. If you are just using the standard unittest test runner then this detail doesn't matter, but if you are a framework author it may be relevant.

setUpClass and tearDownClass

These must be implemented as class methods:

```
import unittest

class Test(unittest.TestCase):
    @classmethod
    def setUpClass(cls):
        cls._connection = createExpensiveConnectionObject()

    @classmethod
    def tearDownClass(cls):
        cls._connection.destroy()
```

If you want the `setUpClass` and `tearDownClass` on base classes called then you must call up to them yourself. The implementations in `TestCase` are empty.

If an exception is raised during a `setUpClass` then the tests in the class are not run and the `tearDownClass` is not run. Skipped classes will not have `setUpClass` or `tearDownClass` run. If the exception is a `SkipTest` exception then the class will be reported as having been skipped instead of as an error.

setUpModule and tearDownModule

These should be implemented as functions:

```
def setUpModule():
    createConnection()

def tearDownModule():
    closeConnection()
```

If an exception is raised in a `setUpModule` then none of the tests in the module will be run and the `tearDownModule` will not be run. If the exception is a `SkipTest` exception then the module will be reported as having been skipped instead of as an error.

26.4.10 Signal Handling

3.2 新版功能.

The `-c/--catch` command-line option to unittest, along with the `catchbreak` parameter to `unittest.main()`, provide more friendly handling of control-C during a test run. With catch break behavior enabled control-C will allow the currently running test to complete, and the test run will then end and report all the results so far. A second control-c will raise a `KeyboardInterrupt` in the usual way.

The control-c handling signal handler attempts to remain compatible with code or tests that install their own `signal.SIGINT` handler. If the unittest handler is called but *isn't* the installed `signal.SIGINT` handler, i.e. it has been replaced by the system under test and delegated to, then it calls the default handler. This will normally be the expected behavior by code that replaces an installed handler and delegates to it. For individual tests that need unittest control-c handling disabled the `removeHandler()` decorator can be used.

There are a few utility functions for framework authors to enable control-c handling functionality within test frameworks.

`unittest.installHandler()`

Install the control-c handler. When a `signal.SIGINT` is received (usually in response to the user pressing control-c) all registered results have `stop()` called.

`unittest.registerResult(result)`

Register a `TestResult` object for control-c handling. Registering a result stores a weak reference to it, so it doesn't prevent the result from being garbage collected.

Registering a `TestResult` object has no side-effects if control-c handling is not enabled, so test frameworks can unconditionally register all results they create independently of whether or not handling is enabled.

`unittest.removeResult(result)`

Remove a registered result. Once a result has been removed then `stop()` will no longer be called on that result object in response to a control-c.

`unittest.removeHandler(function=None)`

When called without arguments this function removes the control-c handler if it has been installed. This function can also be used as a test decorator to temporarily remove the handler while the test is being executed:

```
@unittest.removeHandler
def test_signal_handling(self):
    ...
```

26.5 unittest.mock — mock object library

3.3 新版功能.

Source code: [Lib/unittest/mock.py](#)

`unittest.mock` is a library for testing in Python. It allows you to replace parts of your system under test with mock objects and make assertions about how they have been used.

`unittest.mock` provides a core `Mock` class removing the need to create a host of stubs throughout your test suite. After performing an action, you can make assertions about which methods / attributes were used and arguments they were called with. You can also specify return values and set needed attributes in the normal way.

Additionally, mock provides a `patch()` decorator that handles patching module and class level attributes within the scope of a test, along with `sentinel` for creating unique objects. See the [quick guide](#) for some examples of how to use `Mock`, `MagicMock` and `patch()`.

Mock is very easy to use and is designed for use with `unittest`. Mock is based on the ‘action -> assertion’ pattern instead of ‘record -> replay’ used by many mocking frameworks.

There is a backport of `unittest.mock` for earlier versions of Python, available as [mock on PyPI](#).

26.5.1 Quick Guide

`Mock` and `MagicMock` objects create all attributes and methods as you access them and store details of how they have been used. You can configure them, to specify return values or limit what attributes are available, and then make assertions about how they have been used:

```
>>> from unittest.mock import MagicMock
>>> thing = ProductionClass()
>>> thing.method = MagicMock(return_value=3)
>>> thing.method(3, 4, 5, key='value')
3
>>> thing.method.assert_called_with(3, 4, 5, key='value')
```

`side_effect` allows you to perform side effects, including raising an exception when a mock is called:

```
>>> mock = Mock(side_effect=KeyError('foo'))
>>> mock()
Traceback (most recent call last):
...
KeyError: 'foo'
```

```
>>> values = {'a': 1, 'b': 2, 'c': 3}
>>> def side_effect(arg):
...     return values[arg]
...
>>> mock.side_effect = side_effect
>>> mock('a'), mock('b'), mock('c')
(1, 2, 3)
>>> mock.side_effect = [5, 4, 3, 2, 1]
>>> mock(), mock(), mock()
(5, 4, 3)
```

Mock has many other ways you can configure it and control its behaviour. For example the *spec* argument configures the mock to take its specification from another object. Attempting to access attributes or methods on the mock that don't exist on the spec will fail with an *AttributeError*.

The *patch()* decorator / context manager makes it easy to mock classes or objects in a module under test. The object you specify will be replaced with a mock (or other object) during the test and restored when the test ends:

```
>>> from unittest.mock import patch
>>> @patch('module.ClassName2')
... @patch('module.ClassName1')
... def test(MockClass1, MockClass2):
...     module.ClassName1()
...     module.ClassName2()
...     assert MockClass1 is module.ClassName1
...     assert MockClass2 is module.ClassName2
...     assert MockClass1.called
...     assert MockClass2.called
...
>>> test()
```

注解: When you nest patch decorators the mocks are passed in to the decorated function in the same order they applied (the normal *python* order that decorators are applied). This means from the bottom up, so in the example above the mock for `module.ClassName1` is passed in first.

With *patch()* it matters that you patch objects in the namespace where they are looked up. This is normally straightforward, but for a quick guide read *where to patch*.

As well as a decorator *patch()* can be used as a context manager in a with statement:

```
>>> with patch.object(ProductionClass, 'method', return_value=None) as mock_method:
...     thing = ProductionClass()
...     thing.method(1, 2, 3)
...
>>> mock_method.assert_called_once_with(1, 2, 3)
```

There is also `patch.dict()` for setting values in a dictionary just during a scope and restoring the dictionary to its original state when the test ends:

```
>>> foo = {'key': 'value'}
>>> original = foo.copy()
>>> with patch.dict(foo, {'newkey': 'newvalue'}, clear=True):
...     assert foo == {'newkey': 'newvalue'}
...
>>> assert foo == original
```

Mock supports the mocking of Python *magic methods*. The easiest way of using magic methods is with the *MagicMock* class. It allows you to do things like:

```
>>> mock = MagicMock()
>>> mock.__str__.return_value = 'foobarbaz'
>>> str(mock)
'foobarbaz'
>>> mock.__str__.assert_called_with()
```

Mock allows you to assign functions (or other Mock instances) to magic methods and they will be called appropriately. The *MagicMock* class is just a Mock variant that has all of the magic methods pre-created for you (well, all the useful ones anyway).

The following is an example of using magic methods with the ordinary Mock class:

```
>>> mock = Mock()
>>> mock.__str__ = Mock(return_value='whewheee')
>>> str(mock)
'whewheee'
```

For ensuring that the mock objects in your tests have the same api as the objects they are replacing, you can use *auto-specing*. Auto-specing can be done through the *autospec* argument to patch, or the `create_autospec()` function. Auto-specing creates mock objects that have the same attributes and methods as the objects they are replacing, and any functions and methods (including constructors) have the same call signature as the real object.

This ensures that your mocks will fail in the same way as your production code if they are used incorrectly:

```
>>> from unittest.mock import create_autospec
>>> def function(a, b, c):
...     pass
...
>>> mock_function = create_autospec(function, return_value='fishy')
>>> mock_function(1, 2, 3)
'fishy'
>>> mock_function.assert_called_once_with(1, 2, 3)
>>> mock_function('wrong arguments')
Traceback (most recent call last):
...
TypeError: <lambda>() takes exactly 3 arguments (1 given)
```

`create_autospec()` can also be used on classes, where it copies the signature of the `__init__` method, and on callable objects where it copies the signature of the `__call__` method.

26.5.2 The Mock Class

Mock is a flexible mock object intended to replace the use of stubs and test doubles throughout your code. Mocks are callable and create attributes as new mocks when you access them¹. Accessing the same attribute will always return the same mock. Mocks record how you use them, allowing you to make assertions about what your code has done to them.

MagicMock is a subclass of *Mock* with all the magic methods pre-created and ready to use. There are also non-callable variants, useful when you are mocking out objects that aren't callable: *NonCallableMock* and *NonCallableMagicMock*.

The *patch()* decorator makes it easy to temporarily replace classes in a particular module with a *Mock* object. By default *patch()* will create a *MagicMock* for you. You can specify an alternative class of *Mock* using the *new_callable* argument to *patch()*.

```
class unittest.mock.Mock(spec=None, side_effect=None, return_value=DEFAULT, wraps=None,
                           name=None, spec_set=None, unsafe=False, **kwargs)
```

Create a new *Mock* object. *Mock* takes several optional arguments that specify the behaviour of the Mock object:

- *spec*: This can be either a list of strings or an existing object (a class or instance) that acts as the specification for the mock object. If you pass in an object then a list of strings is formed by calling *dir* on the object (excluding unsupported magic attributes and methods). Accessing any attribute not in this list will raise an *AttributeError*.

If *spec* is an object (rather than a list of strings) then *__class__* returns the class of the spec object. This allows mocks to pass *isinstance()* tests.

- *spec_set*: A stricter variant of *spec*. If used, attempting to *set* or get an attribute on the mock that isn't on the object passed as *spec_set* will raise an *AttributeError*.
- *side_effect*: A function to be called whenever the Mock is called. See the *side_effect* attribute. Useful for raising exceptions or dynamically changing return values. The function is called with the same arguments as the mock, and unless it returns *DEFAULT*, the return value of this function is used as the return value.

Alternatively *side_effect* can be an exception class or instance. In this case the exception will be raised when the mock is called.

If *side_effect* is an iterable then each call to the mock will return the next value from the iterable.

A *side_effect* can be cleared by setting it to *None*.

- *return_value*: The value returned when the mock is called. By default this is a new Mock (created on first access). See the *return_value* attribute.
- *unsafe*: By default if any attribute starts with *assert* or *assert* will raise an *AttributeError*. Passing *unsafe=True* will allow access to these attributes.

3.5 新版功能.

- *wraps*: Item for the mock object to wrap. If *wraps* is not *None* then calling the Mock will pass the call through to the wrapped object (returning the real result). Attribute access on the mock will return a Mock object that wraps the corresponding attribute of the wrapped object (so attempting to access an attribute that doesn't exist will raise an *AttributeError*).

If the mock has an explicit *return_value* set then calls are not passed to the wrapped object and the *return_value* is returned instead.

- *name*: If the mock has a name then it will be used in the repr of the mock. This can be useful for debugging. The name is propagated to child mocks.

¹ The only exceptions are magic methods and attributes (those that have leading and trailing double underscores). Mock doesn't create these but instead raises an *AttributeError*. This is because the interpreter will often implicitly request these methods, and gets very confused to get a new Mock object when it expects a magic method. If you need magic method support see *magic methods*.

Mocks can also be called with arbitrary keyword arguments. These will be used to set attributes on the mock after it is created. See the `configure_mock()` method for details.

assert_called(*args, **kwargs)

Assert that the mock was called at least once.

```
>>> mock = Mock()
>>> mock.method()
<Mock name='mock.method()' id='...'>
>>> mock.method.assert_called()
```

3.6 新版功能.

assert_called_once(*args, **kwargs)

Assert that the mock was called exactly once.

```
>>> mock = Mock()
>>> mock.method()
<Mock name='mock.method()' id='...'>
>>> mock.method.assert_called_once()
>>> mock.method()
<Mock name='mock.method()' id='...'>
>>> mock.method.assert_called_once()
Traceback (most recent call last):
...
AssertionError: Expected 'method' to have been called once. Called 2 times.
```

3.6 新版功能.

assert_called_with(*args, **kwargs)

This method is a convenient way of asserting that calls are made in a particular way:

```
>>> mock = Mock()
>>> mock.method(1, 2, 3, test='wow')
<Mock name='mock.method()' id='...'>
>>> mock.method.assert_called_with(1, 2, 3, test='wow')
```

assert_called_once_with(*args, **kwargs)

Assert that the mock was called exactly once and that that call was with the specified arguments.

```
>>> mock = Mock(return_value=None)
>>> mock('foo', bar='baz')
>>> mock.assert_called_once_with('foo', bar='baz')
>>> mock('other', bar='values')
>>> mock.assert_called_once_with('other', bar='values')
Traceback (most recent call last):
...
AssertionError: Expected 'mock' to be called once. Called 2 times.
```

assert_any_call(*args, **kwargs)

assert the mock has been called with the specified arguments.

The assert passes if the mock has *ever* been called, unlike `assert_called_with()` and `assert_called_once_with()` that only pass if the call is the most recent one, and in the case of `assert_called_once_with()` it must also be the only call.

```
>>> mock = Mock(return_value=None)
>>> mock(1, 2, arg='thing')
```

(下页继续)

(续上页)

```
>>> mock('some', 'thing', 'else')
>>> mock.assert_any_call(1, 2, arg='thing')
```

assert_has_calls (*calls*, *any_order=False*)

assert the mock has been called with the specified calls. The *mock_calls* list is checked for the calls.

If *any_order* is false (the default) then the calls must be sequential. There can be extra calls before or after the specified calls.

If *any_order* is true then the calls can be in any order, but they must all appear in *mock_calls*.

```
>>> mock = Mock(return_value=None)
>>> mock(1)
>>> mock(2)
>>> mock(3)
>>> mock(4)
>>> calls = [call(2), call(3)]
>>> mock.assert_has_calls(calls)
>>> calls = [call(4), call(2), call(3)]
>>> mock.assert_has_calls(calls, any_order=True)
```

assert_not_called()

Assert the mock was never called.

```
>>> m = Mock()
>>> m.hello.assert_not_called()
>>> obj = m.hello()
>>> m.hello.assert_not_called()
Traceback (most recent call last):
...
AssertionError: Expected 'hello' to not have been called. Called 1 times.
```

3.5 新版功能.

reset_mock (*, *return_value=False*, *side_effect=False*)

The *reset_mock* method resets all the call attributes on a mock object:

```
>>> mock = Mock(return_value=None)
>>> mock('hello')
>>> mock.called
True
>>> mock.reset_mock()
>>> mock.called
False
```

在 3.6 版更改: Added two keyword only argument to the *reset_mock* function.

This can be useful where you want to make a series of assertions that reuse the same object. Note that *reset_mock()* *doesn't* clear the return value, *side_effect* or any child attributes you have set using normal assignment by default. In case you want to reset *return_value* or *side_effect*, then pass the corresponding parameter as *True*. Child mocks and the return value mock (if any) are reset as well.

注解: *return_value*, and *side_effect* are keyword only argument.

mock_add_spec (*spec*, *spec_set=False*)

Add a spec to a mock. *spec* can either be an object or a list of strings. Only attributes on the *spec* can be

fetches attributes from the mock.

If `spec_set` is true then only attributes on the spec can be set.

attach_mock (*mock*, *attribute*)

Attach a mock as an attribute of this one, replacing its name and parent. Calls to the attached mock will be recorded in the `method_calls` and `mock_calls` attributes of this one.

configure_mock (***kwargs*)

Set attributes on the mock through keyword arguments.

Attributes plus return values and side effects can be set on child mocks using standard dot notation and unpacking a dictionary in the method call:

```
>>> mock = Mock()
>>> attrs = {'method.return_value': 3, 'other.side_effect': KeyError}
>>> mock.configure_mock(**attrs)
>>> mock.method()
3
>>> mock.other()
Traceback (most recent call last):
...
KeyError
```

The same thing can be achieved in the constructor call to mocks:

```
>>> attrs = {'method.return_value': 3, 'other.side_effect': KeyError}
>>> mock = Mock(some_attribute='eggs', **attrs)
>>> mock.some_attribute
'eggs'
>>> mock.method()
3
>>> mock.other()
Traceback (most recent call last):
...
KeyError
```

`configure_mock()` exists to make it easier to do configuration after the mock has been created.

__dir__ ()

`Mock` objects limit the results of `dir(some_mock)` to useful results. For mocks with a `spec` this includes all the permitted attributes for the mock.

See `FILTER_DIR` for what this filtering does, and how to switch it off.

__get_child_mock (***kw*)

Create the child mocks for attributes and return value. By default child mocks will be the same type as the parent. Subclasses of `Mock` may want to override this to customize the way child mocks are made.

For non-callable mocks the callable variant will be used (rather than any custom subclass).

called

A boolean representing whether or not the mock object has been called:

```
>>> mock = Mock(return_value=None)
>>> mock.called
False
>>> mock()
>>> mock.called
True
```

call_count

An integer telling you how many times the mock object has been called:

```
>>> mock = Mock(return_value=None)
>>> mock.call_count
0
>>> mock()
>>> mock()
>>> mock.call_count
2
```

return_value

Set this to configure the value returned by calling the mock:

```
>>> mock = Mock()
>>> mock.return_value = 'fish'
>>> mock()
'fish'
```

The default return value is a mock object and you can configure it in the normal way:

```
>>> mock = Mock()
>>> mock.return_value.attribute = sentinel.Attribute
>>> mock.return_value()
<Mock name='mock()' id='...'>
>>> mock.return_value.assert_called_with()
```

return_value can also be set in the constructor:

```
>>> mock = Mock(return_value=3)
>>> mock.return_value
3
>>> mock()
3
```

side_effect

This can either be a function to be called when the mock is called, an iterable or an exception (class or instance) to be raised.

If you pass in a function it will be called with same arguments as the mock and unless the function returns the *DEFAULT* singleton the call to the mock will then return whatever the function returns. If the function returns *DEFAULT* then the mock will return its normal value (from the *return_value*).

If you pass in an iterable, it is used to retrieve an iterator which must yield a value on every call. This value can either be an exception instance to be raised, or a value to be returned from the call to the mock (*DEFAULT* handling is identical to the function case).

An example of a mock that raises an exception (to test exception handling of an API):

```
>>> mock = Mock()
>>> mock.side_effect = Exception('Boom!')
>>> mock()
Traceback (most recent call last):
...
Exception: Boom!
```

Using *side_effect* to return a sequence of values:

```
>>> mock = Mock()
>>> mock.side_effect = [3, 2, 1]
>>> mock(), mock(), mock()
(3, 2, 1)
```

Using a callable:

```
>>> mock = Mock(return_value=3)
>>> def side_effect(*args, **kwargs):
...     return DEFAULT
...
>>> mock.side_effect = side_effect
>>> mock()
3
```

`side_effect` can be set in the constructor. Here's an example that adds one to the value the mock is called with and returns it:

```
>>> side_effect = lambda value: value + 1
>>> mock = Mock(side_effect=side_effect)
>>> mock(3)
4
>>> mock(-8)
-7
```

Setting `side_effect` to `None` clears it:

```
>>> m = Mock(side_effect=KeyError, return_value=3)
>>> m()
Traceback (most recent call last):
...
KeyError
>>> m.side_effect = None
>>> m()
3
```

call_args

This is either `None` (if the mock hasn't been called), or the arguments that the mock was last called with. This will be in the form of a tuple: the first member is any ordered arguments the mock was called with (or an empty tuple) and the second member is any keyword arguments (or an empty dictionary).

```
>>> mock = Mock(return_value=None)
>>> print(mock.call_args)
None
>>> mock()
>>> mock.call_args
call()
>>> mock.call_args == ()
True
>>> mock(3, 4)
>>> mock.call_args
call(3, 4)
>>> mock.call_args == ((3, 4),)
True
>>> mock(3, 4, 5, key='fish', next='w00t!')
>>> mock.call_args
call(3, 4, 5, key='fish', next='w00t!')
```

`call_args`, along with members of the lists `call_args_list`, `method_calls` and `mock_calls` are *call* objects. These are tuples, so they can be unpacked to get at the individual arguments and make more complex assertions. See *calls as tuples*.

`call_args_list`

This is a list of all the calls made to the mock object in sequence (so the length of the list is the number of times it has been called). Before any calls have been made it is an empty list. The *call* object can be used for conveniently constructing lists of calls to compare with `call_args_list`.

```
>>> mock = Mock(return_value=None)
>>> mock()
>>> mock(3, 4)
>>> mock(key='fish', next='w00t!')
>>> mock.call_args_list
[call(), call(3, 4), call(key='fish', next='w00t!')]
>>> expected = [(), ((3, 4),), ({'key': 'fish', 'next': 'w00t!',})]
>>> mock.call_args_list == expected
True
```

Members of `call_args_list` are *call* objects. These can be unpacked as tuples to get at the individual arguments. See *calls as tuples*.

`method_calls`

As well as tracking calls to themselves, mocks also track calls to methods and attributes, and *their* methods and attributes:

```
>>> mock = Mock()
>>> mock.method()
<Mock name='mock.method()' id='...'>
>>> mock.property.method.attribute()
<Mock name='mock.property.method.attribute()' id='...'>
>>> mock.method_calls
[call.method(), call.property.method.attribute()]
```

Members of `method_calls` are *call* objects. These can be unpacked as tuples to get at the individual arguments. See *calls as tuples*.

`mock_calls`

`mock_calls` records *all* calls to the mock object, its methods, magic methods *and* return value mocks.

```
>>> mock = MagicMock()
>>> result = mock(1, 2, 3)
>>> mock.first(a=3)
<MagicMock name='mock.first()' id='...'>
>>> mock.second()
<MagicMock name='mock.second()' id='...'>
>>> int(mock)
1
>>> result(1)
<MagicMock name='mock()' id='...'>
>>> expected = [call(1, 2, 3), call.first(a=3), call.second(),
... call.__int__(), call()(1)]
>>> mock.mock_calls == expected
True
```

Members of `mock_calls` are *call* objects. These can be unpacked as tuples to get at the individual arguments. See *calls as tuples*.

注解: The way `mock_calls` are recorded means that where nested calls are made, the parameters of ancestor calls are not recorded and so will always compare equal:

```
>>> mock = MagicMock()
>>> mock.top(a=3).bottom()
<MagicMock name='mock.top().bottom()' id='...'>
>>> mock.mock_calls
[call.top(a=3), call.top().bottom()]
>>> mock.mock_calls[-1] == call.top(a=-1).bottom()
True
```

`__class__`

Normally the `__class__` attribute of an object will return its type. For a mock object with a spec, `__class__` returns the spec class instead. This allows mock objects to pass `isinstance()` tests for the object they are replacing / masquerading as:

```
>>> mock = Mock(spec=3)
>>> isinstance(mock, int)
True
```

`__class__` is assignable to, this allows a mock to pass an `isinstance()` check without forcing you to use a spec:

```
>>> mock = Mock()
>>> mock.__class__ = dict
>>> isinstance(mock, dict)
True
```

class `unittest.mock.NonCallableMock` (*spec=None, wraps=None, name=None, spec_set=None, **kwargs*)

A non-callable version of `Mock`. The constructor parameters have the same meaning of `Mock`, with the exception of `return_value` and `side_effect` which have no meaning on a non-callable mock.

Mock objects that use a class or an instance as a spec or spec_set are able to pass `isinstance()` tests:

```
>>> mock = Mock(spec=SomeClass)
>>> isinstance(mock, SomeClass)
True
>>> mock = Mock(spec_set=SomeClass())
>>> isinstance(mock, SomeClass)
True
```

The `Mock` classes have support for mocking magic methods. See [magic methods](#) for the full details.

The mock classes and the `patch()` decorators all take arbitrary keyword arguments for configuration. For the `patch()` decorators the keywords are passed to the constructor of the mock being created. The keyword arguments are for configuring attributes of the mock:

```
>>> m = MagicMock(attribute=3, other='fish')
>>> m.attribute
3
>>> m.other
'fish'
```

The return value and side effect of child mocks can be set in the same way, using dotted notation. As you can't use dotted names directly in a call you have to create a dictionary and unpack it using `**`:

```
>>> attrs = {'method.return_value': 3, 'other.side_effect': KeyError}
>>> mock = Mock(some_attribute='eggs', **attrs)
>>> mock.some_attribute
'eggs'
>>> mock.method()
3
>>> mock.other()
Traceback (most recent call last):
...
KeyError
```

A callable mock which was created with a *spec* (or a *spec_set*) will introspect the specification object's signature when matching calls to the mock. Therefore, it can match the actual call's arguments regardless of whether they were passed positionally or by name:

```
>>> def f(a, b, c): pass
...
>>> mock = Mock(spec=f)
>>> mock(1, 2, c=3)
<Mock name='mock()' id='140161580456576'>
>>> mock.assert_called_with(1, 2, 3)
>>> mock.assert_called_with(a=1, b=2, c=3)
```

This applies to *assert_called_with()*, *assert_called_once_with()*, *assert_has_calls()* and *assert_any_call()*. When *Autospeccing*, it will also apply to method calls on the mock object.

在 3.4 版更改: Added signature introspection on specced and autospecced mock objects.

class `unittest.mock.PropertyMock` (*args, **kwargs)

A mock intended to be used as a property, or other descriptor, on a class. *PropertyMock* provides *__get__()* and *__set__()* methods so you can specify a return value when it is fetched.

Fetching a *PropertyMock* instance from an object calls the mock, with no args. Setting it calls the mock with the value being set.

```
>>> class Foo:
...     @property
...     def foo(self):
...         return 'something'
...     @foo.setter
...     def foo(self, value):
...         pass
...
>>> with patch('__main__.Foo.foo', new_callable=PropertyMock) as mock_foo:
...     mock_foo.return_value = 'mockity-mock'
...     this_foo = Foo()
...     print(this_foo.foo)
...     this_foo.foo = 6
...
mockity-mock
>>> mock_foo.mock_calls
[call(), call(6)]
```

Because of the way mock attributes are stored you can't directly attach a *PropertyMock* to a mock object. Instead you can attach it to the mock type object:

```
>>> m = MagicMock()
>>> p = PropertyMock(return_value=3)
```

(下页继续)

(续上页)

```
>>> type(m).foo = p
>>> m.foo
3
>>> p.assert_called_once_with()
```

Calling

Mock objects are callable. The call will return the value set as the `return_value` attribute. The default return value is a new Mock object; it is created the first time the return value is accessed (either explicitly or by calling the Mock) - but it is stored and the same one returned each time.

Calls made to the object will be recorded in the attributes like `call_args` and `call_args_list`.

If `side_effect` is set then it will be called after the call has been recorded, so if `side_effect` raises an exception the call is still recorded.

The simplest way to make a mock raise an exception when called is to make `side_effect` an exception class or instance:

```
>>> m = MagicMock(side_effect=IndexError)
>>> m(1, 2, 3)
Traceback (most recent call last):
...
IndexError
>>> m.mock_calls
[call(1, 2, 3)]
>>> m.side_effect = KeyError('Bang!')
>>> m('two', 'three', 'four')
Traceback (most recent call last):
...
KeyError: 'Bang!'
>>> m.mock_calls
[call(1, 2, 3), call('two', 'three', 'four')]
```

If `side_effect` is a function then whatever that function returns is what calls to the mock return. The `side_effect` function is called with the same arguments as the mock. This allows you to vary the return value of the call dynamically, based on the input:

```
>>> def side_effect(value):
...     return value + 1
...
>>> m = MagicMock(side_effect=side_effect)
>>> m(1)
2
>>> m(2)
3
>>> m.mock_calls
[call(1), call(2)]
```

If you want the mock to still return the default return value (a new mock), or any set return value, then there are two ways of doing this. Either return `mock.return_value` from inside `side_effect`, or return `DEFAULT`:

```
>>> m = MagicMock()
>>> def side_effect(*args, **kwargs):
...     return m.return_value
```

(下页继续)

(续上页)

```

...
>>> m.side_effect = side_effect
>>> m.return_value = 3
>>> m()
3
>>> def side_effect(*args, **kwargs):
...     return DEFAULT
...
>>> m.side_effect = side_effect
>>> m()
3

```

To remove a `side_effect`, and return to the default behaviour, set the `side_effect` to `None`:

```

>>> m = MagicMock(return_value=6)
>>> def side_effect(*args, **kwargs):
...     return 3
...
>>> m.side_effect = side_effect
>>> m()
3
>>> m.side_effect = None
>>> m()
6

```

The `side_effect` can also be any iterable object. Repeated calls to the mock will return values from the iterable (until the iterable is exhausted and a *StopIteration* is raised):

```

>>> m = MagicMock(side_effect=[1, 2, 3])
>>> m()
1
>>> m()
2
>>> m()
3
>>> m()
Traceback (most recent call last):
...
StopIteration

```

If any members of the iterable are exceptions they will be raised instead of returned:

```

>>> iterable = (33, ValueError, 66)
>>> m = MagicMock(side_effect=iterable)
>>> m()
33
>>> m()
Traceback (most recent call last):
...
ValueError
>>> m()
66

```

Deleting Attributes

Mock objects create attributes on demand. This allows them to pretend to be objects of any type.

You may want a mock object to return `False` to a `hasattr()` call, or raise an `AttributeError` when an attribute is fetched. You can do this by providing an object as a spec for a mock, but that isn't always convenient.

You “block” attributes by deleting them. Once deleted, accessing an attribute will raise an `AttributeError`.

```
>>> mock = MagicMock()
>>> hasattr(mock, 'm')
True
>>> del mock.m
>>> hasattr(mock, 'm')
False
>>> del mock.f
>>> mock.f
Traceback (most recent call last):
...
AttributeError: f
```

Mock names and the name attribute

Since “name” is an argument to the `Mock` constructor, if you want your mock object to have a “name” attribute you can't just pass it in at creation time. There are two alternatives. One option is to use `configure_mock()`:

```
>>> mock = MagicMock()
>>> mock.configure_mock(name='my_name')
>>> mock.name
'my_name'
```

A simpler option is to simply set the “name” attribute after mock creation:

```
>>> mock = MagicMock()
>>> mock.name = "foo"
```

Attaching Mocks as Attributes

When you attach a mock as an attribute of another mock (or as the return value) it becomes a “child” of that mock. Calls to the child are recorded in the `method_calls` and `mock_calls` attributes of the parent. This is useful for configuring child mocks and then attaching them to the parent, or for attaching mocks to a parent that records all calls to the children and allows you to make assertions about the order of calls between mocks:

```
>>> parent = MagicMock()
>>> child1 = MagicMock(return_value=None)
>>> child2 = MagicMock(return_value=None)
>>> parent.child1 = child1
>>> parent.child2 = child2
>>> child1(1)
>>> child2(2)
>>> parent.mock_calls
[call.child1(1), call.child2(2)]
```

The exception to this is if the mock has a name. This allows you to prevent the “parenting” if for some reason you don't want it to happen.

```
>>> mock = MagicMock()
>>> not_a_child = MagicMock(name='not-a-child')
```

(下页继续)

(续上页)

```
>>> mock.attribute = not_a_child
>>> mock.attribute()
<MagicMock name='not-a-child()' id='...'>
>>> mock.mock_calls
[]
```

Mocks created for you by `patch()` are automatically given names. To attach mocks that have names to a parent you use the `attach_mock()` method:

```
>>> thing1 = object()
>>> thing2 = object()
>>> parent = MagicMock()
>>> with patch('__main__.thing1', return_value=None) as child1:
...     with patch('__main__.thing2', return_value=None) as child2:
...         parent.attach_mock(child1, 'child1')
...         parent.attach_mock(child2, 'child2')
...         child1('one')
...         child2('two')
...
>>> parent.mock_calls
[call.child1('one'), call.child2('two')]
```

26.5.3 The patchers

The patch decorators are used for patching objects only within the scope of the function they decorate. They automatically handle the unpatching for you, even if exceptions are raised. All of these functions can also be used in with statements or as class decorators.

patch

注解: `patch()` is straightforward to use. The key is to do the patching in the right namespace. See the section [where to patch](#).

`unittest.mock.patch(target, new=DEFAULT, spec=None, create=False, spec_set=None, autospec=None, new_callable=None, **kwargs)`

`patch()` acts as a function decorator, class decorator or a context manager. Inside the body of the function or with statement, the `target` is patched with a `new` object. When the function/with statement exits the patch is undone.

If `new` is omitted, then the target is replaced with a `MagicMock`. If `patch()` is used as a decorator and `new` is omitted, the created mock is passed in as an extra argument to the decorated function. If `patch()` is used as a context manager the created mock is returned by the context manager.

`target` should be a string in the form `'package.module.ClassName'`. The `target` is imported and the specified object replaced with the `new` object, so the `target` must be importable from the environment you are calling `patch()` from. The target is imported when the decorated function is executed, not at decoration time.

The `spec` and `spec_set` keyword arguments are passed to the `MagicMock` if patch is creating one for you.

In addition you can pass `spec=True` or `spec_set=True`, which causes patch to pass in the object being mocked as the `spec/spec_set` object.

`new_callable` allows you to specify a different class, or callable object, that will be called to create the `new` object. By default `MagicMock` is used.

A more powerful form of *spec* is *autospec*. If you set `autospec=True` then the mock will be created with a spec from the object being replaced. All attributes of the mock will also have the spec of the corresponding attribute of the object being replaced. Methods and functions being mocked will have their arguments checked and will raise a *TypeError* if they are called with the wrong signature. For mocks replacing a class, their return value (the 'instance') will have the same spec as the class. See the `create_autospec()` function and *Autospeccing*.

Instead of `autospec=True` you can pass `autospec=some_object` to use an arbitrary object as the spec instead of the one being replaced.

By default `patch()` will fail to replace attributes that don't exist. If you pass in `create=True`, and the attribute doesn't exist, patch will create the attribute for you when the patched function is called, and delete it again afterwards. This is useful for writing tests against attributes that your production code creates at runtime. It is off by default because it can be dangerous. With it switched on you can write passing tests against APIs that don't actually exist!

注解： 在 3.5 版更改: If you are patching builtins in a module then you don't need to pass `create=True`, it will be added by default.

Patch can be used as a `TestCase` class decorator. It works by decorating each test method in the class. This reduces the boilerplate code when your test methods share a common patchings set. `patch()` finds tests by looking for method names that start with `patch.TEST_PREFIX`. By default this is 'test', which matches the way *unittest* finds tests. You can specify an alternative prefix by setting `patch.TEST_PREFIX`.

Patch can be used as a context manager, with the `with` statement. Here the patching applies to the indented block after the `with` statement. If you use "as" then the patched object will be bound to the name after the "as"; very useful if `patch()` is creating a mock object for you.

`patch()` takes arbitrary keyword arguments. These will be passed to the *Mock* (or *new_callable*) on construction.

`patch.dict(...)`, `patch.multiple(...)` and `patch.object(...)` are available for alternate use-cases.

`patch()` as function decorator, creating the mock for you and passing it into the decorated function:

```
>>> @patch('__main__.SomeClass')
... def function(normal_argument, mock_class):
...     print(mock_class is SomeClass)
...
>>> function(None)
True
```

Patching a class replaces the class with a *MagicMock* instance. If the class is instantiated in the code under test then it will be the *return_value* of the mock that will be used.

If the class is instantiated multiple times you could use *side_effect* to return a new mock each time. Alternatively you can set the *return_value* to be anything you want.

To configure return values on methods of *instances* on the patched class you must do this on the *return_value*. For example:

```
>>> class Class:
...     def method(self):
...         pass
...
>>> with patch('__main__.Class') as MockClass:
...     instance = MockClass.return_value
...     instance.method.return_value = 'foo'
...     assert Class() is instance
```

(下页继续)

(续上页)

```
...     assert Class().method() == 'foo'
...
```

If you use *spec* or *spec_set* and *patch()* is replacing a *class*, then the return value of the created mock will have the same spec.

```
>>> Original = Class
>>> patcher = patch('__main__.Class', spec=True)
>>> MockClass = patcher.start()
>>> instance = MockClass()
>>> assert isinstance(instance, Original)
>>> patcher.stop()
```

The *new_callable* argument is useful where you want to use an alternative class to the default *MagicMock* for the created mock. For example, if you wanted a *NonCallableMock* to be used:

```
>>> thing = object()
>>> with patch('__main__.thing', new_callable=NonCallableMock) as mock_thing:
...     assert thing is mock_thing
...     thing()
...
Traceback (most recent call last):
...
TypeError: 'NonCallableMock' object is not callable
```

Another use case might be to replace an object with an *io.StringIO* instance:

```
>>> from io import StringIO
>>> def foo():
...     print('Something')
...
>>> @patch('sys.stdout', new_callable=StringIO)
... def test(mock_stdout):
...     foo()
...     assert mock_stdout.getvalue() == 'Something\n'
...
>>> test()
```

When *patch()* is creating a mock for you, it is common that the first thing you need to do is to configure the mock. Some of that configuration can be done in the call to *patch*. Any arbitrary keywords you pass into the call will be used to set attributes on the created mock:

```
>>> patcher = patch('__main__.thing', first='one', second='two')
>>> mock_thing = patcher.start()
>>> mock_thing.first
'one'
>>> mock_thing.second
'two'
```

As well as attributes on the created mock attributes, like the *return_value* and *side_effect*, of child mocks can also be configured. These aren't syntactically valid to pass in directly as keyword arguments, but a dictionary with these as keys can still be expanded into a *patch()* call using ****:

```
>>> config = {'method.return_value': 3, 'other.side_effect': KeyError}
>>> patcher = patch('__main__.thing', **config)
>>> mock_thing = patcher.start()
```

(下页继续)

(续上页)

```
>>> mock_thing.method()
3
>>> mock_thing.other()
Traceback (most recent call last):
...
KeyError
```

patch.object

`patch.object(target, attribute, new=DEFAULT, spec=None, create=False, spec_set=None, autospec=None, new_callable=None, **kwargs)`

patch the named member (*attribute*) on an object (*target*) with a mock object.

`patch.object()` can be used as a decorator, class decorator or a context manager. Arguments *new*, *spec*, *create*, *spec_set*, *autospec* and *new_callable* have the same meaning as for `patch()`. Like `patch()`, `patch.object()` takes arbitrary keyword arguments for configuring the mock object it creates.

When used as a class decorator `patch.object()` honours `patch.TEST_PREFIX` for choosing which methods to wrap.

You can either call `patch.object()` with three arguments or two arguments. The three argument form takes the object to be patched, the attribute name and the object to replace the attribute with.

When calling with the two argument form you omit the replacement object, and a mock is created for you and passed in as an extra argument to the decorated function:

```
>>> @patch.object(SomeClass, 'class_method')
... def test(mock_method):
...     SomeClass.class_method(3)
...     mock_method.assert_called_with(3)
...
>>> test()
```

spec, *create* and the other arguments to `patch.object()` have the same meaning as they do for `patch()`.

patch.dict

`patch.dict(in_dict, values=(), clear=False, **kwargs)`

Patch a dictionary, or dictionary like object, and restore the dictionary to its original state after the test.

in_dict can be a dictionary or a mapping like container. If it is a mapping then it must at least support getting, setting and deleting items plus iterating over keys.

in_dict can also be a string specifying the name of the dictionary, which will then be fetched by importing it.

values can be a dictionary of values to set in the dictionary. *values* can also be an iterable of (*key*, *value*) pairs.

If *clear* is true then the dictionary will be cleared before the new values are set.

`patch.dict()` can also be called with arbitrary keyword arguments to set values in the dictionary.

`patch.dict()` can be used as a context manager, decorator or class decorator. When used as a class decorator `patch.dict()` honours `patch.TEST_PREFIX` for choosing which methods to wrap.

`patch.dict()` can be used to add members to a dictionary, or simply let a test change a dictionary, and ensure the dictionary is restored when the test ends.

```
>>> foo = {}
>>> with patch.dict(foo, {'newkey': 'newvalue'}):
...     assert foo == {'newkey': 'newvalue'}
...
>>> assert foo == {}
```

```
>>> import os
>>> with patch.dict('os.environ', {'newkey': 'newvalue'}):
...     print(os.environ['newkey'])
...
newvalue
>>> assert 'newkey' not in os.environ
```

Keywords can be used in the `patch.dict()` call to set values in the dictionary:

```
>>> mymodule = MagicMock()
>>> mymodule.function.return_value = 'fish'
>>> with patch.dict('sys.modules', mymodule=mymodule):
...     import mymodule
...     mymodule.function('some', 'args')
...
'fish'
```

`patch.dict()` can be used with dictionary like objects that aren't actually dictionaries. At the very minimum they must support item getting, setting, deleting and either iteration or membership test. This corresponds to the magic methods `__getitem__()`, `__setitem__()`, `__delitem__()` and either `__iter__()` or `__contains__()`.

```
>>> class Container:
...     def __init__(self):
...         self.values = {}
...     def __getitem__(self, name):
...         return self.values[name]
...     def __setitem__(self, name, value):
...         self.values[name] = value
...     def __delitem__(self, name):
...         del self.values[name]
...     def __iter__(self):
...         return iter(self.values)
...
>>> thing = Container()
>>> thing['one'] = 1
>>> with patch.dict(thing, one=2, two=3):
...     assert thing['one'] == 2
...     assert thing['two'] == 3
...
>>> assert thing['one'] == 1
>>> assert list(thing) == ['one']
```

patch.multiple

`patch.multiple(target, spec=None, create=False, spec_set=None, autospec=None, new_callable=None, **kwargs)`

Perform multiple patches in a single call. It takes the object to be patched (either as an object or a string to fetch the object by importing) and keyword arguments for the patches:

```
with patch.multiple(settings, FIRST_PATCH='one', SECOND_PATCH='two'):
    ...
```

Use `DEFAULT` as the value if you want `patch.multiple()` to create mocks for you. In this case the created mocks are passed into a decorated function by keyword, and a dictionary is returned when `patch.multiple()` is used as a context manager.

`patch.multiple()` can be used as a decorator, class decorator or a context manager. The arguments `spec`, `spec_set`, `create`, `autospec` and `new_callable` have the same meaning as for `patch()`. These arguments will be applied to *all* patches done by `patch.multiple()`.

When used as a class decorator `patch.multiple()` honours `patch.TEST_PREFIX` for choosing which methods to wrap.

If you want `patch.multiple()` to create mocks for you, then you can use `DEFAULT` as the value. If you use `patch.multiple()` as a decorator then the created mocks are passed into the decorated function by keyword.

```
>>> thing = object()
>>> other = object()
```

```
>>> @patch.multiple('__main__', thing=DEFAULT, other=DEFAULT)
... def test_function(thing, other):
...     assert isinstance(thing, MagicMock)
...     assert isinstance(other, MagicMock)
...
>>> test_function()
```

`patch.multiple()` can be nested with other patch decorators, but put arguments passed by keyword *after* any of the standard arguments created by `patch()`:

```
>>> @patch('sys.exit')
... @patch.multiple('__main__', thing=DEFAULT, other=DEFAULT)
... def test_function(mock_exit, other, thing):
...     assert 'other' in repr(other)
...     assert 'thing' in repr(thing)
...     assert 'exit' in repr(mock_exit)
...
>>> test_function()
```

If `patch.multiple()` is used as a context manager, the value returned by the context manager is a dictionary where created mocks are keyed by name:

```
>>> with patch.multiple('__main__', thing=DEFAULT, other=DEFAULT) as values:
...     assert 'other' in repr(values['other'])
...     assert 'thing' in repr(values['thing'])
...     assert values['thing'] is thing
...     assert values['other'] is other
...
>>>
```

patch methods: start and stop

All the patchers have `start()` and `stop()` methods. These make it simpler to do patching in `setUp` methods or where you want to do multiple patches without nesting decorators or with statements.

To use them call `patch()`, `patch.object()` or `patch.dict()` as normal and keep a reference to the returned patcher object. You can then call `start()` to put the patch in place and `stop()` to undo it.

If you are using `patch()` to create a mock for you then it will be returned by the call to `patcher.start`.

```
>>> patcher = patch('package.module.ClassName')
>>> from package import module
>>> original = module.ClassName
>>> new_mock = patcher.start()
>>> assert module.ClassName is not original
>>> assert module.ClassName is new_mock
>>> patcher.stop()
>>> assert module.ClassName is original
>>> assert module.ClassName is not new_mock
```

A typical use case for this might be for doing multiple patches in the `setUp` method of a `TestCase`:

```
>>> class MyTest(TestCase):
...     def setUp(self):
...         self.patcher1 = patch('package.module.Class1')
...         self.patcher2 = patch('package.module.Class2')
...         self.MockClass1 = self.patcher1.start()
...         self.MockClass2 = self.patcher2.start()
...
...     def tearDown(self):
...         self.patcher1.stop()
...         self.patcher2.stop()
...
...     def test_something(self):
...         assert package.module.Class1 is self.MockClass1
...         assert package.module.Class2 is self.MockClass2
...
>>> MyTest('test_something').run()
```

警告: If you use this technique you must ensure that the patching is “undone” by calling `stop`. This can be fiddlier than you might think, because if an exception is raised in the `setUp` then `tearDown` is not called. `unittest.TestCase.addCleanup()` makes this easier:

```
>>> class MyTest(TestCase):
...     def setUp(self):
...         patcher = patch('package.module.Class')
...         self.MockClass = patcher.start()
...         self.addCleanup(patcher.stop)
...
...     def test_something(self):
...         assert package.module.Class is self.MockClass
...
... 
```

As an added bonus you no longer need to keep a reference to the `patcher` object.

It is also possible to stop all patches which have been started by using `patch.stopall()`.

`patch.stopall()`

Stop all active patches. Only stops patches started with `start`.

patch builtins

You can patch any builtins within a module. The following example patches builtin `ord()`:

```
>>> @patch('__main__.ord')
... def test(mock_ord):
...     mock_ord.return_value = 101
...     print(ord('c'))
...
>>> test()
101
```

TEST_PREFIX

All of the patchers can be used as class decorators. When used in this way they wrap every test method on the class. The patchers recognise methods that start with 'test' as being test methods. This is the same way that the *unittest.TestLoader* finds test methods by default.

It is possible that you want to use a different prefix for your tests. You can inform the patchers of the different prefix by setting `patch.TEST_PREFIX`:

```
>>> patch.TEST_PREFIX = 'foo'
>>> value = 3
>>>
>>> @patch('__main__.value', 'not three')
... class Thing:
...     def foo_one(self):
...         print(value)
...     def foo_two(self):
...         print(value)
...
>>>
>>> Thing().foo_one()
not three
>>> Thing().foo_two()
not three
>>> value
3
```

Nesting Patch Decorators

If you want to perform multiple patches then you can simply stack up the decorators.

You can stack up multiple patch decorators using this pattern:

```
>>> @patch.object(SomeClass, 'class_method')
... @patch.object(SomeClass, 'static_method')
... def test(mock1, mock2):
...     assert SomeClass.static_method is mock1
...     assert SomeClass.class_method is mock2
...     SomeClass.static_method('foo')
...     SomeClass.class_method('bar')
...     return mock1, mock2
...
>>> mock1, mock2 = test()
>>> mock1.assert_called_once_with('foo')
>>> mock2.assert_called_once_with('bar')
```

Note that the decorators are applied from the bottom upwards. This is the standard way that Python applies decorators. The order of the created mocks passed into your test function matches this order.

Where to patch

`patch()` works by (temporarily) changing the object that a *name* points to with another one. There can be many names pointing to any individual object, so for patching to work you must ensure that you patch the name used by the system under test.

The basic principle is that you patch where an object is *looked up*, which is not necessarily the same place as where it is defined. A couple of examples will help to clarify this.

Imagine we have a project that we want to test with the following structure:

```
a.py
-> Defines SomeClass

b.py
-> from a import SomeClass
-> some_function instantiates SomeClass
```

Now we want to test `some_function` but we want to mock out `SomeClass` using `patch()`. The problem is that when we import module `b`, which we will have to do then it imports `SomeClass` from module `a`. If we use `patch()` to mock out `a.SomeClass` then it will have no effect on our test; module `b` already has a reference to the *real* `SomeClass` and it looks like our patching had no effect.

The key is to patch out `SomeClass` where it is used (or where it is looked up). In this case `some_function` will actually look up `SomeClass` in module `b`, where we have imported it. The patching should look like:

```
@patch('b.SomeClass')
```

However, consider the alternative scenario where instead of `from a import SomeClass` module `b` does `import a` and `some_function` uses `a.SomeClass`. Both of these import forms are common. In this case the class we want to patch is being looked up in the module and so we have to patch `a.SomeClass` instead:

```
@patch('a.SomeClass')
```

Patching Descriptors and Proxy Objects

Both `patch` and `patch.object` correctly patch and restore descriptors: class methods, static methods and properties. You should patch these on the *class* rather than an instance. They also work with *some* objects that proxy attribute access, like the `django settings` object.

26.5.4 MagicMock and magic method support

Mocking Magic Methods

`Mock` supports mocking the Python protocol methods, also known as “magic methods”. This allows mock objects to replace containers or other objects that implement Python protocols.

Because magic methods are looked up differently from normal methods², this support has been specially implemented. This means that only specific magic methods are supported. The supported list includes *almost* all of them. If there are

² Magic methods *should* be looked up on the class rather than the instance. Different versions of Python are inconsistent about applying this rule. The supported protocol methods should work with all supported versions of Python.

any missing that you need please let us know.

You mock magic methods by setting the method you are interested in to a function or a mock instance. If you are using a function then it *must* take `self` as the first argument³.

```
>>> def __str__(self):
...     return 'fooble'
...
>>> mock = Mock()
>>> mock.__str__ = __str__
>>> str(mock)
'fooble'
```

```
>>> mock = Mock()
>>> mock.__str__ = Mock()
>>> mock.__str__.return_value = 'fooble'
>>> str(mock)
'fooble'
```

```
>>> mock = Mock()
>>> mock.__iter__ = Mock(return_value=iter([]))
>>> list(mock)
[]
```

One use case for this is for mocking objects used as context managers in a `with` statement:

```
>>> mock = Mock()
>>> mock.__enter__ = Mock(return_value='foo')
>>> mock.__exit__ = Mock(return_value=False)
>>> with mock as m:
...     assert m == 'foo'
...
>>> mock.__enter__.assert_called_with()
>>> mock.__exit__.assert_called_with(None, None, None)
```

Calls to magic methods do not appear in `method_calls`, but they are recorded in `mock_calls`.

注解: If you use the `spec` keyword argument to create a mock then attempting to set a magic method that isn't in the spec will raise an `AttributeError`.

The full list of supported magic methods is:

- `__hash__`, `__sizeof__`, `__repr__` and `__str__`
- `__dir__`, `__format__` and `__subclasses__`
- `__floor__`, `__trunc__` and `__ceil__`
- Comparisons: `__lt__`, `__gt__`, `__le__`, `__ge__`, `__eq__` and `__ne__`
- Container methods: `__getitem__`, `__setitem__`, `__delitem__`, `__contains__`, `__len__`, `__iter__`, `__reversed__` and `__missing__`
- Context manager: `__enter__` and `__exit__`
- Unary numeric methods: `__neg__`, `__pos__` and `__invert__`

³ The function is basically hooked up to the class, but each `Mock` instance is kept isolated from the others.

- The numeric methods (including right hand and in-place variants): `__add__`, `__sub__`, `__mul__`, `__matmul__`, `__div__`, `__truediv__`, `__floordiv__`, `__mod__`, `__divmod__`, `__lshift__`, `__rshift__`, `__and__`, `__xor__`, `__or__`, and `__pow__`
- Numeric conversion methods: `__complex__`, `__int__`, `__float__` and `__index__`
- Descriptor methods: `__get__`, `__set__` and `__delete__`
- Pickling: `__reduce__`, `__reduce_ex__`, `__getinitargs__`, `__getnewargs__`, `__getstate__` and `__setstate__`

The following methods exist but are *not* supported as they are either in use by mock, can't be set dynamically, or can cause problems:

- `__getattr__`, `__setattr__`, `__init__` and `__new__`
- `__prepare__`, `__instancecheck__`, `__subclasscheck__`, `__del__`

Magic Mock

There are two `MagicMock` variants: `MagicMock` and `NonCallableMagicMock`.

class `unittest.mock.MagicMock(*args, **kw)`

`MagicMock` is a subclass of `Mock` with default implementations of most of the magic methods. You can use `MagicMock` without having to configure the magic methods yourself.

The constructor parameters have the same meaning as for `Mock`.

If you use the `spec` or `spec_set` arguments then *only* magic methods that exist in the spec will be created.

class `unittest.mock.NonCallableMagicMock(*args, **kw)`

A non-callable version of `MagicMock`.

The constructor parameters have the same meaning as for `MagicMock`, with the exception of `return_value` and `side_effect` which have no meaning on a non-callable mock.

The magic methods are setup with `MagicMock` objects, so you can configure them and use them in the usual way:

```
>>> mock = MagicMock()
>>> mock[3] = 'fish'
>>> mock.__setitem__.assert_called_with(3, 'fish')
>>> mock.__getitem__.return_value = 'result'
>>> mock[2]
'result'
```

By default many of the protocol methods are required to return objects of a specific type. These methods are preconfigured with a default return value, so that they can be used without you having to do anything if you aren't interested in the return value. You can still *set* the return value manually if you want to change the default.

Methods and their defaults:

- `__lt__`: `NotImplemented`
- `__gt__`: `NotImplemented`
- `__le__`: `NotImplemented`
- `__ge__`: `NotImplemented`
- `__int__`: `1`
- `__contains__`: `False`
- `__len__`: `0`

- `__iter__`: `iter([])`
- `__exit__`: `False`
- `__complex__`: `1j`
- `__float__`: `1.0`
- `__bool__`: `True`
- `__index__`: `1`
- `__hash__`: default hash for the mock
- `__str__`: default str for the mock
- `__sizeof__`: default sizeof for the mock

例如:

```
>>> mock = MagicMock()
>>> int(mock)
1
>>> len(mock)
0
>>> list(mock)
[]
>>> object() in mock
False
```

The two equality methods, `__eq__()` and `__ne__()`, are special. They do the default equality comparison on identity, using the *side_effect* attribute, unless you change their return value to return something else:

```
>>> MagicMock() == 3
False
>>> MagicMock() != 3
True
>>> mock = MagicMock()
>>> mock.__eq__.return_value = True
>>> mock == 3
True
```

The return value of `MagicMock.__iter__()` can be any iterable object and isn't required to be an iterator:

```
>>> mock = MagicMock()
>>> mock.__iter__.return_value = ['a', 'b', 'c']
>>> list(mock)
['a', 'b', 'c']
>>> list(mock)
['a', 'b', 'c']
```

If the return value is an iterator, then iterating over it once will consume it and subsequent iterations will result in an empty list:

```
>>> mock.__iter__.return_value = iter(['a', 'b', 'c'])
>>> list(mock)
['a', 'b', 'c']
>>> list(mock)
[]
```

`MagicMock` has all of the supported magic methods configured except for some of the obscure and obsolete ones. You can still set these up if you want.

Magic methods that are supported but not setup by default in `MagicMock` are:

- `__subclasses__`
- `__dir__`
- `__format__`
- `__get__`, `__set__` and `__delete__`
- `__reversed__` and `__missing__`
- `__reduce__`, `__reduce_ex__`, `__getinitargs__`, `__getnewargs__`, `__getstate__` and `__setstate__`
- `__getformat__` and `__setformat__`

26.5.5 Helpers

sentinel

`unittest.mock.sentinel`

The `sentinel` object provides a convenient way of providing unique objects for your tests.

Attributes are created on demand when you access them by name. Accessing the same attribute will always return the same object. The objects returned have a sensible `repr` so that test failure messages are readable.

The `sentinel` attributes don't preserve their identity when they are *copied* or *pickled*.

Sometimes when testing you need to test that a specific object is passed as an argument to another method, or returned. It can be common to create named sentinel objects to test this. `sentinel` provides a convenient way of creating and testing the identity of objects like this.

In this example we monkey patch method to return `sentinel.some_object`:

```
>>> real = ProductionClass()
>>> real.method = Mock(name="method")
>>> real.method.return_value = sentinel.some_object
>>> result = real.method()
>>> assert result is sentinel.some_object
>>> sentinel.some_object
sentinel.some_object
```

DEFAULT

`unittest.mock.DEFAULT`

The `DEFAULT` object is a pre-created sentinel (actually `sentinel.DEFAULT`). It can be used by *side_effect* functions to indicate that the normal return value should be used.

call

`unittest.mock.call(*args, **kwargs)`

`call()` is a helper object for making simpler assertions, for comparing with `call_args`, `call_args_list`, `mock_calls` and `method_calls`. `call()` can also be used with `assert_has_calls()`.

```
>>> m = MagicMock(return_value=None)
>>> m(1, 2, a='foo', b='bar')
>>> m()
>>> m.call_args_list == [call(1, 2, a='foo', b='bar'), call()]
True
```

`call.call_list()`

For a call object that represents multiple calls, `call_list()` returns a list of all the intermediate calls as well as the final call.

`call_list` is particularly useful for making assertions on “chained calls”. A chained call is multiple calls on a single line of code. This results in multiple entries in `mock_calls` on a mock. Manually constructing the sequence of calls can be tedious.

`call_list()` can construct the sequence of calls from the same chained call:

```
>>> m = MagicMock()
>>> m(1).method(arg='foo').other('bar')(2.0)
<MagicMock name='mock().method().other()' id='...'>
>>> kall = call(1).method(arg='foo').other('bar')(2.0)
>>> kall.call_list()
[call(1),
 call().method(arg='foo'),
 call().method().other('bar'),
 call().method().other()(2.0)]
>>> m.mock_calls == kall.call_list()
True
```

A call object is either a tuple of (positional args, keyword args) or (name, positional args, keyword args) depending on how it was constructed. When you construct them yourself this isn’t particularly interesting, but the call objects that are in the `Mock.call_args`, `Mock.call_args_list` and `Mock.mock_calls` attributes can be introspected to get at the individual arguments they contain.

The call objects in `Mock.call_args` and `Mock.call_args_list` are two-tuples of (positional args, keyword args) whereas the call objects in `Mock.mock_calls`, along with ones you construct yourself, are three-tuples of (name, positional args, keyword args).

You can use their “tupleness” to pull out the individual arguments for more complex introspection and assertions. The positional arguments are a tuple (an empty tuple if there are no positional arguments) and the keyword arguments are a dictionary:

```
>>> m = MagicMock(return_value=None)
>>> m(1, 2, 3, arg='one', arg2='two')
>>> kall = m.call_args
>>> args, kwargs = kall
>>> args
(1, 2, 3)
>>> kwargs
{'arg2': 'two', 'arg': 'one'}
>>> args is kall[0]
True
>>> kwargs is kall[1]
True
```

```
>>> m = MagicMock()
>>> m.foo(4, 5, 6, arg='two', arg2='three')
<MagicMock name='mock.foo()' id='...'>
```

(下页继续)

(续上页)

```
>>> kall = m.mock_calls[0]
>>> name, args, kwargs = kall
>>> name
'foo'
>>> args
(4, 5, 6)
>>> kwargs
{'arg2': 'three', 'arg': 'two'}
>>> name is m.mock_calls[0][0]
True
```

create_autospec

`unittest.mock.create_autospec(spec, spec_set=False, instance=False, **kwargs)`

Create a mock object using another object as a spec. Attributes on the mock will use the corresponding attribute on the *spec* object as their spec.

Functions or methods being mocked will have their arguments checked to ensure that they are called with the correct signature.

If *spec_set* is `True` then attempting to set attributes that don't exist on the spec object will raise an *AttributeError*.

If a class is used as a spec then the return value of the mock (the instance of the class) will have the same spec. You can use a class as the spec for an instance object by passing *instance=True*. The returned mock will only be callable if instances of the mock are callable.

create_autospec() also takes arbitrary keyword arguments that are passed to the constructor of the created mock.

See *Autospeccing* for examples of how to use auto-speccing with *create_autospec()* and the *autospec* argument to *patch()*.

ANY

`unittest.mock.ANY`

Sometimes you may need to make assertions about *some* of the arguments in a call to mock, but either not care about some of the arguments or want to pull them individually out of *call_args* and make more complex assertions on them.

To ignore certain arguments you can pass in objects that compare equal to *everything*. Calls to *assert_called_with()* and *assert_called_once_with()* will then succeed no matter what was passed in.

```
>>> mock = Mock(return_value=None)
>>> mock('foo', bar=object())
>>> mock.assert_called_once_with('foo', bar=ANY)
```

ANY can also be used in comparisons with call lists like *mock_calls*:

```
>>> m = MagicMock(return_value=None)
>>> m(1)
>>> m(1, 2)
>>> m(object())
>>> m.mock_calls == [call(1), call(1, 2), ANY]
True
```

FILTER_DIR

`unittest.mock.FILTER_DIR`

`FILTER_DIR` is a module level variable that controls the way mock objects respond to `dir()` (only for Python 2.6 or more recent). The default is `True`, which uses the filtering described below, to only show useful members. If you dislike this filtering, or need to switch it off for diagnostic purposes, then set `mock.FILTER_DIR = False`.

With filtering on, `dir(some_mock)` shows only useful attributes and will include any dynamically created attributes that wouldn't normally be shown. If the mock was created with a *spec* (or *autospec* of course) then all the attributes from the original are shown, even if they haven't been accessed yet:

```
>>> dir(Mock())
['assert_any_call',
 'assert_called_once_with',
 'assert_called_with',
 'assert_has_calls',
 'attach_mock',
 ...
>>> from urllib import request
>>> dir(Mock(spec=request))
['AbstractBasicAuthHandler',
 'AbstractDigestAuthHandler',
 'AbstractHTTPHandler',
 'BaseHandler',
 ...]
```

Many of the not-very-useful (private to *Mock* rather than the thing being mocked) underscore and double underscore prefixed attributes have been filtered from the result of calling `dir()` on a *Mock*. If you dislike this behaviour you can switch it off by setting the module level switch `FILTER_DIR`:

```
>>> from unittest import mock
>>> mock.FILTER_DIR = False
>>> dir(mock.Mock())
['_NonCallableMock__get_return_value',
 '_NonCallableMock__get_side_effect',
 '_NonCallableMock__return_value_doc',
 '_NonCallableMock__set_return_value',
 '_NonCallableMock__set_side_effect',
 '__call__',
 '__class__',
 ...]
```

Alternatively you can just use `vars(my_mock)` (instance members) and `dir(type(my_mock))` (type members) to bypass the filtering irrespective of `mock.FILTER_DIR`.

mock_open

`unittest.mock.mock_open(mock=None, read_data=None)`

A helper function to create a mock to replace the use of `open()`. It works for `open()` called directly or used as a context manager.

The *mock* argument is the mock object to configure. If `None` (the default) then a *MagicMock* will be created for you, with the API limited to methods or attributes available on standard file handles.

read_data is a string for the `read()`, `readline()`, and `readlines()` methods of the file handle to return. Calls to those methods will take data from *read_data* until it is depleted. The mock of these methods is pretty simplistic: every time the *mock* is called, the *read_data* is rewound to the start. If you need more control over

the data that you are feeding to the tested code you will need to customize this mock for yourself. When that is insufficient, one of the in-memory filesystem packages on [PyPI](#) can offer a realistic filesystem for testing.

在 3.4 版更改: Added `readline()` and `readlines()` support. The mock of `read()` changed to consume `read_data` rather than returning it on each call.

在 3.5 版更改: `read_data` is now reset on each call to the `mock`.

Using `open()` as a context manager is a great way to ensure your file handles are closed properly and is becoming common:

```
with open('/some/path', 'w') as f:
    f.write('something')
```

The issue is that even if you mock out the call to `open()` it is the *returned object* that is used as a context manager (and has `__enter__()` and `__exit__()` called).

Mocking context managers with a `MagicMock` is common enough and fiddly enough that a helper function is useful.

```
>>> m = mock_open()
>>> with patch('__main__.open', m):
...     with open('foo', 'w') as h:
...         h.write('some stuff')
...
>>> m.mock_calls
[call('foo', 'w'),
 call().__enter__(),
 call().write('some stuff'),
 call().__exit__(None, None, None)]
>>> m.assert_called_once_with('foo', 'w')
>>> handle = m()
>>> handle.write.assert_called_once_with('some stuff')
```

And for reading files:

```
>>> with patch('__main__.open', mock_open(read_data='bibble')) as m:
...     with open('foo') as h:
...         result = h.read()
...
>>> m.assert_called_once_with('foo')
>>> assert result == 'bibble'
```

Autospecping

Autospecping is based on the existing `spec` feature of `mock`. It limits the api of mocks to the api of an original object (the `spec`), but it is recursive (implemented lazily) so that attributes of mocks only have the same api as the attributes of the `spec`. In addition mocked functions / methods have the same call signature as the original so they raise a `TypeError` if they are called incorrectly.

Before I explain how auto-specping works, here's why it is needed.

`Mock` is a very powerful and flexible object, but it suffers from two flaws when used to mock out objects from a system under test. One of these flaws is specific to the `Mock` api and the other is a more general problem with using mock objects.

First the problem specific to `Mock`. `Mock` has two assert methods that are extremely handy: `assert_called_with()` and `assert_called_once_with()`.

```
>>> mock = Mock(name='Thing', return_value=None)
>>> mock(1, 2, 3)
>>> mock.assert_called_once_with(1, 2, 3)
>>> mock(1, 2, 3)
>>> mock.assert_called_once_with(1, 2, 3)
Traceback (most recent call last):
...
AssertionError: Expected 'mock' to be called once. Called 2 times.
```

Because mocks auto-create attributes on demand, and allow you to call them with arbitrary arguments, if you misspell one of these assert methods then your assertion is gone:

```
>>> mock = Mock(name='Thing', return_value=None)
>>> mock(1, 2, 3)
>>> mock.assret_called_once_with(4, 5, 6)
```

Your tests can pass silently and incorrectly because of the typo.

The second issue is more general to mocking. If you refactor some of your code, rename members and so on, any tests for code that is still using the *old api* but uses mocks instead of the real objects will still pass. This means your tests can all pass even though your code is broken.

Note that this is another reason why you need integration tests as well as unit tests. Testing everything in isolation is all fine and dandy, but if you don't test how your units are “wired together” there is still lots of room for bugs that tests might have caught.

mock already provides a feature to help with this, called speccing. If you use a class or instance as the *spec* for a mock then you can only access attributes on the mock that exist on the real class:

```
>>> from urllib import request
>>> mock = Mock(spec=request.Request)
>>> mock.assret_called_with
Traceback (most recent call last):
...
AttributeError: Mock object has no attribute 'assret_called_with'
```

The spec only applies to the mock itself, so we still have the same issue with any methods on the mock:

```
>>> mock.has_data()
<mock.Mock object at 0x...>
>>> mock.has_data.assret_called_with()
```

Auto-speccing solves this problem. You can either pass `autospec=True` to `patch()` / `patch.object()` or use the `create_autospec()` function to create a mock with a spec. If you use the `autospec=True` argument to `patch()` then the object that is being replaced will be used as the spec object. Because the speccing is done “lazily” (the spec is created as attributes on the mock are accessed) you can use it with very complex or deeply nested objects (like modules that import modules that import modules) without a big performance hit.

Here's an example of it in use:

```
>>> from urllib import request
>>> patcher = patch('__main__.request', autospec=True)
>>> mock_request = patcher.start()
>>> request is mock_request
True
>>> mock_request.Request
<MagicMock name='request.Request' spec='Request' id='...>
```

You can see that `request.Request` has a spec. `request.Request` takes two arguments in the constructor (one of which is *self*). Here's what happens if we try to call it incorrectly:

```
>>> req = request.Request()
Traceback (most recent call last):
...
TypeError: <lambda>() takes at least 2 arguments (1 given)
```

The spec also applies to instantiated classes (i.e. the return value of specced mocks):

```
>>> req = request.Request('foo')
>>> req
<NonCallableMagicMock name='request.Request()' spec='Request' id='... '>
```

`Request` objects are not callable, so the return value of instantiating our mocked out `request.Request` is a non-callable mock. With the spec in place any typos in our asserts will raise the correct error:

```
>>> req.add_header('spam', 'eggs')
<MagicMock name='request.Request().add_header()' id='... '>
>>> req.add_header.assert_called_with
Traceback (most recent call last):
...
AttributeError: Mock object has no attribute 'assert_called_with'
>>> req.add_header.assert_called_with('spam', 'eggs')
```

In many cases you will just be able to add `autospec=True` to your existing `patch()` calls and then be protected against bugs due to typos and api changes.

As well as using *autospec* through `patch()` there is a `create_autospec()` for creating autospecced mocks directly:

```
>>> from urllib import request
>>> mock_request = create_autospec(request)
>>> mock_request.Request('foo', 'bar')
<NonCallableMagicMock name='mock.Request()' spec='Request' id='... '>
```

This isn't without caveats and limitations however, which is why it is not the default behaviour. In order to know what attributes are available on the spec object, *autospec* has to introspect (access attributes) the spec. As you traverse attributes on the mock a corresponding traversal of the original object is happening under the hood. If any of your specced objects have properties or descriptors that can trigger code execution then you may not be able to use *autospec*. On the other hand it is much better to design your objects so that introspection is safe⁴.

A more serious problem is that it is common for instance attributes to be created in the `__init__()` method and not to exist on the class at all. *autospec* can't know about any dynamically created attributes and restricts the api to visible attributes.

```
>>> class Something:
...     def __init__(self):
...         self.a = 33
...
>>> with patch('__main__.Something', autospec=True):
...     thing = Something()
...     thing.a
...
Traceback (most recent call last):
...
AttributeError: Mock object has no attribute 'a'
```

⁴ This only applies to classes or already instantiated objects. Calling a mocked class to create a mock instance *does not* create a real instance. It is only attribute lookups - along with calls to `dir()` - that are done.

There are a few different ways of resolving this problem. The easiest, but not necessarily the least annoying, way is to simply set the required attributes on the mock after creation. Just because *autospec* doesn't allow you to fetch attributes that don't exist on the spec it doesn't prevent you setting them:

```
>>> with patch('__main__.Something', autospec=True):
...     thing = Something()
...     thing.a = 33
... 
```

There is a more aggressive version of both *spec* and *autospec* that *does* prevent you setting non-existent attributes. This is useful if you want to ensure your code only *sets* valid attributes too, but obviously it prevents this particular scenario:

```
>>> with patch('__main__.Something', autospec=True, spec_set=True):
...     thing = Something()
...     thing.a = 33
...
Traceback (most recent call last):
...
AttributeError: Mock object has no attribute 'a'
```

Probably the best way of solving the problem is to add class attributes as default values for instance members initialised in `__init__()`. Note that if you are only setting default attributes in `__init__()` then providing them via class attributes (shared between instances of course) is faster too. e.g.

```
class Something:
    a = 33
```

This brings up another issue. It is relatively common to provide a default value of `None` for members that will later be an object of a different type. `None` would be useless as a spec because it wouldn't let you access *any* attributes or methods on it. As `None` is *never* going to be useful as a spec, and probably indicates a member that will normally of some other type, *autospec* doesn't use a spec for members that are set to `None`. These will just be ordinary mocks (well - *MagicMocks*):

```
>>> class Something:
...     member = None
...
>>> mock = create_autospec(Something)
>>> mock.member.foo.bar.baz()
<MagicMock name='mock.member.foo.bar.baz()' id='...'>
```

If modifying your production classes to add defaults isn't to your liking then there are more options. One of these is simply to use an instance as the spec rather than the class. The other is to create a subclass of the production class and add the defaults to the subclass without affecting the production class. Both of these require you to use an alternative object as the spec. Thankfully *patch()* supports this - you can simply pass the alternative object as the *autospec* argument:

```
>>> class Something:
...     def __init__(self):
...         self.a = 33
...
>>> class SomethingForTest(Something):
...     a = 33
...
>>> p = patch('__main__.Something', autospec=SomethingForTest)
>>> mock = p.start()
>>> mock.a
<NonCallableMagicMock name='Something.a' spec='int' id='...'>
```

26.6 unittest.mock 上手指南

3.3 新版功能.

26.6.1 使用 mock

模拟方法调用

使用 *Mock* 的常见场景：

- 模拟函数调用
- 记录“对象上的方法调用”

你可能需要替换一个对象上的方法，用于确认此方法被系统中的其他部分调用过，并且调用时使用了正确的参数。

```
>>> real = SomeClass()
>>> real.method = MagicMock(name='method')
>>> real.method(3, 4, 5, key='value')
<MagicMock name='method()' id='...'>
```

使用了 *mock*（本例中的 `real.method`）之后，它有方法和属性可以让你针对它是被如何使用的下断言。

注解：在多数示例中，*Mock* 与 *MagicMock* 两个类可以相互替换，而 *MagicMock* 是一个更适用的类，通常情况下，使用它就可以了。

如果 *mock* 被调用，它的 *called* 属性就会变成 `True`，更重要的是，我们可以使用 `assert_called_with()` 或者 `assert_called_once_with()` 方法来确认它在被调用时使用了正确的参数。

在如下的测试示例中，验证对于 `ProductionClass().method` 的调用会导致 `something` 的调用。

```
>>> class ProductionClass:
...     def method(self):
...         self.something(1, 2, 3)
...     def something(self, a, b, c):
...         pass
...
>>> real = ProductionClass()
>>> real.something = MagicMock()
>>> real.method()
>>> real.something.assert_called_once_with(1, 2, 3)
```

对象上的方法调用的 mock

In the last example we patched a method directly on an object to check that it was called correctly. Another common use case is to pass an object into a method (or some part of the system under test) and then check that it is used in the correct way.

The simple `ProductionClass` below has a `closer` method. If it is called with an object then it calls `close` on it.

```
>>> class ProductionClass:
...     def closer(self, something):
...         something.close()
... 
```

So to test it we need to pass in an object with a `close` method and check that it was called correctly.

```
>>> real = ProductionClass()
>>> mock = Mock()
>>> real.closer(mock)
>>> mock.close.assert_called_with()
```

We don't have to do any work to provide the `'close'` method on our mock. Accessing `close` creates it. So, if `'close'` hasn't already been called then accessing it in the test will create it, but `assert_called_with()` will raise a failure exception.

Mocking Classes

A common use case is to mock out classes instantiated by your code under test. When you patch a class, then that class is replaced with a mock. Instances are created by *calling the class*. This means you access the “mock instance” by looking at the return value of the mocked class.

In the example below we have a function `some_function` that instantiates `Foo` and calls a method on it. The call to `patch()` replaces the class `Foo` with a mock. The `Foo` instance is the result of calling the mock, so it is configured by modifying the mock `return_value`.

```
>>> def some_function():
...     instance = module.Foo()
...     return instance.method()
...
>>> with patch('module.Foo') as mock:
...     instance = mock.return_value
...     instance.method.return_value = 'the result'
...     result = some_function()
...     assert result == 'the result'
```

Naming your mocks

It can be useful to give your mocks a name. The name is shown in the repr of the mock and can be helpful when the mock appears in test failure messages. The name is also propagated to attributes or methods of the mock:

```
>>> mock = MagicMock(name='foo')
>>> mock
<MagicMock name='foo' id='...'>
>>> mock.method
<MagicMock name='foo.method' id='...'>
```

Tracking all Calls

Often you want to track more than a single call to a method. The `mock_calls` attribute records all calls to child attributes of the mock - and also to their children.

```
>>> mock = MagicMock()
>>> mock.method()
<MagicMock name='mock.method()' id='...'>
>>> mock.attribute.method(10, x=53)
<MagicMock name='mock.attribute.method()' id='...'>
>>> mock.mock_calls
[call.method(), call.attribute.method(10, x=53)]
```

If you make an assertion about `mock_calls` and any unexpected methods have been called, then the assertion will fail. This is useful because as well as asserting that the calls you expected have been made, you are also checking that they were made in the right order and with no additional calls:

You use the `call` object to construct lists for comparing with `mock_calls`:

```
>>> expected = [call.method(), call.attribute.method(10, x=53)]
>>> mock.mock_calls == expected
True
```

However, parameters to calls that return mocks are not recorded, which means it is not possible to track nested calls where the parameters used to create ancestors are important:

```
>>> m = Mock()
>>> m.factory(important=True).deliver()
<Mock name='mock.factory().deliver()' id='...'>
>>> m.mock_calls[-1] == call.factory(important=False).deliver()
True
```

Setting Return Values and Attributes

Setting the return values on a mock object is trivially easy:

```
>>> mock = Mock()
>>> mock.return_value = 3
>>> mock()
3
```

Of course you can do the same for methods on the mock:

```
>>> mock = Mock()
>>> mock.method.return_value = 3
>>> mock.method()
3
```

The return value can also be set in the constructor:

```
>>> mock = Mock(return_value=3)
>>> mock()
3
```

If you need an attribute setting on your mock, just do it:

```
>>> mock = Mock()
>>> mock.x = 3
>>> mock.x
3
```

Sometimes you want to mock up a more complex situation, like for example `mock.connection.cursor().execute("SELECT 1")`. If we wanted this call to return a list, then we have to configure the result of the nested call.

We can use `call` to construct the set of calls in a “chained call” like this for easy assertion afterwards:

```
>>> mock = Mock()
>>> cursor = mock.connection.cursor.return_value
>>> cursor.execute.return_value = ['foo']
>>> mock.connection.cursor().execute("SELECT 1")
['foo']
>>> expected = call.connection.cursor().execute("SELECT 1").call_list()
>>> mock.mock_calls
[call.connection.cursor(), call.connection.cursor().execute('SELECT 1')]
>>> mock.mock_calls == expected
True
```

It is the call to `.call_list()` that turns our call object into a list of calls representing the chained calls.

Raising exceptions with mocks

A useful attribute is `side_effect`. If you set this to an exception class or instance then the exception will be raised when the mock is called.

```
>>> mock = Mock(side_effect=Exception('Boom!'))
>>> mock()
Traceback (most recent call last):
...
Exception: Boom!
```

Side effect functions and iterables

`side_effect` can also be set to a function or an iterable. The use case for `side_effect` as an iterable is where your mock is going to be called several times, and you want each call to return a different value. When you set `side_effect` to an iterable every call to the mock returns the next value from the iterable:

```
>>> mock = MagicMock(side_effect=[4, 5, 6])
>>> mock()
4
>>> mock()
5
>>> mock()
6
```

For more advanced use cases, like dynamically varying the return values depending on what the mock is called with, `side_effect` can be a function. The function will be called with the same arguments as the mock. Whatever the function returns is what the call returns:

```
>>> vals = {(1, 2): 1, (2, 3): 2}
>>> def side_effect(*args):
...     return vals[args]
...
>>> mock = MagicMock(side_effect=side_effect)
>>> mock(1, 2)
1
```

(下页继续)

(续上页)

```
>>> mock(2, 3)
2
```

Creating a Mock from an Existing Object

One problem with over use of mocking is that it couples your tests to the implementation of your mocks rather than your real code. Suppose you have a class that implements `some_method`. In a test for another class, you provide a mock of this object that *also* provides `some_method`. If later you refactor the first class, so that it no longer has `some_method` - then your tests will continue to pass even though your code is now broken!

`Mock` allows you to provide an object as a specification for the mock, using the `spec` keyword argument. Accessing methods / attributes on the mock that don't exist on your specification object will immediately raise an attribute error. If you change the implementation of your specification, then tests that use that class will start failing immediately without you having to instantiate the class in those tests.

```
>>> mock = Mock(spec=SomeClass)
>>> mock.old_method()
Traceback (most recent call last):
...
AttributeError: object has no attribute 'old_method'
```

Using a specification also enables a smarter matching of calls made to the mock, regardless of whether some parameters were passed as positional or named arguments:

```
>>> def f(a, b, c): pass
...
>>> mock = Mock(spec=f)
>>> mock(1, 2, 3)
<Mock name='mock()' id='140161580456576'>
>>> mock.assert_called_with(a=1, b=2, c=3)
```

If you want this smarter matching to also work with method calls on the mock, you can use *auto-specing*.

If you want a stronger form of specification that prevents the setting of arbitrary attributes as well as the getting of them then you can use `spec_set` instead of `spec`.

26.6.2 Patch Decorators

注解: With `patch()` it matters that you patch objects in the namespace where they are looked up. This is normally straightforward, but for a quick guide read *where to patch*.

A common need in tests is to patch a class attribute or a module attribute, for example patching a builtin or patching a class in a module to test that it is instantiated. Modules and classes are effectively global, so patching on them has to be undone after the test or the patch will persist into other tests and cause hard to diagnose problems.

`mock` provides three convenient decorators for this: `patch()`, `patch.object()` and `patch.dict()`. `patch` takes a single string, of the form `package.module.Class.attribute` to specify the attribute you are patching. It also optionally takes a value that you want the attribute (or class or whatever) to be replaced with. 'patch.object' takes an object and the name of the attribute you would like patched, plus optionally the value to patch it with.

`patch.object`:

```
>>> original = SomeClass.attribute
>>> @patch.object(SomeClass, 'attribute', sentinel.attribute)
... def test():
...     assert SomeClass.attribute == sentinel.attribute
...
>>> test()
>>> assert SomeClass.attribute == original
```

```
>>> @patch('package.module.attribute', sentinel.attribute)
... def test():
...     from package.module import attribute
...     assert attribute is sentinel.attribute
...
>>> test()
```

If you are patching a module (including *builtins*) then use *patch()* instead of *patch.object()*:

```
>>> mock = MagicMock(return_value=sentinel.file_handle)
>>> with patch('builtins.open', mock):
...     handle = open('filename', 'r')
...
>>> mock.assert_called_with('filename', 'r')
>>> assert handle == sentinel.file_handle, "incorrect file handle returned"
```

The module name can be ‘dotted’, in the form *package.module* if needed:

```
>>> @patch('package.module.ClassName.attribute', sentinel.attribute)
... def test():
...     from package.module import ClassName
...     assert ClassName.attribute == sentinel.attribute
...
>>> test()
```

A nice pattern is to actually decorate test methods themselves:

```
>>> class MyTest(unittest.TestCase):
...     @patch.object(SomeClass, 'attribute', sentinel.attribute)
...     def test_something(self):
...         self.assertEqual(SomeClass.attribute, sentinel.attribute)
...
>>> original = SomeClass.attribute
>>> MyTest('test_something').test_something()
>>> assert SomeClass.attribute == original
```

If you want to patch with a Mock, you can use *patch()* with only one argument (or *patch.object()* with two arguments). The mock will be created for you and passed into the test function / method:

```
>>> class MyTest(unittest.TestCase):
...     @patch.object(SomeClass, 'static_method')
...     def test_something(self, mock_method):
...         SomeClass.static_method()
...         mock_method.assert_called_with()
...
>>> MyTest('test_something').test_something()
```

You can stack up multiple patch decorators using this pattern:

```
>>> class MyTest(unittest.TestCase):
...     @patch('package.module.ClassName1')
...     @patch('package.module.ClassName2')
...     def test_something(self, MockClass2, MockClass1):
...         self.assertIs(package.module.ClassName1, MockClass1)
...         self.assertIs(package.module.ClassName2, MockClass2)
...
>>> MyTest('test_something').test_something()
```

When you nest patch decorators the mocks are passed in to the decorated function in the same order they applied (the normal *python* order that decorators are applied). This means from the bottom up, so in the example above the mock for `test_module.ClassName2` is passed in first.

There is also `patch.dict()` for setting values in a dictionary just during a scope and restoring the dictionary to its original state when the test ends:

```
>>> foo = {'key': 'value'}
>>> original = foo.copy()
>>> with patch.dict(foo, {'newkey': 'newvalue'}, clear=True):
...     assert foo == {'newkey': 'newvalue'}
...
>>> assert foo == original
```

`patch`, `patch.object` and `patch.dict` can all be used as context managers.

Where you use `patch()` to create a mock for you, you can get a reference to the mock using the “as” form of the with statement:

```
>>> class ProductionClass:
...     def method(self):
...         pass
...
>>> with patch.object(ProductionClass, 'method') as mock_method:
...     mock_method.return_value = None
...     real = ProductionClass()
...     real.method(1, 2, 3)
...
>>> mock_method.assert_called_with(1, 2, 3)
```

As an alternative `patch`, `patch.object` and `patch.dict` can be used as class decorators. When used in this way it is the same as applying the decorator individually to every method whose name starts with “test” .

26.6.3 Further Examples

Here are some more examples for some slightly more advanced scenarios.

Mocking chained calls

Mocking chained calls is actually straightforward with mock once you understand the `return_value` attribute. When a mock is called for the first time, or you fetch its `return_value` before it has been called, a new *Mock* is created.

This means that you can see how the object returned from a call to a mocked object has been used by interrogating the `return_value` mock:

```
>>> mock = Mock()
>>> mock().foo(a=2, b=3)
<Mock name='mock().foo()' id='...'>
>>> mock.return_value.foo.assert_called_with(a=2, b=3)
```

From here it is a simple step to configure and then make assertions about chained calls. Of course another alternative is writing your code in a more testable way in the first place...

So, suppose we have some code that looks a little bit like this:

```
>>> class Something:
...     def __init__(self):
...         self.backend = BackendProvider()
...     def method(self):
...         response = self.backend.get_endpoint('foobar').create_call('spam', 'eggs'
↪).start_call()
...         # more code
```

Assuming that `BackendProvider` is already well tested, how do we test `method()`? Specifically, we want to test that the code section `# more code` uses the response object in the correct way.

As this chain of calls is made from an instance attribute we can monkey patch the `backend` attribute on a `Something` instance. In this particular case we are only interested in the return value from the final call to `start_call` so we don't have much configuration to do. Let's assume the object it returns is 'file-like', so we'll ensure that our response object uses the builtin `open()` as its spec.

To do this we create a mock instance as our mock backend and create a mock response object for it. To set the response as the return value for that final `start_call` we could do this:

```
mock_backend.get_endpoint.return_value.create_call.return_value.start_call.return_
↪value = mock_response
```

We can do that in a slightly nicer way using the `configure_mock()` method to directly set the return value for us:

```
>>> something = Something()
>>> mock_response = Mock(spec=open)
>>> mock_backend = Mock()
>>> config = {'get_endpoint.return_value.create_call.return_value.start_call.return_
↪value': mock_response}
>>> mock_backend.configure_mock(**config)
```

With these we monkey patch the "mock backend" in place and can make the real call:

```
>>> something.backend = mock_backend
>>> something.method()
```

Using `mock_calls` we can check the chained call with a single assert. A chained call is several calls in one line of code, so there will be several entries in `mock_calls`. We can use `call.call_list()` to create this list of calls for us:

```
>>> chained = call.get_endpoint('foobar').create_call('spam', 'eggs').start_call()
>>> call_list = chained.call_list()
>>> assert mock_backend.mock_calls == call_list
```

Partial mocking

In some tests I wanted to mock out a call to `datetime.date.today()` to return a known date, but I didn't want to prevent the code under test from creating new date objects. Unfortunately `datetime.date` is written in C, and so I couldn't just monkey-patch out the static `date.today()` method.

I found a simple way of doing this that involved effectively wrapping the date class with a mock, but passing through calls to the constructor to the real class (and returning real instances).

The `patch decorator` is used here to mock out the `date` class in the module under test. The `side_effect` attribute on the mock date class is then set to a lambda function that returns a real date. When the mock date class is called a real date will be constructed and returned by `side_effect`.

```
>>> from datetime import date
>>> with patch('mymodule.date') as mock_date:
...     mock_date.today.return_value = date(2010, 10, 8)
...     mock_date.side_effect = lambda *args, **kw: date(*args, **kw)
...
...     assert mymodule.date.today() == date(2010, 10, 8)
...     assert mymodule.date(2009, 6, 8) == date(2009, 6, 8)
... 
```

Note that we don't patch `datetime.date` globally, we patch `date` in the module that *uses* it. See [where to patch](#).

When `date.today()` is called a known date is returned, but calls to the `date(...)` constructor still return normal dates. Without this you can find yourself having to calculate an expected result using exactly the same algorithm as the code under test, which is a classic testing anti-pattern.

Calls to the date constructor are recorded in the `mock_date` attributes (`call_count` and `friends`) which may also be useful for your tests.

An alternative way of dealing with mocking dates, or other builtin classes, is discussed in [this blog entry](#).

Mocking a Generator Method

A Python generator is a function or method that uses the `yield` statement to return a series of values when iterated over¹.

A generator method / function is called to return the generator object. It is the generator object that is then iterated over. The protocol method for iteration is `__iter__()`, so we can mock this using a `MagicMock`.

Here's an example class with an "iter" method implemented as a generator:

```
>>> class Foo:
...     def iter(self):
...         for i in [1, 2, 3]:
...             yield i
...
>>> foo = Foo()
>>> list(foo.iter())
[1, 2, 3]
```

How would we mock this class, and in particular its "iter" method?

To configure the values returned from the iteration (implicit in the call to `list`), we need to configure the object returned by the call to `foo.iter()`.

¹ There are also generator expressions and more [advanced uses](#) of generators, but we aren't concerned about them here. A very good introduction to generators and how powerful they are is: [Generator Tricks for Systems Programmers](#).

```
>>> mock_foo = MagicMock()
>>> mock_foo.iter.return_value = iter([1, 2, 3])
>>> list(mock_foo.iter())
[1, 2, 3]
```

Applying the same patch to every test method

If you want several patches in place for multiple test methods the obvious way is to apply the patch decorators to every method. This can feel like unnecessary repetition. For Python 2.6 or more recent you can use `patch()` (in all its various forms) as a class decorator. This applies the patches to all test methods on the class. A test method is identified by methods whose names start with `test`:

```
>>> @patch('mymodule.SomeClass')
... class MyTest(TestCase):
...     def test_one(self, MockSomeClass):
...         self.assertIs(mymodule.SomeClass, MockSomeClass)
...     def test_two(self, MockSomeClass):
...         self.assertIs(mymodule.SomeClass, MockSomeClass)
...     def not_a_test(self):
...         return 'something'
...
>>> MyTest('test_one').test_one()
>>> MyTest('test_two').test_two()
>>> MyTest('test_two').not_a_test()
'something'
```

An alternative way of managing patches is to use the *patch methods: start and stop*. These allow you to move the patching into your `setUp` and `tearDown` methods.

```
>>> class MyTest(TestCase):
...     def setUp(self):
...         self.patcher = patch('mymodule.foo')
...         self.mock_foo = self.patcher.start()
...     def test_foo(self):
...         self.assertIs(mymodule.foo, self.mock_foo)
...     def tearDown(self):
...         self.patcher.stop()
...
>>> MyTest('test_foo').run()
```

If you use this technique you must ensure that the patching is “undone” by calling `stop`. This can be fiddlier than you might think, because if an exception is raised in the `setUp` then `tearDown` is not called. `unittest.TestCase.addCleanup()` makes this easier:

```
>>> class MyTest(TestCase):
...     def setUp(self):
...         patcher = patch('mymodule.foo')
...         self.addCleanup(patcher.stop)
...         self.mock_foo = patcher.start()
...     def test_foo(self):
```

(下页继续)

(续上页)

```

...     def test_foo(self):
...         self.assertIs(mymodule.foo, self.mock_foo)
...
>>> MyTest('test_foo').run()

```

Mocking Unbound Methods

Whilst writing tests today I needed to patch an *unbound method* (patching the method on the class rather than on the instance). I needed self to be passed in as the first argument because I want to make asserts about which objects were calling this particular method. The issue is that you can't patch with a mock for this, because if you replace an unbound method with a mock it doesn't become a bound method when fetched from the instance, and so it doesn't get self passed in. The workaround is to patch the unbound method with a real function instead. The `patch()` decorator makes it so simple to patch out methods with a mock that having to create a real function becomes a nuisance.

If you pass `autospec=True` to patch then it does the patching with a *real* function object. This function object has the same signature as the one it is replacing, but delegates to a mock under the hood. You still get your mock auto-created in exactly the same way as before. What it means though, is that if you use it to patch out an unbound method on a class the mocked function will be turned into a bound method if it is fetched from an instance. It will have `self` passed in as the first argument, which is exactly what I wanted:

```

>>> class Foo:
...     def foo(self):
...         pass
...
>>> with patch.object(Foo, 'foo', autospec=True) as mock_foo:
...     mock_foo.return_value = 'foo'
...     foo = Foo()
...     foo.foo()
...
'foo'
>>> mock_foo.assert_called_once_with(foo)

```

If we don't use `autospec=True` then the unbound method is patched out with a Mock instance instead, and isn't called with `self`.

Checking multiple calls with mock

mock has a nice API for making assertions about how your mock objects are used.

```

>>> mock = Mock()
>>> mock.foo_bar.return_value = None
>>> mock.foo_bar('baz', spam='eggs')
>>> mock.foo_bar.assert_called_with('baz', spam='eggs')

```

If your mock is only being called once you can use the `assert_called_once_with()` method that also asserts that the `call_count` is one.

```

>>> mock.foo_bar.assert_called_once_with('baz', spam='eggs')
>>> mock.foo_bar()
>>> mock.foo_bar.assert_called_once_with('baz', spam='eggs')
Traceback (most recent call last):
...
AssertionError: Expected to be called once. Called 2 times.

```

Both `assert_called_with` and `assert_called_once_with` make assertions about the *most recent* call. If your mock is going to be called several times, and you want to make assertions about *all* those calls you can use `call_args_list`:

```
>>> mock = Mock(return_value=None)
>>> mock(1, 2, 3)
>>> mock(4, 5, 6)
>>> mock()
>>> mock.call_args_list
[call(1, 2, 3), call(4, 5, 6), call()]
```

The `call` helper makes it easy to make assertions about these calls. You can build up a list of expected calls and compare it to `call_args_list`. This looks remarkably similar to the repr of the `call_args_list`:

```
>>> expected = [call(1, 2, 3), call(4, 5, 6), call()]
>>> mock.call_args_list == expected
True
```

Coping with mutable arguments

Another situation is rare, but can bite you, is when your mock is called with mutable arguments. `call_args` and `call_args_list` store *references* to the arguments. If the arguments are mutated by the code under test then you can no longer make assertions about what the values were when the mock was called.

Here's some example code that shows the problem. Imagine the following functions defined in 'mymodule' :

```
def frob(val):
    pass

def grob(val):
    "First frob and then clear val"
    frob(val)
    val.clear()
```

When we try to test that `grob` calls `frob` with the correct argument look what happens:

```
>>> with patch('mymodule.frob') as mock_frob:
...     val = {6}
...     mymodule.grob(val)
...
>>> val
set()
>>> mock_frob.assert_called_with({6})
Traceback (most recent call last):
...
AssertionError: Expected: (({6},), {})
Called with: ((set(),), {})
```

One possibility would be for mock to copy the arguments you pass in. This could then cause problems if you do assertions that rely on object identity for equality.

Here's one solution that uses the `side_effect` functionality. If you provide a `side_effect` function for a mock then `side_effect` will be called with the same args as the mock. This gives us an opportunity to copy the arguments and store them for later assertions. In this example I'm using *another* mock to store the arguments so that I can use the mock methods for doing the assertion. Again a helper function sets this up for me.

```

>>> from copy import deepcopy
>>> from unittest.mock import Mock, patch, DEFAULT
>>> def copy_call_args(mock):
...     new_mock = Mock()
...     def side_effect(*args, **kwargs):
...         args = deepcopy(args)
...         kwargs = deepcopy(kwargs)
...         new_mock(*args, **kwargs)
...         return DEFAULT
...     mock.side_effect = side_effect
...     return new_mock
...
>>> with patch('mymodule.frob') as mock_frob:
...     new_mock = copy_call_args(mock_frob)
...     val = {6}
...     mymodule.grob(val)
...
>>> new_mock.assert_called_with({6})
>>> new_mock.call_args
call({6})

```

`copy_call_args` is called with the mock that will be called. It returns a new mock that we do the assertion on. The `side_effect` function makes a copy of the args and calls our `new_mock` with the copy.

注解: If your mock is only going to be used once there is an easier way of checking arguments at the point they are called. You can simply do the checking inside a `side_effect` function.

```

>>> def side_effect(arg):
...     assert arg == {6}
...
>>> mock = Mock(side_effect=side_effect)
>>> mock({6})
>>> mock(set())
Traceback (most recent call last):
...
AssertionError

```

An alternative approach is to create a subclass of `Mock` or `MagicMock` that copies (using `copy.deepcopy()`) the arguments. Here's an example implementation:

```

>>> from copy import deepcopy
>>> class CopyingMock(MagicMock):
...     def __call__(self, *args, **kwargs):
...         args = deepcopy(args)
...         kwargs = deepcopy(kwargs)
...         return super(CopyingMock, self).__call__(*args, **kwargs)
...
>>> c = CopyingMock(return_value=None)
>>> arg = set()
>>> c(arg)
>>> arg.add(1)
>>> c.assert_called_with(set())
>>> c.assert_called_with(arg)
Traceback (most recent call last):
...

```

(下页继续)

(续上页)

```

AssertionError: Expected call: mock({1})
Actual call: mock(set())
>>> c.foo
<CopyingMock name='mock.foo' id='...'>

```

When you subclass `Mock` or `MagicMock` all dynamically created attributes, and the `return_value` will use your subclass automatically. That means all children of a `CopyingMock` will also have the type `CopyingMock`.

Nesting Patches

Using `patch` as a context manager is nice, but if you do multiple patches you can end up with nested `with` statements indenting further and further to the right:

```

>>> class MyTest(TestCase):
...     def test_foo(self):
...         with patch('mymodule.Foo') as mock_foo:
...             with patch('mymodule.Bar') as mock_bar:
...                 with patch('mymodule.Spam') as mock_spam:
...                     assert mymodule.Foo is mock_foo
...                     assert mymodule.Bar is mock_bar
...                     assert mymodule.Spam is mock_spam
...
>>> original = mymodule.Foo
>>> MyTest('test_foo').test_foo()
>>> assert mymodule.Foo is original

```

With unittest cleanup functions and the *patch methods: start and stop* we can achieve the same effect without the nested indentation. A simple helper method, `create_patch`, puts the patch in place and returns the created mock for us:

```

>>> class MyTest(TestCase):
...     def create_patch(self, name):
...         patcher = patch(name)
...         thing = patcher.start()
...         self.addCleanup(patcher.stop)
...         return thing
...
...     def test_foo(self):
...         mock_foo = self.create_patch('mymodule.Foo')
...         mock_bar = self.create_patch('mymodule.Bar')
...         mock_spam = self.create_patch('mymodule.Spam')
...
...         assert mymodule.Foo is mock_foo
...         assert mymodule.Bar is mock_bar
...         assert mymodule.Spam is mock_spam
...
>>> original = mymodule.Foo
>>> MyTest('test_foo').run()
>>> assert mymodule.Foo is original

```

Mocking a dictionary with MagicMock

You may want to mock a dictionary, or other container object, recording all access to it whilst having it still behave like a dictionary.

We can do this with *MagicMock*, which will behave like a dictionary, and using *side_effect* to delegate dictionary access to a real underlying dictionary that is under our control.

When the `__getitem__()` and `__setitem__()` methods of our *MagicMock* are called (normal dictionary access) then *side_effect* is called with the key (and in the case of `__setitem__` the value too). We can also control what is returned.

After the *MagicMock* has been used we can use attributes like *call_args_list* to assert about how the dictionary was used:

```
>>> my_dict = {'a': 1, 'b': 2, 'c': 3}
>>> def getitem(name):
...     return my_dict[name]
...
>>> def setitem(name, val):
...     my_dict[name] = val
...
>>> mock = MagicMock()
>>> mock.__getitem__.side_effect = getitem
>>> mock.__setitem__.side_effect = setitem
```

注解: An alternative to using *MagicMock* is to use *Mock* and *only* provide the magic methods you specifically want:

```
>>> mock = Mock()
>>> mock.__getitem__ = Mock(side_effect=getitem)
>>> mock.__setitem__ = Mock(side_effect=setitem)
```

A *third* option is to use *MagicMock* but passing in *dict* as the *spec* (or *spec_set*) argument so that the *MagicMock* created only has dictionary magic methods available:

```
>>> mock = MagicMock(spec_set=dict)
>>> mock.__getitem__.side_effect = getitem
>>> mock.__setitem__.side_effect = setitem
```

With these side effect functions in place, the *mock* will behave like a normal dictionary but recording the access. It even raises a *KeyError* if you try to access a key that doesn't exist.

```
>>> mock['a']
1
>>> mock['c']
3
>>> mock['d']
Traceback (most recent call last):
...
KeyError: 'd'
>>> mock['b'] = 'fish'
>>> mock['d'] = 'eggs'
>>> mock['b']
'fish'
>>> mock['d']
'eggs'
```

After it has been used you can make assertions about the access using the normal mock methods and attributes:

```
>>> mock.__getitem__.call_args_list
[call('a'), call('c'), call('d'), call('b'), call('d')]
>>> mock.__setitem__.call_args_list
[call('b', 'fish'), call('d', 'eggs')]
>>> my_dict
{'a': 1, 'c': 3, 'b': 'fish', 'd': 'eggs'}
```

Mock subclasses and their attributes

There are various reasons why you might want to subclass `Mock`. One reason might be to add helper methods. Here's a silly example:

```
>>> class MyMock(MagicMock):
...     def has_been_called(self):
...         return self.called
...
>>> mymock = MyMock(return_value=None)
>>> mymock
<MyMock id='...'>
>>> mymock.has_been_called()
False
>>> mymock()
>>> mymock.has_been_called()
True
```

The standard behaviour for `Mock` instances is that attributes and the return value mocks are of the same type as the mock they are accessed on. This ensures that `Mock` attributes are `Mocks` and `MagicMock` attributes are `MagicMocks`². So if you're subclassing to add helper methods then they'll also be available on the attributes and return value mock of instances of your subclass.

```
>>> mymock.foo
<MyMock name='mock.foo' id='...'>
>>> mymock.foo.has_been_called()
False
>>> mymock.foo()
<MyMock name='mock.foo()' id='...'>
>>> mymock.foo.has_been_called()
True
```

Sometimes this is inconvenient. For example, [one user](#) is subclassing mock to create a [Twisted adaptor](#). Having this applied to attributes too actually causes errors.

`Mock` (in all its flavours) uses a method called `_get_child_mock` to create these “sub-mocks” for attributes and return values. You can prevent your subclass being used for attributes by overriding this method. The signature is that it takes arbitrary keyword arguments (`**kwargs`) which are then passed onto the mock constructor:

```
>>> class Subclass(MagicMock):
...     def _get_child_mock(self, **kwargs):
...         return MagicMock(**kwargs)
...
>>> mymock = Subclass()
```

(下页继续)

² An exception to this rule are the non-callable mocks. Attributes use the callable variant because otherwise non-callable mocks couldn't have callable methods.

```
>>> mymock.foo
<MagicMock name='mock.foo' id='...'>
>>> assert isinstance(mymock, Subclass)
>>> assert not isinstance(mymock.foo, Subclass)
>>> assert not isinstance(mymock(), Subclass)
```

(续上页)

```
<Mock name='mock.module.fooble()' id='... '>
>>> mock.module.fooble.assert_called_once_with()
```

Tracking order of calls and less verbose call assertions

The `Mock` class allows you to track the *order* of method calls on your mock objects through the `method_calls` attribute. This doesn't allow you to track the order of calls between separate mock objects, however we can use `mock_calls` to achieve the same effect.

Because mocks track calls to child mocks in `mock_calls`, and accessing an arbitrary attribute of a mock creates a child mock, we can create our separate mocks from a parent one. Calls to those child mock will then all be recorded, in order, in the `mock_calls` of the parent:

```
>>> manager = Mock()
>>> mock_foo = manager.foo
>>> mock_bar = manager.bar
```

```
>>> mock_foo.something()
<Mock name='mock.foo.something()' id='... '>
>>> mock_bar.other.thing()
<Mock name='mock.bar.other.thing()' id='... '>
```

```
>>> manager.mock_calls
[call.foo.something(), call.bar.other.thing()]
```

We can then assert about the calls, including the order, by comparing with the `mock_calls` attribute on the manager mock:

```
>>> expected_calls = [call.foo.something(), call.bar.other.thing()]
>>> manager.mock_calls == expected_calls
True
```

If `patch` is creating, and putting in place, your mocks then you can attach them to a manager mock using the `attach_mock()` method. After attaching calls will be recorded in `mock_calls` of the manager.

```
>>> manager = MagicMock()
>>> with patch('mymodule.Class1') as MockClass1:
...     with patch('mymodule.Class2') as MockClass2:
...         manager.attach_mock(MockClass1, 'MockClass1')
...         manager.attach_mock(MockClass2, 'MockClass2')
...         MockClass1().foo()
...         MockClass2().bar()
...
<MagicMock name='mock.MockClass1().foo()' id='... '>
<MagicMock name='mock.MockClass2().bar()' id='... '>
>>> manager.mock_calls
[call.MockClass1(),
 call.MockClass1().foo(),
 call.MockClass2(),
 call.MockClass2().bar()]
```

If many calls have been made, but you're only interested in a particular sequence of them then an alternative is to use the `assert_has_calls()` method. This takes a list of calls (constructed with the `call` object). If that sequence of calls are in `mock_calls` then the assert succeeds.

```
>>> m = MagicMock()
>>> m().foo().bar().baz()
<MagicMock name='mock().foo().bar().baz()' id='...'>
>>> m.one().two().three()
<MagicMock name='mock.one().two().three()' id='...'>
>>> calls = call.one().two().three().call_list()
>>> m.assert_has_calls(calls)
```

Even though the chained call `m.one().two().three()` aren't the only calls that have been made to the mock, the assert still succeeds.

Sometimes a mock may have several calls made to it, and you are only interested in asserting about *some* of those calls. You may not even care about the order. In this case you can pass `any_order=True` to `assert_has_calls`:

```
>>> m = MagicMock()
>>> m(1), m.two(2, 3), m.seven(7), m.fifty('50')
(...)
>>> calls = [call.fifty('50'), call(1), call.seven(7)]
>>> m.assert_has_calls(calls, any_order=True)
```

More complex argument matching

Using the same basic concept as [ANY](#) we can implement matchers to do more complex assertions on objects used as arguments to mocks.

Suppose we expect some object to be passed to a mock that by default compares equal based on object identity (which is the Python default for user defined classes). To use `assert_called_with()` we would need to pass in the exact same object. If we are only interested in some of the attributes of this object then we can create a matcher that will check these attributes for us.

You can see in this example how a ‘standard’ call to `assert_called_with` isn't sufficient:

```
>>> class Foo:
...     def __init__(self, a, b):
...         self.a, self.b = a, b
...
>>> mock = Mock(return_value=None)
>>> mock(Foo(1, 2))
>>> mock.assert_called_with(Foo(1, 2))
Traceback (most recent call last):
...
AssertionError: Expected: call(<__main__.Foo object at 0x...>)
Actual call: call(<__main__.Foo object at 0x...>)
```

A comparison function for our `Foo` class might look something like this:

```
>>> def compare(self, other):
...     if not type(self) == type(other):
...         return False
...     if self.a != other.a:
...         return False
...     if self.b != other.b:
...         return False
...     return True
...
```

And a matcher object that can use comparison functions like this for its equality operation would look something like this:

```
>>> class Matcher:
...     def __init__(self, compare, some_obj):
...         self.compare = compare
...         self.some_obj = some_obj
...     def __eq__(self, other):
...         return self.compare(self.some_obj, other)
... 
```

Putting all this together:

```
>>> match_foo = Matcher(compare, Foo(1, 2))
>>> mock.assert_called_with(match_foo)
```

The `Matcher` is instantiated with our `compare` function and the `Foo` object we want to compare against. In `assert_called_with` the `Matcher` equality method will be called, which compares the object the mock was called with against the one we created our matcher with. If they match then `assert_called_with` passes, and if they don't an `AssertionError` is raised:

```
>>> match_wrong = Matcher(compare, Foo(3, 4))
>>> mock.assert_called_with(match_wrong)
Traceback (most recent call last):
...
AssertionError: Expected: ((<Matcher object at 0x...>,), {})
Called with: ((<Foo object at 0x...>,), {})
```

With a bit of tweaking you could have the comparison function raise the `AssertionError` directly and provide a more useful failure message.

As of version 1.5, the Python testing library `PyHamcrest` provides similar functionality, that may be useful here, in the form of its equality matcher (`hamcrest.library.integration.match_equality`).

26.7 2to3 - 自动将 Python 2 代码转为 Python 3 代码

`2to3` 是一个 Python 程序，它可以用来读取 Python 2.x 版本的代码，并使用一系列的 * 修复器 `fixer` 来将其转换为合法的 Python 3.x 代码。标准库中已经包含了丰富的修复器，这足以处理绝大多数代码。不过 `2to3` 的支持库 `lib2to3` 是一个很灵活通用的库，所以你也可以为 `2to3` 编写你自己的修复器。`lib2to3` 也可以用在那些需要自动处理 Python 代码的应用中。

26.7.1 使用 2to3

`2to3` 通常会作为脚本和 Python 解释器一起安装，你可以在 Python 根目录的 `Tools/scripts` 文件夹下找到它。

`2to3` 的基本调用参数是一个需要转换的文件或目录列表。对于目录，会递归地寻找其中的 Python 源码。

这里有一个 Python 2.x 的源码文件，`example.py`：

```
def greet(name):
    print "Hello, {0}!".format(name)
print "What's your name?"
name = raw_input()
greet(name)
```

它可以在命令行中使用 `2to3` 转换成 Python 3.x 版本的代码：

```
$ 2to3 example.py
```

这个命令会打印出和源文件的区别。通过传入 `-w` 参数，`2to3` 也可以把需要的修改写回到原文件中（除非传入了 `-n` 参数，否则会为原始文件创建一个副本）：

```
$ 2to3 -w example.py
```

在转换完成后，`example.py` 看起来像是这样：

```
def greet(name):
    print("Hello, {0}!".format(name))
print("What's your name?")
name = input()
greet(name)
```

注释和缩进都会在转换过程中保持不变。

默认情况下，`2to3` 会执行预定义修复器的集合。使用 `-l` 参数可以列出所有可用的修复器。使用 `-f` 参数可以明确指定需要使用的修复器集合。而使用 `-x` 参数则可以明确指定不使用的修复器。下面的例子会只使用 `imports` 和 `has_key` 修复器运行：

```
$ 2to3 -f imports -f has_key example.py
```

这个命令会执行除了 `apply` 之外的所有修复器：

```
$ 2to3 -x apply example.py
```

有一些修复器是需要 * 显式指定 * 的，它们默认不会执行，必须在命令行中列出才会执行。比如下面的例子，除了默认的修复器以外，还会执行 `idioms` 修复器：

```
$ 2to3 -f all -f idioms example.py
```

注意这里使用 `all` 来启用所有默认的修复器。

有些情况下 `2to3` 会找到源码中有一些需要修改，但是无法自动处理的代码。在这种情况下，`2to3` 会在差异处下面打印一个警告信息。你应该定位到相应的代码并对其进行修改，以使其兼容 `Python 3.x`。

`2to3` 也可以重构 `doctests`。使用 `-d` 开启这个模式。需要注意 * 只有 * `doctests` 会被重构。这种模式下不需要文件是合法的 `Python` 代码。举例来说，`reST` 文档中类似 `doctests` 的示例也可以使用这个选项进行重构。

`-v` 选项可以输出更多转换程序的详细信息。

由于某些 `print` 语句可被解读为函数调用或是语句，`2to3` 并不是总能读取包含 `print` 函数的文件。当 `2to3` 检测到存在 `from __future__ import print_function` 编译器指令时，会修改其内部语法将 `print()` 解读为函数。这一变动也可以使用 `-p` 选项手动开启。使用 `-p` 来为已经转换过 `print` 语句的代码运行修复器。

`-o` 或 `--output-dir` 选项可以指定将转换后的文件写入其他目录中。由于这种情况下不会覆写原始文件，所以创建副本文件毫无意义，因此也需要使用 `-n` 选项来禁用创建副本。

3.2.3 新版功能: 增加了 `-o` 选项。

`-W` 或 `--write-unchanged-files` 选项用来告诉 `2to3` 始终需要输出文件，即使没有任何改动。这在使用 `-o` 参数时十分有用，这样就可以将整个 `Python` 源码包完整地转换到另一个目录。这个选项隐含了 `-w` 选项，否则等于没有作用。

3.2.3 新版功能: 增加了 `-w` 选项。

`--add-suffix` 选项接受一个字符串，用来作为后缀附加在输出文件名后面的后面。由于写入的文件名与原始文件不同，所以没有必要创建副本，因此 `-n` 选项也是必要的。举个例子：

```
$ 2to3 -n -W --add-suffix=3 example.py
```

这样会把转换后的文件写入 `example.py3` 文件。

3.2.3 新版功能: 增加了 `--add-suffix` 选项。

将整个项目从一个目录转换到另一个目录可以用这样的命令:

```
$ 2to3 --output-dir=python3-version/mycode -W -n python2-version/mycode
```

26.7.2 修复器

转换代码的每一个步骤都封装在修复器中。可以使用 `2to3 -l` 来列出可用的修复器。之前已经提到, 每个修复器都可以独立地打开或是关闭。下面会对各个修复器做更详细的描述。

apply

移除对 `apply()` 的使用, 举例来说, `apply(function, *args, **kwargs)` 会被转换成 `function(*args, **kwargs)`。

asserts

将已弃用的 `unittest` 方法替换为正确的。

Python 2.x	Python 3.x
<code>failUnlessEqual(a, b)</code>	<code>assertEqual(a, b)</code>
<code>assertEquals(a, b)</code>	<code>assertEqual(a, b)</code>
<code>failIfEqual(a, b)</code>	<code>assertNotEqual(a, b)</code>
<code>assertNotEquals(a, b)</code>	<code>assertNotEqual(a, b)</code>
<code>failUnless(a)</code>	<code>assertTrue(a)</code>
<code>assert_(a)</code>	<code>assertTrue(a)</code>
<code>failIf(a)</code>	<code>assertFalse(a)</code>
<code>failUnlessRaises(exc, cal)</code>	<code>assertRaises(exc, cal)</code>
<code>failUnlessAlmostEqual(a, b)</code>	<code>assertAlmostEqual(a, b)</code>
<code>assertAlmostEquals(a, b)</code>	<code>assertAlmostEqual(a, b)</code>
<code>failIfAlmostEqual(a, b)</code>	<code>assertNotAlmostEqual(a, b)</code>
<code>assertNotAlmostEquals(a, b)</code>	<code>assertNotAlmostEqual(a, b)</code>

basestring

将 `basestring` 转换为 `str`。

buffer

将 `buffer` 转换为 `memoryview`。这个修复器是可选的, 因为 `memoryview` API 和 `buffer` 很相似, 但不完全一样。

dict

修复字典迭代方法。`dict.iteritems()` 会转换成 `dict.items()`, `dict.iterkeys()` 会转换成 `dict.keys()`, `dict.itervalues()` 会转换成 `dict.values()`。类似的, `dict.viewitems()`, `dict.viewkeys()` 和 `dict.viewvalues()` 会分别转换成 `dict.items()`, `dict.keys()` 和 `dict.values()`。另外也会将原有的 `dict.items()`, `dict.keys()` 和 `dict.values()` 方法调用用 `list` 包装一层。

except

将 `except X, T` 转换为 `except X as T`。

exec

将 `exec` 语句转换为 `exec()` 函数调用。

execfile

移除 `execfile()` 的使用。`execfile()` 的实参会使用 `open()`、`compile()` 和 `exec()` 包装。

exitfunc

将对 `sys.exitfunc` 的赋值改为使用 `atexit` 模块代替。

filter

将 `filter()` 函数用 `list` 包装一层。

funcattrs

修复已经重命名的函数属性。比如 `my_function.func_closure` 会被转换为 `my_function.__closure__`。

future

移除 `from __future__ import new_feature` 语句。

getcwdu

将 `os.getcwdu()` 重命名为 `os.getcwd()`。

has_key

将 `dict.has_key(key)` 转换为 `key in dict`。

idioms

这是一个可选的修复器，会进行多种转换，将 Python 代码变成更加常见的写法。类似 `type(x) is SomeClass` 和 `type(x) == SomeClass` 的类型对比会被转换成 `isinstance(x, SomeClass)`。`while 1` 转换成 `while True`。这个修复器还会在合适的地方使用 `sorted()` 函数。举个例子，这样的代码块：

```
L = list(some_iterable)
L.sort()
```

会被转换为：

```
L = sorted(some_iterable)
```

import

检测 sibling imports，并将其转换成相对 import。

imports

处理标准库模块的重命名。

imports2

处理标准库中其他模块的重命名。这个修复器由于一些技术上的限制，因此和 `imports` 拆分开了。

input

将 `input(prompt)` 转换为 `eval(input(prompt))`。

intern

将 `intern()` 转换为 `sys.intern()`。

isinstance

修复 `isinstance()` 函数第二个实参中重复的类型。举例来说，`isinstance(x, (int, int))` 会转换为 `isinstance(x, int)`，`isinstance(x, (int, float, int))` 会转换为 `isinstance(x, (int, float))`。

itertools_imports

移除 `itertools.ifilter()`、`itertools.izip()` 以及 `itertools.imap()` 的 import。对 `itertools.ifilterfalse()` 的 import 也会替换成 `itertools.filterfalse()`。

itertools

修改 `itertools.ifilter()`、`itertools.izip()` 和 `itertools.imap()` 的调用为对应的内建实现。`itertools.ifilterfalse()` 会替换成 `itertools.filterfalse()`。

long

将 long 重命名为 *int*。

map

用 *list* 包装 *map()*。同时也会将 `map(None, x)` 替换为 `list(x)`。使用 `from future_builtins import map` 禁用这个修复器。

metaclass

将老的元类语法 (类体中的 `__metaclass__ = Meta`) 替换为新的 (`class X(metaclass=Meta)`)。

methodattrs

修复老的方法属性名。例如 `meth.im_func` 会被转换为 `meth.__func__`。

ne

转换老的不等语法, 将 `<>` 转为 `!=`。

next

将迭代器的 `next()` 方法调用转为 *next()* 函数。也会将 *next()* 方法重命名为 `__next__()`。

nonzero

将 `__nonzero__()` 转换为 `__bool__()`。

numliterals

将八进制字面量转为新的语法。

operator

Converts calls to various functions in the *operator* module to other, but equivalent, function calls. When needed, the appropriate import statements are added, e.g. `import collections`. The following mapping are made:

Python 2.x	Python 3.x
<code>operator.isCallable(obj)</code>	<code>hasattr(obj, '__call__')</code>
<code>operator.sequenceIncludes(obj)</code>	<code>operator.contains(obj)</code>
<code>operator.isSequenceType(obj)</code>	<code>isinstance(obj, collections.Sequence)</code>
<code>operator.isMappingType(obj)</code>	<code>isinstance(obj, collections.Mapping)</code>
<code>operator.isNumberType(obj)</code>	<code>isinstance(obj, numbers.Number)</code>
<code>operator.repeat(obj, n)</code>	<code>operator.mul(obj, n)</code>
<code>operator.irepeat(obj, n)</code>	<code>operator.imul(obj, n)</code>

paren

在列表生成式中增加必须的括号。例如将 `[x for x in 1, 2]` 转换为 `[x for x in (1, 2)]`。

print

将 `print` 语句转换为 *print()* 函数。

raise

将 `raise E, V` 转换为 `raise E(V)`, 将 `raise E, V, T` 转换为 `raise E(V).with_traceback(T)`。如果 `E` 是元组, 这样的转换是不正确的, 因为用元组代替异常的做法在 3.0 中已经移除了。

raw_input

将 `raw_input()` 转换为 *input()*。

reduce

将 `reduce()` 转换为 *functools.reduce()*。

reload

Converts `reload()` to *imp.reload()*。

renames

将 `sys.maxint` 转换为 *sys.maxsize*。

repr

将反引号 `repr` 表达式替换为 `repr()` 函数。

set_literal

将 `set` 构造函数替换为 `set literals` 写法。这个修复器是可选的。

standarderror

将 `StandardError` 重命名为 `Exception`。

sys_exc

将弃用的 `sys.exc_value`, `sys.exc_type`, `sys.exc_traceback` 替换为 `sys.exc_info()` 的用法。

throw

修复生成器的 `throw()` 方法的 API 变更。

tuple_params

移除隐式的元组参数解包。这个修复器会插入临时变量。

types

修复 `type` 模块中一些成员的移除引起的代码问题。

unicode

将 `unicode` 重命名为 `str`。

urllib

将 `urllib` 和 `urllib2` 重命名为 `urllib` 包。

ws_comma

移除逗号分隔的元素之间多余的空白。这个修复器是可选的。

xrange

将 `xrange()` 重命名为 `range()`，并用 `list` 包装原有的 `range()`。

xreadlines

将 `for x in file.xreadlines()` 转换为 `for x in file`。

zip

用 `list` 包装 `zip()`。如果使用了 `from future_builtins import zip` 的话会禁用。

26.7.3 lib2to3 —— 2to3 支持库

源代码: [Lib/lib2to3/](#)

注解: `lib2to3` API 并不稳定，并可能在未来大幅修改。

26.8 test — Regression tests package for Python

注解: The `test` package is meant for internal use by Python only. It is documented for the benefit of the core developers of Python. Any use of this package outside of Python's standard library is discouraged as code mentioned here can change or be removed without notice between releases of Python.

The `test` package contains all regression tests for Python as well as the modules `test.support` and `test.regrtest`. `test.support` is used to enhance your tests while `test.regrtest` drives the testing suite.

Each module in the `test` package whose name starts with `test_` is a testing suite for a specific module or feature. All new tests should be written using the `unittest` or `doctest` module. Some older tests are written using a “traditional” testing style that compares output printed to `sys.stdout`; this style of test is considered deprecated.

参见:

Module `unittest` Writing PyUnit regression tests.

`unittest` — 文档测试 Tests embedded in documentation strings.

26.8.1 Writing Unit Tests for the `test` package

It is preferred that tests that use the `unittest` module follow a few guidelines. One is to name the test module by starting it with `test_` and end it with the name of the module being tested. The test methods in the test module should start with `test_` and end with a description of what the method is testing. This is needed so that the methods are recognized by the test driver as test methods. Also, no documentation string for the method should be included. A comment (such as `# Tests function returns only True or False`) should be used to provide documentation for test methods. This is done because documentation strings get printed out if they exist and thus what test is being run is not stated.

A basic boilerplate is often used:

```
import unittest
from test import support

class MyTestCase1(unittest.TestCase):

    # Only use setUp() and tearDown() if necessary

    def setUp(self):
        ... code to execute in preparation for tests ...

    def tearDown(self):
        ... code to execute to clean up after tests ...

    def test_feature_one(self):
        # Test feature one.
        ... testing code ...

    def test_feature_two(self):
        # Test feature two.
        ... testing code ...

    ... more test methods ...

class MyTestCase2(unittest.TestCase):
    ... same structure as MyTestCase1 ...

... more test classes ...

if __name__ == '__main__':
    unittest.main()
```

This code pattern allows the testing suite to be run by `test.regrtest`, on its own as a script that supports the `unittest` CLI, or via the `python -m unittest` CLI.

The goal for regression testing is to try to break code. This leads to a few guidelines to be followed:

- The testing suite should exercise all classes, functions, and constants. This includes not just the external API that is to be presented to the outside world but also “private” code.
- Whitebox testing (examining the code being tested when the tests are being written) is preferred. Blackbox testing (testing only the published user interface) is not complete enough to make sure all boundary and edge cases are tested.
- Make sure all possible values are tested including invalid ones. This makes sure that not only all valid values are acceptable but also that improper values are handled correctly.
- Exhaust as many code paths as possible. Test where branching occurs and thus tailor input to make sure as many different paths through the code are taken.
- Add an explicit test for any bugs discovered for the tested code. This will make sure that the error does not crop up again if the code is changed in the future.
- Make sure to clean up after your tests (such as close and remove all temporary files).
- If a test is dependent on a specific condition of the operating system then verify the condition already exists before attempting the test.
- Import as few modules as possible and do it as soon as possible. This minimizes external dependencies of tests and also minimizes possible anomalous behavior from side-effects of importing a module.
- Try to maximize code reuse. On occasion, tests will vary by something as small as what type of input is used. Minimize code duplication by subclassing a basic test class with a class that specifies the input:

```
class TestFuncAcceptsSequencesMixin:

    func = mySuperWhammyFunction

    def test_func(self):
        self.func(self.arg)

class AcceptLists(TestFuncAcceptsSequencesMixin, unittest.TestCase):
    arg = [1, 2, 3]

class AcceptStrings(TestFuncAcceptsSequencesMixin, unittest.TestCase):
    arg = 'abc'

class AcceptTuples(TestFuncAcceptsSequencesMixin, unittest.TestCase):
    arg = (1, 2, 3)
```

When using this pattern, remember that all classes that inherit from `unittest.TestCase` are run as tests. The Mixin class in the example above does not have any data and so can't be run by itself, thus it does not inherit from `unittest.TestCase`.

参见:

Test Driven Development A book by Kent Beck on writing tests before code.

26.8.2 Running tests using the command-line interface

The `test` package can be run as a script to drive Python's regression test suite, thanks to the `-m` option: **python -m test**. Under the hood, it uses `test.regrtest`; the call **python -m test.regrtest** used in previous Python versions still works. Running the script by itself automatically starts running all regression tests in the `test` package. It does this by finding all modules in the package whose name starts with `test_`, importing them, and executing the function `test_main()` if present or loading the tests via `unittest.TestLoader.loadTestsFromModule` if `test_main`

does not exist. The names of tests to execute may also be passed to the script. Specifying a single regression test (**python -m test test_spam**) will minimize output and only print whether the test passed or failed.

Running *test* directly allows what resources are available for tests to use to be set. You do this by using the `-u` command-line option. Specifying `all` as the value for the `-u` option enables all possible resources: **python -m test -uall**. If all but one resource is desired (a more common case), a comma-separated list of resources that are not desired may be listed after `all`. The command **python -m test -uall,-audio,-largefile** will run *test* with all resources except the `audio` and `largefile` resources. For a list of all resources and more command-line options, run **python -m test -h**.

Some other ways to execute the regression tests depend on what platform the tests are being executed on. On Unix, you can run **make test** at the top-level directory where Python was built. On Windows, executing **rt.bat** from your PCBuild directory will run all regression tests.

26.9 test.support —Utilities for the Python test suite

The *test.support* module provides support for Python's regression test suite.

注解: *test.support* is not a public module. It is documented here to help Python developers write tests. The API of this module is subject to change without backwards compatibility concerns between releases.

This module defines the following exceptions:

exception *test.support.TestFailed*

Exception to be raised when a test fails. This is deprecated in favor of *unittest*-based tests and *unittest.TestCase*'s assertion methods.

exception *test.support.ResourceDenied*

Subclass of *unittest.SkipTest*. Raised when a resource (such as a network connection) is not available. Raised by the *requires()* function.

The *test.support* module defines the following constants:

test.support.verbose

True when verbose output is enabled. Should be checked when more detailed information is desired about a running test. *verbose* is set by *test.regrtest*.

test.support.is_jython

True if the running interpreter is Jython.

test.support.TESTFN

Set to a name that is safe to use as the name of a temporary file. Any temporary file that is created should be closed and unlinked (removed).

The *test.support* module defines the following functions:

test.support.forget (*module_name*)

Remove the module named *module_name* from *sys.modules* and delete any byte-compiled files of the module.

test.support.is_resource_enabled (*resource*)

Return True if *resource* is enabled and available. The list of available resources is only set when *test.regrtest* is executing the tests.

test.support.requires (*resource*, *msg=None*)

Raise *ResourceDenied* if *resource* is not available. *msg* is the argument to *ResourceDenied* if it is raised. Always returns True if called by a function whose `__name__` is `'__main__'`. Used when tests are executed by *test.regrtest*.

`test.support.findfile(filename, subdir=None)`

Return the path to the file named *filename*. If no match is found *filename* is returned. This does not equal a failure since it could be the path to the file.

Setting *subdir* indicates a relative path to use to find the file rather than looking directly in the path directories.

`test.support.run_unittest(*classes)`

Execute `unittest.TestCase` subclasses passed to the function. The function scans the classes for methods starting with the prefix `test_` and executes the tests individually.

It is also legal to pass strings as parameters; these should be keys in `sys.modules`. Each associated module will be scanned by `unittest.TestLoader.loadTestsFromModule()`. This is usually seen in the following `test_main()` function:

```
def test_main():
    support.run_unittest(__name__)
```

This will run all tests defined in the named module.

`test.support.run_doctest(module, verbosity=None)`

Run `doctest.testmod()` on the given *module*. Return (failure_count, test_count).

If *verbosity* is None, `doctest.testmod()` is run with verbosity set to *verbose*. Otherwise, it is run with verbosity set to None.

`test.support.check_warnings(*filters, quiet=True)`

A convenience wrapper for `warnings.catch_warnings()` that makes it easier to test that a warning was correctly raised. It is approximately equivalent to calling `warnings.catch_warnings(record=True)` with `warnings.simplefilter()` set to always and with the option to automatically validate the results that are recorded.

`check_warnings` accepts 2-tuples of the form ("message regexp", `WarningCategory`) as positional arguments. If one or more *filters* are provided, or if the optional keyword argument *quiet* is `False`, it checks to make sure the warnings are as expected: each specified filter must match at least one of the warnings raised by the enclosed code or the test fails, and if any warnings are raised that do not match any of the specified filters the test fails. To disable the first of these checks, set *quiet* to `True`.

If no arguments are specified, it defaults to:

```
check_warnings(("", Warning), quiet=True)
```

In this case all warnings are caught and no errors are raised.

On entry to the context manager, a `WarningRecorder` instance is returned. The underlying warnings list from `catch_warnings()` is available via the recorder object's *warnings* attribute. As a convenience, the attributes of the object representing the most recent warning can also be accessed directly through the recorder object (see example below). If no warning has been raised, then any of the attributes that would otherwise be expected on an object representing a warning will return `None`.

The recorder object also has a `reset()` method, which clears the warnings list.

The context manager is designed to be used like this:

```
with check_warnings(("assertion is always true", SyntaxWarning),
                    ("", UserWarning)):
    exec('assert(False, "Hey!")')
    warnings.warn(UserWarning("Hide me!"))
```

In this case if either warning was not raised, or some other warning was raised, `check_warnings()` would raise an error.

When a test needs to look more deeply into the warnings, rather than just checking whether or not they occurred, code like this can be used:

```
with check_warnings(quiet=True) as w:
    warnings.warn("foo")
    assert str(w.args[0]) == "foo"
    warnings.warn("bar")
    assert str(w.args[0]) == "bar"
    assert str(w.warnings[0].args[0]) == "foo"
    assert str(w.warnings[1].args[0]) == "bar"
    w.reset()
    assert len(w.warnings) == 0
```

Here all warnings will be caught, and the test code tests the captured warnings directly.

在 3.2 版更改: New optional arguments *filters* and *quiet*.

```
test.support.captured_stdin()
test.support.captured_stdout()
test.support.captured_stderr()
```

A context managers that temporarily replaces the named stream with *io.StringIO* object.

Example use with output streams:

```
with captured_stdout() as stdout, captured_stderr() as stderr:
    print("hello")
    print("error", file=sys.stderr)
assert stdout.getvalue() == "hello\n"
assert stderr.getvalue() == "error\n"
```

Example use with input stream:

```
with captured_stdin() as stdin:
    stdin.write('hello\n')
    stdin.seek(0)
    # call test code that consumes from sys.stdin
    captured = input()
self.assertEqual(captured, "hello")
```

```
test.support.temp_dir(path=None, quiet=False)
```

A context manager that creates a temporary directory at *path* and yields the directory.

If *path* is None, the temporary directory is created using *tempfile.mkdtemp()*. If *quiet* is False, the context manager raises an exception on error. Otherwise, if *path* is specified and cannot be created, only a warning is issued.

```
test.support.change_cwd(path, quiet=False)
```

A context manager that temporarily changes the current working directory to *path* and yields the directory.

If *quiet* is False, the context manager raises an exception on error. Otherwise, it issues only a warning and keeps the current working directory the same.

```
test.support.temp_cwd(name='tempcwd', quiet=False)
```

A context manager that temporarily creates a new directory and changes the current working directory (CWD).

The context manager creates a temporary directory in the current directory with name *name* before temporarily changing the current working directory. If *name* is None, the temporary directory is created using *tempfile.mkdtemp()*.

If *quiet* is False and it is not possible to create or change the CWD, an error is raised. Otherwise, only a warning is raised and the original CWD is used.

`test.support.temp_umask(umask)`
A context manager that temporarily sets the process umask.

`test.support.can_symlink()`
Return True if the OS supports symbolic links, False otherwise.

`@test.support.skip_unless_symlink`
A decorator for running tests that require support for symbolic links.

`@test.support.anticipate_failure(condition)`
A decorator to conditionally mark tests with `unittest.expectedFailure()`. Any use of this decorator should have an associated comment identifying the relevant tracker issue.

`@test.support.run_with_locale(catstr, *locales)`
A decorator for running a function in a different locale, correctly resetting it after it has finished. `catstr` is the locale category as a string (for example "LC_ALL"). The `locales` passed will be tried sequentially, and the first valid locale will be used.

`test.support.make_bad_fd()`
Create an invalid file descriptor by opening and closing a temporary file, and returning its descriptor.

`test.support.import_module(name, deprecated=False)`
This function imports and returns the named module. Unlike a normal import, this function raises `unittest.SkipTest` if the module cannot be imported.

Module and package deprecation messages are suppressed during this import if `deprecated` is True.

3.1 新版功能.

`test.support.import_fresh_module(name, fresh=(), blocked=(), deprecated=False)`
This function imports and returns a fresh copy of the named Python module by removing the named module from `sys.modules` before doing the import. Note that unlike `reload()`, the original module is not affected by this operation.

`fresh` is an iterable of additional module names that are also removed from the `sys.modules` cache before doing the import.

`blocked` is an iterable of module names that are replaced with None in the module cache during the import to ensure that attempts to import them raise `ImportError`.

The named module and any modules named in the `fresh` and `blocked` parameters are saved before starting the import and then reinserted into `sys.modules` when the fresh import is complete.

Module and package deprecation messages are suppressed during this import if `deprecated` is True.

This function will raise `ImportError` if the named module cannot be imported.

Example use:

```
# Get copies of the warnings module for testing without affecting the
# version being used by the rest of the test suite. One copy uses the
# C implementation, the other is forced to use the pure Python fallback
# implementation
py_warnings = import_fresh_module('warnings', blocked=['_warnings'])
c_warnings = import_fresh_module('warnings', fresh=['_warnings'])
```

3.1 新版功能.

`test.support.bind_port(sock, host=HOST)`
Bind the socket to a free port and return the port number. Relies on ephemeral ports in order to ensure we are using an unbound port. This is important as many tests may be running simultaneously, especially in a buildbot environment. This method raises an exception if the `sock.family` is `AF_INET` and `sock.type` is `SOCK_STREAM`,

and the socket has `SO_REUSEADDR` or `SO_REUSEPORT` set on it. Tests should never set these socket options for TCP/IP sockets. The only case for setting these options is testing multicasting via multiple UDP sockets.

Additionally, if the `SO_EXCLUSIVEADDRUSE` socket option is available (i.e. on Windows), it will be set on the socket. This will prevent anyone else from binding to our host/port for the duration of the test.

`test.support.find_unused_port(family=socket.AF_INET, socktype=socket.SOCK_STREAM)`

Returns an unused port that should be suitable for binding. This is achieved by creating a temporary socket with the same family and type as the `sock` parameter (default is `AF_INET`, `SOCK_STREAM`), and binding it to the specified host address (defaults to `0.0.0.0`) with the port set to 0, eliciting an unused ephemeral port from the OS. The temporary socket is then closed and deleted, and the ephemeral port is returned.

Either this method or `bind_port()` should be used for any tests where a server socket needs to be bound to a particular port for the duration of the test. Which one to use depends on whether the calling code is creating a python socket, or if an unused port needs to be provided in a constructor or passed to an external program (i.e. the `-accept` argument to openssl's `s_server` mode). Always prefer `bind_port()` over `find_unused_port()` where possible. Using a hard coded port is discouraged since it can make multiple instances of the test impossible to run simultaneously, which is a problem for buildbots.

`test.support.load_package_tests(pkg_dir, loader, standard_tests, pattern)`

Generic implementation of the `unittest` `load_tests` protocol for use in test packages. `pkg_dir` is the root directory of the package; `loader`, `standard_tests`, and `pattern` are the arguments expected by `load_tests`. In simple cases, the test package's `__init__.py` can be the following:

```
import os
from test.support import load_package_tests

def load_tests(*args):
    return load_package_tests(os.path.dirname(__file__), *args)
```

`test.support.detect_api_mismatch(ref_api, other_api, *, ignore=())`

Returns the set of attributes, functions or methods of `ref_api` not found on `other_api`, except for a defined list of items to be ignored in this check specified in `ignore`.

By default this skips private attributes beginning with `'_'` but includes all magic methods, i.e. those starting and ending in `'__'`.

3.5 新版功能.

`test.support.check_all(test_case, module, name_of_module=None, extra=(), blacklist=())`

Assert that the `__all__` variable of `module` contains all public names.

The module's public names (its API) are detected automatically based on whether they match the public name convention and were defined in `module`.

The `name_of_module` argument can specify (as a string or tuple thereof) what module(s) an API could be defined in in order to be detected as a public API. One case for this is when `module` imports part of its public API from other modules, possibly a C backend (like `csv` and its `_csv`).

The `extra` argument can be a set of names that wouldn't otherwise be automatically detected as "public", like objects without a proper `__module__` attribute. If provided, it will be added to the automatically detected ones.

The `blacklist` argument can be a set of names that must not be treated as part of the public API even though their names indicate otherwise.

Example use:

```
import bar
import foo
import unittest
```

(下页继续)

(续上页)

```

from test import support

class MiscTestCase(unittest.TestCase):
    def test__all__(self):
        support.check__all__(self, foo)

class OtherTestCase(unittest.TestCase):
    def test__all__(self):
        extra = {'BAR_CONST', 'FOO_CONST'}
        blacklist = {'baz'} # Undocumented name.
        # bar imports part of its API from _bar.
        support.check__all__(self, bar, ('bar', '_bar'),
                                extra=extra, blacklist=blacklist)

```

3.6 新版功能.

The `test.support` module defines the following classes:

class `test.support.TransientResource` (*exc*, ***kwargs*)

Instances are a context manager that raises `ResourceDenied` if the specified exception type is raised. Any keyword arguments are treated as attribute/value pairs to be compared against any exception raised within the `with` statement. Only if all pairs match properly against attributes on the exception is `ResourceDenied` raised.

class `test.support.EnvironmentVarGuard`

Class used to temporarily set or unset environment variables. Instances can be used as a context manager and have a complete dictionary interface for querying/modifying the underlying `os.environ`. After exit from the context manager all changes to environment variables done through this instance will be rolled back.

在 3.1 版更改: Added dictionary interface.

`EnvironmentVarGuard.set` (*envvar*, *value*)

Temporarily set the environment variable *envvar* to the value of *value*.

`EnvironmentVarGuard.unset` (*envvar*)

Temporarily unset the environment variable *envvar*.

class `test.support.SuppressCrashReport`

A context manager used to try to prevent crash dialog popups on tests that are expected to crash a subprocess.

On Windows, it disables Windows Error Reporting dialogs using `SetErrorMode`.

On UNIX, `resource.setrlimit()` is used to set `resource.RLIMIT_CORE`'s soft limit to 0 to prevent core dump file creation.

On both platforms, the old value is restored by `__exit__()`.

class `test.support.WarningsRecorder`

Class used to record warnings for unit tests. See documentation of `check_warnings()` above for more details.

class `test.support.FakePath` (*path*)

Simple *path-like object*. It implements the `__fspath__()` method which just returns the *path* argument. If *path* is an exception, it will be raised in `__fspath__()`.

这些库可以帮助你进行 Python 开发：调试器使你能够逐步执行代码，分析堆栈帧并设置断点等，而分析器运行代码并为你提供执行时间的详细分类，从而使你能够找出你程序中的瓶颈。

27.1 bdb —Debugger framework

Source code: [Lib/bdb.py](#)

The `bdb` module handles basic debugger functions, like setting breakpoints or managing execution via the debugger.

定义了以下异常：

exception `bdb.BdbQuit`

Exception raised by the `Bdb` class for quitting the debugger.

The `bdb` module also defines two classes:

class `bdb.Breakpoint` (*self, file, line, temporary=0, cond=None, funcname=None*)

This class implements temporary breakpoints, ignore counts, disabling and (re-)enabling, and conditionals.

Breakpoints are indexed by number through a list called `bpybnumber` and by (`file`, `line`) pairs through `bplist`. The former points to a single instance of class `Breakpoint`. The latter points to a list of such instances since there may be more than one breakpoint per line.

When creating a breakpoint, its associated filename should be in canonical form. If a *funcname* is defined, a breakpoint hit will be counted when the first line of that function is executed. A conditional breakpoint always counts a hit.

`Breakpoint` instances have the following methods:

deleteMe ()

Delete the breakpoint from the list associated to a file/line. If it is the last breakpoint in that position, it also deletes the entry for the file/line.

enable()

Mark the breakpoint as enabled.

disable()

Mark the breakpoint as disabled.

bpformat()

Return a string with all the information about the breakpoint, nicely formatted:

- The breakpoint number.
- If it is temporary or not.
- Its file, line position.
- The condition that causes a break.
- If it must be ignored the next N times.
- The breakpoint hit count.

3.2 新版功能.

bpprint (*out=None*)

Print the output of *bpformat()* to the file *out*, or if it is *None*, to standard output.

class `bdb.Bdb` (*skip=None*)

The *Bdb* class acts as a generic Python debugger base class.

This class takes care of the details of the trace facility; a derived class should implement user interaction. The standard debugger class (*pdb.Pdb*) is an example.

The *skip* argument, if given, must be an iterable of glob-style module name patterns. The debugger will not step into frames that originate in a module that matches one of these patterns. Whether a frame is considered to originate in a certain module is determined by the `__name__` in the frame globals.

3.1 新版功能: The *skip* argument.

The following methods of *Bdb* normally don't need to be overridden.

canonic (*filename*)

Auxiliary method for getting a filename in a canonical form, that is, as a case-normalized (on case-insensitive filesystems) absolute path, stripped of surrounding angle brackets.

reset ()

Set the *botframe*, *stopframe*, *returnframe* and *quitting* attributes with values ready to start debugging.

trace_dispatch (*frame, event, arg*)

This function is installed as the trace function of debugged frames. Its return value is the new trace function (in most cases, that is, itself).

The default implementation decides how to dispatch a frame, depending on the type of event (passed as a string) that is about to be executed. *event* can be one of the following:

- "line": A new line of code is going to be executed.
- "call": A function is about to be called, or another code block entered.
- "return": A function or other code block is about to return.
- "exception": An exception has occurred.
- "c_call": A C function is about to be called.
- "c_return": A C function has returned.

- `"c_exception"`: A C function has raised an exception.

For the Python events, specialized functions (see below) are called. For the C events, no action is taken.

The `arg` parameter depends on the previous event.

See the documentation for `sys.settrace()` for more information on the trace function. For more information on code and frame objects, refer to types.

dispatch_line (*frame*)

If the debugger should stop on the current line, invoke the `user_line()` method (which should be overridden in subclasses). Raise a `BdbQuit` exception if the `Bdb.quitting` flag is set (which can be set from `user_line()`). Return a reference to the `trace_dispatch()` method for further tracing in that scope.

dispatch_call (*frame*, *arg*)

If the debugger should stop on this function call, invoke the `user_call()` method (which should be overridden in subclasses). Raise a `BdbQuit` exception if the `Bdb.quitting` flag is set (which can be set from `user_call()`). Return a reference to the `trace_dispatch()` method for further tracing in that scope.

dispatch_return (*frame*, *arg*)

If the debugger should stop on this function return, invoke the `user_return()` method (which should be overridden in subclasses). Raise a `BdbQuit` exception if the `Bdb.quitting` flag is set (which can be set from `user_return()`). Return a reference to the `trace_dispatch()` method for further tracing in that scope.

dispatch_exception (*frame*, *arg*)

If the debugger should stop at this exception, invokes the `user_exception()` method (which should be overridden in subclasses). Raise a `BdbQuit` exception if the `Bdb.quitting` flag is set (which can be set from `user_exception()`). Return a reference to the `trace_dispatch()` method for further tracing in that scope.

Normally derived classes don't override the following methods, but they may if they want to redefine the definition of stopping and breakpoints.

stop_here (*frame*)

This method checks if the *frame* is somewhere below `botframe` in the call stack. `botframe` is the frame in which debugging started.

break_here (*frame*)

This method checks if there is a breakpoint in the filename and line belonging to *frame* or, at least, in the current function. If the breakpoint is a temporary one, this method deletes it.

break_anywhere (*frame*)

This method checks if there is a breakpoint in the filename of the current frame.

Derived classes should override these methods to gain control over debugger operation.

user_call (*frame*, *argument_list*)

This method is called from `dispatch_call()` when there is the possibility that a break might be necessary anywhere inside the called function.

user_line (*frame*)

This method is called from `dispatch_line()` when either `stop_here()` or `break_here()` yields True.

user_return (*frame*, *return_value*)

This method is called from `dispatch_return()` when `stop_here()` yields True.

user_exception (*frame*, *exc_info*)

This method is called from `dispatch_exception()` when `stop_here()` yields True.

do_clear (*arg*)

Handle how a breakpoint must be removed when it is a temporary one.

This method must be implemented by derived classes.

Derived classes and clients can call the following methods to affect the stepping state.

set_step ()

Stop after one line of code.

set_next (*frame*)

Stop on the next line in or below the given frame.

set_return (*frame*)

Stop when returning from the given frame.

set_until (*frame*)

Stop when the line with the line no greater than the current one is reached or when returning from current frame.

set_trace ([*frame*])

Start debugging from *frame*. If *frame* is not specified, debugging starts from caller's frame.

set_continue ()

Stop only at breakpoints or when finished. If there are no breakpoints, set the system trace function to None.

set_quit ()

Set the `quitting` attribute to True. This raises `BdbQuit` in the next call to one of the `dispatch_*` () methods.

Derived classes and clients can call the following methods to manipulate breakpoints. These methods return a string containing an error message if something went wrong, or None if all is well.

set_break (*filename*, *lineno*, *temporary=0*, *cond*, *funcname*)

Set a new breakpoint. If the *lineno* line doesn't exist for the *filename* passed as argument, return an error message. The *filename* should be in canonical form, as described in the `canonic` () method.

clear_break (*filename*, *lineno*)

Delete the breakpoints in *filename* and *lineno*. If none were set, an error message is returned.

clear_bpbynumber (*arg*)

Delete the breakpoint which has the index *arg* in the `Breakpoint.bpbynumber`. If *arg* is not numeric or out of range, return an error message.

clear_all_file_breaks (*filename*)

Delete all breakpoints in *filename*. If none were set, an error message is returned.

clear_all_breaks ()

Delete all existing breakpoints.

get_bpbynumber (*arg*)

Return a breakpoint specified by the given number. If *arg* is a string, it will be converted to a number. If *arg* is a non-numeric string, if the given breakpoint never existed or has been deleted, a `ValueError` is raised.

3.2 新版功能.

get_break (*filename*, *lineno*)

Check if there is a breakpoint for *lineno* of *filename*.

get_breaks (*filename*, *lineno*)

Return all breakpoints for *lineno* in *filename*, or an empty list if none are set.

get_file_breaks (*filename*)

Return all breakpoints in *filename*, or an empty list if none are set.

get_all_breaks()

Return all breakpoints that are set.

Derived classes and clients can call the following methods to get a data structure representing a stack trace.

get_stack(f, t)

Get a list of records for a frame and all higher (calling) and lower frames, and the size of the higher part.

format_stack_entry(frame_lineno, lprefix=':')

Return a string with information about a stack entry, identified by a (frame, lineno) tuple:

- The canonical form of the filename which contains the frame.
- The function name, or "<lambda>".
- The input arguments.
- The return value.
- The line of code (if it exists).

The following two methods can be called by clients to use a debugger to debug a *statement*, given as a string.

run(cmd, globals=None, locals=None)

Debug a statement executed via the `exec()` function. *globals* defaults to `__main__.__dict__`, *locals* defaults to *globals*.

runeval(expr, globals=None, locals=None)

Debug an expression executed via the `eval()` function. *globals* and *locals* have the same meaning as in `run()`.

runctx(cmd, globals, locals)

For backwards compatibility. Calls the `run()` method.

runcall(func, *args, **kwargs)

Debug a single function call, and return its result.

Finally, the module defines the following functions:

`bdb.checkfuncname(b, frame)`

Check whether we should break here, depending on the way the breakpoint *b* was set.

If it was set via line number, it checks if `b.line` is the same as the one in the frame also passed as argument. If the breakpoint was set via function name, we have to check we are in the right frame (the right function) and if we are in its first executable line.

`bdb.effective(file, line, frame)`

Determine if there is an effective (active) breakpoint at this line of code. Return a tuple of the breakpoint and a boolean that indicates if it is ok to delete a temporary breakpoint. Return `(None, None)` if there is no matching breakpoint.

`bdb.set_trace()`

Start debugging with a `Bdb` instance from caller's frame.

27.2 faulthandler —Dump the Python traceback

3.3 新版功能.

This module contains functions to dump Python tracebacks explicitly, on a fault, after a timeout, or on a user signal. Call `faulthandler.enable()` to install fault handlers for the SIGSEGV, SIGFPE, SIGABRT, SIGBUS, and SIGILL

signals. You can also enable them at startup by setting the `PYTHONFAULTHANDLER` environment variable or by using the `-X faulthandler` command line option.

The fault handler is compatible with system fault handlers like Apport or the Windows fault handler. The module uses an alternative stack for signal handlers if the `sigaltstack()` function is available. This allows it to dump the traceback even on a stack overflow.

The fault handler is called on catastrophic cases and therefore can only use signal-safe functions (e.g. it cannot allocate memory on the heap). Because of this limitation traceback dumping is minimal compared to normal Python tracebacks:

- Only ASCII is supported. The `backslashreplace` error handler is used on encoding.
- Each string is limited to 500 characters.
- Only the filename, the function name and the line number are displayed. (no source code)
- It is limited to 100 frames and 100 threads.
- The order is reversed: the most recent call is shown first.

By default, the Python traceback is written to `sys.stderr`. To see tracebacks, applications must be run in the terminal. A log file can alternatively be passed to `faulthandler.enable()`.

The module is implemented in C, so tracebacks can be dumped on a crash or when Python is deadlocked.

27.2.1 Dumping the traceback

`faulthandler.dump_traceback (file=sys.stderr, all_threads=True)`

Dump the tracebacks of all threads into *file*. If *all_threads* is `False`, dump only the current thread.

在 3.5 版更改: Added support for passing file descriptor to this function.

27.2.2 Fault handler state

`faulthandler.enable (file=sys.stderr, all_threads=True)`

Enable the fault handler: install handlers for the `SIGSEGV`, `SIGFPE`, `SIGABRT`, `SIGBUS` and `SIGILL` signals to dump the Python traceback. If *all_threads* is `True`, produce tracebacks for every running thread. Otherwise, dump only the current thread.

The *file* must be kept open until the fault handler is disabled: see *issue with file descriptors*.

在 3.5 版更改: Added support for passing file descriptor to this function.

在 3.6 版更改: On Windows, a handler for Windows exception is also installed.

`faulthandler.disable ()`

Disable the fault handler: uninstall the signal handlers installed by `enable()`.

`faulthandler.is_enabled ()`

Check if the fault handler is enabled.

27.2.3 Dumping the tracebacks after a timeout

`faulthandler.dump_traceback_later (timeout, repeat=False, file=sys.stderr, exit=False)`

Dump the tracebacks of all threads, after a timeout of *timeout* seconds, or every *timeout* seconds if *repeat* is `True`. If *exit* is `True`, call `_exit()` with status=1 after dumping the tracebacks. (Note `_exit()` exits the process immediately, which means it doesn't do any cleanup like flushing file buffers.) If the function is called twice, the new call replaces previous parameters and resets the timeout. The timer has a sub-second resolution.

The *file* must be kept open until the traceback is dumped or `cancel_dump_traceback_later()` is called: see [issue with file descriptors](#).

This function is implemented using a watchdog thread and therefore is not available if Python is compiled with threads disabled.

在 3.5 版更改: Added support for passing file descriptor to this function.

```
faulthandler.cancel_dump_traceback_later()
Cancel the last call to dump_traceback_later().
```

27.2.4 Dumping the traceback on a user signal

```
faulthandler.register(signum, file=sys.stderr, all_threads=True, chain=False)
```

Register a user signal: install a handler for the *signum* signal to dump the traceback of all threads, or of the current thread if *all_threads* is `False`, into *file*. Call the previous handler if *chain* is `True`.

The *file* must be kept open until the signal is unregistered by `unregister()`: see [issue with file descriptors](#).

Not available on Windows.

在 3.5 版更改: Added support for passing file descriptor to this function.

```
faulthandler.unregister(signum)
```

Unregister a user signal: uninstall the handler of the *signum* signal installed by `register()`. Return `True` if the signal was registered, `False` otherwise.

Not available on Windows.

27.2.5 Issue with file descriptors

`enable()`, `dump_traceback_later()` and `register()` keep the file descriptor of their *file* argument. If the file is closed and its file descriptor is reused by a new file, or if `os.dup2()` is used to replace the file descriptor, the traceback will be written into a different file. Call these functions again each time that the file is replaced.

27.2.6 示例

Example of a segmentation fault on Linux with and without enabling the fault handler:

```
$ python3 -c "import ctypes; ctypes.string_at(0)"
Segmentation fault

$ python3 -q -X faulthandler
>>> import ctypes
>>> ctypes.string_at(0)
Fatal Python error: Segmentation fault

Current thread 0x00007fb899f39700 (most recent call first):
  File "/home/python/cpython/Lib/ctypes/__init__.py", line 486 in string_at
  File "<stdin>", line 1 in <module>
Segmentation fault
```

27.3 pdb —Python 的调试器

源代码: [Lib/pdb.py](#)

The module `pdb` defines an interactive source code debugger for Python programs. It supports setting (conditional) breakpoints and single stepping at the source line level, inspection of stack frames, source code listing, and evaluation of arbitrary Python code in the context of any stack frame. It also supports post-mortem debugging and can be called under program control.

The debugger is extensible – it is actually defined as the class `Pdb`. This is currently undocumented but easily understood by reading the source. The extension interface uses the modules `bdb` and `cmd`.

The debugger's prompt is `(Pdb)`. Typical usage to run a program under control of the debugger is:

```
>>> import pdb
>>> import mymodule
>>> pdb.run('mymodule.test()')
> <string>(0)?()
(Pdb) continue
> <string>(1)?()
(Pdb) continue
NameError: 'spam'
> <string>(1)?()
(Pdb)
```

在 3.3 版更改: Tab-completion via the `readline` module is available for commands and command arguments, e.g. the current global and local names are offered as arguments of the `p` command.

`pdb.py` can also be invoked as a script to debug other scripts. For example:

```
python3 -m pdb myscript.py
```

When invoked as a script, `pdb` will automatically enter post-mortem debugging if the program being debugged exits abnormally. After post-mortem debugging (or after normal exit of the program), `pdb` will restart the program. Automatic restarting preserves `pdb`'s state (such as breakpoints) and in most cases is more useful than quitting the debugger upon program's exit.

3.2 新版功能: `pdb.py` now accepts a `-c` option that executes commands as if given in a `.pdbrc` file, see [Debugger Commands](#).

The typical usage to break into the debugger from a running program is to insert

```
import pdb; pdb.set_trace()
```

at the location you want to break into the debugger. You can then step through the code following this statement, and continue running without the debugger using the `continue` command.

The typical usage to inspect a crashed program is:

```
>>> import pdb
>>> import mymodule
>>> mymodule.test()
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
  File "./mymodule.py", line 4, in test
    test2()
  File "./mymodule.py", line 3, in test2
```

(下页继续)

(续上页)

```

    print(spam)
NameError: spam
>>> pdb.pm()
> ./mymodule.py(3)test2()
-> print(spam)
(Pdb)

```

The module defines the following functions; each enters the debugger in a slightly different way:

`pdb.run(statement, globals=None, locals=None)`

Execute the *statement* (given as a string or a code object) under debugger control. The debugger prompt appears before any code is executed; you can set breakpoints and type *continue*, or you can step through the statement using *step* or *next* (all these commands are explained below). The optional *globals* and *locals* arguments specify the environment in which the code is executed; by default the dictionary of the module `__main__` is used. (See the explanation of the built-in `exec()` or `eval()` functions.)

`pdb.runeval(expression, globals=None, locals=None)`

Evaluate the *expression* (given as a string or a code object) under debugger control. When `runeval()` returns, it returns the value of the expression. Otherwise this function is similar to `run()`.

`pdb.runcall(function, *args, **kwargs)`

Call the *function* (a function or method object, not a string) with the given arguments. When `runcall()` returns, it returns whatever the function call returned. The debugger prompt appears as soon as the function is entered.

`pdb.set_trace()`

Enter the debugger at the calling stack frame. This is useful to hard-code a breakpoint at a given point in a program, even if the code is not otherwise being debugged (e.g. when an assertion fails).

`pdb.post_mortem(traceback=None)`

Enter post-mortem debugging of the given *traceback* object. If no *traceback* is given, it uses the one of the exception that is currently being handled (an exception must be being handled if the default is to be used).

`pdb.pm()`

Enter post-mortem debugging of the traceback found in `sys.last_traceback`.

The `run*` functions and `set_trace()` are aliases for instantiating the `Pdb` class and calling the method of the same name. If you want to access further features, you have to do this yourself:

class `pdb.Pdb(completekey='tab', stdin=None, stdout=None, skip=None, nosigint=False, readrc=True)`

Pdb is the debugger class.

The *completekey*, *stdin* and *stdout* arguments are passed to the underlying `cmd.Cmd` class; see the description there.

The *skip* argument, if given, must be an iterable of glob-style module name patterns. The debugger will not step into frames that originate in a module that matches one of these patterns.¹

By default, `Pdb` sets a handler for the SIGINT signal (which is sent when the user presses `Ctrl-C` on the console) when you give a *continue* command. This allows you to break into the debugger again by pressing `Ctrl-C`. If you want `Pdb` not to touch the SIGINT handler, set *nosigint* to true.

The *readrc* argument defaults to true and controls whether `Pdb` will load `.pdbrc` files from the filesystem.

Example call to enable tracing with *skip*:

```
import pdb; pdb.Pdb(skip=['django.*']).set_trace()
```

3.1 新版功能: The *skip* argument.

3.2 新版功能: The *nosigint* argument. Previously, a SIGINT handler was never set by `Pdb`.

¹ Whether a frame is considered to originate in a certain module is determined by the `__name__` in the frame `globals`.

在 3.6 版更改: The *readrc* argument.

```
run (statement, globals=None, locals=None)
runeval (expression, globals=None, locals=None)
runcall (function, *args, **kws)
set_trace ()
```

See the documentation for the functions explained above.

27.3.1 Debugger Commands

The commands recognized by the debugger are listed below. Most commands can be abbreviated to one or two letters as indicated; e.g. `h(elp)` means that either `h` or `help` can be used to enter the help command (but not `he` or `hel`, nor `H` or `Help` or `HELP`). Arguments to commands must be separated by whitespace (spaces or tabs). Optional arguments are enclosed in square brackets (`[]`) in the command syntax; the square brackets must not be typed. Alternatives in the command syntax are separated by a vertical bar (`|`).

Entering a blank line repeats the last command entered. Exception: if the last command was a *list* command, the next 11 lines are listed.

Commands that the debugger doesn't recognize are assumed to be Python statements and are executed in the context of the program being debugged. Python statements can also be prefixed with an exclamation point (`!`). This is a powerful way to inspect the program being debugged; it is even possible to change a variable or call a function. When an exception occurs in such a statement, the exception name is printed but the debugger's state is not changed.

The debugger supports *aliases*. Aliases can have parameters which allows one a certain level of adaptability to the context under examination.

Multiple commands may be entered on a single line, separated by `;;`. (A single `;` is not used as it is the separator for multiple commands in a line that is passed to the Python parser.) No intelligence is applied to separating the commands; the input is split at the first `;;` pair, even if it is in the middle of a quoted string.

If a file `.pdbrc` exists in the user's home directory or in the current directory, it is read in and executed as if it had been typed at the debugger prompt. This is particularly useful for aliases. If both files exist, the one in the home directory is read first and aliases defined there can be overridden by the local file.

在 3.2 版更改: `.pdbrc` can now contain commands that continue debugging, such as *continue* or *next*. Previously, these commands had no effect.

h(elp) [*command*]

Without argument, print the list of available commands. With a *command* as argument, print help about that command. `help pdb` displays the full documentation (the docstring of the *pdb* module). Since the *command* argument must be an identifier, `help exec` must be entered to get help on the `!` command.

w(here)

Print a stack trace, with the most recent frame at the bottom. An arrow indicates the current frame, which determines the context of most commands.

d(own) [*count*]

Move the current frame *count* (default one) levels down in the stack trace (to a newer frame).

u(p) [*count*]

Move the current frame *count* (default one) levels up in the stack trace (to an older frame).

b(reak) [(*filename*:*lineno* | *function*) [, *condition*]]

With a *lineno* argument, set a break there in the current file. With a *function* argument, set a break at the first executable statement within that function. The line number may be prefixed with a filename and a colon, to specify a breakpoint in another file (probably one that hasn't been loaded yet). The file is searched on *sys.path*. Note that each breakpoint is assigned a number to which all the other breakpoint commands refer.

If a second argument is present, it is an expression which must evaluate to true before the breakpoint is honored.

Without argument, list all breaks, including for each breakpoint, the number of times that breakpoint has been hit, the current ignore count, and the associated condition if any.

tbreak [(filename:]lineno | function) [, condition]

Temporary breakpoint, which is removed automatically when it is first hit. The arguments are the same as for *break*.

cl(ear) [filename:]lineno | bnumber [bnumber ...]

With a *filename:lineno* argument, clear all the breakpoints at this line. With a space separated list of breakpoint numbers, clear those breakpoints. Without argument, clear all breaks (but first ask confirmation).

disable [bnumber [bnumber ...]]

Disable the breakpoints given as a space separated list of breakpoint numbers. Disabling a breakpoint means it cannot cause the program to stop execution, but unlike clearing a breakpoint, it remains in the list of breakpoints and can be (re-)enabled.

enable [bnumber [bnumber ...]]

Enable the breakpoints specified.

ignore bnumber [count]

Set the ignore count for the given breakpoint number. If count is omitted, the ignore count is set to 0. A breakpoint becomes active when the ignore count is zero. When non-zero, the count is decremented each time the breakpoint is reached and the breakpoint is not disabled and any associated condition evaluates to true.

condition bnumber [condition]

Set a new *condition* for the breakpoint, an expression which must evaluate to true before the breakpoint is honored. If *condition* is absent, any existing condition is removed; i.e., the breakpoint is made unconditional.

commands [bnumber]

Specify a list of commands for breakpoint number *bnumber*. The commands themselves appear on the following lines. Type a line containing just *end* to terminate the commands. An example:

```
(Pdb) commands 1
(com) p some_variable
(com) end
(Pdb)
```

To remove all commands from a breakpoint, type *commands* and follow it immediately with *end*; that is, give no commands.

With no *bnumber* argument, *commands* refers to the last breakpoint set.

You can use breakpoint commands to start your program up again. Simply use the *continue* command, or *step*, or any other command that resumes execution.

Specifying any command resuming execution (currently *continue*, *step*, *next*, *return*, *jump*, *quit* and their abbreviations) terminates the command list (as if that command was immediately followed by *end*). This is because any time you resume execution (even with a simple *next* or *step*), you may encounter another breakpoint—which could have its own command list, leading to ambiguities about which list to execute.

If you use the ‘*silent*’ command in the command list, the usual message about stopping at a breakpoint is not printed. This may be desirable for breakpoints that are to print a specific message and then *continue*. If none of the other commands print anything, you see no sign that the breakpoint was reached.

s(step)

Execute the current line, stop at the first possible occasion (either in a function that is called or on the next line in the current function).

n(ext)

Continue execution until the next line in the current function is reached or it returns. (The difference between *next* and *step* is that *step* stops inside a called function, while *next* executes called functions at (nearly) full speed, only stopping at the next line in the current function.)

unt(il) [*lineno*]

Without argument, continue execution until the line with a number greater than the current one is reached.

With a line number, continue execution until a line with a number greater or equal to that is reached. In both cases, also stop when the current frame returns.

在 3.2 版更改: Allow giving an explicit line number.

r(eturn)

Continue execution until the current function returns.

c(ontinue)

Continue execution, only stop when a breakpoint is encountered.

j(ump) *lineno*

Set the next line that will be executed. Only available in the bottom-most frame. This lets you jump back and execute code again, or jump forward to skip code that you don't want to run.

It should be noted that not all jumps are allowed – for instance it is not possible to jump into the middle of a *for* loop or out of a *finally* clause.

l(list) [*first*[, *last*]]

List source code for the current file. Without arguments, list 11 lines around the current line or continue the previous listing. With *.* as argument, list 11 lines around the current line. With one argument, list 11 lines around at that line. With two arguments, list the given range; if the second argument is less than the first, it is interpreted as a count.

The current line in the current frame is indicated by *->*. If an exception is being debugged, the line where the exception was originally raised or propagated is indicated by *>>*, if it differs from the current line.

3.2 新版功能: The *>>* marker.

ll | *longlist*

List all source code for the current function or frame. Interesting lines are marked as for *list*.

3.2 新版功能.

a(rgs)

Print the argument list of the current function.

p *expression*

Evaluate the *expression* in the current context and print its value.

注解: *print()* can also be used, but is not a debugger command —this executes the Python *print()* function.

pp *expression*

Like the *p* command, except the value of the expression is pretty-printed using the *pprint* module.

whatis *expression*

Print the type of the *expression*.

source *expression*

Try to get source code for the given object and display it.

3.2 新版功能.

display [expression]

Display the value of the expression if it changed, each time execution stops in the current frame.

Without expression, list all display expressions for the current frame.

3.2 新版功能.

undisplay [expression]

Do not display the expression any more in the current frame. Without expression, clear all display expressions for the current frame.

3.2 新版功能.

interact

Start an interactive interpreter (using the [code](#) module) whose global namespace contains all the (global and local) names found in the current scope.

3.2 新版功能.

alias [name [command]]

Create an alias called *name* that executes *command*. The command must *not* be enclosed in quotes. Replaceable parameters can be indicated by %1, %2, and so on, while %* is replaced by all the parameters. If no command is given, the current alias for *name* is shown. If no arguments are given, all aliases are listed.

Aliases may be nested and can contain anything that can be legally typed at the pdb prompt. Note that internal pdb commands *can* be overridden by aliases. Such a command is then hidden until the alias is removed. Aliasing is recursively applied to the first word of the command line; all other words in the line are left alone.

As an example, here are two useful aliases (especially when placed in the `.pdbrc` file):

```
# Print instance variables (usage "pi classInst")
alias pi for k in %1.__dict__.keys(): print("%1.", k, "=", %1.__dict__[k])
# Print instance variables in self
alias ps pi self
```

unalias name

Delete the specified alias.

! statement

Execute the (one-line) *statement* in the context of the current stack frame. The exclamation point can be omitted unless the first word of the statement resembles a debugger command. To set a global variable, you can prefix the assignment command with a `global` statement on the same line, e.g.:

```
(Pdb) global list_options; list_options = ['-l']
(Pdb)
```

run [args ...]**restart** [args ...]

Restart the debugged Python program. If an argument is supplied, it is split with [shlex](#) and the result is used as the new `sys.argv`. History, breakpoints, actions and debugger options are preserved. *restart* is an alias for *run*.

q(uit)

Quit from the debugger. The program being executed is aborted.

27.4 The Python Profilers

Source code: [Lib/profile.py](#) and [Lib/pstats.py](#)

27.4.1 Introduction to the profilers

`cProfile` and `profile` provide *deterministic profiling* of Python programs. A *profile* is a set of statistics that describes how often and for how long various parts of the program executed. These statistics can be formatted into reports via the `pstats` module.

The Python standard library provides two different implementations of the same profiling interface:

1. `cProfile` is recommended for most users; it's a C extension with reasonable overhead that makes it suitable for profiling long-running programs. Based on `lsprof`, contributed by Brett Rosen and Ted Czottter.
2. `profile`, a pure Python module whose interface is imitated by `cProfile`, but which adds significant overhead to profiled programs. If you're trying to extend the profiler in some way, the task might be easier with this module. Originally designed and written by Jim Roskind.

注解： The profiler modules are designed to provide an execution profile for a given program, not for benchmarking purposes (for that, there is `timeit` for reasonably accurate results). This particularly applies to benchmarking Python code against C code: the profilers introduce overhead for Python code, but not for C-level functions, and so the C code would seem faster than any Python one.

27.4.2 Instant User's Manual

This section is provided for users that “don't want to read the manual.” It provides a very brief overview, and allows a user to rapidly perform profiling on an existing application.

To profile a function that takes a single argument, you can do:

```
import cProfile
import re
cProfile.run('re.compile("foo|bar")')
```

(Use `profile` instead of `cProfile` if the latter is not available on your system.)

The above action would run `re.compile()` and print profile results like the following:

```
197 function calls (192 primitive calls) in 0.002 seconds

Ordered by: standard name

ncalls  tottime  percall  cumtime  percall filename:lineno(function)
   1      0.000   0.000    0.001    0.001 <string>:1(<module>)
   1      0.000   0.000    0.001    0.001 re.py:212(compile)
   1      0.000   0.000    0.001    0.001 re.py:268(_compile)
   1      0.000   0.000    0.000    0.000 sre_compile.py:172(_compile_charset)
   1      0.000   0.000    0.000    0.000 sre_compile.py:201(_optimize_charset)
   4      0.000   0.000    0.000    0.000 sre_compile.py:25(_identityfunction)
  3/1      0.000   0.000    0.000    0.000 sre_compile.py:33(_compile)
```

The first line indicates that 197 calls were monitored. Of those calls, 192 were *primitive*, meaning that the call was not induced via recursion. The next line: `Ordered by: standard name`, indicates that the text string in the far right column was used to sort the output. The column headings include:

ncalls for the number of calls.

tottime for the total time spent in the given function (and excluding time made in calls to sub-functions)

percall is the quotient of `tottime` divided by `ncalls`

cumtime is the cumulative time spent in this and all subfunctions (from invocation till exit). This figure is accurate *even* for recursive functions.

percall is the quotient of `cumtime` divided by primitive calls

filename:lineno(function) provides the respective data of each function

When there are two numbers in the first column (for example 3/1), it means that the function recursed. The second value is the number of primitive calls and the former is the total number of calls. Note that when the function does not recurse, these two values are the same, and only the single figure is printed.

Instead of printing the output at the end of the profile run, you can save the results to a file by specifying a filename to the `run()` function:

```
import cProfile
import re
cProfile.run('re.compile("foo|bar")', 'restats')
```

The `pstats.Stats` class reads profile results from a file and formats them in various ways.

The file `cProfile` can also be invoked as a script to profile another script. For example:

```
python -m cProfile [-o output_file] [-s sort_order] myscript.py
```

`-o` writes the profile results to a file instead of to stdout

`-s` specifies one of the `sort_stats()` sort values to sort the output by. This only applies when `-o` is not supplied.

The `pstats` module's `Stats` class has a variety of methods for manipulating and printing the data saved into a profile results file:

```
import pstats
p = pstats.Stats('restats')
p.strip_dirs().sort_stats(-1).print_stats()
```

The `strip_dirs()` method removed the extraneous path from all the module names. The `sort_stats()` method sorted all the entries according to the standard module/line/name string that is printed. The `print_stats()` method printed out all the statistics. You might try the following sort calls:

```
p.sort_stats('name')
p.print_stats()
```

The first call will actually sort the list by function name, and the second call will print out the statistics. The following are some interesting calls to experiment with:

```
p.sort_stats('cumulative').print_stats(10)
```

This sorts the profile by cumulative time in a function, and then only prints the ten most significant lines. If you want to understand what algorithms are taking time, the above line is what you would use.

If you were looking to see what functions were looping a lot, and taking a lot of time, you would do:

```
p.sort_stats('time').print_stats(10)
```

to sort according to time spent within each function, and then print the statistics for the top ten functions.

You might also try:

```
p.sort_stats('file').print_stats('__init__')
```

This will sort all the statistics by file name, and then print out statistics for only the class init methods (since they are spelled with `__init__` in them). As one final example, you could try:

```
p.sort_stats('time', 'cumulative').print_stats(.5, 'init')
```

This line sorts statistics with a primary key of time, and a secondary key of cumulative time, and then prints out some of the statistics. To be specific, the list is first culled down to 50% (re: `.5`) of its original size, then only lines containing `init` are maintained, and that sub-sub-list is printed.

If you wondered what functions called the above functions, you could now (`p` is still sorted according to the last criteria) do:

```
p.print_callers(.5, 'init')
```

and you would get a list of callers for each of the listed functions.

If you want more functionality, you're going to have to read the manual, or guess what the following functions do:

```
p.print_callees()
p.add('restats')
```

Invoked as a script, the `pstats` module is a statistics browser for reading and examining profile dumps. It has a simple line-oriented interface (implemented using `cmd`) and interactive help.

27.4.3 profile and cProfile Module Reference

Both the `profile` and `cProfile` modules provide the following functions:

`profile.run(command, filename=None, sort=-1)`

This function takes a single argument that can be passed to the `exec()` function, and an optional file name. In all cases this routine executes:

```
exec(command, __main__.__dict__, __main__.__dict__)
```

and gathers profiling statistics from the execution. If no file name is present, then this function automatically creates a `Stats` instance and prints a simple profiling report. If the sort value is specified, it is passed to this `Stats` instance to control how the results are sorted.

`profile.runcx(command, globals, locals, filename=None, sort=-1)`

This function is similar to `run()`, with added arguments to supply the globals and locals dictionaries for the `command` string. This routine executes:

```
exec(command, globals, locals)
```

and gathers profiling statistics as in the `run()` function above.

class `profile.Profile(timer=None, timeunit=0.0, subcalls=True, builtins=True)`

This class is normally only used if more precise control over profiling is needed than what the `cProfile.run()` function provides.

A custom timer can be supplied for measuring how long code takes to run via the `timer` argument. This must be a function that returns a single number representing the current time. If the number is an integer, the `timeunit` specifies a multiplier that specifies the duration of each unit of time. For example, if the timer returns times measured in thousands of seconds, the time unit would be `.001`.

Directly using the `Profile` class allows formatting profile results without writing the profile data to a file:

```
import cProfile, pstats, io
pr = cProfile.Profile()
pr.enable()
# ... do something ...
pr.disable()
s = io.StringIO()
sortby = 'cumulative'
ps = pstats.Stats(pr, stream=s).sort_stats(sortby)
ps.print_stats()
print(s.getvalue())
```

enable()

Start collecting profiling data.

disable()

Stop collecting profiling data.

create_stats()

Stop collecting profiling data and record the results internally as the current profile.

print_stats(sort=-1)

Create a *Stats* object based on the current profile and print the results to stdout.

dump_stats(filename)

Write the results of the current profile to *filename*.

run(cmd)

Profile the cmd via *exec()*.

runctx(cmd, globals, locals)

Profile the cmd via *exec()* with the specified global and local environment.

runcall(func, *args, **kwargs)

Profile func(*args, **kwargs)

Note that profiling will only work if the called command/function actually returns. If the interpreter is terminated (e.g. via a *sys.exit()* call during the called command/function execution) no profiling results will be printed.

27.4.4 The Stats Class

Analysis of the profiler data is done using the *Stats* class.

class pstats.Stats(*filenames or profile, stream=sys.stdout)

This class constructor creates an instance of a “statistics object” from a *filename* (or list of filenames) or from a *Profile* instance. Output will be printed to the stream specified by *stream*.

The file selected by the above constructor must have been created by the corresponding version of *profile* or *cProfile*. To be specific, there is *no* file compatibility guaranteed with future versions of this profiler, and there is no compatibility with files produced by other profilers, or the same profiler run on a different operating system. If several files are provided, all the statistics for identical functions will be coalesced, so that an overall view of several processes can be considered in a single report. If additional files need to be combined with data in an existing *Stats* object, the *add()* method can be used.

Instead of reading the profile data from a file, a *cProfile.Profile* or *profile.Profile* object can be used as the profile data source.

Stats objects have the following methods:

strip_dirs()

This method for the *Stats* class removes all leading path information from file names. It is very useful in

reducing the size of the printout to fit within (close to) 80 columns. This method modifies the object, and the stripped information is lost. After performing a strip operation, the object is considered to have its entries in a “random” order, as it was just after object initialization and loading. If `strip_dirs()` causes two function names to be indistinguishable (they are on the same line of the same filename, and have the same function name), then the statistics for these two entries are accumulated into a single entry.

add (*filenames)

This method of the `Stats` class accumulates additional profiling information into the current profiling object. Its arguments should refer to filenames created by the corresponding version of `profile.run()` or `cProfile.run()`. Statistics for identically named (re: file, line, name) functions are automatically accumulated into single function statistics.

dump_stats (filename)

Save the data loaded into the `Stats` object to a file named *filename*. The file is created if it does not exist, and is overwritten if it already exists. This is equivalent to the method of the same name on the `profile.Profile` and `cProfile.Profile` classes.

sort_stats (*keys)

This method modifies the `Stats` object by sorting it according to the supplied criteria. The argument is typically a string identifying the basis of a sort (example: 'time' or 'name').

When more than one key is provided, then additional keys are used as secondary criteria when there is equality in all keys selected before them. For example, `sort_stats('name', 'file')` will sort all the entries according to their function name, and resolve all ties (identical function names) by sorting by file name.

Abbreviations can be used for any key names, as long as the abbreviation is unambiguous. The following are the keys currently defined:

Valid Arg	意义
'calls'	call count
'cumulative'	cumulative time
'cumtime'	cumulative time
'file'	file name
'filename'	file name
'module'	file name
'ncalls'	call count
'pcalls'	primitive call count
'line'	line number
'name'	function name
'nfl'	name/file/line
'stdname'	standard name
'time'	internal time
'tottime'	internal time

Note that all sorts on statistics are in descending order (placing most time consuming items first), where as name, file, and line number searches are in ascending order (alphabetical). The subtle distinction between 'nfl' and 'stdname' is that the standard name is a sort of the name as printed, which means that the embedded line numbers get compared in an odd way. For example, lines 3, 20, and 40 would (if the file names were the same) appear in the string order 20, 3 and 40. In contrast, 'nfl' does a numeric compare of the line numbers. In fact, `sort_stats('nfl')` is the same as `sort_stats('name', 'file', 'line')`.

For backward-compatibility reasons, the numeric arguments -1, 0, 1, and 2 are permitted. They are interpreted as 'stdname', 'calls', 'time', and 'cumulative' respectively. If this old style format (numeric) is used, only one sort key (the numeric key) will be used, and additional arguments will be silently ignored.

reverse_order()

This method for the *Stats* class reverses the ordering of the basic list within the object. Note that by default ascending vs descending order is properly selected based on the sort key of choice.

print_stats(*restrictions)

This method for the *Stats* class prints out a report as described in the *profile.run()* definition.

The order of the printing is based on the last *sort_stats()* operation done on the object (subject to caveats in *add()* and *strip_dirs()*).

The arguments provided (if any) can be used to limit the list down to the significant entries. Initially, the list is taken to be the complete set of profiled functions. Each restriction is either an integer (to select a count of lines), or a decimal fraction between 0.0 and 1.0 inclusive (to select a percentage of lines), or a string that will interpreted as a regular expression (to pattern match the standard name that is printed). If several restrictions are provided, then they are applied sequentially. For example:

```
print_stats(.1, 'foo:')
```

would first limit the printing to first 10% of list, and then only print functions that were part of filename *.foo:*. In contrast, the command:

```
print_stats('foo:', .1)
```

would limit the list to all functions having file names *.foo:*, and then proceed to only print the first 10% of them.

print_callers(*restrictions)

This method for the *Stats* class prints a list of all functions that called each function in the profiled database. The ordering is identical to that provided by *print_stats()*, and the definition of the restricting argument is also identical. Each caller is reported on its own line. The format differs slightly depending on the profiler that produced the stats:

- With *profile*, a number is shown in parentheses after each caller to show how many times this specific call was made. For convenience, a second non-parenthesized number repeats the cumulative time spent in the function at the right.
- With *cProfile*, each caller is preceded by three numbers: the number of times this specific call was made, and the total and cumulative times spent in the current function while it was invoked by this specific caller.

print_callees(*restrictions)

This method for the *Stats* class prints a list of all function that were called by the indicated function. Aside from this reversal of direction of calls (re: called vs was called by), the arguments and ordering are identical to the *print_callers()* method.

27.4.5 What Is Deterministic Profiling?

Deterministic profiling is meant to reflect the fact that all *function call*, *function return*, and *exception* events are monitored, and precise timings are made for the intervals between these events (during which time the user's code is executing). In contrast, *statistical profiling* (which is not done by this module) randomly samples the effective instruction pointer, and deduces where time is being spent. The latter technique traditionally involves less overhead (as the code does not need to be instrumented), but provides only relative indications of where time is being spent.

In Python, since there is an interpreter active during execution, the presence of instrumented code is not required to do deterministic profiling. Python automatically provides a *hook* (optional callback) for each event. In addition, the interpreted nature of Python tends to add so much overhead to execution, that deterministic profiling tends to only add small processing overhead in typical applications. The result is that deterministic profiling is not that expensive, yet provides extensive run time statistics about the execution of a Python program.

Call count statistics can be used to identify bugs in code (surprising counts), and to identify possible inline-expansion points (high call counts). Internal time statistics can be used to identify “hot loops” that should be carefully optimized. Cumulative time statistics should be used to identify high level errors in the selection of algorithms. Note that the unusual handling of cumulative times in this profiler allows statistics for recursive implementations of algorithms to be directly compared to iterative implementations.

27.4.6 Limitations

One limitation has to do with accuracy of timing information. There is a fundamental problem with deterministic profilers involving accuracy. The most obvious restriction is that the underlying “clock” is only ticking at a rate (typically) of about .001 seconds. Hence no measurements will be more accurate than the underlying clock. If enough measurements are taken, then the “error” will tend to average out. Unfortunately, removing this first error induces a second source of error.

The second problem is that it “takes a while” from when an event is dispatched until the profiler’s call to get the time actually *gets* the state of the clock. Similarly, there is a certain lag when exiting the profiler event handler from the time that the clock’s value was obtained (and then squirreled away), until the user’s code is once again executing. As a result, functions that are called many times, or call many functions, will typically accumulate this error. The error that accumulates in this fashion is typically less than the accuracy of the clock (less than one clock tick), but it *can* accumulate and become very significant.

The problem is more important with `profile` than with the lower-overhead `cProfile`. For this reason, `profile` provides a means of calibrating itself for a given platform so that this error can be probabilistically (on the average) removed. After the profiler is calibrated, it will be more accurate (in a least square sense), but it will sometimes produce negative numbers (when call counts are exceptionally low, and the gods of probability work against you :-).) Do *not* be alarmed by negative numbers in the profile. They should *only* appear if you have calibrated your profiler, and the results are actually better than without calibration.

27.4.7 Calibration

The profiler of the `profile` module subtracts a constant from each event handling time to compensate for the overhead of calling the time function, and socking away the results. By default, the constant is 0. The following procedure can be used to obtain a better constant for a given platform (see [Limitations](#)).

```
import profile
pr = profile.Profile()
for i in range(5):
    print(pr.calibrate(10000))
```

The method executes the number of Python calls given by the argument, directly and again under the profiler, measuring the time for both. It then computes the hidden overhead per profiler event, and returns that as a float. For example, on a 1.8Ghz Intel Core i5 running Mac OS X, and using Python’s `time.clock()` as the timer, the magical number is about 4.04e-6.

The object of this exercise is to get a fairly consistent result. If your computer is *very* fast, or your timer function has poor resolution, you might have to pass 100000, or even 1000000, to get consistent results.

When you have a consistent answer, there are three ways you can use it:

```
import profile

# 1. Apply computed bias to all Profile instances created hereafter.
profile.Profile.bias = your_computed_bias
```

(下页继续)

(续上页)

```
# 2. Apply computed bias to a specific Profile instance.
pr = profile.Profile()
pr.bias = your_computed_bias

# 3. Specify computed bias in instance constructor.
pr = profile.Profile(bias=your_computed_bias)
```

If you have a choice, you are better off choosing a smaller constant, and then your results will “less often” show up as negative in profile statistics.

27.4.8 Using a custom timer

If you want to change how current time is determined (for example, to force use of wall-clock time or elapsed process time), pass the timing function you want to the `Profile` class constructor:

```
pr = profile.Profile(your_time_func)
```

The resulting profiler will then call `your_time_func`. Depending on whether you are using `profile.Profile` or `cProfile.Profile`, `your_time_func`'s return value will be interpreted differently:

`profile.Profile` `your_time_func` should return a single number, or a list of numbers whose sum is the current time (like what `os.times()` returns). If the function returns a single time number, or the list of returned numbers has length 2, then you will get an especially fast version of the dispatch routine.

Be warned that you should calibrate the profiler class for the timer function that you choose (see [Calibration](#)). For most machines, a timer that returns a lone integer value will provide the best results in terms of low overhead during profiling. (`os.times()` is *pretty* bad, as it returns a tuple of floating point values). If you want to substitute a better timer in the cleanest fashion, derive a class and hardwire a replacement dispatch method that best handles your timer call, along with the appropriate calibration constant.

`cProfile.Profile` `your_time_func` should return a single number. If it returns integers, you can also invoke the class constructor with a second argument specifying the real duration of one unit of time. For example, if `your_integer_time_func` returns times measured in thousands of seconds, you would construct the `Profile` instance as follows:

```
pr = cProfile.Profile(your_integer_time_func, 0.001)
```

As the `cProfile.Profile` class cannot be calibrated, custom timer functions should be used with care and should be as fast as possible. For the best results with a custom timer, it might be necessary to hard-code it in the C source of the internal `_lsprof` module.

Python 3.3 adds several new functions in `time` that can be used to make precise measurements of process or wall-clock time. For example, see `time.perf_counter()`.

27.5 timeit — 测量小代码片段的执行时间

源码: [Lib/timeit.py](#)

该模块提供了一种简单的方法来计算一小段 Python 代码的耗时。它有命令行界面 以及一个可调用 方法。它避免了许多用于测量执行时间的常见陷阱。另见 Tim Peters 对 O'Reilly 出版的 *Python Cookbook* 中“算法”章节的介绍。

27.5.1 基本示例

以下示例显示了如何使用命令行界面 来比较三个不同的表达式：

```
$ python3 -m timeit '"-".join(str(n) for n in range(100))'
10000 loops, best of 3: 30.2 usec per loop
$ python3 -m timeit '"-".join([str(n) for n in range(100)])'
10000 loops, best of 3: 27.5 usec per loop
$ python3 -m timeit '"-".join(map(str, range(100)))'
10000 loops, best of 3: 23.2 usec per loop
```

这可以通过Python 接口 实现

```
>>> import timeit
>>> timeit.timeit('"-".join(str(n) for n in range(100))', number=10000)
0.3018611848820001
>>> timeit.timeit('"-".join([str(n) for n in range(100)])', number=10000)
0.2727368790656328
>>> timeit.timeit('"-".join(map(str, range(100)))', number=10000)
0.23702679807320237
```

Note however that `timeit` will automatically determine the number of repetitions only when the command-line interface is used. In the 示例 section you can find more advanced examples.

27.5.2 Python 接口

该模块定义了三个便利函数和一个公共类：

`timeit.timeit(stmt='pass', setup='pass', timer=<default timer>, number=1000000, globals=None)`

使用给定语句、`setup` 代码和 `timer` 函数创建一个 `Timer` 实例，并执行 `number` 次其 `timeit()` 方法。可选的 `globals` 参数指定用于执行代码的命名空间。

在 3.5 版更改：添加可选参数 `globals`。

`timeit.repeat(stmt='pass', setup='pass', timer=<default timer>, repeat=3, number=1000000, globals=None)`

使用给定语句、`setup` 代码和 `timer` 函数创建一个 `Timer` 实例，并使用给定的 `repeat` 计数和 `number` 执行运行其 `repeat()` 方法。可选的 `globals` 参数指定用于执行代码的命名空间。

在 3.5 版更改：添加可选参数 `globals`。

`timeit.default_timer()`

默认的计时器，总是 `time.perf_counter()`。

在 3.3 版更改：`time.perf_counter()` 现在是默认计时器。

`class timeit.Timer(stmt='pass', setup='pass', timer=<timer function>, globals=None)`

用于小代码片段的计数执行速度的类。

构造函数接受一个将计时的语句、一个用于设置的附加语句和一个定时器函数。两个语句都默认为 `'pass'`；计时器函数与平台有关（请参阅模块文档字符串）。`stmt` 和 `setup` 也可能包含多个以分号分隔的语句，只要它们不包含多行字符串文字即可。该语句默认在 `timeit` 的命名空间内执行；可以通过将命名空间传递给 `globals` 来控制此行为。

要测量第一个语句的执行时间，请使用 `timeit()` 方法。`repeat()` 和 `autorange()` 方法是方便的方法来调用 `timeit()` 多次。

`setup` 的执行时间从总体计时执行中排除。

`stmt` 和 `setup` 参数也可以使用不带参数的可调用对象。这将在一个计时器函数中嵌入对它们的调用，然后由 `timeit()` 执行。请注意，由于额外的函数调用，在这种情况下，计时开销会略大一些。

在 3.5 版更改：添加可选参数 `globals`。

timeit (`number=1000000`)

执行 `number` 次主要语句。这将执行一次 `setup` 语句，然后返回执行主语句多次所需的时间，以秒为单位测量为浮点数。参数是通过循环的次数，默认为一百万。要使用的主语句、`setup` 语句和 `timer` 函数将传递给构造函数。

注解： By default, `timeit()` temporarily turns off *garbage collection* during the timing. The advantage of this approach is that it makes independent timings more comparable. This disadvantage is that GC may be an important component of the performance of the function being measured. If so, GC can be re-enabled as the first statement in the `setup` string. For example:

```
timeit.Timer('for i in range(10): oct(i)', 'gc.enable()').timeit()
```

autorange (`callback=None`)

自动决定调用多少次 `timeit()`。

This is a convenience function that calls `timeit()` repeatedly so that the total time ≥ 0.2 second, returning the eventual (number of loops, time taken for that number of loops). It calls `timeit()` with `number` set to successive powers of ten (10, 100, 1000, ...) up to a maximum of one billion, until the time taken is at least 0.2 second, or the maximum is reached.

如果给出 `callback` 并且不是 `None`，则在每次试验后将使用两个参数调用它：`callback(number, time_taken)`。

3.6 新版功能。

repeat (`repeat=3, number=1000000`)

调用 `timeit()` 几次。

这是一个方便的函数，它反复调用 `timeit()`，返回结果列表。第一个参数指定调用 `timeit()` 的次数。第二个参数指定 `timeit()` 的 `number` 参数。

注解： 从结果向量计算并报告平均值和标准差这些是很诱人的。但是，这不是很有用。在典型情况下，最低值给出了机器运行给定代码段的速度下限；结果向量中较高的值通常不是由 Python 的速度变化引起的，而是由于其他过程干扰你的计时准确性。所以结果的 `min()` 可能是你应该感兴趣的唯一数字。之后，你应该看看整个向量并应用常识而不是统计。

print_exc (`file=None`)

帮助程序从计时代码中打印回溯。

典型使用：

```
t = Timer(...)          # outside the try/except
try:
    t.timeit(...)        # or t.repeat(...)
except Exception:
    t.print_exc()
```

与标准回溯相比，优势在于将显示已编译模板中的源行。可选的 `file` 参数指向发送回溯的位置；它默认为 `sys.stderr`。

27.5.3 命令行界面

从命令行调用程序时，使用以下表单：

```
python -m timeit [-n N] [-r N] [-u U] [-s S] [-t] [-c] [-h] [statement ...]
```

如果了解以下选项：

- n N, --number=N**
执行‘语句’多少次
- r N, --repeat=N**
how many times to repeat the timer (default 3)
- s S, --setup=S**
最初要执行一次的语句（默认为 pass）
- p, --process**
测量进程时间，而不是 wallclock 时间，使用 `time.process_time()` 而不是 `time.perf_counter()`，这是默认值
3.3 新版功能.
- t, --time**
use `time.time()` (deprecated)
- u, --unit=U**
specify a time unit for timer output; can select usec, msec, or sec
3.5 新版功能.
- c, --clock**
use `time.clock()` (deprecated)
- v, --verbose**
打印原始计时结果；重复更多位数精度
- h, --help**
打印一条简短的使用信息并退出

可以通过将每一行指定为单独的语句参数来给出多行语句；通过在引号中包含参数并使用前导空格可以缩进行。多个 `-s` 选项的处理方式相似。

如果 `-n` 未给出，则通过尝试 10 的连续幂次来计算合适数量的循环，直到总时间至少为 0.2 秒。

`default_timer()` measurements can be affected by other programs running on the same machine, so the best thing to do when accurate timing is necessary is to repeat the timing a few times and use the best time. The `-r` option is good for this; the default of 3 repetitions is probably enough in most cases. You can use `time.process_time()` to measure CPU time.

注解： 执行 pass 语句会产生一定的基线开销。这里的代码不会试图隐藏它，但你应该知道它。可以通过不带参数调用程序来测量基线开销，并且 Python 版本之间可能会有所不同。

27.5.4 示例

可以提供在一个在开头只执行一次的 setup 语句：

```
$ python -m timeit -s 'text = "sample string"; char = "g"' 'char in text'
10000000 loops, best of 3: 0.0877 usec per loop
$ python -m timeit -s 'text = "sample string"; char = "g"' 'text.find(char)'
1000000 loops, best of 3: 0.342 usec per loop
```

```
>>> import timeit
>>> timeit.timeit('char in text', setup='text = "sample string"; char = "g"')
0.41440500499993504
>>> timeit.timeit('text.find(char)', setup='text = "sample string"; char = "g"')
1.7246671520006203
```

使用 `Timer` 类及其方法可以完成同样的操作:

```
>>> import timeit
>>> t = timeit.Timer('char in text', setup='text = "sample string"; char = "g"')
>>> t.timeit()
0.3955516149999312
>>> t.repeat()
[0.40193588800002544, 0.3960157959998014, 0.39594301399984033]
```

以下示例显示如何计算包含多行的表达式。在这里我们对比使用 `hasattr()` 与 `try/except` 的开销来测试缺失与提供对象属性:

```
$ python -m timeit 'try: ' ' str.__bool__' 'except AttributeError: ' ' pass'
100000 loops, best of 3: 15.7 usec per loop
$ python -m timeit 'if hasattr(str, "__bool__"): pass'
100000 loops, best of 3: 4.26 usec per loop

$ python -m timeit 'try: ' ' int.__bool__' 'except AttributeError: ' ' pass'
1000000 loops, best of 3: 1.43 usec per loop
$ python -m timeit 'if hasattr(int, "__bool__"): pass'
100000 loops, best of 3: 2.23 usec per loop
```

```
>>> import timeit
>>> # attribute is missing
>>> s = """\
... try:
...     str.__bool__
... except AttributeError:
...     pass
... """
>>> timeit.timeit(stmt=s, number=100000)
0.9138244460009446
>>> s = "if hasattr(str, '__bool__'): pass"
>>> timeit.timeit(stmt=s, number=100000)
0.5829014980008651
>>>
>>> # attribute is present
>>> s = """\
... try:
...     int.__bool__
... except AttributeError:
...     pass
... """
>>> timeit.timeit(stmt=s, number=100000)
0.04215312199994514
```

(下页继续)

(续上页)

```
>>> s = "if hasattr(int, '__bool__'): pass"
>>> timeit.timeit(stmt=s, number=100000)
0.08588060699912603
```

要让 `timeit` 模块访问你定义的函数，你可以传递一个包含 `import` 语句的 `setup` 参数：

```
def test():
    """Stupid test function"""
    L = [i for i in range(100)]

if __name__ == '__main__':
    import timeit
    print(timeit.timeit("test()", setup="from __main__ import test"))
```

另一种选择是将 `globals()` 传递给 `globals` 参数，这将导致代码在当前的全局命名空间中执行。这比单独指定 `import` 更方便

```
def f(x):
    return x**2
def g(x):
    return x**4
def h(x):
    return x**8

import timeit
print(timeit.timeit('[func(42) for func in (f,g,h)]', globals=globals()))
```

27.6 `trace` — Trace or track Python statement execution

Source code: [Lib/trace.py](#)

The `trace` module allows you to trace program execution, generate annotated statement coverage listings, print caller/callee relationships and list functions executed during a program run. It can be used in another program or from the command line.

参见：

Coverage.py A popular third-party coverage tool that provides HTML output along with advanced features such as branch coverage.

27.6.1 Command-Line Usage

The `trace` module can be invoked from the command line. It can be as simple as

```
python -m trace --count -C . somefile.py ...
```

The above will execute `somefile.py` and generate annotated listings of all Python modules imported during the execution into the current directory.

--help

Display usage and exit.

--version

Display the version of the module and exit.

Main options

At least one of the following options must be specified when invoking *trace*. The *--listfuncs* option is mutually exclusive with the *--trace* and *--count* options. When *--listfuncs* is provided, neither *--count* nor *--trace* are accepted, and vice versa.

-c, --count

Produce a set of annotated listing files upon program completion that shows how many times each statement was executed. See also *--coverdir*, *--file* and *--no-report* below.

-t, --trace

Display lines as they are executed.

-l, --listfuncs

Display the functions executed by running the program.

-r, --report

Produce an annotated list from an earlier program run that used the *--count* and *--file* option. This does not execute any code.

-T, --trackcalls

Display the calling relationships exposed by running the program.

Modifiers**-f, --file=<file>**

Name of a file to accumulate counts over several tracing runs. Should be used with the *--count* option.

-C, --coverdir=<dir>

Directory where the report files go. The coverage report for *package.module* is written to file *dir/package/module.cover*.

-m, --missing

When generating annotated listings, mark lines which were not executed with >>>>>.

-s, --summary

When using *--count* or *--report*, write a brief summary to stdout for each file processed.

-R, --no-report

Do not generate annotated listings. This is useful if you intend to make several runs with *--count*, and then produce a single set of annotated listings at the end.

-g, --timing

Prefix each line with the time since the program started. Only used while tracing.

Filters

These options may be repeated multiple times.

--ignore-module=<mod>

Ignore each of the given module names and its submodules (if it is a package). The argument can be a list of names separated by a comma.

--ignore-dir=<dir>

Ignore all modules and packages in the named directory and subdirectories. The argument can be a list of directories separated by `os.pathsep`.

27.6.2 Programmatic Interface

class `trace.Trace` (*count=1, trace=1, countfuncs=0, countcallers=0, ignoremods=(), ignoredirs=(), infile=None, outfile=None, timing=False*)

Create an object to trace execution of a single statement or expression. All parameters are optional. *count* enables counting of line numbers. *trace* enables line execution tracing. *countfuncs* enables listing of the functions called during the run. *countcallers* enables call relationship tracking. *ignoremods* is a list of modules or packages to ignore. *ignoredirs* is a list of directories whose modules or packages should be ignored. *infile* is the name of the file from which to read stored count information. *outfile* is the name of the file in which to write updated count information. *timing* enables a timestamp relative to when tracing was started to be displayed.

run (*cmd*)

Execute the command and gather statistics from the execution with the current tracing parameters. *cmd* must be a string or code object, suitable for passing into `exec()`.

runtcx (*cmd, globals=None, locals=None*)

Execute the command and gather statistics from the execution with the current tracing parameters, in the defined global and local environments. If not defined, *globals* and *locals* default to empty dictionaries.

runfunc (*func, *args, **kwargs*)

Call *func* with the given arguments under control of the `Trace` object with the current tracing parameters.

results ()

Return a `CoverageResults` object that contains the cumulative results of all previous calls to `run`, `runtcx` and `runfunc` for the given `Trace` instance. Does not reset the accumulated trace results.

class `trace.CoverageResults`

A container for coverage results, created by `Trace.results()`. Should not be created directly by the user.

update (*other*)

Merge in data from another `CoverageResults` object.

write_results (*show_missing=True, summary=False, coverdir=None*)

Write coverage results. Set *show_missing* to show lines that had no hits. Set *summary* to include in the output the coverage summary per module. *coverdir* specifies the directory into which the coverage result files will be output. If `None`, the results for each source file are placed in its directory.

A simple example demonstrating the use of the programmatic interface:

```
import sys
import trace

# create a Trace object, telling it what to ignore, and whether to
# do tracing or line-counting or both.
tracer = trace.Trace(
    ignoredirs=[sys.prefix, sys.exec_prefix],
    trace=0,
    count=1)
```

(下页继续)

(续上页)

```
# run the new command using the given tracer
tracer.run('main() ')

# make a report, placing output in the current directory
r = tracer.results()
r.write_results(show_missing=True, coverdir=".")
```

27.7 tracemalloc —Trace memory allocations

3.4 新版功能.

Source code: [Lib/tracemalloc.py](#)

The tracemalloc module is a debug tool to trace memory blocks allocated by Python. It provides the following information:

- Traceback where an object was allocated
- Statistics on allocated memory blocks per filename and per line number: total size, number and average size of allocated memory blocks
- Compute the differences between two snapshots to detect memory leaks

To trace most memory blocks allocated by Python, the module should be started as early as possible by setting the `PYTHONTRACEMALLOC` environment variable to 1, or by using `-X tracemalloc` command line option. The `tracemalloc.start()` function can be called at runtime to start tracing Python memory allocations.

By default, a trace of an allocated memory block only stores the most recent frame (1 frame). To store 25 frames at startup: set the `PYTHONTRACEMALLOC` environment variable to 25, or use the `-X tracemalloc=25` command line option.

27.7.1 示例

Display the top 10

Display the 10 files allocating the most memory:

```
import tracemalloc

tracemalloc.start()

# ... run your application ...

snapshot = tracemalloc.take_snapshot()
top_stats = snapshot.statistics('lineno')

print("[ Top 10 ]")
for stat in top_stats[:10]:
    print(stat)
```

Example of output of the Python test suite:

```
[ Top 10 ]
<frozen importlib._bootstrap>:716: size=4855 KiB, count=39328, average=126 B
<frozen importlib._bootstrap>:284: size=521 KiB, count=3199, average=167 B
/usr/lib/python3.4/collections/__init__.py:368: size=244 KiB, count=2315, average=108 B
↪B
/usr/lib/python3.4/unittest/case.py:381: size=185 KiB, count=779, average=243 B
/usr/lib/python3.4/unittest/case.py:402: size=154 KiB, count=378, average=416 B
/usr/lib/python3.4/abc.py:133: size=88.7 KiB, count=347, average=262 B
<frozen importlib._bootstrap>:1446: size=70.4 KiB, count=911, average=79 B
<frozen importlib._bootstrap>:1454: size=52.0 KiB, count=25, average=2131 B
<string>:5: size=49.7 KiB, count=148, average=344 B
/usr/lib/python3.4/sysconfig.py:411: size=48.0 KiB, count=1, average=48.0 KiB
```

We can see that Python loaded 4855 KiB data (bytecode and constants) from modules and that the `collections` module allocated 244 KiB to build `namedtuple` types.

See `Snapshot.statistics()` for more options.

Compute differences

Take two snapshots and display the differences:

```
import tracemalloc
tracemalloc.start()
# ... start your application ...

snapshot1 = tracemalloc.take_snapshot()
# ... call the function leaking memory ...
snapshot2 = tracemalloc.take_snapshot()

top_stats = snapshot2.compare_to(snapshot1, 'lineno')

print("[ Top 10 differences ]")
for stat in top_stats[:10]:
    print(stat)
```

Example of output before/after running some tests of the Python test suite:

```
[ Top 10 differences ]
<frozen importlib._bootstrap>:716: size=8173 KiB (+4428 KiB), count=71332 (+39369), ↪
↪average=117 B
/usr/lib/python3.4/linecache.py:127: size=940 KiB (+940 KiB), count=8106 (+8106), ↪
↪average=119 B
/usr/lib/python3.4/unittest/case.py:571: size=298 KiB (+298 KiB), count=589 (+589), ↪
↪average=519 B
<frozen importlib._bootstrap>:284: size=1005 KiB (+166 KiB), count=7423 (+1526), ↪
↪average=139 B
/usr/lib/python3.4/mimetypes.py:217: size=112 KiB (+112 KiB), count=1334 (+1334), ↪
↪average=86 B
/usr/lib/python3.4/http/server.py:848: size=96.0 KiB (+96.0 KiB), count=1 (+1), ↪
↪average=96.0 KiB
/usr/lib/python3.4/inspect.py:1465: size=83.5 KiB (+83.5 KiB), count=109 (+109), ↪
↪average=784 B
/usr/lib/python3.4/unittest/mock.py:491: size=77.7 KiB (+77.7 KiB), count=143 (+143), ↪
↪average=557 B
/usr/lib/python3.4/urllib/parse.py:476: size=71.8 KiB (+71.8 KiB), count=969 (+969), ↪
↪average=76 B
```

(下页继续)

(续上页)

```
/usr/lib/python3.4/contextlib.py:38: size=67.2 KiB (+67.2 KiB), count=126 (+126), ↵
↵average=546 B
```

We can see that Python has loaded 8173 KiB of module data (bytecode and constants), and that this is 4428 KiB more than had been loaded before the tests, when the previous snapshot was taken. Similarly, the `linecache` module has cached 940 KiB of Python source code to format tracebacks, all of it since the previous snapshot.

If the system has little free memory, snapshots can be written on disk using the `Snapshot.dump()` method to analyze the snapshot offline. Then use the `Snapshot.load()` method reload the snapshot.

Get the traceback of a memory block

Code to display the traceback of the biggest memory block:

```
import tracemalloc

# Store 25 frames
tracemalloc.start(25)

# ... run your application ...

snapshot = tracemalloc.take_snapshot()
top_stats = snapshot.statistics('traceback')

# pick the biggest memory block
stat = top_stats[0]
print("%s memory blocks: %.1f KiB" % (stat.count, stat.size / 1024))
for line in stat.traceback.format():
    print(line)
```

Example of output of the Python test suite (traceback limited to 25 frames):

```
903 memory blocks: 870.1 KiB
File "<frozen importlib._bootstrap>", line 716
File "<frozen importlib._bootstrap>", line 1036
File "<frozen importlib._bootstrap>", line 934
File "<frozen importlib._bootstrap>", line 1068
File "<frozen importlib._bootstrap>", line 619
File "<frozen importlib._bootstrap>", line 1581
File "<frozen importlib._bootstrap>", line 1614
File "/usr/lib/python3.4/doctest.py", line 101
import pdb
File "<frozen importlib._bootstrap>", line 284
File "<frozen importlib._bootstrap>", line 938
File "<frozen importlib._bootstrap>", line 1068
File "<frozen importlib._bootstrap>", line 619
File "<frozen importlib._bootstrap>", line 1581
File "<frozen importlib._bootstrap>", line 1614
File "/usr/lib/python3.4/test/support/__init__.py", line 1728
import doctest
File "/usr/lib/python3.4/test/test_pickletools.py", line 21
    support.run_doctest(pickletools)
File "/usr/lib/python3.4/test/regtest.py", line 1276
    test_runner()
File "/usr/lib/python3.4/test/regtest.py", line 976
```

(下页继续)

(续上页)

```

display_failure=not verbose)
File "/usr/lib/python3.4/test/regtest.py", line 761
    match_tests=ns.match_tests)
File "/usr/lib/python3.4/test/regtest.py", line 1563
    main()
File "/usr/lib/python3.4/test/__main__.py", line 3
    regtest.main_in_temp_cwd()
File "/usr/lib/python3.4/runpy.py", line 73
    exec(code, run_globals)
File "/usr/lib/python3.4/runpy.py", line 160
    "__main__", fname, loader, pkg_name)

```

We can see that the most memory was allocated in the `importlib` module to load data (bytecode and constants) from modules: 870.1 KiB. The traceback is where the `importlib` loaded data most recently: on the `import pdb` line of the `doctest` module. The traceback may change if a new module is loaded.

Pretty top

Code to display the 10 lines allocating the most memory with a pretty output, ignoring `<frozen importlib._bootstrap>` and `<unknown>` files:

```

import linecache
import os
import tracemalloc

def display_top(snapshot, key_type='lineno', limit=10):
    snapshot = snapshot.filter_traces((
        tracemalloc.Filter(False, "<frozen importlib._bootstrap>"),
        tracemalloc.Filter(False, "<unknown>"),
    ))
    top_stats = snapshot.statistics(key_type)

    print("Top %s lines" % limit)
    for index, stat in enumerate(top_stats[:limit], 1):
        frame = stat.traceback[0]
        # replace "/path/to/module/file.py" with "module/file.py"
        filename = os.sep.join(frame.filename.split(os.sep)[-2:])
        print("#%s: %s: %s: %.1f KiB"
              % (index, filename, frame.lineno, stat.size / 1024))
        line = linecache.getline(frame.filename, frame.lineno).strip()
        if line:
            print('    %s' % line)

    other = top_stats[limit:]
    if other:
        size = sum(stat.size for stat in other)
        print("%s other: %.1f KiB" % (len(other), size / 1024))
    total = sum(stat.size for stat in top_stats)
    print("Total allocated size: %.1f KiB" % (total / 1024))

tracemalloc.start()

# ... run your application ...

snapshot = tracemalloc.take_snapshot()

```

(下页继续)

(续上页)

```
display_top(snapshot)
```

Example of output of the Python test suite:

```
Top 10 lines
#1: Lib/base64.py:414: 419.8 KiB
    _b85chars2 = [(a + b) for a in _b85chars for b in _b85chars]
#2: Lib/base64.py:306: 419.8 KiB
    _a85chars2 = [(a + b) for a in _a85chars for b in _a85chars]
#3: collections/__init__.py:368: 293.6 KiB
    exec(class_definition, namespace)
#4: Lib/abc.py:133: 115.2 KiB
    cls = super().__new__(mcls, name, bases, namespace)
#5: unittest/case.py:574: 103.1 KiB
    testMethod()
#6: Lib/linecache.py:127: 95.4 KiB
    lines = fp.readlines()
#7: urllib/parse.py:476: 71.8 KiB
    for a in _hexdig for b in _hexdig}
#8: <string>:5: 62.0 KiB
#9: Lib/_weakrefset.py:37: 60.0 KiB
    self.data = set()
#10: Lib/base64.py:142: 59.8 KiB
    _b32tab2 = [a + b for a in _b32tab for b in _b32tab]
6220 other: 3602.8 KiB
Total allocated size: 5303.1 KiB
```

See `Snapshot.statistics()` for more options.

27.7.2 API

函数

`tracemalloc.clear_traces()`

Clear traces of memory blocks allocated by Python.

See also `stop()`.

`tracemalloc.get_object_traceback(obj)`

Get the traceback where the Python object *obj* was allocated. Return a `Traceback` instance, or `None` if the `tracemalloc` module is not tracing memory allocations or did not trace the allocation of the object.

See also `gc.get_referrers()` and `sys.getsizeof()` functions.

`tracemalloc.get_traceback_limit()`

Get the maximum number of frames stored in the traceback of a trace.

The `tracemalloc` module must be tracing memory allocations to get the limit, otherwise an exception is raised.

The limit is set by the `start()` function.

`tracemalloc.get_traced_memory()`

Get the current size and peak size of memory blocks traced by the `tracemalloc` module as a tuple: (current: int, peak: int).

`tracemalloc.get_tracemalloc_memory()`

Get the memory usage in bytes of the `tracemalloc` module used to store traces of memory blocks. Return an `int`.

`tracemalloc.is_tracing()`

True if the `tracemalloc` module is tracing Python memory allocations, False otherwise.

See also `start()` and `stop()` functions.

`tracemalloc.start(nframe: int=1)`

Start tracing Python memory allocations: install hooks on Python memory allocators. Collected tracebacks of traces will be limited to `nframe` frames. By default, a trace of a memory block only stores the most recent frame: the limit is 1. `nframe` must be greater or equal to 1.

Storing more than 1 frame is only useful to compute statistics grouped by 'traceback' or to compute cumulative statistics: see the `Snapshot.compare_to()` and `Snapshot.statistics()` methods.

Storing more frames increases the memory and CPU overhead of the `tracemalloc` module. Use the `get_tracemalloc_memory()` function to measure how much memory is used by the `tracemalloc` module.

The `PYTHONTRACEMALLOC` environment variable (`PYTHONTRACEMALLOC=NFRAME`) and the `-X tracemalloc=NFRAME` command line option can be used to start tracing at startup.

See also `stop()`, `is_tracing()` and `get_traceback_limit()` functions.

`tracemalloc.stop()`

Stop tracing Python memory allocations: uninstall hooks on Python memory allocators. Also clears all previously collected traces of memory blocks allocated by Python.

Call `take_snapshot()` function to take a snapshot of traces before clearing them.

See also `start()`, `is_tracing()` and `clear_traces()` functions.

`tracemalloc.take_snapshot()`

Take a snapshot of traces of memory blocks allocated by Python. Return a new `Snapshot` instance.

The snapshot does not include memory blocks allocated before the `tracemalloc` module started to trace memory allocations.

Tracebacks of traces are limited to `get_traceback_limit()` frames. Use the `nframe` parameter of the `start()` function to store more frames.

The `tracemalloc` module must be tracing memory allocations to take a snapshot, see the `start()` function.

See also the `get_object_traceback()` function.

DomainFilter

class `tracemalloc.DomainFilter` (*inclusive: bool, domain: int*)

Filter traces of memory blocks by their address space (domain).

3.6 新版功能.

inclusive

If *inclusive* is True (include), match memory blocks allocated in the address space *domain*.

If *inclusive* is False (exclude), match memory blocks not allocated in the address space *domain*.

domain

Address space of a memory block (int). Read-only property.

Filter

class tracemalloc.**Filter**(*inclusive: bool, filename_pattern: str, lineno: int=None, all_frames: bool=False, domain: int=None*)

Filter on traces of memory blocks.

See the `fnmatch.fnmatch()` function for the syntax of `filename_pattern`. The `'.pyc'` file extension is replaced with `'.py'`.

示例:

- `Filter(True, subprocess.__file__)` only includes traces of the `subprocess` module
- `Filter(False, tracemalloc.__file__)` excludes traces of the `tracemalloc` module
- `Filter(False, "<unknown>")` excludes empty tracebacks

在 3.5 版更改: The `'.pyo'` file extension is no longer replaced with `'.py'`.

在 3.6 版更改: Added the `domain` attribute.

domain

Address space of a memory block (int or None).

inclusive

If `inclusive` is `True` (include), only match memory blocks allocated in a file with a name matching `filename_pattern` at line number `lineno`.

If `inclusive` is `False` (exclude), ignore memory blocks allocated in a file with a name matching `filename_pattern` at line number `lineno`.

lineno

Line number (int) of the filter. If `lineno` is `None`, the filter matches any line number.

filename_pattern

Filename pattern of the filter (str). Read-only property.

all_frames

If `all_frames` is `True`, all frames of the traceback are checked. If `all_frames` is `False`, only the most recent frame is checked.

This attribute has no effect if the traceback limit is 1. See the `get_traceback_limit()` function and `Snapshot.traceback_limit` attribute.

Frame

class tracemalloc.**Frame**

Frame of a traceback.

The `Traceback` class is a sequence of `Frame` instances.

filename

Filename (str).

lineno

Line number (int).

Snapshot

class tracemalloc.**Snapshot**

Snapshot of traces of memory blocks allocated by Python.

The `take_snapshot()` function creates a snapshot instance.

compare_to (*old_snapshot: Snapshot, key_type: str, cumulative: bool=False*)

Compute the differences with an old snapshot. Get statistics as a sorted list of `StatisticDiff` instances grouped by `key_type`.

See the `Snapshot.statistics()` method for `key_type` and `cumulative` parameters.

The result is sorted from the biggest to the smallest by: absolute value of `StatisticDiff.size_diff`, `StatisticDiff.size`, absolute value of `StatisticDiff.count_diff`, `Statistic.count` and then by `StatisticDiff.traceback`.

dump (*filename*)

Write the snapshot into a file.

Use `load()` to reload the snapshot.

filter_traces (*filters*)

Create a new `Snapshot` instance with a filtered `traces` sequence, *filters* is a list of `DomainFilter` and `Filter` instances. If *filters* is an empty list, return a new `Snapshot` instance with a copy of the traces.

All inclusive filters are applied at once, a trace is ignored if no inclusive filters match it. A trace is ignored if at least one exclusive filter matches it.

在 3.6 版更改: `DomainFilter` instances are now also accepted in *filters*.

classmethod load (*filename*)

Load a snapshot from a file.

See also `dump()`.

statistics (*key_type: str, cumulative: bool=False*)

Get statistics as a sorted list of `Statistic` instances grouped by `key_type`:

key_type	描述
'filename'	filename
'lineno'	filename and line number
'traceback'	traceback

If `cumulative` is `True`, cumulate size and count of memory blocks of all frames of the traceback of a trace, not only the most recent frame. The cumulative mode can only be used with `key_type` equals to 'filename' and 'lineno'.

The result is sorted from the biggest to the smallest by: `Statistic.size`, `Statistic.count` and then by `Statistic.traceback`.

traceback_limit

Maximum number of frames stored in the traceback of `traces`: result of the `get_traceback_limit()` when the snapshot was taken.

traces

Traces of all memory blocks allocated by Python: sequence of `Trace` instances.

The sequence has an undefined order. Use the `Snapshot.statistics()` method to get a sorted list of statistics.

Statistic

class tracemalloc.Statistic

Statistic on memory allocations.

Snapshot.statistics() returns a list of *Statistic* instances.

See also the *StatisticDiff* class.

count

Number of memory blocks (int).

size

Total size of memory blocks in bytes (int).

traceback

Traceback where the memory block was allocated, *Traceback* instance.

StatisticDiff

class tracemalloc.**StatisticDiff**

Statistic difference on memory allocations between an old and a new *Snapshot* instance.

Snapshot.compare_to() returns a list of *StatisticDiff* instances. See also the *Statistic* class.

count

Number of memory blocks in the new snapshot (int): 0 if the memory blocks have been released in the new snapshot.

count_diff

Difference of number of memory blocks between the old and the new snapshots (int): 0 if the memory blocks have been allocated in the new snapshot.

size

Total size of memory blocks in bytes in the new snapshot (int): 0 if the memory blocks have been released in the new snapshot.

size_diff

Difference of total size of memory blocks in bytes between the old and the new snapshots (int): 0 if the memory blocks have been allocated in the new snapshot.

traceback

Traceback where the memory blocks were allocated, *Traceback* instance.

Trace

class tracemalloc.**Trace**

Trace of a memory block.

The *Snapshot.traces* attribute is a sequence of *Trace* instances.

size

Size of the memory block in bytes (int).

traceback

Traceback where the memory block was allocated, *Traceback* instance.

Traceback

class tracemalloc.**Traceback**

Sequence of *Frame* instances sorted from the most recent frame to the oldest frame.

A traceback contains at least 1 frame. If the tracemalloc module failed to get a frame, the filename "<unknown>" at line number 0 is used.

When a snapshot is taken, tracebacks of traces are limited to `get_traceback_limit()` frames. See the `take_snapshot()` function.

The `Trace.traceback` attribute is an instance of `Traceback` instance.

format (*limit=None*)

Format the traceback as a list of lines with newlines. Use the `linecache` module to retrieve lines from the source code. If *limit* is set, only format the *limit* most recent frames.

Similar to the `traceback.format_tb()` function, except that `format()` does not include newlines.

示例:

```
print("Traceback (most recent call first):")
for line in traceback:
    print(line)
```

输出:

```
Traceback (most recent call first):
  File "test.py", line 9
    obj = Object()
  File "test.py", line 12
    tb = tracemalloc.get_object_traceback(f())
```

这些库可帮助你发布和安装 Python 软件。虽然这些模块设计为与 ‘Python 包索引’ <<https://pypi.org>> 结合使用，但它们也可以与本地索引服务器一起使用，或者根本不使用任何索引服务器。

28.1 `distutils` — 构建和安装 Python 模块

The `distutils` package provides support for building and installing additional modules into a Python installation. The new modules may be either 100%-pure Python, or may be extension modules written in C, or may be collections of Python packages which include modules coded in both Python and C.

Most Python users will *not* want to use this module directly, but instead use the cross-version tools maintained by the Python Packaging Authority. In particular, `setuptools` is an enhanced alternative to `distutils` that provides:

- support for declaring project dependencies
- additional mechanisms for configuring which files to include in source releases (including plugins for integration with version control systems)
- the ability to declare project “entry points”, which can be used as the basis for application plugin systems
- the ability to automatically generate Windows command line executables at installation time rather than needing to prebuild them
- consistent behaviour across all supported Python versions

The recommended `pip` installer runs all `setup.py` scripts with `setuptools`, even if the script itself only imports `distutils`. Refer to the [Python Packaging User Guide](#) for more information.

For the benefits of packaging tool authors and users seeking a deeper understanding of the details of the current packaging and distribution system, the legacy `distutils` based user documentation and API reference remain available:

- [install-index](#)
- [distutils-index](#)

28.2 ensurepip — Bootstrapping the pip installer

3.4 新版功能.

The *ensurepip* package provides support for bootstrapping the `pip` installer into an existing Python installation or virtual environment. This bootstrapping approach reflects the fact that `pip` is an independent project with its own release cycle, and the latest available stable version is bundled with maintenance and feature releases of the CPython reference interpreter.

In most cases, end users of Python shouldn't need to invoke this module directly (as `pip` should be bootstrapped by default), but it may be needed if installing `pip` was skipped when installing Python (or when creating a virtual environment) or after explicitly uninstalling `pip`.

注解: This module *does not* access the internet. All of the components needed to bootstrap `pip` are included as internal parts of the package.

参见:

installing-index The end user guide for installing Python packages

PEP 453: Explicit bootstrapping of pip in Python installations The original rationale and specification for this module.

28.2.1 Command line interface

The command line interface is invoked using the interpreter's `-m` switch.

The simplest possible invocation is:

```
python -m ensurepip
```

This invocation will install `pip` if it is not already installed, but otherwise does nothing. To ensure the installed version of `pip` is at least as recent as the one bundled with `ensurepip`, pass the `--upgrade` option:

```
python -m ensurepip --upgrade
```

By default, `pip` is installed into the current virtual environment (if one is active) or into the system site packages (if there is no active virtual environment). The installation location can be controlled through two additional command line options:

- `--root <dir>`: Installs `pip` relative to the given root directory rather than the root of the currently active virtual environment (if any) or the default root for the current Python installation.
- `--user`: Installs `pip` into the user site packages directory rather than globally for the current Python installation (this option is not permitted inside an active virtual environment).

By default, the scripts `pipX` and `pipX.Y` will be installed (where `X.Y` stands for the version of Python used to invoke `ensurepip`). The scripts installed can be controlled through two additional command line options:

- `--altinstall`: if an alternate installation is requested, the `pipX` script will *not* be installed.
- `--default-pip`: if a “default pip” installation is requested, the `pip` script will be installed in addition to the two regular scripts.

Providing both of the script selection options will trigger an exception.

在 3.6.3 版更改: The exit status is non-zero if the command fails.

28.2.2 Module API

`ensurepip` exposes two functions for programmatic use:

`ensurepip.version()`

Returns a string specifying the bundled version of pip that will be installed when bootstrapping an environment.

`ensurepip.bootstrap(root=None, upgrade=False, user=False, altinstall=False, default_pip=False, verbosity=0)`

Bootstraps pip into the current or designated environment.

root specifies an alternative root directory to install relative to. If *root* is `None`, then installation uses the default install location for the current environment.

upgrade indicates whether or not to upgrade an existing installation of an earlier version of pip to the bundled version.

user indicates whether to use the user scheme rather than installing globally.

By default, the scripts `pipX` and `pipX.Y` will be installed (where `X.Y` stands for the current version of Python).

If *altinstall* is set, then `pipX` will *not* be installed.

If *default_pip* is set, then `pip` will be installed in addition to the two regular scripts.

Setting both *altinstall* and *default_pip* will trigger `ValueError`.

verbosity controls the level of output to `sys.stdout` from the bootstrapping operation.

注解： The bootstrapping process has side effects on both `sys.path` and `os.environ`. Invoking the command line interface in a subprocess instead allows these side effects to be avoided.

注解： The bootstrapping process may install additional modules required by pip, but other software should not assume those dependencies will always be present by default (as the dependencies may be removed in a future version of pip).

28.3 venv — 创建虚拟环境

3.3 新版功能.

源码: [Lib/venv/](#)

The `venv` module provides support for creating lightweight “virtual environments” with their own site directories, optionally isolated from system site directories. Each virtual environment has its own Python binary (allowing creation of environments with various Python versions) and can have its own independent set of installed Python packages in its site directories.

有关 Python 虚拟环境的更多信息，请参阅 [PEP 405](#)。

注解： 从 Python 3.6 开始，不推荐使用 `pyvenv` 脚本，而是使用 `python3 -m venv` 来帮助防止任何关于虚拟环境将基于哪个 Python 解释器的混淆。

28.3.1 创建虚拟环境

通过执行 `venv` 指令来创建一个虚拟环境:

```
python3 -m venv /path/to/new/virtual/environment
```

Running this command creates the target directory (creating any parent directories that don't exist already) and places a `pyvenv.cfg` file in it with a `home` key pointing to the Python installation from which the command was run. It also creates a `bin` (or `Scripts` on Windows) subdirectory containing a copy/symlink of the Python binary/binaries (as appropriate for the platform or arguments used at environment creation time). It also creates an (initially empty) `lib/pythonX.Y/site-packages` subdirectory (on Windows, this is `Lib\site-packages`). If an existing directory is specified, it will be re-used.

3.6 版后已移除: `pyvenv` 是 Python 3.3 和 3.4 中创建虚拟环境的推荐工具, 不过在 Python 3.6 中已弃用。

在 3.5 版更改: 现在推荐使用 `venv` 来创建虚拟环境。

参见:

Python 打包用户指南: 创建和使用虚拟环境

On Windows, invoke the `venv` command as follows:

```
c:\>c:\Python35\python -m venv c:\path\to\myenv
```

Alternatively, if you configured the `PATH` and `PATHEXT` variables for your Python installation:

```
c:\>python -m venv c:\path\to\myenv
```

The command, if run with `-h`, will show the available options:

```
usage: venv [-h] [--system-site-packages] [--symlinks | --copies] [--clear]
           [--upgrade] [--without-pip]
           ENV_DIR [ENV_DIR ...]

Creates virtual Python environments in one or more target directories.

positional arguments:
  ENV_DIR                A directory to create the environment in.

optional arguments:
  -h, --help            show this help message and exit
  --system-site-packages
                        Give the virtual environment access to the system
                        site-packages dir.
  --symlinks            Try to use symlinks rather than copies, when symlinks
                        are not the default for the platform.
  --copies              Try to use copies rather than symlinks, even when
                        symlinks are the default for the platform.
  --clear               Delete the contents of the environment directory if it
                        already exists, before environment creation.
  --upgrade             Upgrade the environment directory to use this version
                        of Python, assuming Python has been upgraded in-place.
  --without-pip         Skips installing or upgrading pip in the virtual
                        environment (pip is bootstrapped by default)

Once an environment has been created, you may wish to activate it, e.g. by
sourcing an activate script in its bin directory.
```

在 3.4 版更改: Installs `pip` by default, added the `--without-pip` and `--copies` options

在 3.4 版更改: In earlier versions, if the target directory already existed, an error was raised, unless the `--clear` or `--upgrade` option was provided.

The created `pyvenv.cfg` file also includes the `include-system-site-packages` key, set to `true` if `venv` is run with the `--system-site-packages` option, `false` otherwise.

Unless the `--without-pip` option is given, *ensurepip* will be invoked to bootstrap `pip` into the virtual environment.

Multiple paths can be given to `venv`, in which case an identical virtual environment will be created, according to the given options, at each provided path.

Once a virtual environment has been created, it can be “activated” using a script in the virtual environment’s binary directory. The invocation of the script is platform-specific (<venv> must be replaced by the path of the directory containing the virtual environment):

平台	Shell	用于激活虚拟环境的命令
Posix	bash/zsh	\$ source <venv>/bin/activate
	fish	\$. <venv>/bin/activate.fish
	csh/tcsh	\$ source <venv>/bin/activate.csh
Windows	cmd.exe	C:\> <venv>\Scripts\activate.bat
	PowerShell	PS C:\> <venv>\Scripts\Activate.ps1

You don’t specifically *need* to activate an environment; activation just prepends the virtual environment’s binary directory to your path, so that “python” invokes the virtual environment’s Python interpreter and you can run installed scripts without having to use their full path. However, all scripts installed in a virtual environment should be runnable without activating it, and run with the virtual environment’s Python automatically.

You can deactivate a virtual environment by typing “deactivate” in your shell. The exact mechanism is platform-specific: for example, the Bash activation script defines a “deactivate” function, whereas on Windows there are separate scripts called `deactivate.bat` and `Deactivate.ps1` which are installed when the virtual environment is created.

3.4 新版功能: `fish` and `csh` activation scripts.

注解: A virtual environment is a Python environment such that the Python interpreter, libraries and scripts installed into it are isolated from those installed in other virtual environments, and (by default) any libraries installed in a “system” Python, i.e., one which is installed as part of your operating system.

A virtual environment is a directory tree which contains Python executable files and other files which indicate that it is a virtual environment.

Common installation tools such as `Setuptools` and `pip` work as expected with virtual environments. In other words, when a virtual environment is active, they install Python packages into the virtual environment without needing to be told to do so explicitly.

When a virtual environment is active (i.e., the virtual environment’s Python interpreter is running), the attributes `sys.prefix` and `sys.exec_prefix` point to the base directory of the virtual environment, whereas `sys.base_prefix` and `sys.base_exec_prefix` point to the non-virtual environment Python installation which was used to create the virtual environment. If a virtual environment is not active, then `sys.prefix` is the same as `sys.base_prefix` and `sys.exec_prefix` is the same as `sys.base_exec_prefix` (they all point to a non-virtual environment Python installation).

When a virtual environment is active, any options that change the installation path will be ignored from all distutils configuration files to prevent projects being inadvertently installed outside of the virtual environment.

When working in a command shell, users can make a virtual environment active by running an `activate` script in the virtual environment’s executables directory (the precise filename is shell-dependent), which prepends the virtual environment’s directory for executables to the `PATH` environment variable for the running shell. There should be no

need in other circumstances to activate a virtual environment—scripts installed into virtual environments have a “shebang” line which points to the virtual environment’s Python interpreter. This means that the script will run with that interpreter regardless of the value of `PATH`. On Windows, “shebang” line processing is supported if you have the Python Launcher for Windows installed (this was added to Python in 3.3 - see [PEP 397](#) for more details). Thus, double-clicking an installed script in a Windows Explorer window should run the script with the correct interpreter without there needing to be any reference to its virtual environment in `PATH`.

28.3.2 API

The high-level method described above makes use of a simple API which provides mechanisms for third-party virtual environment creators to customize environment creation according to their needs, the `EnvBuilder` class.

class `venv.EnvBuilder` (*system_site_packages=False, clear=False, symlinks=False, upgrade=False, with_pip=False, prompt=None*)

The `EnvBuilder` class accepts the following keyword arguments on instantiation:

- `system_site_packages` – a Boolean value indicating that the system Python site-packages should be available to the environment (defaults to `False`).
- `clear` – a Boolean value which, if true, will delete the contents of any existing target directory, before creating the environment.
- `symlinks` – a Boolean value indicating whether to attempt to symlink the Python binary (and any necessary DLLs or other binaries, e.g. `pythonw.exe`), rather than copying.
- `upgrade` – a Boolean value which, if true, will upgrade an existing environment with the running Python - for use when that Python has been upgraded in-place (defaults to `False`).
- `with_pip` – a Boolean value which, if true, ensures pip is installed in the virtual environment. This uses `ensurepip` with the `--default-pip` option.
- `prompt` – a String to be used after virtual environment is activated (defaults to `None` which means directory name of the environment would be used).

在 3.4 版更改: Added the `with_pip` parameter

3.6 新版功能: Added the `prompt` parameter

Creators of third-party virtual environment tools will be free to use the provided `EnvBuilder` class as a base class.

The returned env-builder is an object which has a method, `create`:

create (*env_dir*)

This method takes as required argument the path (absolute or relative to the current directory) of the target directory which is to contain the virtual environment. The `create` method will either create the environment in the specified directory, or raise an appropriate exception.

The `create` method of the `EnvBuilder` class illustrates the hooks available for subclass customization:

```
def create(self, env_dir):
    """
    Create a virtualized Python environment in a directory.
    env_dir is the target directory to create an environment in.
    """
    env_dir = os.path.abspath(env_dir)
    context = self.ensure_directories(env_dir)
    self.create_configuration(context)
    self.setup_python(context)
```

(下页继续)

(续上页)

```
self.setup_scripts(context)
self.post_setup(context)
```

Each of the methods `ensure_directories()`, `create_configuration()`, `setup_python()`, `setup_scripts()` and `post_setup()` can be overridden.

ensure_directories (*env_dir*)

Creates the environment directory and all necessary directories, and returns a context object. This is just a holder for attributes (such as paths), for use by the other methods. The directories are allowed to exist already, as long as either `clear` or `upgrade` were specified to allow operating on an existing environment directory.

create_configuration (*context*)

Creates the `pyenvv.cfg` configuration file in the environment.

setup_python (*context*)

Creates a copy of the Python executable (and, under Windows, DLLs) in the environment. On a POSIX system, if a specific executable `python3.x` was used, symlinks to `python` and `python3` will be created pointing to that executable, unless files with those names already exist.

setup_scripts (*context*)

Installs activation scripts appropriate to the platform into the virtual environment.

post_setup (*context*)

A placeholder method which can be overridden in third party implementations to pre-install packages in the virtual environment or perform other post-creation steps.

In addition, `EnvBuilder` provides this utility method that can be called from `setup_scripts()` or `post_setup()` in subclasses to assist in installing custom scripts into the virtual environment.

install_scripts (*context*, *path*)

path is the path to a directory that should contain subdirectories “common”, “posix”, “nt”, each containing scripts destined for the bin directory in the environment. The contents of “common” and the directory corresponding to `os.name` are copied after some text replacement of placeholders:

- `__VENV_DIR__` is replaced with the absolute path of the environment directory.
- `__VENV_NAME__` is replaced with the environment name (final path segment of environment directory).
- `__VENV_PROMPT__` is replaced with the prompt (the environment name surrounded by parentheses and with a following space)
- `__VENV_BIN_NAME__` is replaced with the name of the bin directory (either `bin` or `Scripts`).
- `__VENV_PYTHON__` is replaced with the absolute path of the environment’s executable.

The directories are allowed to exist (for when an existing environment is being upgraded).

有一个方便实用的模块级别的函数:

```
venv.create(env_dir, system_site_packages=False, clear=False, symlinks=False, with_pip=False)
```

通过关键词参数来创建一个 `EnvBuilder`, 并且使用 `env_dir` 参数来调用它的 `create()` 方法。

在 3.4 版更改: Added the `with_pip` parameter

28.3.3 一个扩展 `EnvBuilder` 的例子

下面的脚本展示了如何通过实现一个子类来扩展 `EnvBuilder`。这个子类会安装 `setuptools` 和 `pip` 的到被创建的虚拟环境中。

```

import os
import os.path
from subprocess import Popen, PIPE
import sys
from threading import Thread
from urllib.parse import urlparse
from urllib.request import urlretrieve
import venv

class ExtendedEnvBuilder(venv.EnvBuilder):
    """
    This builder installs setuptools and pip so that you can pip or
    easy_install other packages into the created virtual environment.

    :param nodist: If True, setuptools and pip are not installed into the
        created virtual environment.
    :param nopip: If True, pip is not installed into the created
        virtual environment.
    :param progress: If setuptools or pip are installed, the progress of the
        installation can be monitored by passing a progress
        callable. If specified, it is called with two
        arguments: a string indicating some progress, and a
        context indicating where the string is coming from.
        The context argument can have one of three values:
        'main', indicating that it is called from virtualize()
        itself, and 'stdout' and 'stderr', which are obtained
        by reading lines from the output streams of a subprocess
        which is used to install the app.

        If a callable is not specified, default progress
        information is output to sys.stderr.
    """

    def __init__(self, *args, **kwargs):
        self.nodist = kwargs.pop('nodist', False)
        self.nopip = kwargs.pop('nopip', False)
        self.progress = kwargs.pop('progress', None)
        self.verbose = kwargs.pop('verbose', False)
        super().__init__(*args, **kwargs)

    def post_setup(self, context):
        """
        Set up any packages which need to be pre-installed into the
        virtual environment being created.

        :param context: The information for the virtual environment
            creation request being processed.
        """
        os.environ['VIRTUAL_ENV'] = context.env_dir
        if not self.nodist:
            self.install_setuptools(context)
            # Can't install pip without setuptools
        if not self.nopip and not self.nodist:
            self.install_pip(context)

    def reader(self, stream, context):
        """

```

(下页继续)

(续上页)

```

Read lines from a subprocess' output stream and either pass to a progress
callable (if specified) or write progress information to sys.stderr.
"""
progress = self.progress
while True:
    s = stream.readline()
    if not s:
        break
    if progress is not None:
        progress(s, context)
    else:
        if not self.verbose:
            sys.stderr.write('.')
        else:
            sys.stderr.write(s.decode('utf-8'))
        sys.stderr.flush()
stream.close()

def install_script(self, context, name, url):
    _, _, path, _, _, _ = urlparse(url)
    fn = os.path.split(path)[-1]
    binpath = context.bin_path
    distpath = os.path.join(binpath, fn)
    # Download script into the virtual environment's binaries folder
    urlretrieve(url, distpath)
    progress = self.progress
    if self.verbose:
        term = '\n'
    else:
        term = ''
    if progress is not None:
        progress('Installing %s ...%s' % (name, term), 'main')
    else:
        sys.stderr.write('Installing %s ...%s' % (name, term))
        sys.stderr.flush()
    # Install in the virtual environment
    args = [context.env_exe, fn]
    p = Popen(args, stdout=PIPE, stderr=PIPE, cwd=binpath)
    t1 = Thread(target=self.reader, args=(p.stdout, 'stdout'))
    t1.start()
    t2 = Thread(target=self.reader, args=(p.stderr, 'stderr'))
    t2.start()
    p.wait()
    t1.join()
    t2.join()
    if progress is not None:
        progress('done.', 'main')
    else:
        sys.stderr.write('done.\n')
    # Clean up - no longer needed
    os.unlink(distpath)

def install_setuptools(self, context):
    """
    Install setuptools in the virtual environment.

```

(下页继续)

(续上页)

```

:param context: The information for the virtual environment
                  creation request being processed.
"""
url = 'https://bitbucket.org/pypa/setuptools/downloads/ez_setup.py'
self.install_script(context, 'setuptools', url)
# clear up the setuptools archive which gets downloaded
pred = lambda o: o.startswith('setuptools-') and o.endswith('.tar.gz')
files = filter(pred, os.listdir(context.bin_path))
for f in files:
    f = os.path.join(context.bin_path, f)
    os.unlink(f)

def install_pip(self, context):
    """
    Install pip in the virtual environment.

    :param context: The information for the virtual environment
                    creation request being processed.
    """
    url = 'https://raw.githubusercontent.com/pypa/pip/master/contrib/get-pip.py'
    self.install_script(context, 'pip', url)

def main(args=None):
    compatible = True
    if sys.version_info < (3, 3):
        compatible = False
    elif not hasattr(sys, 'base_prefix'):
        compatible = False
    if not compatible:
        raise ValueError('This script is only for use with '
                          'Python 3.3 or later')
    else:
        import argparse

        parser = argparse.ArgumentParser(prog=__name__,
                                         description='Creates virtual Python '
                                                         'environments in one or '
                                                         'more target '
                                                         'directories.')

        parser.add_argument('dirs', metavar='ENV_DIR', nargs='+',
                             help='A directory in which to create the '
                                   'virtual environment.')
        parser.add_argument('--no-setuptools', default=False,
                             action='store_true', dest='nodist',
                             help="Don't install setuptools or pip in the "
                                   "virtual environment.")
        parser.add_argument('--no-pip', default=False,
                             action='store_true', dest='nopip',
                             help="Don't install pip in the virtual "
                                   "environment.")
        parser.add_argument('--system-site-packages', default=False,
                             action='store_true', dest='system_site',
                             help='Give the virtual environment access to the '
                                   'system site-packages dir.')

        if os.name == 'nt':
            use_symlinks = False

```

(下页继续)

(续上页)

```

else:
    use_symlinks = True
parser.add_argument('--symlinks', default=use_symlinks,
                    action='store_true', dest='symlinks',
                    help='Try to use symlinks rather than copies, '
                         'when symlinks are not the default for '
                         'the platform.')
parser.add_argument('--clear', default=False, action='store_true',
                    dest='clear', help='Delete the contents of the '
                                       'virtual environment '
                                       'directory if it already '
                                       'exists, before virtual '
                                       'environment creation.')
parser.add_argument('--upgrade', default=False, action='store_true',
                    dest='upgrade', help='Upgrade the virtual '
                                       'environment directory to '
                                       'use this version of '
                                       'Python, assuming Python '
                                       'has been upgraded '
                                       'in-place.')
parser.add_argument('--verbose', default=False, action='store_true',
                    dest='verbose', help='Display the output '
                                       'from the scripts which '
                                       'install setuptools and pip.')

options = parser.parse_args(args)
if options.upgrade and options.clear:
    raise ValueError('you cannot supply --upgrade and --clear together.')
builder = ExtendedEnvBuilder(system_site_packages=options.system_site,
                             clear=options.clear,
                             symlinks=options.symlinks,
                             upgrade=options.upgrade,
                             nodist=options.nodist,
                             nopip=options.nopip,
                             verbose=options.verbose)

for d in options.dirs:
    builder.create(d)

if __name__ == '__main__':
    rc = 1
    try:
        main()
        rc = 0
    except Exception as e:
        print('Error: %s' % e, file=sys.stderr)
    sys.exit(rc)

```

This script is also available for download [online](#).

28.4 zipapp —Manage executable python zip archives

3.5 新版功能.

Source code: [Lib/zipapp.py](#)

This module provides tools to manage the creation of zip files containing Python code, which can be executed directly by the Python interpreter. The module provides both a 命令行界面 and a *Python API*.

28.4.1 Basic Example

The following example shows how the 命令行界面 can be used to create an executable archive from a directory containing Python code. When run, the archive will execute the `main` function from the module `myapp` in the archive.

```
$ python -m zipapp myapp -m "myapp:main"
$ python myapp.pyz
<output from myapp>
```

28.4.2 命令行界面

When called as a program from the command line, the following form is used:

```
$ python -m zipapp source [options]
```

If *source* is a directory, this will create an archive from the contents of *source*. If *source* is a file, it should be an archive, and it will be copied to the target archive (or the contents of its shebang line will be displayed if the `-info` option is specified).

The following options are understood:

- o <output>, --output=<output>**
Write the output to a file named *output*. If this option is not specified, the output filename will be the same as the input *source*, with the extension `.pyz` added. If an explicit filename is given, it is used as is (so a `.pyz` extension should be included if required).

An output filename must be specified if the *source* is an archive (and in that case, *output* must not be the same as *source*).
- p <interpreter>, --python=<interpreter>**
Add a `#!` line to the archive specifying *interpreter* as the command to run. Also, on POSIX, make the archive executable. The default is to write no `#!` line, and not make the file executable.
- m <mainfn>, --main=<mainfn>**
Write a `__main__.py` file to the archive that executes *mainfn*. The *mainfn* argument should have the form “`pkg.mod:fn`”, where “`pkg.mod`” is a package/module in the archive, and “`fn`” is a callable in the given module. The `__main__.py` file will execute that callable.

`--main` cannot be specified when copying an archive.
- info**
Display the interpreter embedded in the archive, for diagnostic purposes. In this case, any other options are ignored and *SOURCE* must be an archive, not a directory.
- h, --help**
Print a short usage message and exit.

28.4.3 Python API

The module defines two convenience functions:

`zipapp.create_archive(source, target=None, interpreter=None, main=None)`
Create an application archive from *source*. The source can be any of the following:

- The name of a directory, or a `pathlib.Path` object referring to a directory, in which case a new application archive will be created from the content of that directory.
- The name of an existing application archive file, or a `pathlib.Path` object referring to such a file, in which case the file is copied to the target (modifying it to reflect the value given for the *interpreter* argument). The file name should include the `.pyz` extension, if required.
- A file object open for reading in bytes mode. The content of the file should be an application archive, and the file object is assumed to be positioned at the start of the archive.

The *target* argument determines where the resulting archive will be written:

- If it is the name of a file, or a `pathlib.Path` object, the archive will be written to that file.
- If it is an open file object, the archive will be written to that file object, which must be open for writing in bytes mode.
- If the target is omitted (or `None`), the source must be a directory and the target will be a file with the same name as the source, with a `.pyz` extension added.

The *interpreter* argument specifies the name of the Python interpreter with which the archive will be executed. It is written as a “shebang” line at the start of the archive. On POSIX, this will be interpreted by the OS, and on Windows it will be handled by the Python launcher. Omitting the *interpreter* results in no shebang line being written. If an interpreter is specified, and the target is a filename, the executable bit of the target file will be set.

The *main* argument specifies the name of a callable which will be used as the main program for the archive. It can only be specified if the source is a directory, and the source does not already contain a `__main__.py` file. The *main* argument should take the form “`pkg.module:callable`” and the archive will be run by importing “`pkg.module`” and executing the given callable with no arguments. It is an error to omit *main* if the source is a directory and does not contain a `__main__.py` file, as otherwise the resulting archive would not be executable.

If a file object is specified for *source* or *target*, it is the caller’s responsibility to close it after calling `create_archive`.

When copying an existing archive, file objects supplied only need `read` and `readline`, or `write` methods. When creating an archive from a directory, if the target is a file object it will be passed to the `zipfile.ZipFile` class, and must supply the methods needed by that class.

`zipapp.get_interpreter (archive)`

Return the interpreter specified in the `#!` line at the start of the archive. If there is no `#!` line, return `None`. The *archive* argument can be a filename or a file-like object open for reading in bytes mode. It is assumed to be at the start of the archive.

28.4.4 示例

Pack up a directory into an archive, and run it.

```
$ python -m zipapp myapp
$ python myapp.pyz
<output from myapp>
```

The same can be done using the `create_archive()` function:

```
>>> import zipapp
>>> zipapp.create_archive('myapp.pyz', 'myapp')
```

To make the application directly executable on POSIX, specify an interpreter to use.

```
$ python -m zipapp myapp -p "/usr/bin/env python"
$ ./myapp.pyz
<output from myapp>
```

To replace the shebang line on an existing archive, create a modified archive using the `create_archive()` function:

```
>>> import zipapp
>>> zipapp.create_archive('old_archive.pyz', 'new_archive.pyz', '/usr/bin/python3')
```

To update the file in place, do the replacement in memory using a `BytesIO` object, and then overwrite the source afterwards. Note that there is a risk when overwriting a file in place that an error will result in the loss of the original file. This code does not protect against such errors, but production code should do so. Also, this method will only work if the archive fits in memory:

```
>>> import zipapp
>>> import io
>>> temp = io.BytesIO()
>>> zipapp.create_archive('myapp.pyz', temp, '/usr/bin/python2')
>>> with open('myapp.pyz', 'wb') as f:
>>>     f.write(temp.getvalue())
```

28.4.5 Specifying the Interpreter

Note that if you specify an interpreter and then distribute your application archive, you need to ensure that the interpreter used is portable. The Python launcher for Windows supports most common forms of POSIX `#!` line, but there are other issues to consider:

- If you use `“/usr/bin/env python”` (or other forms of the `“python”` command, such as `“/usr/bin/python”`), you need to consider that your users may have either Python 2 or Python 3 as their default, and write your code to work under both versions.
- If you use an explicit version, for example `“/usr/bin/env python3”` your application will not work for users who do not have that version. (This may be what you want if you have not made your code Python 2 compatible).
- There is no way to say `“python X.Y or later”`, so be careful of using an exact version like `“/usr/bin/env python3.4”` as you will need to change your shebang line for users of Python 3.5, for example.

Typically, you should use an `“/usr/bin/env python2”` or `“/usr/bin/env python3”`, depending on whether your code is written for Python 2 or 3.

28.4.6 Creating Standalone Applications with zipapp

Using the `zipapp` module, it is possible to create self-contained Python programs, which can be distributed to end users who only need to have a suitable version of Python installed on their system. The key to doing this is to bundle all of the application’s dependencies into the archive, along with the application code.

The steps to create a standalone archive are as follows:

1. Create your application in a directory as normal, so you have a `myapp` directory containing a `__main__.py` file, and any supporting application code.
2. Install all of your application’s dependencies into the `myapp` directory, using `pip`:

```
$ python -m pip install -r requirements.txt --target myapp
```

(this assumes you have your project requirements in a `requirements.txt` file - if not, you can just list the dependencies manually on the pip command line).

3. Optionally, delete the `.dist-info` directories created by pip in the `myapp` directory. These hold metadata for pip to manage the packages, and as you won't be making any further use of pip they aren't required - although it won't do any harm if you leave them.

4. Package the application using:

```
$ python -m zipapp -p "interpreter" myapp
```

This will produce a standalone executable, which can be run on any machine with the appropriate interpreter available. See *Specifying the Interpreter* for details. It can be shipped to users as a single file.

On Unix, the `myapp.pyz` file is executable as it stands. You can rename the file to remove the `.pyz` extension if you prefer a “plain” command name. On Windows, the `myapp.pyz[w]` file is executable by virtue of the fact that the Python interpreter registers the `.pyz` and `.pyzw` file extensions when installed.

Making a Windows executable

On Windows, registration of the `.pyz` extension is optional, and furthermore, there are certain places that don't recognise registered extensions “transparently” (the simplest example is that `subprocess.run(['myapp'])` won't find your application - you need to explicitly specify the extension).

On Windows, therefore, it is often preferable to create an executable from the zipapp. This is relatively easy, although it does require a C compiler. The basic approach relies on the fact that zipfiles can have arbitrary data prepended, and Windows exe files can have arbitrary data appended. So by creating a suitable launcher and tacking the `.pyz` file onto the end of it, you end up with a single-file executable that runs your application.

A suitable launcher can be as simple as the following:

```
#define Py_LIMITED_API 1
#include "Python.h"

#define WIN32_LEAN_AND_MEAN
#include <windows.h>

#ifdef WINDOWS
int WINAPI wWinMain(
    HINSTANCE hInstance,      /* handle to current instance */
    HINSTANCE hPrevInstance,  /* handle to previous instance */
    LPWSTR lpCmdLine,         /* pointer to command line */
    int nCmdShow              /* show state of window */
)
#else
int wmain()
#endif
{
    wchar_t **myargv = _alloca((__argc + 1) * sizeof(wchar_t*));
    myargv[0] = __wargv[0];
    memcpy(myargv + 1, __wargv, __argc * sizeof(wchar_t *));
    return Py_Main(__argc+1, myargv);
}
```

If you define the `WINDOWS` preprocessor symbol, this will generate a GUI executable, and without it, a console executable.

To compile the executable, you can either just use the standard MSVC command line tools, or you can take advantage of the fact that `distutils` knows how to compile Python source:

```

>>> from distutils.ccompiler import new_compiler
>>> import distutils.sysconfig
>>> import sys
>>> import os
>>> from pathlib import Path

>>> def compile(src):
>>>     src = Path(src)
>>>     cc = new_compiler()
>>>     exe = src.stem
>>>     cc.add_include_dir(distutils.sysconfig.get_python_inc())
>>>     cc.add_library_dir(os.path.join(sys.base_exec_prefix, 'libs'))
>>>     # First the CLI executable
>>>     objs = cc.compile([str(src)])
>>>     cc.link_executable(objs, exe)
>>>     # Now the GUI executable
>>>     cc.define_macro('WINDOWS')
>>>     objs = cc.compile([str(src)])
>>>     cc.link_executable(objs, exe + '.w')

>>> if __name__ == "__main__":
>>>     compile("zastub.c")

```

The resulting launcher uses the “Limited ABI”, so it will run unchanged with any version of Python 3.x. All it needs is for Python (`python3.dll`) to be on the user’s PATH.

For a fully standalone distribution, you can distribute the launcher with your application appended, bundled with the Python “embedded” distribution. This will run on any PC with the appropriate architecture (32 bit or 64 bit).

Caveats

There are some limitations to the process of bundling your application into a single file. In most, if not all, cases they can be addressed without needing major changes to your application.

1. If your application depends on a package that includes a C extension, that package cannot be run from a zip file (this is an OS limitation, as executable code must be present in the filesystem for the OS loader to load it). In this case, you can exclude that dependency from the zipfile, and either require your users to have it installed, or ship it alongside your zipfile and add code to your `__main__.py` to include the directory containing the unzipped module in `sys.path`. In this case, you will need to make sure to ship appropriate binaries for your target architecture(s) (and potentially pick the correct version to add to `sys.path` at runtime, based on the user’s machine).
2. If you are shipping a Windows executable as described above, you either need to ensure that your users have `python3.dll` on their PATH (which is not the default behaviour of the installer) or you should bundle your application with the embedded distribution.
3. The suggested launcher above uses the Python embedding API. This means that in your application, `sys.executable` will be your application, and *not* a conventional Python interpreter. Your code and its dependencies need to be prepared for this possibility. For example, if your application uses the `multiprocessing` module, it will need to call `multiprocessing.set_executable()` to let the module know where to find the standard Python interpreter.

28.4.7 The Python Zip Application Archive Format

Python has been able to execute zip files which contain a `__main__.py` file since version 2.6. In order to be executed by Python, an application archive simply has to be a standard zip file containing a `__main__.py` file which will be run

as the entry point for the application. As usual for any Python script, the parent of the script (in this case the zip file) will be placed on `sys.path` and thus further modules can be imported from the zip file.

The zip file format allows arbitrary data to be prepended to a zip file. The zip application format uses this ability to prepend a standard POSIX “shebang” line to the file (`#!/path/to/interpreter`).

Formally, the Python zip application format is therefore:

1. An optional shebang line, containing the characters `b'#!'` followed by an interpreter name, and then a new-line (`b'\n'`) character. The interpreter name can be anything acceptable to the OS “shebang” processing, or the Python launcher on Windows. The interpreter should be encoded in UTF-8 on Windows, and in `sys.getfilesystemencoding()` on POSIX.
2. Standard zipfile data, as generated by the `zipfile` module. The zipfile content *must* include a file called `__main__.py` (which must be in the “root” of the zipfile - i.e., it cannot be in a subdirectory). The zipfile data can be compressed or uncompressed.

If an application archive has a shebang line, it may have the executable bit set on POSIX systems, to allow it to be executed directly.

There is no requirement that the tools in this module are used to create application archives - the module is a convenience, but archives in the above format created by any means are acceptable to Python.

本章里描述的模块提供了和 Python 解释器及其环境交互相关的广泛服务。以下是综述：

29.1 sys — 系统相关的参数和函数

该模块提供了一些变量和函数。这些变量可能被解释器使用，也可能由解释器提供。这些函数会影响解释器。本模块总是可用的。

`sys.abiflags`

在 POSIX 系统上，以标准的 `configure` 脚本构建的 Python 中，这个变量会包含 [PEP 3149](#) 中定义的 ABI 标签。

3.2 新版功能.

`sys.argv`

一个列表，其中包含了被传递给 Python 脚本的命令行参数。`argv[0]` 为脚本的名称（是否是完整的路径名取决于操作系统）。如果是通过 Python 解释器的命令行参数 `-c` 来执行的，`argv[0]` 会被设置成字符串 `'-c'`。如果没有脚本名被传递给 Python 解释器，`argv[0]` 为空字符串。

为了遍历标准输入，或者通过命令行传递的文件列表，参照 `fileinput` 模块

`sys.base_exec_prefix`

在 `site.py` 运行之前，Python 启动的时候被设置为跟 `exec_prefix` 同样的值。如果不是运行在虚拟环境中，两个值会保持相同；如果 `site.py` 发现处于一个虚拟环境中，`prefix` 和 `exec_prefix` 将会指向虚拟环境。然而 `base_prefix` 和 `base_exec_prefix` 将仍然会指向基础的 Python 环境（用来创建虚拟环境的 Python 环境）

3.3 新版功能.

`sys.base_prefix`

在 `site.py` 运行之前，Python 启动的时候被设置为跟 `prefix` 同样的值。如果不是运行在虚拟环境中，两个值会保持相同；如果 `site.py` 发现处于一个虚拟环境中，`prefix` 和 `exec_prefix` 将会指

向虚拟环境。然而`base_prefix`和`base_exec_prefix`将仍然会指向基础的 Python 环境（用来创建虚拟环境的 Python 环境）

3.3 新版功能.

`sys.byteorder`

本地字节顺序的指示符。在大端序（最高有效位优先）操作系统上值为 'big'，在小端序（最低有效位优先）操作系统上为 'little'。

`sys.builtin_module_names`

一个元素为字符串的元组。包含了所有的被编译进 Python 解释器的模块。（这个信息无法通过其他的办法获取，`modules.keys()` 只包括被导入过的模块。）

`sys.call_tracing(func, args)`

在启用跟踪时调用 `func(*args)` 来保存跟踪状态，然后恢复跟踪状态。这将从检查点的调试器调用，以便递归地调试其他的一些代码。

`sys.copyright`

一个字符串，包含了 Python 解释器有关的版权信息

`sys._clear_type_cache()`

清除内部的类型缓存。类型缓存是为了加速查找方法和属性的。在调试引用泄漏的时候调用这个函数只会清除不必要的引用。

这个函数应该只在内部为了一些特定的目的使用。

`sys._current_frames()`

返回一个字典，将每个线程的标识符映射到调用函数时该线程中当前活动的最顶层堆栈帧。注意`traceback`模块中的函数可以在给定帧的情况下构建调用堆栈。

This is most useful for debugging deadlock: this function does not require the deadlocked threads' cooperation, and such threads' call stacks are frozen for as long as they remain deadlocked. The frame returned for a non-deadlocked thread may bear no relationship to that thread' s current activity by the time calling code examines the frame.

这个函数应该只在内部为了一些特定的目的使用。

`sys._debugmallocstats()`

Print low-level information to stderr about the state of CPython' s memory allocator.

If Python is configured `--with-pydebug`, it also performs some expensive internal consistency checks.

3.3 新版功能.

CPython implementation detail: This function is specific to CPython. The exact output format is not defined here, and may change.

`sys.dllhandle`

Integer specifying the handle of the Python DLL. Availability: Windows.

`sys.displayhook(value)`

If `value` is not None, this function prints `repr(value)` to `sys.stdout`, and saves `value` in `builtins._`. If `repr(value)` is not encodable to `sys.stdout.encoding` with `sys.stdout.errors` error handler (which is probably 'strict'), encode it to `sys.stdout.encoding` with 'backslashreplace' error handler.

`sys.displayhook` is called on the result of evaluating an *expression* entered in an interactive Python session. The display of these values can be customized by assigning another one-argument function to `sys.displayhook`.

伪代码:

```
def displayhook(value):
    if value is None:
        return
    # Set '_' to None to avoid recursion
    builtins._ = None
    text = repr(value)
    try:
        sys.stdout.write(text)
    except UnicodeEncodeError:
        bytes = text.encode(sys.stdout.encoding, 'backslashreplace')
        if hasattr(sys.stdout, 'buffer'):
            sys.stdout.buffer.write(bytes)
        else:
            text = bytes.decode(sys.stdout.encoding, 'strict')
            sys.stdout.write(text)
    sys.stdout.write("\n")
    builtins._ = value
```

在 3.2 版更改: Use 'backslashreplace' error handler on *UnicodeEncodeError*.

`sys.dont_write_bytecode`

If this is true, Python won't try to write .pyc files on the import of source modules. This value is initially set to True or False depending on the -B command line option and the PYTHONDONTWRITEBYTECODE environment variable, but you can set it yourself to control bytecode file generation.

`sys.excepthook (type, value, traceback)`

This function prints out a given traceback and exception to `sys.stderr`.

When an exception is raised and uncaught, the interpreter calls `sys.excepthook` with three arguments, the exception class, exception instance, and a traceback object. In an interactive session this happens just before control is returned to the prompt; in a Python program this happens just before the program exits. The handling of such top-level exceptions can be customized by assigning another three-argument function to `sys.excepthook`.

`sys.__displayhook__`

`sys.__excepthook__`

These objects contain the original values of `displayhook` and `excepthook` at the start of the program. They are saved so that `displayhook` and `excepthook` can be restored in case they happen to get replaced with broken objects.

`sys.exc_info()`

This function returns a tuple of three values that give information about the exception that is currently being handled. The information returned is specific both to the current thread and to the current stack frame. If the current stack frame is not handling an exception, the information is taken from the calling stack frame, or its caller, and so on until a stack frame is found that is handling an exception. Here, “handling an exception” is defined as “executing an except clause.” For any stack frame, only information about the exception being currently handled is accessible.

If no exception is being handled anywhere on the stack, a tuple containing three `None` values is returned. Otherwise, the values returned are `(type, value, traceback)`. Their meaning is: *type* gets the type of the exception being handled (a subclass of *BaseException*); *value* gets the exception instance (an instance of the exception type); *traceback* gets a traceback object (see the Reference Manual) which encapsulates the call stack at the point where the exception originally occurred.

`sys.exec_prefix`

A string giving the site-specific directory prefix where the platform-dependent Python files are installed; by default, this is also `'/usr/local'`. This can be set at build time with the `--exec-prefix` argument to the **configure** script. Specifically, all configuration files (e.g. the `pyconfig.h` header file) are installed in the directory `exec_prefix/lib/pythonX.Y/config`, and shared library modules are installed in `exec_prefix/lib/pythonX.Y/lib-dynload`, where *X.Y* is the version number of Python, for example

3.2.

注解: If a *virtual environment* is in effect, this value will be changed in `site.py` to point to the virtual environment. The value for the Python installation will still be available, via `base_exec_prefix`.

sys.executable

A string giving the absolute path of the executable binary for the Python interpreter, on systems where this makes sense. If Python is unable to retrieve the real path to its executable, `sys.executable` will be an empty string or `None`.

sys.exit([arg])

Exit from Python. This is implemented by raising the `SystemExit` exception, so cleanup actions specified by finally clauses of `try` statements are honored, and it is possible to intercept the exit attempt at an outer level.

The optional argument *arg* can be an integer giving the exit status (defaulting to zero), or another type of object. If it is an integer, zero is considered “successful termination” and any nonzero value is considered “abnormal termination” by shells and the like. Most systems require it to be in the range 0–127, and produce undefined results otherwise. Some systems have a convention for assigning specific meanings to specific exit codes, but these are generally underdeveloped; Unix programs generally use 2 for command line syntax errors and 1 for all other kind of errors. If another type of object is passed, `None` is equivalent to passing zero, and any other object is printed to `stderr` and results in an exit code of 1. In particular, `sys.exit("some error message")` is a quick way to exit a program when an error occurs.

Since `exit()` ultimately “only” raises an exception, it will only exit the process when called from the main thread, and the exception is not intercepted.

在 3.6 版更改: If an error occurs in the cleanup after the Python interpreter has caught `SystemExit` (such as an error flushing buffered data in the standard streams), the exit status is changed to 120.

sys.flags

The *struct sequence* `flags` exposes the status of command line flags. The attributes are read only.

属性	flag
<code>debug</code>	<code>-d</code>
<code>inspect</code>	<code>-i</code>
<code>interactive</code>	<code>-i</code>
<code>isolated</code>	<code>-I</code>
<code>optimize</code>	<code>-O</code> or <code>-OO</code>
<code>dont_write_bytecode</code>	<code>-B</code>
<code>no_user_site</code>	<code>-s</code>
<code>no_site</code>	<code>-S</code>
<code>ignore_environment</code>	<code>-E</code>
<code>verbose</code>	<code>-v</code>
<code>bytes_warning</code>	<code>-b</code>
<code>quiet</code>	<code>-q</code>
<code>hash_randomization</code>	<code>-R</code>

在 3.2 版更改: Added `quiet` attribute for the new `-q` flag.

3.2.3 新版功能: The `hash_randomization` attribute.

在 3.3 版更改: Removed obsolete `division_warning` attribute.

在 3.4 版更改: Added `isolated` attribute for `-I` isolated flag.

sys.float_info

A *struct sequence* holding information about the float type. It contains low level information about the precision

and internal representation. The values correspond to the various floating-point constants defined in the standard header file `float.h` for the ‘C’ programming language; see section 5.2.4.2.2 of the 1999 ISO/IEC C standard [C99], ‘Characteristics of floating types’, for details.

属性	float.h macro	explanation
<code>epsilon</code>	<code>DBL_EPSILON</code>	difference between 1 and the least value greater than 1 that is representable as a float
<code>dig</code>	<code>DBL_DIG</code>	maximum number of decimal digits that can be faithfully represented in a float; see below
<code>mant_dig</code>	<code>DBL_MANT_DIG</code>	float precision: the number of base-radix digits in the significand of a float
<code>max</code>	<code>DBL_MAX</code>	maximum representable finite float
<code>max_exp</code>	<code>DBL_MAX_EXP</code>	maximum integer e such that $\text{radix}^{**}(e-1)$ is a representable finite float
<code>max_10_exp</code>	<code>DBL_MAX_10_EXP</code>	maximum integer e such that $10^{**}e$ is in the range of representable finite floats
<code>min</code>	<code>DBL_MIN</code>	minimum positive normalized float
<code>min_exp</code>	<code>DBL_MIN_EXP</code>	minimum integer e such that $\text{radix}^{**}(e-1)$ is a normalized float
<code>min_10_exp</code>	<code>DBL_MIN_10_EXP</code>	minimum integer e such that $10^{**}e$ is a normalized float
<code>radix</code>	<code>FLT_RADIX</code>	radix of exponent representation
<code>rounds</code>	<code>FLT_ROUNDS</code>	integer constant representing the rounding mode used for arithmetic operations. This reflects the value of the system <code>FLT_ROUNDS</code> macro at interpreter startup time. See section 5.2.4.2.2 of the C99 standard for an explanation of the possible values and their meanings.

The attribute `sys.float_info.dig` needs further explanation. If s is any string representing a decimal number with at most `sys.float_info.dig` significant digits, then converting s to a float and back again will recover a string representing the same decimal value:

```
>>> import sys
>>> sys.float_info.dig
15
>>> s = '3.14159265358979'      # decimal string with 15 significant digits
>>> format(float(s), '.15g')    # convert to float and back -> same value
'3.14159265358979'
```

But for strings with more than `sys.float_info.dig` significant digits, this isn't always true:

```
>>> s = '9876543211234567'     # 16 significant digits is too many!
>>> format(float(s), '.16g')    # conversion changes value
'9876543211234568'
```

`sys.float_repr_style`

A string indicating how the `repr()` function behaves for floats. If the string has value `'short'` then for a finite float x , `repr(x)` aims to produce a short string with the property that `float(repr(x)) == x`. This is the usual behaviour in Python 3.1 and later. Otherwise, `float_repr_style` has value `'legacy'` and `repr(x)` behaves in the same way as it did in versions of Python prior to 3.1.

3.1 新版功能.

`sys.getallocatedblocks()`

Return the number of memory blocks currently allocated by the interpreter, regardless of their size. This function is mainly useful for tracking and debugging memory leaks. Because of the interpreter's internal caches, the result

can vary from call to call; you may have to call `_clear_type_cache()` and `gc.collect()` to get more predictable results.

If a Python build or implementation cannot reasonably compute this information, `getallocatedblocks()` is allowed to return 0 instead.

3.4 新版功能.

`sys.getcheckinterval()`
Return the interpreter's "check interval"; see `setcheckinterval()`.

3.2 版后已移除: Use `getswitchinterval()` instead.

`sys.getdefaultencoding()`
Return the name of the current default string encoding used by the Unicode implementation.

`sys.getdlopenflags()`
Return the current value of the flags that are used for `dlopen()` calls. Symbolic names for the flag values can be found in the `os` module (RTLD_XXX constants, e.g. `os.RTLD_LAZY`). Availability: Unix.

`sys.getfilesystemencoding()`
Return the name of the encoding used to convert between Unicode filenames and bytes filenames. For best compatibility, str should be used for filenames in all cases, although representing filenames as bytes is also supported. Functions accepting or returning filenames should support either str or bytes and internally convert to the system's preferred representation.

This encoding is always ASCII-compatible.

`os.fsencode()` and `os.fsdecode()` should be used to ensure that the correct encoding and errors mode are used.

- On Mac OS X, the encoding is 'utf-8'.
- On Unix, the encoding is the locale encoding.
- On Windows, the encoding may be 'utf-8' or 'mbcs', depending on user configuration.

在 3.2 版更改: `getfilesystemencoding()` result cannot be None anymore.

在 3.6 版更改: Windows is no longer guaranteed to return 'mbcs'. See [PEP 529](#) and `_enablelegacywindowsfsencoding()` for more information.

`sys.getfilesystemencodeerrors()`
Return the name of the error mode used to convert between Unicode filenames and bytes filenames. The encoding name is returned from `getfilesystemencoding()`.

`os.fsencode()` and `os.fsdecode()` should be used to ensure that the correct encoding and errors mode are used.

3.6 新版功能.

`sys.getrefcount(object)`
Return the reference count of the *object*. The count returned is generally one higher than you might expect, because it includes the (temporary) reference as an argument to `getrefcount()`.

`sys.getrecursionlimit()`
Return the current value of the recursion limit, the maximum depth of the Python interpreter stack. This limit prevents infinite recursion from causing an overflow of the C stack and crashing Python. It can be set by `setrecursionlimit()`.

`sys.getsizeof(object[, default])`
Return the size of an object in bytes. The object can be any type of object. All built-in objects will return correct results, but this does not have to hold true for third-party extensions as it is implementation specific.

Only the memory consumption directly attributed to the object is accounted for, not the memory consumption of objects it refers to.

If given, *default* will be returned if the object does not provide means to retrieve the size. Otherwise a `TypeError` will be raised.

`getsizeof()` calls the object's `__sizeof__` method and adds an additional garbage collector overhead if the object is managed by the garbage collector.

See [recursive sizeof recipe](#) for an example of using `getsizeof()` recursively to find the size of containers and all their contents.

`sys.getswitchinterval()`

Return the interpreter's "thread switch interval"; see `setswitchinterval()`.

3.2 新版功能.

`sys._getframe([depth])`

Return a frame object from the call stack. If optional integer *depth* is given, return the frame object that many calls below the top of the stack. If that is deeper than the call stack, `ValueError` is raised. The default for *depth* is zero, returning the frame at the top of the call stack.

CPython implementation detail: This function should be used for internal and specialized purposes only. It is not guaranteed to exist in all implementations of Python.

`sys.getprofile()`

Get the profiler function as set by `setprofile()`.

`sys.gettrace()`

Get the trace function as set by `settrace()`.

CPython implementation detail: The `gettrace()` function is intended only for implementing debuggers, profilers, coverage tools and the like. Its behavior is part of the implementation platform, rather than part of the language definition, and thus may not be available in all Python implementations.

`sys.getwindowsversion()`

Return a named tuple describing the Windows version currently running. The named elements are *major*, *minor*, *build*, *platform*, *service_pack*, *service_pack_minor*, *service_pack_major*, *suite_mask*, *product_type* and *platform_version*. *service_pack* contains a string, *platform_version* a 3-tuple and all other values are integers. The components can also be accessed by name, so `sys.getwindowsversion()[0]` is equivalent to `sys.getwindowsversion().major`. For compatibility with prior versions, only the first 5 elements are retrievable by indexing.

platform will be 2 (`VER_PLATFORM_WIN32_NT`).

product_type may be one of the following values:

常数	意义
1 (<code>VER_NT_WORKSTATION</code>)	The system is a workstation.
2 (<code>VER_NT_DOMAIN_CONTROLLER</code>)	The system is a domain controller.
3 (<code>VER_NT_SERVER</code>)	The system is a server, but not a domain controller.

This function wraps the Win32 `GetVersionEx()` function; see the Microsoft documentation on `OSVERSIONINFOEX()` for more information about these fields.

platform_version returns the accurate major version, minor version and build number of the current operating system, rather than the version that is being emulated for the process. It is intended for use in logging rather than for feature detection.

Availability: Windows.

在 3.2 版更改: Changed to a named tuple and added *service_pack_minor*, *service_pack_major*, *suite_mask*, and *product_type*.

在 3.6 版更改: Added *platform_version*

`sys.get_asyncgen_hooks()`

Returns an *asyncgen_hooks* object, which is similar to a *namedtuple* of the form (*firstiter*, *finalizer*), where *firstiter* and *finalizer* are expected to be either `None` or functions which take an *asynchronous generator iterator* as an argument, and are used to schedule finalization of an asynchronous generator by an event loop.

3.6 新版功能: See [PEP 525](#) for more details.

注解: This function has been added on a provisional basis (see [PEP 411](#) for details.)

`sys.get_coroutine_wrapper()`

Returns `None`, or a wrapper set by *set_coroutine_wrapper()*.

3.5 新版功能: See [PEP 492](#) for more details.

注解: This function has been added on a provisional basis (see [PEP 411](#) for details.) Use it only for debugging purposes.

`sys.hash_info`

A *struct sequence* giving parameters of the numeric hash implementation. For more details about hashing of numeric types, see [数字类型的哈希运算](#).

属性	explanation
<code>width</code>	width in bits used for hash values
<code>modulus</code>	prime modulus P used for numeric hash scheme
<code>inf</code>	hash value returned for a positive infinity
<code>nan</code>	hash value returned for a nan
<code>imag</code>	multiplier used for the imaginary part of a complex number
<code>algorithm</code>	name of the algorithm for hashing of str, bytes, and memoryview
<code>hash_bits</code>	internal output size of the hash algorithm
<code>seed_bits</code>	size of the seed key of the hash algorithm

3.2 新版功能.

在 3.4 版更改: Added *algorithm*, *hash_bits* and *seed_bits*

`sys.hexversion`

The version number encoded as a single integer. This is guaranteed to increase with each version, including proper support for non-production releases. For example, to test that the Python interpreter is at least version 1.5.2, use:

```
if sys.hexversion >= 0x010502F0:
    # use some advanced feature
    ...
else:
    # use an alternative implementation or warn the user
    ...
```

This is called `hexversion` since it only really looks meaningful when viewed as the result of passing it to the built-in *hex()* function. The *struct sequence* *sys.version_info* may be used for a more human-friendly encoding of the same information.

More details of `hexversion` can be found at `apiabiversion`.

sys.implementation

An object containing information about the implementation of the currently running Python interpreter. The following attributes are required to exist in all Python implementations.

name is the implementation's identifier, e.g. `'cpython'`. The actual string is defined by the Python implementation, but it is guaranteed to be lower case.

version is a named tuple, in the same format as `sys.version_info`. It represents the version of the Python *implementation*. This has a distinct meaning from the specific version of the Python *language* to which the currently running interpreter conforms, which `sys.version_info` represents. For example, for PyPy 1.8 `sys.implementation.version` might be `sys.version_info(1, 8, 0, 'final', 0)`, whereas `sys.version_info` would be `sys.version_info(2, 7, 2, 'final', 0)`. For CPython they are the same value, since it is the reference implementation.

hexversion is the implementation version in hexadecimal format, like `sys.hexversion`.

cache_tag is the tag used by the import machinery in the filenames of cached modules. By convention, it would be a composite of the implementation's name and version, like `'cpython-33'`. However, a Python implementation may use some other value if appropriate. If *cache_tag* is set to `None`, it indicates that module caching should be disabled.

`sys.implementation` may contain additional attributes specific to the Python implementation. These non-standard attributes must start with an underscore, and are not described here. Regardless of its contents, `sys.implementation` will not change during a run of the interpreter, nor between implementation versions. (It may change between Python language versions, however.) See [PEP 421](#) for more information.

3.3 新版功能.

sys.int_info

A *struct sequence* that holds information about Python's internal representation of integers. The attributes are read only.

属性	解释
<code>bits_per_digit</code>	number of bits held in each digit. Python integers are stored internally in base $2^{**int_info.bits_per_digit}$
<code>sizeof_digit</code>	size in bytes of the C type used to represent a digit

3.1 新版功能.

sys.__interactivehook__

When this attribute exists, its value is automatically called (with no arguments) when the interpreter is launched in interactive mode. This is done after the `PYTHONSTARTUP` file is read, so that you can set this hook there. The `site` module *sets this*.

3.4 新版功能.

sys.intern (string)

Enter *string* in the table of “interned” strings and return the interned string – which is *string* itself or a copy. Interning strings is useful to gain a little performance on dictionary lookup – if the keys in a dictionary are interned, and the lookup key is interned, the key comparisons (after hashing) can be done by a pointer compare instead of a string compare. Normally, the names used in Python programs are automatically interned, and the dictionaries used to hold module, class or instance attributes have interned keys.

Interned strings are not immortal; you must keep a reference to the return value of `intern()` around to benefit from it.

sys.is_finalizing()

Return `True` if the Python interpreter is *shutting down*, `False` otherwise.

3.5 新版功能.

`sys.last_type`
`sys.last_value`
`sys.last_traceback`

These three variables are not always defined; they are set when an exception is not handled and the interpreter prints an error message and a stack traceback. Their intended use is to allow an interactive user to import a debugger module and engage in post-mortem debugging without having to re-execute the command that caused the error. (Typical use is `import pdb; pdb.pm()` to enter the post-mortem debugger; see [pdb](#) module for more information.)

The meaning of the variables is the same as that of the return values from `exc_info()` above.

`sys.maxsize`

An integer giving the maximum value a variable of type `Py_ssize_t` can take. It's usually $2^{*}31 - 1$ on a 32-bit platform and $2^{*}63 - 1$ on a 64-bit platform.

`sys.maxunicode`

An integer giving the value of the largest Unicode code point, i.e. 1114111 (`0x10FFFF` in hexadecimal).

在 3.3 版更改: Before [PEP 393](#), `sys.maxunicode` used to be either `0xFFFF` or `0x10FFFF`, depending on the configuration option that specified whether Unicode characters were stored as UCS-2 or UCS-4.

`sys.meta_path`

A list of *meta path finder* objects that have their `find_spec()` methods called to see if one of the objects can find the module to be imported. The `find_spec()` method is called with at least the absolute name of the module being imported. If the module to be imported is contained in a package, then the parent package's `__path__` attribute is passed in as a second argument. The method returns a *module spec*, or `None` if the module cannot be found.

参见:

`importlib.abc.MetaPathFinder` The abstract base class defining the interface of finder objects on `meta_path`.

`importlib.machinery.ModuleSpec` The concrete class which `find_spec()` should return instances of.

在 3.4 版更改: *Module specs* were introduced in Python 3.4, by [PEP 451](#). Earlier versions of Python looked for a method called `find_module()`. This is still called as a fallback if a `meta_path` entry doesn't have a `find_spec()` method.

`sys.modules`

This is a dictionary that maps module names to modules which have already been loaded. This can be manipulated to force reloading of modules and other tricks. However, replacing the dictionary will not necessarily work as expected and deleting essential items from the dictionary may cause Python to fail.

`sys.path`

A list of strings that specifies the search path for modules. Initialized from the environment variable `PYTHONPATH`, plus an installation-dependent default.

As initialized upon program startup, the first item of this list, `path[0]`, is the directory containing the script that was used to invoke the Python interpreter. If the script directory is not available (e.g. if the interpreter is invoked interactively or if the script is read from standard input), `path[0]` is the empty string, which directs Python to search modules in the current directory first. Notice that the script directory is inserted *before* the entries inserted as a result of `PYTHONPATH`.

A program is free to modify this list for its own purposes. Only strings and bytes should be added to `sys.path`; all other data types are ignored during import.

参见:

Module `site` This describes how to use `.pth` files to extend `sys.path`.

`sys.path_hooks`

A list of callables that take a path argument to try to create a *finder* for the path. If a finder can be created, it is to be returned by the callable, else raise `ImportError`.

Originally specified in **PEP 302**.

`sys.path_importer_cache`

A dictionary acting as a cache for *finder* objects. The keys are paths that have been passed to `sys.path_hooks` and the values are the finders that are found. If a path is a valid file system path but no finder is found on `sys.path_hooks` then `None` is stored.

Originally specified in **PEP 302**.

在 3.3 版更改: `None` is stored instead of `imp.NullImporter` when no finder is found.

`sys.platform`

This string contains a platform identifier that can be used to append platform-specific components to `sys.path`, for instance.

For Unix systems, except on Linux, this is the lowercased OS name as returned by `uname -s` with the first part of the version as returned by `uname -r` appended, e.g. `'sunos5'` or `'freebsd8'`, *at the time when Python was built*. Unless you want to test for a specific system version, it is therefore recommended to use the following idiom:

```
if sys.platform.startswith('freebsd'):
    # FreeBSD-specific code here...
elif sys.platform.startswith('linux'):
    # Linux-specific code here...
```

For other systems, the values are:

System	platform value
Linux	<code>'linux'</code>
Windows	<code>'win32'</code>
Windows/Cygwin	<code>'cygwin'</code>
Mac OS X	<code>'darwin'</code>

在 3.3 版更改: On Linux, `sys.platform` doesn't contain the major version anymore. It is always `'linux'`, instead of `'linux2'` or `'linux3'`. Since older Python versions include the version number, it is recommended to always use the `startswith` idiom presented above.

参见:

`os.name` has a coarser granularity. `os.uname()` gives system-dependent version information.

`platform` 模块提供了对系统标识更详细的检查。

`sys.prefix`

A string giving the site-specific directory prefix where the platform independent Python files are installed; by default, this is the string `'/usr/local'`. This can be set at build time with the `--prefix` argument to the **configure** script. The main collection of Python library modules is installed in the directory `prefix/lib/pythonX.Y` while the platform independent header files (all except `pyconfig.h`) are stored in `prefix/include/pythonX.Y`, where `X.Y` is the version number of Python, for example `3.2`.

注解: If a *virtual environment* is in effect, this value will be changed in `site.py` to point to the virtual environment. The value for the Python installation will still be available, via `base_prefix`.

`sys.ps1`

`sys.ps2`

Strings specifying the primary and secondary prompt of the interpreter. These are only defined if the interpreter is in interactive mode. Their initial values in this case are `'>>> '` and `'... '`. If a non-string object is assigned to either variable, its `str()` is re-evaluated each time the interpreter prepares to read a new interactive command; this can be used to implement a dynamic prompt.

`sys.setcheckinterval(interval)`

Set the interpreter's "check interval". This integer value determines how often the interpreter checks for periodic things such as thread switches and signal handlers. The default is 100, meaning the check is performed every 100 Python virtual instructions. Setting it to a larger value may increase performance for programs using threads. Setting it to a value ≤ 0 checks every virtual instruction, maximizing responsiveness as well as overhead.

3.2 版后已移除: This function doesn't have an effect anymore, as the internal logic for thread switching and asynchronous tasks has been rewritten. Use `setswitchinterval()` instead.

`sys.setdlopenflags(n)`

Set the flags used by the interpreter for `dlopen()` calls, such as when the interpreter loads extension modules. Among other things, this will enable a lazy resolving of symbols when importing a module, if called as `sys.setdlopenflags(0)`. To share symbols across extension modules, call as `sys.setdlopenflags(os.RTLD_GLOBAL)`. Symbolic names for the flag values can be found in the `os` module (RTLD_XXX constants, e.g. `os.RTLD_LAZY`).

Availability: Unix.

`sys.setprofile(profilefunc)`

Set the system's profile function, which allows you to implement a Python source code profiler in Python. See chapter *The Python Profilers* for more information on the Python profiler. The system's profile function is called similarly to the system's trace function (see `settrace()`), but it is called with different events, for example it isn't called for each executed line of code (only on call and return, but the return event is reported even when an exception has been set). The function is thread-specific, but there is no way for the profiler to know about context switches between threads, so it does not make sense to use this in the presence of multiple threads. Also, its return value is not used, so it can simply return `None`.

Profile functions should have three arguments: *frame*, *event*, and *arg*. *frame* is the current stack frame. *event* is a string: `'call'`, `'return'`, `'c_call'`, `'c_return'`, or `'c_exception'`. *arg* depends on the event type.

The events have the following meaning:

'call' A function is called (or some other code block entered). The profile function is called; *arg* is `None`.

'return' A function (or other code block) is about to return. The profile function is called; *arg* is the value that will be returned, or `None` if the event is caused by an exception being raised.

'c_call' A C function is about to be called. This may be an extension function or a built-in. *arg* is the C function object.

'c_return' A C function has returned. *arg* is the C function object.

'c_exception' A C function has raised an exception. *arg* is the C function object.

`sys.setrecursionlimit(limit)`

Set the maximum depth of the Python interpreter stack to *limit*. This limit prevents infinite recursion from causing an overflow of the C stack and crashing Python.

The highest possible limit is platform-dependent. A user may need to set the limit higher when they have a program that requires deep recursion and a platform that supports a higher limit. This should be done with care, because a too-high limit can lead to a crash.

If the new limit is too low at the current recursion depth, a `RecursionError` exception is raised.

在 3.5.1 版更改: A `RecursionError` exception is now raised if the new limit is too low at the current recursion depth.

`sys.setswitchinterval(interval)`

Set the interpreter's thread switch interval (in seconds). This floating-point value determines the ideal duration of the “timeslices” allocated to concurrently running Python threads. Please note that the actual value can be higher, especially if long-running internal functions or methods are used. Also, which thread becomes scheduled at the end of the interval is the operating system's decision. The interpreter doesn't have its own scheduler.

3.2 新版功能.

`sys.settrace(tracefunc)`

Set the system's trace function, which allows you to implement a Python source code debugger in Python. The function is thread-specific; for a debugger to support multiple threads, it must be registered using `settrace()` for each thread being debugged.

Trace functions should have three arguments: *frame*, *event*, and *arg*. *frame* is the current stack frame. *event* is a string: 'call', 'line', 'return' or 'exception'. *arg* depends on the event type.

The trace function is invoked (with *event* set to 'call') whenever a new local scope is entered; it should return a reference to a local trace function to be used that scope, or `None` if the scope shouldn't be traced.

The local trace function should return a reference to itself (or to another function for further tracing in that scope), or `None` to turn off tracing in that scope.

The events have the following meaning:

'call' A function is called (or some other code block entered). The global trace function is called; *arg* is `None`; the return value specifies the local trace function.

'line' The interpreter is about to execute a new line of code or re-execute the condition of a loop. The local trace function is called; *arg* is `None`; the return value specifies the new local trace function. See `Objects/lnotab_notes.txt` for a detailed explanation of how this works.

'return' A function (or other code block) is about to return. The local trace function is called; *arg* is the value that will be returned, or `None` if the event is caused by an exception being raised. The trace function's return value is ignored.

'exception' An exception has occurred. The local trace function is called; *arg* is a tuple (*exception*, *value*, *traceback*); the return value specifies the new local trace function.

Note that as an exception is propagated down the chain of callers, an 'exception' event is generated at each level.

For more information on code and frame objects, refer to types.

CPython implementation detail: The `settrace()` function is intended only for implementing debuggers, profilers, coverage tools and the like. Its behavior is part of the implementation platform, rather than part of the language definition, and thus may not be available in all Python implementations.

`sys.set_asyncgen_hooks(firstiter, finalizer)`

Accepts two optional keyword arguments which are callables that accept an *asynchronous generator iterator* as an argument. The *firstiter* callable will be called when an asynchronous generator is iterated for the first time. The *finalizer* will be called when an asynchronous generator is about to be garbage collected.

3.6 新版功能: See [PEP 525](#) for more details, and for a reference example of a *finalizer* method see the implementation of `asyncio.Loop.shutdown_asyncgens` in [Lib/asyncio/base_events.py](#)

注解: This function has been added on a provisional basis (see [PEP 411](#) for details.)

`sys.set_coroutine_wrapper(wrapper)`

Allows intercepting creation of *coroutine* objects (only ones that are created by an `async def` function; generators decorated with `types.coroutine()` or `asyncio.coroutine()` will not be intercepted).

The *wrapper* argument must be either:

- a callable that accepts one argument (a coroutine object);
- None, to reset the wrapper.

If called twice, the new wrapper replaces the previous one. The function is thread-specific.

The *wrapper* callable cannot define new coroutines directly or indirectly:

```
def wrapper(coro):
    async def wrap(coro):
        return await coro
    return wrap(coro)
sys.set_coroutine_wrapper(wrapper)

async def foo():
    pass

# The following line will fail with a RuntimeError, because
# ``wrapper`` creates a ``wrap(coro)`` coroutine:
foo()
```

See also `get_coroutine_wrapper()`.

3.5 新版功能: See [PEP 492](#) for more details.

注解: This function has been added on a provisional basis (see [PEP 411](#) for details.) Use it only for debugging purposes.

`sys._enablelegacywindowsfsencoding()`

Changes the default filesystem encoding and errors mode to ‘mbcs’ and ‘replace’ respectively, for consistency with versions of Python prior to 3.6.

This is equivalent to defining the `PYTHONLEGACYWINDOWSFSENCODING` environment variable before launching Python.

Availability: Windows

3.6 新版功能: 有关更多详细信息, 请参阅 [PEP 529](#)。

`sys.stdin`

`sys.stdout`

`sys.stderr`

File objects used by the interpreter for standard input, output and errors:

- `stdin` is used for all interactive input (including calls to `input()`);
- `stdout` is used for the output of `print()` and *expression* statements and for the prompts of `input()`;
- The interpreter’s own prompts and its error messages go to `stderr`.

These streams are regular *text files* like those returned by the `open()` function. Their parameters are chosen as follows:

- The character encoding is platform-dependent. Under Windows, if the stream is interactive (that is, if its `isatty()` method returns `True`), the console codepage is used, otherwise the ANSI code page. Under other platforms, the locale encoding is used (see `locale.getpreferredencoding()`).

Under all platforms though, you can override this value by setting the `PYTHONIOENCODING` environment variable before starting Python.

- When interactive, standard streams are line-buffered. Otherwise, they are block-buffered like regular text files. You can override this value with the `-u` command-line option.

注解: To write or read binary data from/to the standard streams, use the underlying binary `buffer` object. For example, to write bytes to `stdout`, use `sys.stdout.buffer.write(b'abc')`.

However, if you are writing a library (and do not control in which context its code will be executed), be aware that the standard streams may be replaced with file-like objects like `io.StringIO` which do not support the `buffer` attribute.

`sys.__stdin__`
`sys.__stdout__`
`sys.__stderr__`

These objects contain the original values of `stdin`, `stderr` and `stdout` at the start of the program. They are used during finalization, and could be useful to print to the actual standard stream no matter if the `sys.std*` object has been redirected.

It can also be used to restore the actual files to known working file objects in case they have been overwritten with a broken object. However, the preferred way to do this is to explicitly save the previous stream before replacing it, and restore the saved object.

注解: Under some conditions `stdin`, `stdout` and `stderr` as well as the original values `__stdin__`, `__stdout__` and `__stderr__` can be `None`. It is usually the case for Windows GUI apps that aren't connected to a console and Python apps started with `pythonw`.

`sys.thread_info`

A *struct sequence* holding information about the thread implementation.

属性	解释
<code>name</code>	Name of the thread implementation: <ul style="list-style-type: none">• <code>'nt'</code>: Windows threads• <code>'pthread'</code>: POSIX threads• <code>'solaris'</code>: Solaris threads
<code>lock</code>	Name of the lock implementation: <ul style="list-style-type: none">• <code>'semaphore'</code>: a lock uses a semaphore• <code>'mutex+cond'</code>: a lock uses a mutex and a condition variable• <code>None</code> if this information is unknown
<code>version</code>	Name and version of the thread library. It is a string, or <code>None</code> if these informations are unknown.

3.3 新版功能.

`sys.tracebacklimit`

When this variable is set to an integer value, it determines the maximum number of levels of traceback information

printed when an unhandled exception occurs. The default is 1000. When set to 0 or less, all traceback information is suppressed and only the exception type and value are printed.

`sys.version`

A string containing the version number of the Python interpreter plus additional information on the build number and compiler used. This string is displayed when the interactive interpreter is started. Do not extract version information out of it, rather, use `version_info` and the functions provided by the `platform` module.

`sys.api_version`

The C API version for this interpreter. Programmers may find this useful when debugging version conflicts between Python and extension modules.

`sys.version_info`

A tuple containing the five components of the version number: *major*, *minor*, *micro*, *releaselevel*, and *serial*. All values except *releaselevel* are integers; the release level is 'alpha', 'beta', 'candidate', or 'final'. The `version_info` value corresponding to the Python version 2.0 is (2, 0, 0, 'final', 0). The components can also be accessed by name, so `sys.version_info[0]` is equivalent to `sys.version_info.major` and so on.

在 3.1 版更改: Added named component attributes.

`sys.warnoptions`

This is an implementation detail of the warnings framework; do not modify this value. Refer to the `warnings` module for more information on the warnings framework.

`sys.winver`

The version number used to form registry keys on Windows platforms. This is stored as string resource 1000 in the Python DLL. The value is normally the first three characters of `version`. It is provided in the `sys` module for informational purposes; modifying this value has no effect on the registry keys used by Python. Availability: Windows.

`sys._xoptions`

A dictionary of the various implementation-specific flags passed through the `-X` command-line option. Option names are either mapped to their values, if given explicitly, or to `True`. Example:

```
$ ./python -Xa=b -Xc
Python 3.2a3+ (py3k, Oct 16 2010, 20:14:50)
[GCC 4.4.3] on linux2
Type "help", "copyright", "credits" or "license" for more information.
>>> import sys
>>> sys._xoptions
{'a': 'b', 'c': True}
```

CPython implementation detail: This is a CPython-specific way of accessing options passed through `-X`. Other implementations may export them through other means, or not at all.

3.2 新版功能.

Citations

29.2 sysconfig —Provide access to Python’ s configuration information

3.2 新版功能.

源代码: [Lib/sysconfig.py](#)

The `sysconfig` module provides access to Python's configuration information like the list of installation paths and the configuration variables relevant for the current platform.

29.2.1 配置变量

A Python distribution contains a `Makefile` and a `pyconfig.h` header file that are necessary to build both the Python binary itself and third-party C extensions compiled using `distutils`.

`sysconfig` puts all variables found in these files in a dictionary that can be accessed using `get_config_vars()` or `get_config_var()`.

Notice that on Windows, it's a much smaller set.

`sysconfig.get_config_vars(*args)`

With no arguments, return a dictionary of all configuration variables relevant for the current platform.

With arguments, return a list of values that result from looking up each argument in the configuration variable dictionary.

For each argument, if the value is not found, return `None`.

`sysconfig.get_config_var(name)`

Return the value of a single variable `name`. Equivalent to `get_config_vars().get(name)`.

If `name` is not found, return `None`.

用法示例:

```
>>> import sysconfig
>>> sysconfig.get_config_var('Py_ENABLE_SHARED')
0
>>> sysconfig.get_config_var('LIBDIR')
'/usr/local/lib'
>>> sysconfig.get_config_vars('AR', 'CXX')
['ar', 'g++']
```

29.2.2 安装路径

Python uses an installation scheme that differs depending on the platform and on the installation options. These schemes are stored in `sysconfig` under unique identifiers based on the value returned by `os.name`.

Every new component that is installed using `distutils` or a Distutils-based system will follow the same scheme to copy its file in the right places.

Python currently supports seven schemes:

- `posix_prefix`: scheme for Posix platforms like Linux or Mac OS X. This is the default scheme used when Python or a component is installed.
- `posix_home`: scheme for Posix platforms used when a `home` option is used upon installation. This scheme is used when a component is installed through Distutils with a specific home prefix.
- `posix_user`: scheme for Posix platforms used when a component is installed through Distutils and the `user` option is used. This scheme defines paths located under the user home directory.
- `nt`: scheme for NT platforms like Windows.
- `nt_user`: scheme for NT platforms, when the `user` option is used.

Each scheme is itself composed of a series of paths and each path has a unique identifier. Python currently uses eight paths:

- *stdlib*: directory containing the standard Python library files that are not platform-specific.
- *platstdlib*: directory containing the standard Python library files that are platform-specific.
- *platlib*: directory for site-specific, platform-specific files.
- *purelib*: directory for site-specific, non-platform-specific files.
- *include*: directory for non-platform-specific header files.
- *platinclude*: directory for platform-specific header files.
- *scripts*: directory for script files.
- *data*: directory for data files.

sysconfig provides some functions to determine these paths.

`sysconfig.get_scheme_names()`

Return a tuple containing all schemes currently supported in *sysconfig*.

`sysconfig.get_path_names()`

Return a tuple containing all path names currently supported in *sysconfig*.

`sysconfig.get_path(name[, scheme[, vars[, expand]]])`

Return an installation path corresponding to the path *name*, from the install scheme named *scheme*.

name has to be a value from the list returned by `get_path_names()`.

sysconfig stores installation paths corresponding to each path name, for each platform, with variables to be expanded. For instance the *stdlib* path for the *nt* scheme is: {base}/Lib.

`get_path()` will use the variables returned by `get_config_vars()` to expand the path. All variables have default values for each platform so one may call this function and get the default value.

If *scheme* is provided, it must be a value from the list returned by `get_scheme_names()`. Otherwise, the default scheme for the current platform is used.

If *vars* is provided, it must be a dictionary of variables that will update the dictionary return by `get_config_vars()`.

If *expand* is set to `False`, the path will not be expanded using the variables.

If *name* is not found, return `None`.

`sysconfig.get_paths([scheme[, vars[, expand]]])`

Return a dictionary containing all installation paths corresponding to an installation scheme. See `get_path()` for more information.

If *scheme* is not provided, will use the default scheme for the current platform.

If *vars* is provided, it must be a dictionary of variables that will update the dictionary used to expand the paths.

If *expand* is set to `false`, the paths will not be expanded.

If *scheme* is not an existing scheme, `get_paths()` will raise a *KeyError*.

29.2.3 其他功能

`sysconfig.get_python_version()`

Return the MAJOR.MINOR Python version number as a string. Similar to '%d.%d' % sys.version_info[:2].

`sysconfig.get_platform()`

Return a string that identifies the current platform.

This is used mainly to distinguish platform-specific build directories and platform-specific built distributions. Typically includes the OS name and version and the architecture (as supplied by `os.uname()`), although the exact information included depends on the OS; e.g. for IRIX the architecture isn't particularly important (IRIX only runs on SGI hardware), but for Linux the kernel version isn't particularly important.

返回值的示例：

- linux-i586
- linux-alpha (?)
- solaris-2.6-sun4u
- irix-5.3
- irix64-6.2

Windows 将返回以下之一：

- win-amd64 (在 AMD64, aka x86_64, Intel64, 和 EM64T 上的 64 位 Windows)
- win-ia64 (64bit Windows on Itanium)
- win32 (all others - specifically, `sys.platform` is returned)

Mac OS X 返回：

- macosx-10.6-ppc
- macosx-10.4-ppc64
- macosx-10.3-i386
- macosx-10.4-fat

对于其他非-POSIX 平台，目前只是返回 `sys.platform`。

`sysconfig.is_python_build()`

Return True if the running Python interpreter was built from source and is being run from its built location, and not from a location resulting from e.g. running `make install` or installing via a binary installer.

`sysconfig.parse_config_h(fp[, vars])`

Parse a `config.h`-style file.

`fp` is a file-like object pointing to the `config.h`-like file.

A dictionary containing name/value pairs is returned. If an optional dictionary is passed in as the second argument, it is used instead of a new dictionary, and updated with the values read in the file.

`sysconfig.get_config_h_filename()`

返回 `pyconfig.h` 的目录

`sysconfig.get_makefile_filename()`

返回 `Makefile` 的目录

29.2.4 Using `sysconfig` as a script

You can use `sysconfig` as a script with Python's `-m` option:

```
$ python -m sysconfig
Platform: "macosx-10.4-i386"
Python version: "3.2"
Current installation scheme: "posix_prefix"

Paths:
    data = "/usr/local"
    include = "/Users/tarek/Dev/svn.python.org/py3k/Include"
    platinclude = "."
    platlib = "/usr/local/lib/python3.2/site-packages"
    platstdlib = "/usr/local/lib/python3.2"
    purelib = "/usr/local/lib/python3.2/site-packages"
    scripts = "/usr/local/bin"
    stdlib = "/usr/local/lib/python3.2"

Variables:
    AC_APPLE_UNIVERSAL_BUILD = "0"
    AIX_GENUINE_CPLUSPLUS = "0"
    AR = "ar"
    ARFLAGS = "rc"
    ...
```

This call will print in the standard output the information returned by `get_platform()`, `get_python_version()`, `get_path()` and `get_config_vars()`.

29.3 builtins — 内建对象

该模块提供对 Python 的所有“内置”标识符的直接访问；例如，`builtins.open` 是内置函数的全名 `open()`。请参阅 [内置函数](#) 和 [内置常量](#) 的文档。

大多数应用程序通常不会显式访问此模块，但在提供与内置值同名的对象的模块中可能很有用，但其中还需要内置该名称。例如，在一个想要实现 `open()` 函数的模块中，它包装了内置的 `open()`，这个模块可以直接使用：

```
import builtins

def open(path):
    f = builtins.open(path, 'r')
    return UpperCaser(f)

class UpperCaser:
    '''Wrapper around a file that converts output to upper-case.'''

    def __init__(self, f):
        self._f = f

    def read(self, count=-1):
        return self._f.read(count).upper()

    # ...
```

作为一个实现细节，大多数模块都将名称 `__builtins__` 作为其全局变量的一部分提供。`__builtins__` 的值通常是这个模块或者这个模块的值 `__dict__` 属性。由于这是一个实现细节，因此 Python 的替代实现可能不会使用它。

29.4 `__main__` — 顶层脚本环境

'`__main__`' 是顶层代码执行的作用域的名称。模块的 `__name__` 在通过标准输入、脚本文件或是交互式命令读入的时候会等于 '`__main__`'。

模块可以通过检查自己的 `__name__` 来得知是否运行在 `main` 作用域中，这使得模块可以在作为脚本或是通过 `python -m` 运行时条件性地执行一些代码，而在被 `import` 时不会执行。

```
if __name__ == "__main__":
    # execute only if run as a script
    main()
```

对软件包来说，通过加入 `__main__.py` 模块可以达到同样的效果，当使用 `-m` 运行模块时，其中的代码会被执行。

29.5 `warnings` — Warning control

源代码： `Lib/warnings.py`

Warning messages are typically issued in situations where it is useful to alert the user of some condition in a program, where that condition (normally) doesn't warrant raising an exception and terminating the program. For example, one might want to issue a warning when a program uses an obsolete module.

Python programmers issue warnings by calling the `warn()` function defined in this module. (C programmers use `PyErr_WarnEx()`; see exceptionhandling for details).

Warning messages are normally written to `sys.stderr`, but their disposition can be changed flexibly, from ignoring all warnings to turning them into exceptions. The disposition of warnings can vary based on the warning category (see below), the text of the warning message, and the source location where it is issued. Repetitions of a particular warning for the same source location are typically suppressed.

There are two stages in warning control: first, each time a warning is issued, a determination is made whether a message should be issued or not; next, if a message is to be issued, it is formatted and printed using a user-settable hook.

The determination whether to issue a warning message is controlled by the warning filter, which is a sequence of matching rules and actions. Rules can be added to the filter by calling `filterwarnings()` and reset to its default state by calling `resetwarnings()`.

The printing of warning messages is done by calling `showwarning()`, which may be overridden; the default implementation of this function formats the message by calling `formatwarning()`, which is also available for use by custom implementations.

参见：

`logging.captureWarnings()` allows you to handle all warnings with the standard logging infrastructure.

29.5.1 警告类别

There are a number of built-in exceptions that represent warning categories. This categorization is useful to be able to filter out groups of warnings. The following warnings category classes are currently defined:

Class	描述
<code>Warning</code>	This is the base class of all warning category classes. It is a subclass of <code>Exception</code> .
<code>UserWarning</code>	The default category for <code>warn()</code> .
<code>DeprecationWarning</code>	Base category for warnings about deprecated features (ignored by default).
<code>SyntaxWarning</code>	Base category for warnings about dubious syntactic features.
<code>RuntimeWarning</code>	Base category for warnings about dubious runtime features.
<code>FutureWarning</code>	Base category for warnings about constructs that will change semantically in the future.
<code>PendingDeprecationWarning</code>	Base category for warnings about features that will be deprecated in the future (ignored by default).
<code>ImportWarning</code>	Base category for warnings triggered during the process of importing a module (ignored by default).
<code>UnicodeWarning</code>	Base category for warnings related to Unicode.
<code>BytesWarning</code>	Base category for warnings related to <code>bytes</code> and <code>bytearray</code> .
<code>ResourceWarning</code>	Base category for warnings related to resource usage.

While these are technically built-in exceptions, they are documented here, because conceptually they belong to the warnings mechanism.

User code can define additional warning categories by subclassing one of the standard warning categories. A warning category must always be a subclass of the `Warning` class.

29.5.2 The Warnings Filter

The warnings filter controls whether warnings are ignored, displayed, or turned into errors (raising an exception).

Conceptually, the warnings filter maintains an ordered list of filter specifications; any specific warning is matched against each filter specification in the list in turn until a match is found; the match determines the disposition of the match. Each entry is a tuple of the form *(action, message, category, module, lineno)*, where:

- action* is one of the following strings:

值	处置
"error"	将匹配警告转换为异常
"ignore"	从不打印匹配的警告
"always"	总是打印匹配的警告
"default"	print the first occurrence of matching warnings for each location where the warning is issued
"module"	print the first occurrence of matching warnings for each module where the warning is issued
"once"	无论位置如何，仅打印第一次出现的匹配警告

- message* is a string containing a regular expression that the start of the warning message must match. The expression is compiled to always be case-insensitive.
- category* is a class (a subclass of `Warning`) of which the warning category must be a subclass in order to match.
- module* is a string containing a regular expression that the module name must match. The expression is compiled to be case-sensitive.
- lineno* is an integer that the line number where the warning occurred must match, or 0 to match all line numbers.

Since the `Warning` class is derived from the built-in `Exception` class, to turn a warning into an error we simply raise `category(message)`.

The warnings filter is initialized by `-W` options passed to the Python interpreter command line. The interpreter saves the arguments for all `-W` options without interpretation in `sys.warnoptions`; the `warnings` module parses these when it is first imported (invalid options are ignored, after printing a message to `sys.stderr`).

Default Warning Filters

By default, Python installs several warning filters, which can be overridden by the command-line options passed to `-W` and calls to `filterwarnings()`.

- `DeprecationWarning` and `PendingDeprecationWarning`, and `ImportWarning` are ignored.
- `BytesWarning` is ignored unless the `-b` option is given once or twice; in this case this warning is either printed (`-b`) or turned into an exception (`-bb`).
- `ResourceWarning` is ignored unless Python was built in debug mode.

在 3.2 版更改: `DeprecationWarning` is now ignored by default in addition to `PendingDeprecationWarning`.

29.5.3 暂时禁止警告

If you are using code that you know will raise a warning, such as a deprecated function, but do not want to see the warning, then it is possible to suppress the warning using the `catch_warnings` context manager:

```
import warnings

def fxn():
    warnings.warn("deprecated", DeprecationWarning)

with warnings.catch_warnings():
    warnings.simplefilter("ignore")
    fxn()
```

While within the context manager all warnings will simply be ignored. This allows you to use known-deprecated code without having to see the warning while not suppressing the warning for other code that might not be aware of its use of deprecated code. Note: this can only be guaranteed in a single-threaded application. If two or more threads use the `catch_warnings` context manager at the same time, the behavior is undefined.

29.5.4 测试警告

To test warnings raised by code, use the `catch_warnings` context manager. With it you can temporarily mutate the warnings filter to facilitate your testing. For instance, do the following to capture all raised warnings to check:

```
import warnings

def fxn():
    warnings.warn("deprecated", DeprecationWarning)

with warnings.catch_warnings(record=True) as w:
    # Cause all warnings to always be triggered.
    warnings.simplefilter("always")
    # Trigger a warning.
    fxn()
    # Verify some things
```

(下页继续)

(续上页)

```

assert len(w) == 1
assert isinstance(w[-1].category, DeprecationWarning)
assert "deprecated" in str(w[-1].message)

```

One can also cause all warnings to be exceptions by using `error` instead of `always`. One thing to be aware of is that if a warning has already been raised because of a `once/default` rule, then no matter what filters are set the warning will not be seen again unless the warnings registry related to the warning has been cleared.

Once the context manager exits, the warnings filter is restored to its state when the context was entered. This prevents tests from changing the warnings filter in unexpected ways between tests and leading to indeterminate test results. The `showwarning()` function in the module is also restored to its original value. Note: this can only be guaranteed in a single-threaded application. If two or more threads use the `catch_warnings` context manager at the same time, the behavior is undefined.

When testing multiple operations that raise the same kind of warning, it is important to test them in a manner that confirms each operation is raising a new warning (e.g. set warnings to be raised as exceptions and check the operations raise exceptions, check that the length of the warning list continues to increase after each operation, or else delete the previous entries from the warnings list before each new operation).

29.5.5 Updating Code For New Versions of Python

Warnings that are only of interest to the developer are ignored by default. As such you should make sure to test your code with typically ignored warnings made visible. You can do this from the command-line by passing `-Wd` to the interpreter (this is shorthand for `-W default`). This enables default handling for all warnings, including those that are ignored by default. To change what action is taken for encountered warnings you simply change what argument is passed to `-W`, e.g. `-W error`. See the `-W` flag for more details on what is possible.

To programmatically do the same as `-Wd`, use:

```
warnings.simplefilter('default')
```

Make sure to execute this code as soon as possible. This prevents the registering of what warnings have been raised from unexpectedly influencing how future warnings are treated.

Having certain warnings ignored by default is done to prevent a user from seeing warnings that are only of interest to the developer. As you do not necessarily have control over what interpreter a user uses to run their code, it is possible that a new version of Python will be released between your release cycles. The new interpreter release could trigger new warnings in your code that were not there in an older interpreter, e.g. `DeprecationWarning` for a module that you are using. While you as a developer want to be notified that your code is using a deprecated module, to a user this information is essentially noise and provides no benefit to them.

The `unittest` module has been also updated to use the `'default'` filter while running tests.

29.5.6 Available Functions

`warnings.warn(message, category=None, stacklevel=1, source=None)`

Issue a warning, or maybe ignore it or raise an exception. The `category` argument, if given, must be a warning category class (see above); it defaults to `UserWarning`. Alternatively `message` can be a `Warning` instance, in which case `category` will be ignored and `message.__class__` will be used. In this case the message text will be `str(message)`. This function raises an exception if the particular warning issued is changed into an error by the warnings filter see above. The `stacklevel` argument can be used by wrapper functions written in Python, like this:

```
def deprecation(message):
    warnings.warn(message, DeprecationWarning, stacklevel=2)
```

This makes the warning refer to `deprecation()`'s caller, rather than to the source of `deprecation()` itself (since the latter would defeat the purpose of the warning message).

source, if supplied, is the destroyed object which emitted a *ResourceWarning*.

在 3.6 版更改: Added *source* parameter.

```
warnings.warn_explicit(message, category, filename, lineno, module=None, registry=None, module_globals=None, source=None)
```

This is a low-level interface to the functionality of `warn()`, passing in explicitly the message, category, filename and line number, and optionally the module name and the registry (which should be the `__warningregistry__` dictionary of the module). The module name defaults to the filename with `.py` stripped; if no registry is passed, the warning is never suppressed. *message* must be a string and *category* a subclass of *Warning* or *message* may be a *Warning* instance, in which case *category* will be ignored.

module_globals, if supplied, should be the global namespace in use by the code for which the warning is issued. (This argument is used to support displaying source for modules found in zipfiles or other non-filesystem import sources).

source, if supplied, is the destroyed object which emitted a *ResourceWarning*.

在 3.6 版更改: Add the *source* parameter.

```
warnings.showwarning(message, category, filename, lineno, file=None, line=None)
```

Write a warning to a file. The default implementation calls `formatwarning(message, category, filename, lineno, line)` and writes the resulting string to *file*, which defaults to `sys.stderr`. You may replace this function with any callable by assigning to `warnings.showwarning`. *line* is a line of source code to be included in the warning message; if *line* is not supplied, `showwarning()` will try to read the line specified by *filename* and *lineno*.

```
warnings.formatwarning(message, category, filename, lineno, line=None)
```

Format a warning the standard way. This returns a string which may contain embedded newlines and ends in a newline. *line* is a line of source code to be included in the warning message; if *line* is not supplied, `formatwarning()` will try to read the line specified by *filename* and *lineno*.

```
warnings.filterwarnings(action, message="", category=Warning, module="", lineno=0, append=False)
```

Insert an entry into the list of *warnings filter specifications*. The entry is inserted at the front by default; if *append* is true, it is inserted at the end. This checks the types of the arguments, compiles the *message* and *module* regular expressions, and inserts them as a tuple in the list of warnings filters. Entries closer to the front of the list override entries later in the list, if both match a particular warning. Omitted arguments default to a value that matches everything.

```
warnings.simplefilter(action, category=Warning, lineno=0, append=False)
```

Insert a simple entry into the list of *warnings filter specifications*. The meaning of the function parameters is as for `filterwarnings()`, but regular expressions are not needed as the filter inserted always matches any message in any module as long as the category and line number match.

```
warnings.resetwarnings()
```

Reset the warnings filter. This discards the effect of all previous calls to `filterwarnings()`, including that of the `-W` command line options and calls to `simplefilter()`.

29.5.7 Available Context Managers

```
class warnings.catch_warnings(*, record=False, module=None)
```

A context manager that copies and, upon exit, restores the warnings filter and the `showwarning()` function. If

the *record* argument is *False* (the default) the context manager returns *None* on entry. If *record* is *True*, a list is returned that is progressively populated with objects as seen by a custom *showwarning()* function (which also suppresses output to `sys.stdout`). Each object in the list has attributes with the same names as the arguments to *showwarning()*.

The *module* argument takes a module that will be used instead of the module returned when you import *warnings* whose filter will be protected. This argument exists primarily for testing the *warnings* module itself.

注解: The *catch_warnings* manager works by replacing and then later restoring the module's *showwarning()* function and internal list of filter specifications. This means the context manager is modifying global state and therefore is not thread-safe.

29.6 contextlib — Utilities for with-statement contexts

源代码 `Lib/contextlib.py`

此模块为涉及 `with` 语句的常见任务提供了实用的程序。更多信息请参见[上下文管理器类型](#) 和 `context-managers`。

29.6.1 工具

提供的函数和类：

class `contextlib.AbstractContextManager`

An *abstract base class* for classes that implement `object.__enter__()` and `object.__exit__()`. A default implementation for `object.__enter__()` is provided which returns `self` while `object.__exit__()` is an abstract method which by default returns `None`. See also the definition of [上下文管理器类型](#).

3.6 新版功能.

`@contextlib.contextmanager`

This function is a *decorator* that can be used to define a factory function for `with` statement context managers, without needing to create a class or separate `__enter__()` and `__exit__()` methods.

While many objects natively support use in `with` statements, sometimes a resource needs to be managed that isn't a context manager in its own right, and doesn't implement a `close()` method for use with `contextlib.closing`.

An abstract example would be the following to ensure correct resource management:

```
from contextlib import contextmanager

@contextmanager
def managed_resource(*args, **kwargs):
    # Code to acquire resource, e.g.:
    resource = acquire_resource(*args, **kwargs)
    try:
        yield resource
    finally:
        # Code to release resource, e.g.:
        release_resource(resource)
```

(下页继续)

(续上页)

```
>>> with managed_resource(timeout=3600) as resource:
...     # Resource is released at the end of this block,
...     # even if code in the block raises an exception
```

The function being decorated must return a *generator*-iterator when called. This iterator must yield exactly one value, which will be bound to the targets in the `with` statement's `as` clause, if any.

At the point where the generator yields, the block nested in the `with` statement is executed. The generator is then resumed after the block is exited. If an unhandled exception occurs in the block, it is reraised inside the generator at the point where the yield occurred. Thus, you can use a `try...except...finally` statement to trap the error (if any), or ensure that some cleanup takes place. If an exception is trapped merely in order to log it or to perform some action (rather than to suppress it entirely), the generator must reraise that exception. Otherwise the generator context manager will indicate to the `with` statement that the exception has been handled, and execution will resume with the statement immediately following the `with` statement.

`contextmanager()` uses *ContextDecorator* so the context managers it creates can be used as decorators as well as in `with` statements. When used as a decorator, a new generator instance is implicitly created on each function call (this allows the otherwise “one-shot” context managers created by `contextmanager()` to meet the requirement that context managers support multiple invocations in order to be used as decorators).

在 3.2 版更改: Use of *ContextDecorator*.

`contextlib.closing(thing)`

Return a context manager that closes *thing* upon completion of the block. This is basically equivalent to:

```
from contextlib import contextmanager

@contextmanager
def closing(thing):
    try:
        yield thing
    finally:
        thing.close()
```

And lets you write code like this:

```
from contextlib import closing
from urllib.request import urlopen

with closing(urlopen('http://www.python.org')) as page:
    for line in page:
        print(line)
```

without needing to explicitly close `page`. Even if an error occurs, `page.close()` will be called when the `with` block is exited.

`contextlib.suppress(*exceptions)`

Return a context manager that suppresses any of the specified exceptions if they occur in the body of a `with` statement and then resumes execution with the first statement following the end of the `with` statement.

As with any other mechanism that completely suppresses exceptions, this context manager should be used only to cover very specific errors where silently continuing with program execution is known to be the right thing to do.

例如:

```
from contextlib import suppress
```

(下页继续)

(续上页)

```

with suppress(FileNotFoundError) :
    os.remove('somefile.tmp')

with suppress(FileNotFoundError) :
    os.remove('someotherfile.tmp')

```

This code is equivalent to:

```

try:
    os.remove('somefile.tmp')
except FileNotFoundError:
    pass

try:
    os.remove('someotherfile.tmp')
except FileNotFoundError:
    pass

```

This context manager is *reentrant*.

3.4 新版功能.

`contextlib.redirect_stdout` (*new_target*)

Context manager for temporarily redirecting `sys.stdout` to another file or file-like object.

This tool adds flexibility to existing functions or classes whose output is hardwired to stdout.

For example, the output of `help()` normally is sent to `sys.stdout`. You can capture that output in a string by redirecting the output to an `io.StringIO` object:

```

f = io.StringIO()
with redirect_stdout(f):
    help(pow)
s = f.getvalue()

```

To send the output of `help()` to a file on disk, redirect the output to a regular file:

```

with open('help.txt', 'w') as f:
    with redirect_stdout(f):
        help(pow)

```

To send the output of `help()` to `sys.stderr`:

```

with redirect_stdout(sys.stderr):
    help(pow)

```

Note that the global side effect on `sys.stdout` means that this context manager is not suitable for use in library code and most threaded applications. It also has no effect on the output of subprocesses. However, it is still a useful approach for many utility scripts.

This context manager is *reentrant*.

3.4 新版功能.

`contextlib.redirect_stderr` (*new_target*)

Similar to `redirect_stdout()` but redirecting `sys.stderr` to another file or file-like object.

This context manager is *reentrant*.

3.5 新版功能.

class contextlib.ContextDecorator

A base class that enables a context manager to also be used as a decorator.

Context managers inheriting from `ContextDecorator` have to implement `__enter__` and `__exit__` as normal. `__exit__` retains its optional exception handling even when used as a decorator.

`ContextDecorator` is used by `contextmanager()`, so you get this functionality automatically.

`ContextDecorator` 的示例:

```
from contextlib import ContextDecorator

class mycontext(ContextDecorator):
    def __enter__(self):
        print('Starting')
        return self

    def __exit__(self, *exc):
        print('Finishing')
        return False

>>> @mycontext()
... def function():
...     print('The bit in the middle')
...
>>> function()
Starting
The bit in the middle
Finishing

>>> with mycontext():
...     print('The bit in the middle')
...
Starting
The bit in the middle
Finishing
```

This change is just syntactic sugar for any construct of the following form:

```
def f():
    with cm():
        # Do stuff
```

`ContextDecorator` lets you instead write:

```
@cm()
def f():
    # Do stuff
```

It makes it clear that the `cm` applies to the whole function, rather than just a piece of it (and saving an indentation level is nice, too).

Existing context managers that already have a base class can be extended by using `ContextDecorator` as a mixin class:

```
from contextlib import ContextDecorator

class mycontext(ContextBaseClass, ContextDecorator):
    def __enter__(self):
```

(下页继续)

(续上页)

```

    return self

def __exit__(self, *exc):
    return False

```

注解: As the decorated function must be able to be called multiple times, the underlying context manager must support use in multiple `with` statements. If this is not the case, then the original construct with the explicit `with` statement inside the function should be used.

3.2 新版功能.

class contextlib.ExitStack

A context manager that is designed to make it easy to programmatically combine other context managers and cleanup functions, especially those that are optional or otherwise driven by input data.

For example, a set of files may easily be handled in a single `with` statement as follows:

```

with ExitStack() as stack:
    files = [stack.enter_context(open(fname)) for fname in filenames]
    # All opened files will automatically be closed at the end of
    # the with statement, even if attempts to open files later
    # in the list raise an exception

```

Each instance maintains a stack of registered callbacks that are called in reverse order when the instance is closed (either explicitly or implicitly at the end of a `with` statement). Note that callbacks are *not* invoked implicitly when the context stack instance is garbage collected.

This stack model is used so that context managers that acquire their resources in their `__init__` method (such as file objects) can be handled correctly.

Since registered callbacks are invoked in the reverse order of registration, this ends up behaving as if multiple nested `with` statements had been used with the registered set of callbacks. This even extends to exception handling - if an inner callback suppresses or replaces an exception, then outer callbacks will be passed arguments based on that updated state.

This is a relatively low level API that takes care of the details of correctly unwinding the stack of exit callbacks. It provides a suitable foundation for higher level context managers that manipulate the exit stack in application specific ways.

3.3 新版功能.

enter_context (cm)

Enters a new context manager and adds its `__exit__()` method to the callback stack. The return value is the result of the context manager's own `__enter__()` method.

These context managers may suppress exceptions just as they normally would if used directly as part of a `with` statement.

push (exit)

Adds a context manager's `__exit__()` method to the callback stack.

As `__enter__` is *not* invoked, this method can be used to cover part of an `__enter__()` implementation with a context manager's own `__exit__()` method.

If passed an object that is not a context manager, this method assumes it is a callback with the same signature as a context manager's `__exit__()` method and adds it directly to the callback stack.

By returning true values, these callbacks can suppress exceptions the same way context manager `__exit__()` methods can.

The passed in object is returned from the function, allowing this method to be used as a function decorator.

callback (*callback*, *args, **kws)

Accepts an arbitrary callback function and arguments and adds it to the callback stack.

Unlike the other methods, callbacks added this way cannot suppress exceptions (as they are never passed the exception details).

The passed in callback is returned from the function, allowing this method to be used as a function decorator.

pop_all()

Transfers the callback stack to a fresh `ExitStack` instance and returns it. No callbacks are invoked by this operation - instead, they will now be invoked when the new stack is closed (either explicitly or implicitly at the end of a `with` statement).

For example, a group of files can be opened as an “all or nothing” operation as follows:

```
with ExitStack() as stack:
    files = [stack.enter_context(open(fname)) for fname in filenames]
    # Hold onto the close method, but don't call it yet.
    close_files = stack.pop_all().close
    # If opening any file fails, all previously opened files will be
    # closed automatically. If all files are opened successfully,
    # they will remain open even after the with statement ends.
    # close_files() can then be invoked explicitly to close them all.
```

close()

Immediately unwinds the callback stack, invoking callbacks in the reverse order of registration. For any context managers and exit callbacks registered, the arguments passed in will indicate that no exception occurred.

29.6.2 例子和配方

This section describes some examples and recipes for making effective use of the tools provided by `contextlib`.

Supporting a variable number of context managers

The primary use case for `ExitStack` is the one given in the class documentation: supporting a variable number of context managers and other cleanup operations in a single `with` statement. The variability may come from the number of context managers needed being driven by user input (such as opening a user specified collection of files), or from some of the context managers being optional:

```
with ExitStack() as stack:
    for resource in resources:
        stack.enter_context(resource)
    if need_special_resource():
        special = acquire_special_resource()
        stack.callback(release_special_resource, special)
    # Perform operations that use the acquired resources
```

As shown, `ExitStack` also makes it quite easy to use `with` statements to manage arbitrary resources that don't natively support the context management protocol.

Simplifying support for single optional context managers

In the specific case of a single optional context manager, *ExitStack* instances can be used as a “do nothing” context manager, allowing a context manager to easily be omitted without affecting the overall structure of the source code:

```
def debug_trace(details):
    if __debug__:
        return TraceContext(details)
    # Don't do anything special with the context in release mode
    return ExitStack()

with debug_trace():
    # Suite is traced in debug mode, but runs normally otherwise
```

Catching exceptions from `__enter__` methods

It is occasionally desirable to catch exceptions from an `__enter__` method implementation, *without* inadvertently catching exceptions from the `with` statement body or the context manager's `__exit__` method. By using *ExitStack* the steps in the context management protocol can be separated slightly in order to allow this:

```
stack = ExitStack()
try:
    x = stack.enter_context(cm)
except Exception:
    # handle __enter__ exception
else:
    with stack:
        # Handle normal case
```

Actually needing to do this is likely to indicate that the underlying API should be providing a direct resource management interface for use with `try/except/finally` statements, but not all APIs are well designed in that regard. When a context manager is the only resource management API provided, then *ExitStack* can make it easier to handle various situations that can't be handled directly in a `with` statement.

Cleaning up in an `__enter__` implementation

As noted in the documentation of *ExitStack.push()*, this method can be useful in cleaning up an already allocated resource if later steps in the `__enter__()` implementation fail.

Here's an example of doing this for a context manager that accepts resource acquisition and release functions, along with an optional validation function, and maps them to the context management protocol:

```
from contextlib import contextmanager, AbstractContextManager, ExitStack

class ResourceManager(AbstractContextManager):

    def __init__(self, acquire_resource, release_resource, check_resource_ok=None):
        self.acquire_resource = acquire_resource
        self.release_resource = release_resource
        if check_resource_ok is None:
            def check_resource_ok(resource):
                return True
        self.check_resource_ok = check_resource_ok
```

(下页继续)

(续上页)

```

@contextmanager
def _cleanup_on_error(self):
    with ExitStack() as stack:
        stack.push(self)
        yield
        # The validation check passed and didn't raise an exception
        # Accordingly, we want to keep the resource, and pass it
        # back to our caller
        stack.pop_all()

def __enter__(self):
    resource = self.acquire_resource()
    with self._cleanup_on_error():
        if not self.check_resource_ok(resource):
            msg = "Failed validation for {!r}"
            raise RuntimeError(msg.format(resource))
    return resource

def __exit__(self, *exc_details):
    # We don't need to duplicate any of our resource release logic
    self.release_resource()

```

Replacing any use of try-finally and flag variables

A pattern you will sometimes see is a try-finally statement with a flag variable to indicate whether or not the body of the finally clause should be executed. In its simplest form (that can't already be handled just by using an except clause instead), it looks something like this:

```

cleanup_needed = True
try:
    result = perform_operation()
    if result:
        cleanup_needed = False
finally:
    if cleanup_needed:
        cleanup_resources()

```

As with any try statement based code, this can cause problems for development and review, because the setup code and the cleanup code can end up being separated by arbitrarily long sections of code.

`ExitStack` makes it possible to instead register a callback for execution at the end of a with statement, and then later decide to skip executing that callback:

```

from contextlib import ExitStack

with ExitStack() as stack:
    stack.callback(cleanup_resources)
    result = perform_operation()
    if result:
        stack.pop_all()

```

This allows the intended cleanup up behaviour to be made explicit up front, rather than requiring a separate flag variable.

If a particular application uses this pattern a lot, it can be simplified even further by means of a small helper class:

```
from contextlib import ExitStack

class Callback(ExitStack):
    def __init__(self, callback, *args, **kwds):
        super(Callback, self).__init__()
        self.callback(callback, *args, **kwds)

    def cancel(self):
        self.pop_all()

with Callback(cleanup_resources) as cb:
    result = perform_operation()
    if result:
        cb.cancel()
```

If the resource cleanup isn't already neatly bundled into a standalone function, then it is still possible to use the decorator form of `ExitStack.callback()` to declare the resource cleanup in advance:

```
from contextlib import ExitStack

with ExitStack() as stack:
    @stack.callback
    def cleanup_resources():
        ...
    result = perform_operation()
    if result:
        stack.pop_all()
```

Due to the way the decorator protocol works, a callback function declared this way cannot take any parameters. Instead, any resources to be released must be accessed as closure variables.

Using a context manager as a function decorator

`ContextDecorator` makes it possible to use a context manager in both an ordinary `with` statement and also as a function decorator.

For example, it is sometimes useful to wrap functions or groups of statements with a logger that can track the time of entry and time of exit. Rather than writing both a function decorator and a context manager for the task, inheriting from `ContextDecorator` provides both capabilities in a single definition:

```
from contextlib import ContextDecorator
import logging

logging.basicConfig(level=logging.INFO)

class track_entry_and_exit(ContextDecorator):
    def __init__(self, name):
        self.name = name

    def __enter__(self):
        logging.info('Entering: %s', self.name)

    def __exit__(self, exc_type, exc, exc_tb):
        logging.info('Exiting: %s', self.name)
```

Instances of this class can be used as both a context manager:

```
with track_entry_and_exit('widget loader'):
    print('Some time consuming activity goes here')
    load_widget()
```

And also as a function decorator:

```
@track_entry_and_exit('widget loader')
def activity():
    print('Some time consuming activity goes here')
    load_widget()
```

Note that there is one additional limitation when using context managers as function decorators: there's no way to access the return value of `__enter__()`. If that value is needed, then it is still necessary to use an explicit `with` statement.

参见:

PEP 343 - “with” 语句 Python `with` 语句的规范描述、背景和示例。

29.6.3 Single use, reusable and reentrant context managers

Most context managers are written in a way that means they can only be used effectively in a `with` statement once. These single use context managers must be created afresh each time they're used - attempting to use them a second time will trigger an exception or otherwise not work correctly.

This common limitation means that it is generally advisable to create context managers directly in the header of the `with` statement where they are used (as shown in all of the usage examples above).

Files are an example of effectively single use context managers, since the first `with` statement will close the file, preventing any further IO operations using that file object.

Context managers created using `contextmanager()` are also single use context managers, and will complain about the underlying generator failing to yield if an attempt is made to use them a second time:

```
>>> from contextlib import contextmanager
>>> @contextmanager
... def singleuse():
...     print("Before")
...     yield
...     print("After")
...
>>> cm = singleuse()
>>> with cm:
...     pass
...
Before
After
>>> with cm:
...     pass
...
Traceback (most recent call last):
...
RuntimeError: generator didn't yield
```

Reentrant context managers

More sophisticated context managers may be “reentrant”. These context managers can not only be used in multiple `with` statements, but may also be used *inside* a `with` statement that is already using the same context manager.

`threading.RLock` is an example of a reentrant context manager, as are `suppress()` and `redirect_stdout()`. Here’s a very simple example of reentrant use:

```
>>> from contextlib import redirect_stdout
>>> from io import StringIO
>>> stream = StringIO()
>>> write_to_stream = redirect_stdout(stream)
>>> with write_to_stream:
...     print("This is written to the stream rather than stdout")
...     with write_to_stream:
...         print("This is also written to the stream")
...
>>> print("This is written directly to stdout")
This is written directly to stdout
>>> print(stream.getvalue())
This is written to the stream rather than stdout
This is also written to the stream
```

Real world examples of reentrancy are more likely to involve multiple functions calling each other and hence be far more complicated than this example.

Note also that being reentrant is *not* the same thing as being thread safe. `redirect_stdout()`, for example, is definitely not thread safe, as it makes a global modification to the system state by binding `sys.stdout` to a different stream.

Reusable context managers

Distinct from both single use and reentrant context managers are “reusable” context managers (or, to be completely explicit, “reusable, but not reentrant” context managers, since reentrant context managers are also reusable). These context managers support being used multiple times, but will fail (or otherwise not work correctly) if the specific context manager instance has already been used in a containing `with` statement.

`threading.Lock` is an example of a reusable, but not reentrant, context manager (for a reentrant lock, it is necessary to use `threading.RLock` instead).

Another example of a reusable, but not reentrant, context manager is `ExitStack`, as it invokes *all* currently registered callbacks when leaving any `with` statement, regardless of where those callbacks were added:

```
>>> from contextlib import ExitStack
>>> stack = ExitStack()
>>> with stack:
...     stack.callback(print, "Callback: from first context")
...     print("Leaving first context")
...
Leaving first context
Callback: from first context
>>> with stack:
...     stack.callback(print, "Callback: from second context")
...     print("Leaving second context")
...
Leaving second context
Callback: from second context
```

(下页继续)

(续上页)

```
>>> with stack:
...     stack.callback(print, "Callback: from outer context")
...     with stack:
...         stack.callback(print, "Callback: from inner context")
...         print("Leaving inner context")
...     print("Leaving outer context")
...
Leaving inner context
Callback: from inner context
Callback: from outer context
Leaving outer context
```

As the output from the example shows, reusing a single stack object across multiple with statements works correctly, but attempting to nest them will cause the stack to be cleared at the end of the innermost with statement, which is unlikely to be desirable behaviour.

Using separate `ExitStack` instances instead of reusing a single instance avoids that problem:

```
>>> from contextlib import ExitStack
>>> with ExitStack() as outer_stack:
...     outer_stack.callback(print, "Callback: from outer context")
...     with ExitStack() as inner_stack:
...         inner_stack.callback(print, "Callback: from inner context")
...         print("Leaving inner context")
...     print("Leaving outer context")
...
Leaving inner context
Callback: from inner context
Leaving outer context
Callback: from outer context
```

29.7 abc — 抽象基类

源代码: [Lib/abc.py](#)

该模块提供了在 Python 中定义抽象基类 (ABC) 的组件, 在 [PEP 3119](#) 中已有概述。查看 PEP 文档了解为什么需要在 Python 中增加这个模块。(也可查看 [PEP 3141](#) 以及 `numbers` 模块了解基于 ABC 的数字类型继承关系。)

`collections` 模块中有一些派生自 ABC 的具体类; 当然这些类还可以进一步被派生。此外, `collections.abc` 子模块中有一些 ABC 可被用于测试一个类或实例是否提供特定的接口, 例如它是否可哈希或它是否为映射等。

该模块提供了一个元类 `ABCMeta`, 可以用来定义抽象类, 另外还提供工具类 `ABC`, 可以用它以继承的方式定义抽象基类。

class `abc.ABC`

一个使用 `ABCMeta` 作为元类的工具类。抽象基类可以通过从 `ABC` 派生来简单地创建, 这就避免了在某些情况下会令人混淆的元类用法, 例如:

```
from abc import ABC

class MyABC(ABC):
    pass
```

注意`ABC` 的类型仍然是`ABCMeta`，因此继承`ABC` 仍然需要关注元类使用中的注意事项，比如可能会导致元类冲突的多重继承。当然你也可以直接使用`ABCMeta` 作为元类来定义抽象基类，例如：

```
from abc import ABCMeta

class MyABC(metaclass=ABCMeta):
    pass
```

3.4 新版功能.

class `abc.ABCMeta`

用于定义抽象基类（ABC）的元类。

使用该元类以创建抽象基类。抽象基类可以像 `mix-in` 类一样直接被子类继承。你也可以将不相关的具体类（包括内建类）和抽象基类注册为“抽象子类”——这些类以及它们的子类会被内建函数 `issubclass()` 识别为对应的抽象基类的子类，但是该抽象基类不会出现在其 `MRO`（Method Resolution Order，方法解析顺序）中，抽象基类中实现的方法也不可调用（即使通过 `super()` 调用也不行）。¹

使用`ABCMeta` 作为元类创建的类含有如下方法：

register (*subclass*)

将“子类”注册为该抽象基类的“抽象子类”，例如：

```
from abc import ABC

class MyABC(ABC):
    pass

MyABC.register(tuple)

assert issubclass(tuple, MyABC)
assert isinstance((), MyABC)
```

在 3.3 版更改：返回注册的子类，使其能够作为类装饰器。

在 3.4 版更改：你可以使用 `get_cache_token()` 函数来检测对 `register()` 的调用。

你也可以在虚基类中重载这个方法。

__subclasshook__ (*subclass*)

（必须定义为类方法。）

检查 *subclass* 是否是该抽象基类的子类。也就是说对于那些你希望定义为该抽象基类的子类的类，你不用对每个类都调用 `register()` 方法了，而是可以直接自定义 `issubclass` 的行为。（这个类方法是在抽象基类的 `__subclasscheck__()` 方法中调用的。）

该方法必须返回 `True`, `False` 或是 `NotImplemented`。如果返回 `True`，*subclass* 就会被认为是这个抽象基类的子类。如果返回 `False`，无论正常情况是否应该认为是其子类，统一视为不是。如果返回 `NotImplemented`，子类检查会按照正常机制继续执行。

为了对这些概念做一演示，请看以下定义 `ABC` 的示例：

```
class Foo:
    def __getitem__(self, index):
        ...
    def __len__(self):
        ...
    def get_iterator(self):
```

（下页继续）

¹ C++ 程序员需要注意：Python 中虚基类的概念和 C++ 中的并不相同。

(续上页)

```

        return iter(self)

class MyIterable(ABC):

    @abstractmethod
    def __iter__(self):
        while False:
            yield None

    def get_iterator(self):
        return self.__iter__()

    @classmethod
    def __subclasshook__(cls, C):
        if cls is MyIterable:
            if any("__iter__" in B.__dict__ for B in C.__mro__):
                return True
            return NotImplemented

MyIterable.register(Foo)

```

ABC `MyIterable` 定义了标准的迭代方法 `__iter__()` 作为一个抽象方法。这里给出的实现仍可在子类中被调用。`get_iterator()` 方法也是 `MyIterable` 抽象基类的一部分，但它并非必须被非抽象派生类所重载。

The `__subclasshook__()` class method defined here says that any class that has an `__iter__()` method in its `__dict__` (or in that of one of its base classes, accessed via the `__mro__` list) is considered a `MyIterable` too.

Finally, the last line makes `Foo` a virtual subclass of `MyIterable`, even though it does not define an `__iter__()` method (it uses the old-style iterable protocol, defined in terms of `__len__()` and `__getitem__()`). Note that this will not make `get_iterator` available as a method of `Foo`, so it is provided separately.

此外，`abc` 模块还提供了这些装饰器：

`@abc.abstractmethod`

用于声明抽象方法的装饰器。

使用此装饰器要求类的元类是 `ABCMeta` 或是从该类派生。一个具有派生自 `ABCMeta` 的元类的类不可以被实例化，除非它全部的抽象方法和特征属性均已被重载。抽象方法可通过任何普通的“super”调用机制来调用。`abstractmethod()` 可被用于声明特性属性和描述器的抽象方法。

Dynamically adding abstract methods to a class, or attempting to modify the abstraction status of a method or class once it is created, are not supported. The `abstractmethod()` only affects subclasses derived using regular inheritance; “virtual subclasses” registered with the ABC’s `register()` method are not affected.

When `abstractmethod()` is applied in combination with other method descriptors, it should be applied as the innermost decorator, as shown in the following usage examples:

```

class C(ABC):
    @abstractmethod
    def my_abstract_method(self, ...):
        ...

    @classmethod
    @abstractmethod
    def my_abstract_classmethod(cls, ...):
        ...

```

(下页继续)

(续上页)

```

@staticmethod
@abstractmethod
def my_abstract_staticmethod(...):
    ...

@property
@abstractmethod
def my_abstract_property(self):
    ...
@my_abstract_property.setter
@abstractmethod
def my_abstract_property(self, val):
    ...

@abstractmethod
def _get_x(self):
    ...
@abstractmethod
def _set_x(self, val):
    ...
x = property(_get_x, _set_x)

```

In order to correctly interoperate with the abstract base class machinery, the descriptor must identify itself as abstract using `__isabstractmethod__`. In general, this attribute should be `True` if any of the methods used to compose the descriptor are abstract. For example, Python's built-in `property` does the equivalent of:

```

class Descriptor:
    ...
    @property
    def __isabstractmethod__(self):
        return any(getattr(f, '__isabstractmethod__', False) for
                    f in (self._fget, self._fset, self._fdel))

```

注解： Unlike Java abstract methods, these abstract methods may have an implementation. This implementation can be called via the `super()` mechanism from the class that overrides it. This could be useful as an end-point for a super-call in a framework that uses cooperative multiple-inheritance.

The `abc` module also supports the following legacy decorators:

`@abc.abstractmethod`

3.2 新版功能.

3.3 版后已移除: It is now possible to use `classmethod` with `abstractmethod()`, making this decorator redundant.

A subclass of the built-in `classmethod()`, indicating an abstract classmethod. Otherwise it is similar to `abstractmethod()`.

This special case is deprecated, as the `classmethod()` decorator is now correctly identified as abstract when applied to an abstract method:

```

class C(ABC):
    @classmethod
    @abstractmethod
    def my_abstract_classmethod(cls, ...):
        ...

```

@abc.abstractstaticmethod

3.2 新版功能.

3.3 版后已移除: It is now possible to use `staticmethod` with `abstractmethod()`, making this decorator redundant.

A subclass of the built-in `staticmethod()`, indicating an abstract staticmethod. Otherwise it is similar to `abstractmethod()`.

This special case is deprecated, as the `staticmethod()` decorator is now correctly identified as abstract when applied to an abstract method:

```
class C(ABC):
    @staticmethod
    @abstractmethod
    def my_abstract_staticmethod(...):
        ...
```

@abc.abstractproperty

3.3 版后已移除: It is now possible to use `property`, `property.getter()`, `property.setter()` and `property.deleter()` with `abstractmethod()`, making this decorator redundant.

A subclass of the built-in `property()`, indicating an abstract property.

This special case is deprecated, as the `property()` decorator is now correctly identified as abstract when applied to an abstract method:

```
class C(ABC):
    @property
    @abstractmethod
    def my_abstract_property(self):
        ...
```

The above example defines a read-only property; you can also define a read-write abstract property by appropriately marking one or more of the underlying methods as abstract:

```
class C(ABC):
    @property
    def x(self):
        ...

    @x.setter
    @abstractmethod
    def x(self, val):
        ...
```

If only some components are abstract, only those components need to be updated to create a concrete property in a subclass:

```
class D(C):
    @C.x.setter
    def x(self, val):
        ...
```

`abc` 模块还提供了这些函数:

abc.get_cache_token()

返回当前抽象基类的缓存令牌

The token is an opaque object (that supports equality testing) identifying the current version of the abstract base class cache for virtual subclasses. The token changes with every call to `ABCMeta.register()` on any ABC.

3.4 新版功能.

29.8 atexit — 退出处理器

`atexit` 模块定义了清理函数的注册和反注册函数. 被注册的函数会在解释器正常终止时执行. `atexit` 会按照注册顺序的 * 逆序 * 执行; 如果你注册了 A, B 和 C, 那么在解释器终止时会依序执行 C, B, A.

注意: 通过该模块注册的函数, 在程序被未被 Python 捕获的信号杀死时并不会执行, 在检测到 Python 内部致命错误以及调用了 `os._exit()` 时也不会执行.

`atexit.register(func, *args, **kwargs)`

将 `func` 注册为终止时执行的函数. 任何传给 `func` 的可选的参数都应当作为参数传给 `register()`. 可以多次注册同样的函数及参数.

在正常的程序终止时 (举例来说, 当调用了 `sys.exit()` 或是主模块的执行完成时), 所有注册过的函数都会以后进先出的顺序执行. 这样做是假定更底层的模块通常会比高层模块更早引入, 因此需要更晚清理.

If an exception is raised during execution of the exit handlers, a traceback is printed (unless `SystemExit` is raised) and the exception information is saved. After all exit handlers have had a chance to run the last exception to be raised is re-raised.

这个函数返回 `func` 对象, 可以把它当作装饰器使用。

`atexit.unregister(func)`

Remove `func` from the list of functions to be run at interpreter shutdown. After calling `unregister()`, `func` is guaranteed not to be called when the interpreter shuts down, even if it was registered more than once. `unregister()` silently does nothing if `func` was not previously registered.

参见:

模块 `readline` 使用 `atexit` 读写 `readline` 历史文件的有用的例子。

29.8.1 atexit 示例

The following simple example demonstrates how a module can initialize a counter from a file when it is imported and save the counter's updated value automatically when the program terminates without relying on the application making an explicit call into this module at termination.

```
try:
    with open("counterfile") as infile:
        _count = int(infile.read())
except FileNotFoundError:
    _count = 0

def incrcounter(n):
    global _count
    _count = _count + n

def savecounter():
    with open("counterfile", "w") as outfile:
```

(下页继续)

(续上页)

```

        outfile.write("%d" % _count)

import atexit
atexit.register(savecounter)

```

Positional and keyword arguments may also be passed to `register()` to be passed along to the registered function when it is called:

```

def goodbye(name, adjective):
    print('Goodbye, %s, it was %s to meet you.' % (name, adjective))

import atexit
atexit.register(goodbye, 'Donny', 'nice')

# or:
atexit.register(goodbye, adjective='nice', name='Donny')

```

作为 *decorator* 使用:

```

import atexit

@atexit.register
def goodbye():
    print("You are now leaving the Python sector.")

```

只有在函数不需要任何参数调用时才能工作。

29.9 traceback — 打印或检索堆栈回溯

源代码: [Lib/traceback.py](#)

This module provides a standard interface to extract, format and print stack traces of Python programs. It exactly mimics the behavior of the Python interpreter when it prints a stack trace. This is useful when you want to print stack traces under program control, such as in a “wrapper” around the interpreter. 该模块提供了一个标准接口来提取、格式化和打印 Python 程序的堆栈跟踪结果。它完全模仿 Python 解释器在打印堆栈跟踪结果时的行为。当您想要在程序控制下打印堆栈跟踪结果时，例如在解释器周围的“包装”中，这是非常有用的。

这个模块使用 `traceback` 对象——这是存储在 `sys.last_traceback` 中的对象类型变量，并作为 `sys.exc_info()` 的第三项被返回。

这个模块定义了以下函数:

`traceback.print_tb(tb, limit=None, file=None)`

Print up to *limit* stack trace entries from traceback object *tb* (starting from the caller’s frame) if *limit* is positive. Otherwise, print the last `abs(limit)` entries. If *limit* is omitted or `None`, all entries are printed. If *file* is omitted or `None`, the output goes to `sys.stderr`; otherwise it should be an open file or file-like object to receive the output.

在 3.5 版更改: Added negative *limit* support.

`traceback.print_exception(etype, value, tb, limit=None, file=None, chain=True)`

Print exception information and stack trace entries from traceback object *tb* to *file*. This differs from `print_tb()` in the following ways:

- if *tb* is not `None`, it prints a header `Traceback (most recent call last):`:

- it prints the exception *etype* and *value* after the stack trace
- if *type(value)* is `SyntaxError` and *value* has the appropriate format, it prints the line where the syntax error occurred with a caret indicating the approximate position of the error.

The optional *limit* argument has the same meaning as for `print_tb()`. If *chain* is true (the default), then chained exceptions (the `__cause__` or `__context__` attributes of the exception) will be printed as well, like the interpreter itself does when printing an unhandled exception.

在 3.5 版更改: The *etype* argument is ignored and inferred from the type of *value*.

`traceback.print_exc(limit=None, file=None, chain=True)`

This is a shorthand for `print_exception(*sys.exc_info(), limit, file, chain)`.

`traceback.print_last(limit=None, file=None, chain=True)`

This is a shorthand for `print_exception(sys.last_type, sys.last_value, sys.last_traceback, limit, file, chain)`. In general it will work only after an exception has reached an interactive prompt (see `sys.last_type`).

`traceback.print_stack(f=None, limit=None, file=None)`

Print up to *limit* stack trace entries (starting from the invocation point) if *limit* is positive. Otherwise, print the last `abs(limit)` entries. If *limit* is omitted or `None`, all entries are printed. The optional *f* argument can be used to specify an alternate stack frame to start. The optional *file* argument has the same meaning as for `print_tb()`.

在 3.5 版更改: Added negative *limit* support.

`traceback.extract_tb(tb, limit=None)`

Return a `StackSummary` object representing a list of “pre-processed” stack trace entries extracted from the traceback object *tb*. It is useful for alternate formatting of stack traces. The optional *limit* argument has the same meaning as for `print_tb()`. A “pre-processed” stack trace entry is a `FrameSummary` object containing attributes `filename`, `lineno`, `name`, and `line` representing the information that is usually printed for a stack trace. The `line` is a string with leading and trailing whitespace stripped; if the source is not available it is `None`.

`traceback.extract_stack(f=None, limit=None)`

Extract the raw traceback from the current stack frame. The return value has the same format as for `extract_tb()`. The optional *f* and *limit* arguments have the same meaning as for `print_stack()`.

`traceback.format_list(extracted_list)`

Given a list of tuples or `FrameSummary` objects as returned by `extract_tb()` or `extract_stack()`, return a list of strings ready for printing. Each string in the resulting list corresponds to the item with the same index in the argument list. Each string ends in a newline; the strings may contain internal newlines as well, for those items whose source text line is not `None`.

`traceback.format_exception_only(etype, value)`

Format the exception part of a traceback. The arguments are the exception type and value such as given by `sys.last_type` and `sys.last_value`. The return value is a list of strings, each ending in a newline. Normally, the list contains a single string; however, for `SyntaxError` exceptions, it contains several lines that (when printed) display detailed information about where the syntax error occurred. The message indicating which exception occurred is the always last string in the list.

`traceback.format_exception(etype, value, tb, limit=None, chain=True)`

Format a stack trace and the exception information. The arguments have the same meaning as the corresponding arguments to `print_exception()`. The return value is a list of strings, each ending in a newline and some containing internal newlines. When these lines are concatenated and printed, exactly the same text is printed as does `print_exception()`.

在 3.5 版更改: The *etype* argument is ignored and inferred from the type of *value*.

`traceback.format_exc(limit=None, chain=True)`

This is like `print_exc(limit)` but returns a string instead of printing to a file.

`traceback.format_tb(tb, limit=None)`

A shorthand for `format_list(extract_tb(tb, limit))`.

`traceback.format_stack(f=None, limit=None)`

A shorthand for `format_list(extract_stack(f, limit))`.

`traceback.clear_frames(tb)`

Clears the local variables of all the stack frames in a traceback `tb` by calling the `clear()` method of each frame object.

3.4 新版功能.

`traceback.walk_stack(f)`

Walk a stack following `f.f_back` from the given frame, yielding the frame and line number for each frame. If `f` is `None`, the current stack is used. This helper is used with `StackSummary.extract()`.

3.5 新版功能.

`traceback.walk_tb(tb)`

Walk a traceback following `tb.next` yielding the frame and line number for each frame. This helper is used with `StackSummary.extract()`.

3.5 新版功能.

The module also defines the following classes:

29.9.1 TracebackException Objects

3.5 新版功能.

`TracebackException` objects are created from actual exceptions to capture data for later printing in a lightweight fashion.

class `traceback.TracebackException(exc_type, exc_value, exc_traceback, *, limit=None, lookup_lines=True, capture_locals=False)`

Capture an exception for later rendering. `limit`, `lookup_lines` and `capture_locals` are as for the `StackSummary` class.

Note that when locals are captured, they are also shown in the traceback.

__cause__

A `TracebackException` of the original `__cause__`.

__context__

A `TracebackException` of the original `__context__`.

__suppress_context__

The `__suppress_context__` value from the original exception.

stack

A `StackSummary` representing the traceback.

exc_type

The class of the original traceback.

filename

For syntax errors - the file name where the error occurred.

lineno

For syntax errors - the line number where the error occurred.

text

For syntax errors - the text where the error occurred.

offset

For syntax errors - the offset into the text where the error occurred.

msg

For syntax errors - the compiler error message.

classmethod from_exception (*exc*, *, *limit=None*, *lookup_lines=True*, *capture_locals=False*)

Capture an exception for later rendering. *limit*, *lookup_lines* and *capture_locals* are as for the `StackSummary` class.

Note that when locals are captured, they are also shown in the traceback.

format (*, *chain=True*)

Format the exception.

If *chain* is not `True`, `__cause__` and `__context__` will not be formatted.

The return value is a generator of strings, each ending in a newline and some containing internal newlines. `print_exception()` is a wrapper around this method which just prints the lines to a file.

The message indicating which exception occurred is always the last string in the output.

format_exception_only ()

Format the exception part of the traceback.

The return value is a generator of strings, each ending in a newline.

Normally, the generator emits a single string; however, for `SyntaxError` exceptions, it emits several lines that (when printed) display detailed information about where the syntax error occurred.

The message indicating which exception occurred is always the last string in the output.

29.9.2 StackSummary Objects

3.5 新版功能.

`StackSummary` objects represent a call stack ready for formatting.

class `traceback.StackSummary`**classmethod extract** (*frame_gen*, *, *limit=None*, *lookup_lines=True*, *capture_locals=False*)

Construct a `StackSummary` object from a frame generator (such as is returned by `walk_stack()` or `walk_tb()`).

If *limit* is supplied, only this many frames are taken from *frame_gen*. If *lookup_lines* is `False`, the returned `FrameSummary` objects will not have read their lines in yet, making the cost of creating the `StackSummary` cheaper (which may be valuable if it may not actually get formatted). If *capture_locals* is `True` the local variables in each `FrameSummary` are captured as object representations.

classmethod from_list (*a_list*)

Construct a `StackSummary` object from a supplied list of `FrameSummary` objects or old-style list of tuples. Each tuple should be a 4-tuple with filename, lineno, name, line as the elements.

format ()

Returns a list of strings ready for printing. Each string in the resulting list corresponds to a single frame from the stack. Each string ends in a newline; the strings may contain internal newlines as well, for those items with source text lines.

For long sequences of the same frame and line, the first few repetitions are shown, followed by a summary line stating the exact number of further repetitions.

在 3.6 版更改: Long sequences of repeated frames are now abbreviated.

29.9.3 FrameSummary Objects

3.5 新版功能.

FrameSummary objects represent a single frame in a traceback.

class `traceback.FrameSummary` (*filename, lineno, name, lookup_line=True, locals=None, line=None*)

Represent a single frame in the traceback or stack that is being formatted or printed. It may optionally have a stringified version of the frames locals included in it. If *lookup_line* is `False`, the source code is not looked up until the *FrameSummary* has the *line* attribute accessed (which also happens when casting it to a tuple). *line* may be directly provided, and will prevent line lookups happening at all. *locals* is an optional local variable dictionary, and if supplied the variable representations are stored in the summary for later display.

29.9.4 Traceback Examples

This simple example implements a basic read-eval-print loop, similar to (but less useful than) the standard Python interactive interpreter loop. For a more complete implementation of the interpreter loop, refer to the [code](#) module.

```
import sys, traceback

def run_user_code(envdir):
    source = input(">>> ")
    try:
        exec(source, envdir)
    except Exception:
        print("Exception in user code:")
        print("-"*60)
        traceback.print_exc(file=sys.stdout)
        print("-"*60)

envdir = {}
while True:
    run_user_code(envdir)
```

The following example demonstrates the different ways to print and format the exception and traceback:

```
import sys, traceback

def lumberjack():
    bright_side_of_death()

def bright_side_of_death():
    return tuple()[0]

try:
    lumberjack()
except IndexError:
    exc_type, exc_value, exc_traceback = sys.exc_info()
    print("*** print_tb:")
    traceback.print_tb(exc_traceback, limit=1, file=sys.stdout)
    print("*** print_exception:")
    # exc_type below is ignored on 3.5 and later
    traceback.print_exception(exc_type, exc_value, exc_traceback,
```

(下页继续)

(续上页)

```

limit=2, file=sys.stdout)
print("*** print_exc:")
traceback.print_exc(limit=2, file=sys.stdout)
print("*** format_exc, first and last line:")
formatted_lines = traceback.format_exc().splitlines()
print(formatted_lines[0])
print(formatted_lines[-1])
print("*** format_exception:")
# exc_type below is ignored on 3.5 and later
print(repr(traceback.format_exception(exc_type, exc_value,
                                     exc_traceback)))

print("*** extract_tb:")
print(repr(traceback.extract_tb(exc_traceback)))
print("*** format_tb:")
print(repr(traceback.format_tb(exc_traceback)))
print("*** tb_lineno:", exc_traceback.tb_lineno)

```

The output for the example would look similar to this:

```

*** print_tb:
  File "<doctest...>", line 10, in <module>
    lumberjack()
*** print_exception:
Traceback (most recent call last):
  File "<doctest...>", line 10, in <module>
    lumberjack()
  File "<doctest...>", line 4, in lumberjack
    bright_side_of_death()
IndexError: tuple index out of range
*** print_exc:
Traceback (most recent call last):
  File "<doctest...>", line 10, in <module>
    lumberjack()
  File "<doctest...>", line 4, in lumberjack
    bright_side_of_death()
IndexError: tuple index out of range
*** format_exc, first and last line:
Traceback (most recent call last):
IndexError: tuple index out of range
*** format_exception:
['Traceback (most recent call last):\n',
 '  File "<doctest...>", line 10, in <module>\n    lumberjack()\n',
 '  File "<doctest...>", line 4, in lumberjack\n    bright_side_of_death()\n',
 '  File "<doctest...>", line 7, in bright_side_of_death\n    return tuple()[0]\n',
 'IndexError: tuple index out of range\n']
*** extract_tb:
[<FrameSummary file <doctest...>, line 10 in <module>>,
 <FrameSummary file <doctest...>, line 4 in lumberjack>,
 <FrameSummary file <doctest...>, line 7 in bright_side_of_death>]
*** format_tb:
['  File "<doctest...>", line 10, in <module>\n    lumberjack()\n',
 '  File "<doctest...>", line 4, in lumberjack\n    bright_side_of_death()\n',
 '  File "<doctest...>", line 7, in bright_side_of_death\n    return tuple()[0]\n']
*** tb_lineno: 10

```

The following example shows the different ways to print and format the stack:

```
>>> import traceback
>>> def another_function():
...     lumberstack()
...
>>> def lumberstack():
...     traceback.print_stack()
...     print(repr(traceback.extract_stack()))
...     print(repr(traceback.format_stack()))
...
>>> another_function()
File "<doctest>", line 10, in <module>
    another_function()
File "<doctest>", line 3, in another_function
    lumberstack()
File "<doctest>", line 6, in lumberstack
    traceback.print_stack()
[('<doctest>', 10, '<module>', 'another_function()'),
 ('<doctest>', 3, 'another_function', 'lumberstack()'),
 ('<doctest>', 7, 'lumberstack', 'print(repr(traceback.extract_stack()))')]
[' File "<doctest>", line 10, in <module>\n    another_function()\n',
 ' File "<doctest>", line 3, in another_function\n    lumberstack()\n',
 ' File "<doctest>", line 8, in lumberstack\n    print(repr(traceback.format_
↳ stack()))\n']
```

This last example demonstrates the final few formatting functions:

```
>>> import traceback
>>> traceback.format_list([('spam.py', 3, '<module>', 'spam.eggs()'),
...                       ('eggs.py', 42, 'eggs', 'return "bacon"')])
[' File "spam.py", line 3, in <module>\n    spam.eggs()\n',
 ' File "eggs.py", line 42, in eggs\n    return "bacon"\n']
>>> an_error = IndexError('tuple index out of range')
>>> traceback.format_exception_only(type(an_error), an_error)
['IndexError: tuple index out of range\n']
```

29.10 `__future__` —Future 语句定义

源代码: `Lib/__future__.py`

`__future__` 是一个真正的模块，这主要有 3 个原因：

- 避免混淆已有的分析 `import` 语句并查找 `import` 的模块的工具。
- 确保 `future` 语句在 2.1 之前的版本运行时至少能抛出 `runtime` 异常 (`import __future__` 会失败，因为 2.1 版本之前没有这个模块)。
- 当引入不兼容的修改时，可以记录其引入的时间以及强制使用的时间。这是一种可执行的文档，并且可以通过 `import __future__` 来做程序性的检查。

`__future__.py` 中的每一条语句都是以下格式的：

```
FeatureName = _Feature(OptionalRelease, MandatoryRelease,
                        CompilerFlag)
```

通常 `OptionalRelease` 要比 `MandatoryRelease` 小，并且都是和 `sys.version_info` 格式一致的 5 元素元组。

```
(PY_MAJOR_VERSION, # the 2 in 2.1.0a3; an int
PY_MINOR_VERSION, # the 1; an int
PY_MICRO_VERSION, # the 0; an int
PY_RELEASE_LEVEL, # "alpha", "beta", "candidate" or "final"; string
PY_RELEASE_SERIAL # the 3; an int
)
```

OptionalRelease 记录了一个特性首次发布时的 Python 版本。

在 *MandatoryRelases* 还没有发布时，*MandatoryRelease* 表示该特性会变成语言的一部分的预测时间。

其他情况下，*MandatoryRelease* 用来记录这个特性是何时成为语言的一部分的。从该版本往后，使用该特性将不需要 *future* 语句，不过很多人还是会加上对应的 *import*。

MandatoryRelease 也可能是 *None*, 表示这个特性已经被撤销。

_Feature 类的实例有两个对应的方法，*getOptionalRelease()* 和 *getMandatoryRelease()*。

CompilerFlag 是一个（位）标记，对于动态编译的代码，需要将这个标记作为第四个参数传入内建函数 *compile()* 中以开启对应的特性。这个标记存储在 *_Feature* 类实例的 *compiler_flag* 属性中。

__future__ 中不会删除特性的描述。从 Python 2.1 中首次加入以来，通过这种方式引入了以下特性：

特性	可选版本	强制加入版本	效果
nested_scopes	2.1.0b1	2.2	PEP 227: <i>Statically Nested Scopes</i>
generators	2.2.0a1	2.3	PEP 255: <i>Simple Generators</i>
division	2.2.0a2	3.0	PEP 238: <i>Changing the Division Operator</i>
absolute_import	2.5.0a1	3.0	PEP 328: <i>Imports: Multi-Line and Absolute/Relative</i>
with_statement	2.5.0a1	2.6	PEP 343: <i>The “with” Statement</i>
print_function	2.6.0a2	3.0	PEP 3105: <i>Make print a function</i>
unicode_literals	2.6.0a2	3.0	PEP 3112: <i>Bytes literals in Python 3000</i>
generator_stop	3.5.0b1	3.7	PEP 479: <i>StopIteration handling inside generators</i>

参见：
future 编译器怎样处理 *future import*。

29.11 gc —垃圾回收器接口

此模块提供可选的垃圾回收器的接口，提供的功能包括：关闭收集器、调整收集频率、设置调试选项。它同时提供对回收器找到但是无法释放的不可达对象的访问。由于 Python 使用了带有引用计数的回收器，如果你确定你的程序不会产生循环引用，你可以关闭回收器。可以通过调用 *gc.disable()* 关闭自动垃圾回收。若要调试一个存在内存泄漏的程序，调用 *gc.set_debug(gc.DEBUG_LEAK)*；需要注意的是，它包含 *gc.DEBUG_SAVEALL*，使得被垃圾回收的对象会被存放在 *gc.garbage* 中以待检查。

gc 模块提供下列函数：

gc.enable()
 启用自动垃圾回收

`gc.disable()`

停用自动垃圾回收

`gc.isenabled()`

如果自动回收启用, 则返回 `True`

`gc.collect(generation=2)`

若被调用时不包含参数, 则启动完全的垃圾回收。可选的参数 *generation* 可以是一个整数, 指明需要回收哪一代 (从 0 到 2) 的垃圾。当参数 *generation* 无效时, 会引发 `ValueError` 异常。返回发现的不可达对象的数目。

每当运行完整收集或最高代 (2) 收集时, 为多个内置类型所维护的空闲列表会被清空。由于特定类型特别是 `float` 的实现, 在某些空闲列表中并非所有项都会被释放。

`gc.set_debug(flags)`

设置垃圾回收器的调试标识位。调试信息会被写入 `sys.stderr`。此文档末尾列出了各个标志位及其含义; 可以使用位操作对多个标志位进行设置以控制调试器。

`gc.get_debug()`

返回当前调试标识位。

`gc.get_objects()`

返回一个含有所有被收集器跟踪的对象的列表。返回的列表中不包括该函数返回的对象。

`gc.get_stats()`

返回一个包含三个字典对象的列表, 每个字典分别包含对应代的从解释器开始运行的垃圾回收统计数据。字典的键的数目在将来可能发生改变, 目前每个字典包含以下内容:

- `collections` 是该代被回收的次数;
- `collected` 是该代中被回收的对象总数;
- `uncollectable` 是在这一代中被发现无法收集的对象总数 (因此被移动到 `garbage` 列表中)。

3.4 新版功能.

`gc.set_threshold(threshold0[, threshold1[, threshold2]])`

设置垃圾回收阈值 (收集频率)。将 `threshold0` 设为零会禁用回收。

垃圾回收器把所有对象分类为三代, 取决于对象幸存于多少次垃圾回收。新创建的对象会被放在最年轻代 (第 0 代)。如果一个对象幸存于一次垃圾回收, 则该对象会被放入下一代。第 2 代是最老的一代, 因此这一代的对象幸存于垃圾回收后, 仍会留在第 2 代。为了判定何时需要进行垃圾回收, 垃圾回收器会跟踪上一次回收后, 分配和释放的对象的数目。当分配对象的数量减去释放对象的数量大于阈值 `threshold0` 时, 回收器开始进行垃圾回收。起初只有第 0 代会被检查。当上一次第 1 代被检查后, 第 0 代被检查的次数多于阈值 `threshold1` 时, 第 1 代也会被检查。相似的, `threshold2` 设置了触发第 2 代被垃圾回收的第 1 代被垃圾回收的次数。

`gc.get_count()`

将当前回收计数以形为 `(count0, count1, count2)` 的元组返回。

`gc.get_threshold()`

将当前回收阈值以形为 `(threshold0, threshold1, threshold2)` 的元组返回。

`gc.get_referrers(*objs)`

返回直接引用任意一个 *objs* 的对象列表。这个函数只定位支持垃圾回收的容器; 引用了其它对象但不支持垃圾回收的扩展类型不会被找到。

需要注意的是, 已经解除对 *objs* 引用的对象, 但仍存在于循环引用中未被回收时, 仍然会被作为引用者出现在返回的列表当中。若要获取当前正在引用 *objs* 的对象, 需要调用 `collect()` 然后再调用 `get_referrers()`。

在使用 `get_referrers()` 返回的对象时必须要小心, 因为其中一些对象可能仍在构造中因此处于暂时的无效状态。不要把 `get_referrers()` 用于调试以外的其它目的。

gc.get_referents(*objs)

返回被任意一个参数中的对象直接引用的对象的列表。返回的被引用对象是被参数中的对象的 C 语言级别方法（若存在）`tp_traverse` 访问到的对象，可能不是所有的实际直接可达对象。只有支持垃圾回收的对象支持 `tp_traverse` 方法，并且此方法只会在需要访问涉及循环引用的对象时使用。因此，可以有以下例子：一个整数对其中一个参数是直接可达的，这个整数有可能出现或不出现在返回的结果列表当中。

gc.is_tracked(obj)

当对象正在被垃圾回收器监控时返回 `True`，否则返回 `False`。一般来说，原子类的实例不会被监控，而非原子类（如容器、用户自定义的对象）会被监控。然而，会有一些特定类型的优化以便减少垃圾回收器在简单实例（如只含有原子性的键和值的字典）上的消耗。

```
>>> gc.is_tracked(0)
False
>>> gc.is_tracked("a")
False
>>> gc.is_tracked([])
True
>>> gc.is_tracked({})
False
>>> gc.is_tracked({"a": 1})
False
>>> gc.is_tracked({"a": []})
True
```

3.1 新版功能.

提供以下变量仅供只读访问（你可以修改但不应该重绑定它们）：

gc.garbage

返回一个列表，列表内为回收器找到的不可达而又不能被释放的对象（即不可回收对象）。在 Python 3.4 以后，该列表在大多数情况下都是空的，除非使用了含有非空的 `tp_del` 的 C 扩展类型的实例。

如果设置了 `DEBUG_SAVEALL`，则所有不可访问对象将被添加至该列表而不会被释放。

在 3.2 版更改：当 *interpreter shutdown* 即解释器关闭时，若此列表非空，会产生 `ResourceWarning`，即资源警告，在默认情况下此警告不会被提醒。如果设置了 `DEBUG_UNCOLLECTABLE`，所有无法被回收的对象会被打印。

在 3.4 版更改：根据 **PEP 442**，带有 `__del__()` 方法的对象最终不再会进入 `gc.garbage`。

gc.callbacks

在垃圾回收器开始前和完成后会被调用的一系列回调函数。这些回调函数在被调用时使用两个参数：*phase* 和 *info*。

phase 可为以下两值之一：

“start”：垃圾回收即将开始。

“stop”：垃圾回收已结束。

info is a dict providing more information for the callback. The following keys are currently defined:

“generation”（代）：正在被回收的最久远的一代。

“collected”（已回收的）：当 **phase** 为 “stop” 时，被成功回收的对象的数目。

“uncollectable”（不可回收的）：当 *phase* 为 “stop” 时，不能被回收并被放入 `garbage` 的对象的数目。

应用程序可以把他们自己的回调函数加入此列表。主要的使用场景有：

统计垃圾回收的数据，如：不同代的回收频率、回收所花费的时间。
使应用程序可以识别和清理他们自己的在 *garbage* 中的不可回收类型的对象。

3.3 新版功能.

以下常量被用于 `set_debug()`：

- `gc.DEBUG_STATS`
在回收完成后打印统计信息。当回收频率设置较高时，这些信息会比较有用。
- `gc.DEBUG_COLLECTABLE`
当发现可回收对象时打印信息。
- `gc.DEBUG_UNCOLLECTABLE`
打印找到的不可回收对象的信息（指不能被回收器回收的不可达对象）。这些对象会被添加到 *garbage* 列表中。

在 3.2 版更改：当 *interpreter shutdown* 时，即解释器关闭时，若 *garbage* 列表中存在对象，这些对象也会被打印输出。
- `gc.DEBUG_SAVEALL`
设置后，所有回收器找到的不可达对象会被添加进 *garbage* 而不是直接被释放。这在调试一个内存泄漏的程序时会很有用。
- `gc.DEBUG_LEAK`
调试内存泄漏的程序时，使回收器打印信息的调试标识位。（等价于 `DEBUG_COLLECTABLE | DEBUG_UNCOLLECTABLE | DEBUG_SAVEALL`）。

29.12 inspect — 检查对象

源代码: [Lib/inspect.py](#)

inspect 模块提供了一些有用的函数帮助获取对象的信息，例如模块、类、方法、函数、回溯、帧对象以及代码对象。例如它可以帮助你检查类的内容，获取某个方法的源代码，取得并格式化某个函数的参数列表，或者获取你需要显示的回溯的详细信息。
该模块提供了 4 种主要的功能：类型检查、获取源代码、检查类与函数、检查解释器的调用堆栈。

29.12.1 类型和成员

getmembers() 函数获取对象的成员，例如类或模块。函数名以” is” 开始的函数主要作为 *getmembers()* 的第 2 个参数使用。它们也可用于判定某对象是否有如下的特殊属性：

类型	属性	描述
module 模块	<code>__doc__</code>	文档字符串
	<code>__file__</code>	文件名 (内置模块没有文件名)
	<code>__doc__</code>	文档字符串
class 类	<code>__doc__</code>	文档字符串
	<code>__name__</code>	类定义时所使用的名称
	<code>__qualname__</code>	qualified name – 限定名称
	<code>__module__</code>	该类型被定义时所在的模块的名称
method 方法	<code>__doc__</code>	文档字符串
	<code>__name__</code>	该方法定义时所使用的名称

下页继续

表 1 – 续上页

类型	属性	描述
函数	<code>__qualname__</code>	qualified name – 限定名称
	<code>__func__</code>	实现该方法的函数对象
	<code>__self__</code>	该方法被绑定的实例, 若没有绑定则为 <code>None</code>
	<code>__doc__</code>	文档字符串
	<code>__name__</code>	用于定义此函数的名称
	<code>__qualname__</code>	qualified name – 限定名称
	<code>__code__</code>	包含已编译函数的代码对象 <i>bytecode</i>
	<code>__defaults__</code>	tuple of any default values for positional or keyword parameters
	<code>__kwdefaults__</code>	mapping of any default values for keyword-only parameters
	<code>__globals__</code>	global namespace in which this function was defined
回溯	<code>__annotations__</code>	mapping of parameters names to annotations; "return" key is reserved for return annotations.
	<code>tb_frame</code>	此级别的框架对象
	<code>tb_lasti</code>	index of last attempted instruction in bytecode
	<code>tb_lineno</code>	current line number in Python source code
框架	<code>tb_next</code>	next inner traceback object (called by this level)
	<code>f_back</code>	next outer frame object (this frame's caller)
	<code>f_builtins</code>	builtins namespace seen by this frame
	<code>f_code</code>	code object being executed in this frame
	<code>f_globals</code>	global namespace seen by this frame
	<code>f_lasti</code>	index of last attempted instruction in bytecode
	<code>f_lineno</code>	current line number in Python source code
	<code>f_locals</code>	local namespace seen by this frame
	<code>f_trace</code>	tracing function for this frame, or <code>None</code>
	<code>co_argcount</code>	number of arguments (not including keyword only arguments, * or ** args)
code	<code>co_code</code>	原始编译字节码的字符串
	<code>co_cellvars</code>	单元变量名称的元组 (通过包含作用域引用)
	<code>co_consts</code>	字节码中使用的常量元组
	<code>co_filename</code>	创建此代码对象的文件的名称
	<code>co_firstlineno</code>	number of first line in Python source code
	<code>co_flags</code>	bitmap of CO_* flags, read more here
	<code>co_inotab</code>	编码的行号到字节码索引的映射
	<code>co_freevars</code>	tuple of names of free variables (referenced via a function's closure)
	<code>co_kwonlyargcount</code>	number of keyword only arguments (not including ** arg)
	<code>co_name</code>	定义此代码对象的名称
	<code>co_names</code>	局部变量名称的元组
	<code>co_nlocals</code>	局部变量的数量
	<code>co_stacksize</code>	需要虚拟机堆栈空间
	<code>co_varnames</code>	参数名和局部变量的元组
生成器	<code>__name__</code>	名称
	<code>__qualname__</code>	qualified name – 限定名称
	<code>gi_frame</code>	框架
	<code>gi_running</code>	生成器在运行吗?
	<code>gi_code</code>	code
协程	<code>gi_yieldfrom</code>	object being iterated by <code>yield from</code> , or <code>None</code>
	<code>__name__</code>	名称
	<code>__qualname__</code>	qualified name – 限定名称
	<code>cr_await</code>	object being awaited on, or <code>None</code>
	<code>cr_frame</code>	框架
	<code>cr_running</code>	is the coroutine running?

下页继续

表 1 – 续上页

类型	属性	描述
	<code>cr_code</code>	<code>code</code>
<code>builtin</code>	<code>__doc__</code>	文档字符串
	<code>__name__</code>	此函数或方法的原始名称
	<code>__qualname__</code>	qualified name – 限定名称
	<code>__self__</code>	instance to which a method is bound, or <code>None</code>

在 3.5 版更改: Add `__qualname__` and `gi_yieldfrom` attributes to generators.

The `__name__` attribute of generators is now set from the function name, instead of the code name, and it can now be modified.

`inspect.getmembers(object[, predicate])`
Return all the members of an object in a list of (name, value) pairs sorted by name. If the optional *predicate* argument is supplied, only members for which the predicate returns a true value are included.

注解: `getmembers()` will only return class attributes defined in the metaclass when the argument is a class and those attributes have been listed in the metaclass' custom `__dir__()`.

`inspect.getmodule(path)`
Return the name of the module named by the file *path*, without including the names of enclosing packages. The file extension is checked against all of the entries in `importlib.machinery.all_suffixes()`. If it matches, the final path component is returned with the extension removed. Otherwise, `None` is returned.

Note that this function *only* returns a meaningful name for actual Python modules - paths that potentially refer to Python packages will still return `None`.

在 3.3 版更改: The function is based directly on `importlib`.

`inspect.ismodule(object)`
Return true if the object is a module.

`inspect.isclass(object)`
Return true if the object is a class, whether built-in or created in Python code.

`inspect.ismethod(object)`
Return true if the object is a bound method written in Python.

`inspect.isfunction(object)`
Return true if the object is a Python function, which includes functions created by a *lambda* expression.

`inspect.isgeneratorfunction(object)`
Return true if the object is a Python generator function.

`inspect.isgenerator(object)`
Return true if the object is a generator.

`inspect.iscoroutinefunction(object)`
Return true if the object is a *coroutine function* (a function defined with an `async def` syntax).

3.5 新版功能.

`inspect.iscoroutine(object)`
Return true if the object is a *coroutine* created by an `async def` function.

3.5 新版功能.

`inspect.isawaitable(object)`
Return true if the object can be used in `await` expression.

Can also be used to distinguish generator-based coroutines from regular generators:

```
def gen():
    yield
@types.coroutine
def gen_coro():
    yield

assert not isawaitable(gen())
assert isawaitable(gen_coro())
```

3.5 新版功能.

`inspect.isasyncgenfunction(object)`

Return true if the object is an *asynchronous generator* function, for example:

```
>>> async def agen():
...     yield 1
...
>>> inspect.isasyncgenfunction(agen)
True
```

3.6 新版功能.

`inspect.isasyncgen(object)`

Return true if the object is an *asynchronous generator iterator* created by an *asynchronous generator* function.

3.6 新版功能.

`inspect.istraceback(object)`

Return true if the object is a traceback.

`inspect.isframe(object)`

Return true if the object is a frame.

`inspect.iscode(object)`

Return true if the object is a code.

`inspect.isbuiltin(object)`

Return true if the object is a built-in function or a bound built-in method.

`inspect.isroutine(object)`

Return true if the object is a user-defined or built-in function or method.

`inspect.isabstract(object)`

Return true if the object is an abstract base class.

`inspect.ismethoddescriptor(object)`

Return true if the object is a method descriptor, but not if `ismethod()`, `isclass()`, `isfunction()` or `isbuiltin()` are true.

This, for example, is true of `int.__add__`. An object passing this test has a `__get__()` method but not a `__set__()` method, but beyond that the set of attributes varies. A `__name__` attribute is usually sensible, and `__doc__` often is.

Methods implemented via descriptors that also pass one of the other tests return false from the `ismethoddescriptor()` test, simply because the other tests promise more – you can, e.g., count on having the `__func__` attribute (etc) when an object passes `ismethod()`.

`inspect.isdatadescriptor(object)`

Return true if the object is a data descriptor.

Data descriptors have both a `__get__` and a `__set__` method. Examples are properties (defined in Python), getsets, and members. The latter two are defined in C and there are more specific tests available for those types, which is robust across Python implementations. Typically, data descriptors will also have `__name__` and `__doc__` attributes (properties, getsets, and members have both of these attributes), but this is not guaranteed.

`inspect.isgetsetdescriptor(object)`

Return true if the object is a getset descriptor.

CPython implementation detail: getsets are attributes defined in extension modules via `PyGetSetDef` structures. For Python implementations without such types, this method will always return `False`.

`inspect.ismemberdescriptor(object)`

Return true if the object is a member descriptor.

CPython implementation detail: Member descriptors are attributes defined in extension modules via `PyMemberDef` structures. For Python implementations without such types, this method will always return `False`.

29.12.2 Retrieving source code

`inspect.getdoc(object)`

Get the documentation string for an object, cleaned up with `cleandoc()`. If the documentation string for an object is not provided and the object is a class, a method, a property or a descriptor, retrieve the documentation string from the inheritance hierarchy.

在 3.5 版更改: Documentation strings are now inherited if not overridden.

`inspect.getcomments(object)`

Return in a single string any lines of comments immediately preceding the object's source code (for a class, function, or method), or at the top of the Python source file (if the object is a module). If the object's source code is unavailable, return `None`. This could happen if the object has been defined in C or the interactive shell.

`inspect.getfile(object)`

Return the name of the (text or binary) file in which an object was defined. This will fail with a `TypeError` if the object is a built-in module, class, or function.

`inspect.getmodule(object)`

Try to guess which module an object was defined in.

`inspect.getsourcefile(object)`

Return the name of the Python source file in which an object was defined. This will fail with a `TypeError` if the object is a built-in module, class, or function.

`inspect.getsourcelines(object)`

Return a list of source lines and starting line number for an object. The argument may be a module, class, method, function, traceback, frame, or code object. The source code is returned as a list of the lines corresponding to the object and the line number indicates where in the original source file the first line of code was found. An `OSError` is raised if the source code cannot be retrieved.

在 3.3 版更改: `OSError` is raised instead of `IOError`, now an alias of the former.

`inspect.getsource(object)`

Return the text of the source code for an object. The argument may be a module, class, method, function, traceback, frame, or code object. The source code is returned as a single string. An `OSError` is raised if the source code cannot be retrieved.

在 3.3 版更改: `OSError` is raised instead of `IOError`, now an alias of the former.

`inspect.cleandoc(doc)`

Clean up indentation from docstrings that are indented to line up with blocks of code.

All leading whitespace is removed from the first line. Any leading whitespace that can be uniformly removed from the second line onwards is removed. Empty lines at the beginning and end are subsequently removed. Also, all tabs are expanded to spaces.

29.12.3 Introspecting callables with the Signature object

3.3 新版功能.

The Signature object represents the call signature of a callable object and its return annotation. To retrieve a Signature object, use the `signature()` function.

`inspect.signature(callable, *, follow_wrapped=True)`
Return a `Signature` object for the given callable:

```
>>> from inspect import signature
>>> def foo(a, *, b:int, **kwargs):
...     pass

>>> sig = signature(foo)

>>> str(sig)
'(a, *, b:int, **kwargs)'

>>> str(sig.parameters['b'])
'b:int'

>>> sig.parameters['b'].annotation
<class 'int'>
```

Accepts a wide range of python callables, from plain functions and classes to `functools.partial()` objects.

Raises `ValueError` if no signature can be provided, and `TypeError` if that type of object is not supported.

3.5 新版功能: `follow_wrapped` parameter. Pass `False` to get a signature of callable specifically (callable.`__wrapped__` will not be used to unwrap decorated callables.)

注解: Some callables may not be introspectable in certain implementations of Python. For example, in CPython, some built-in functions defined in C provide no metadata about their arguments.

class `inspect.Signature` (`parameters=None`, *, `return_annotation=Signature.empty`)

A Signature object represents the call signature of a function and its return annotation. For each parameter accepted by the function it stores a `Parameter` object in its `parameters` collection.

The optional `parameters` argument is a sequence of `Parameter` objects, which is validated to check that there are no parameters with duplicate names, and that the parameters are in the right order, i.e. positional-only first, then positional-or-keyword, and that parameters with defaults follow parameters without defaults.

The optional `return_annotation` argument, can be an arbitrary Python object, is the “return” annotation of the callable.

Signature objects are *immutable*. Use `Signature.replace()` to make a modified copy.

在 3.5 版更改: Signature objects are picklable and hashable.

empty

A special class-level marker to specify absence of a return annotation.

parameters

An ordered mapping of parameters' names to the corresponding *Parameter* objects.

return_annotation

The “return” annotation for the callable. If the callable has no “return” annotation, this attribute is set to *Signature.empty*.

bind (*args, **kwargs)

Create a mapping from positional and keyword arguments to parameters. Returns *BoundArguments* if *args and **kwargs match the signature, or raises a *TypeError*.

bind_partial (*args, **kwargs)

Works the same way as *Signature.bind()*, but allows the omission of some required arguments (mimics *functools.partial()* behavior.) Returns *BoundArguments*, or raises a *TypeError* if the passed arguments do not match the signature.

replace (*[, parameters][, return_annotation])

Create a new *Signature* instance based on the instance *replace* was invoked on. It is possible to pass different *parameters* and/or *return_annotation* to override the corresponding properties of the base signature. To remove *return_annotation* from the copied *Signature*, pass in *Signature.empty*.

```
>>> def test(a, b):
...     pass
>>> sig = signature(test)
>>> new_sig = sig.replace(return_annotation="new return anno")
>>> str(new_sig)
"(a, b) -> 'new return anno'"
```

classmethod from_callable (obj, *, follow_wrapped=True)

Return a *Signature* (or its subclass) object for a given callable obj. Pass *follow_wrapped=False* to get a signature of obj without unwrapping its *__wrapped__* chain.

This method simplifies subclassing of *Signature*:

```
class MySignature(Signature):
    pass
sig = MySignature.from_callable(min)
assert isinstance(sig, MySignature)
```

3.5 新版功能.

class inspect.Parameter (name, kind, *, default=Parameter.empty, annotation=Parameter.empty)

Parameter objects are *immutable*. Instead of modifying a *Parameter* object, you can use *Parameter.replace()* to create a modified copy.

在 3.5 版更改: *Parameter* objects are picklable and hashable.

empty

A special class-level marker to specify absence of default values and annotations.

name

The name of the parameter as a string. The name must be a valid Python identifier.

CPython implementation detail: CPython generates implicit parameter names of the form *.0* on the code objects used to implement comprehensions and generator expressions.

在 3.6 版更改: These parameter names are exposed by this module as names like *implicit0*.

default

The default value for the parameter. If the parameter has no default value, this attribute is set to *Parameter.empty*.

annotation

The annotation for the parameter. If the parameter has no annotation, this attribute is set to `Parameter.empty`.

kind

Describes how argument values are bound to the parameter. Possible values (accessible via `Parameter`, like `Parameter.KEYWORD_ONLY`):

名称	意义
<code>POSITIONAL_ONLY</code>	Value must be supplied as a positional argument. Python has no explicit syntax for defining positional-only parameters, but many built-in and extension module functions (especially those that accept only one or two parameters) accept them.
<code>POSITIONAL_OR_KEYWORD</code>	Value may be supplied as either a keyword or positional argument (this is the standard binding behaviour for functions implemented in Python.)
<code>VAR_POSITIONAL</code>	A tuple of positional arguments that aren't bound to any other parameter. This corresponds to a <code>*args</code> parameter in a Python function definition.
<code>KEYWORD_ONLY</code>	Value must be supplied as a keyword argument. Keyword only parameters are those which appear after a <code>*</code> or <code>*args</code> entry in a Python function definition.
<code>VAR_KEYWORD</code>	A dict of keyword arguments that aren't bound to any other parameter. This corresponds to a <code>**kwargs</code> parameter in a Python function definition.

Example: print all keyword-only arguments without default values:

```
>>> def foo(a, b, *, c, d=10):
...     pass

>>> sig = signature(foo)
>>> for param in sig.parameters.values():
...     if (param.kind == param.KEYWORD_ONLY and
...         param.default is param.empty):
...         print('Parameter:', param)
Parameter: c
```

replace (`*, name`][, `kind`][, `default`][, `annotation`])

Create a new `Parameter` instance based on the instance replaced was invoked on. To override a `Parameter` attribute, pass the corresponding argument. To remove a default value or/and an annotation from a `Parameter`, pass `Parameter.empty`.

```
>>> from inspect import Parameter
>>> param = Parameter('foo', Parameter.KEYWORD_ONLY, default=42)
>>> str(param)
'foo=42'

>>> str(param.replace()) # Will create a shallow copy of 'param'
'foo=42'

>>> str(param.replace(default=Parameter.empty, annotation='spam'))
'foo: 'spam''
```

在 3.4 版更改: In Python 3.3 `Parameter` objects were allowed to have `name` set to `None` if their `kind` was

set to `POSITIONAL_ONLY`. This is no longer permitted.

class `inspect.BindArguments`

Result of a `Signature.bind()` or `Signature.bind_partial()` call. Holds the mapping of arguments to the function's parameters.

arguments

An ordered, mutable mapping (`collections.OrderedDict`) of parameters' names to arguments' values. Contains only explicitly bound arguments. Changes in `arguments` will reflect in `args` and `kwargs`.

Should be used in conjunction with `Signature.parameters` for any argument processing purposes.

注解: Arguments for which `Signature.bind()` or `Signature.bind_partial()` relied on a default value are skipped. However, if needed, use `BindArguments.apply_defaults()` to add them.

args

A tuple of positional arguments values. Dynamically computed from the `arguments` attribute.

kwargs

A dict of keyword arguments values. Dynamically computed from the `arguments` attribute.

signature

A reference to the parent `Signature` object.

apply_defaults()

Set default values for missing arguments.

For variable-positional arguments (`*args`) the default is an empty tuple.

For variable-keyword arguments (`**kwargs`) the default is an empty dict.

```
>>> def foo(a, b='ham', *args): pass
>>> ba = inspect.signature(foo).bind('spam')
>>> ba.apply_defaults()
>>> ba.arguments
OrderedDict([('a', 'spam'), ('b', 'ham'), ('args', ())])
```

3.5 新版功能.

The `args` and `kwargs` properties can be used to invoke functions:

```
def test(a, *, b):
    ...

sig = signature(test)
ba = sig.bind(10, b=20)
test(*ba.args, **ba.kwargs)
```

参见:

PEP 362 - Function Signature Object. The detailed specification, implementation details and examples.

29.12.4 类与函数

`inspect.getclasstree(classes, unique=False)`

Arrange the given list of classes into a hierarchy of nested lists. Where a nested list appears, it contains classes derived from the class whose entry immediately precedes the list. Each entry is a 2-tuple containing a class and

a tuple of its base classes. If the *unique* argument is true, exactly one entry appears in the returned structure for each class in the given list. Otherwise, classes using multiple inheritance and their descendants will appear multiple times.

`inspect.getargspec(func)`

Get the names and default values of a Python function's parameters. A *named tuple* `ArgSpec(args, varargs, keywords, defaults)` is returned. *args* is a list of the parameter names. *varargs* and *keywords* are the names of the * and ** parameters or None. *defaults* is a tuple of default argument values or None if there are no default arguments; if this tuple has *n* elements, they correspond to the last *n* elements listed in *args*.

3.0 版后已移除: Use `getfullargspec()` for an updated API that is usually a drop-in replacement, but also correctly handles function annotations and keyword-only parameters.

Alternatively, use `signature()` and *Signature Object*, which provide a more structured introspection API for callables.

`inspect.getfullargspec(func)`

Get the names and default values of a Python function's parameters. A *named tuple* is returned:

```
FullArgSpec(args, varargs, varkw, defaults, kwoonlyargs, kwoonlydefaults,
             annotations)
```

args is a list of the positional parameter names. *varargs* is the name of the * parameter or None if arbitrary positional arguments are not accepted. *varkw* is the name of the ** parameter or None if arbitrary keyword arguments are not accepted. *defaults* is an *n*-tuple of default argument values corresponding to the last *n* positional parameters, or None if there are no such defaults defined. *kwoonlyargs* is a list of keyword-only parameter names. *kwoonlydefaults* is a dictionary mapping parameter names from *kwoonlyargs* to the default values used if no argument is supplied. *annotations* is a dictionary mapping parameter names to annotations. The special key "return" is used to report the function return value annotation (if any).

Note that `signature()` and *Signature Object* provide the recommended API for callable introspection, and support additional behaviours (like positional-only arguments) that are sometimes encountered in extension module APIs. This function is retained primarily for use in code that needs to maintain compatibility with the Python 2 `inspect` module API.

在 3.4 版更改: This function is now based on `signature()`, but still ignores `__wrapped__` attributes and includes the already bound first parameter in the signature output for bound methods.

在 3.6 版更改: This method was previously documented as deprecated in favour of `signature()` in Python 3.5, but that decision has been reversed in order to restore a clearly supported standard interface for single-source Python 2/3 code migrating away from the legacy `getargspec()` API.

`inspect.getargvalues(frame)`

Get information about arguments passed into a particular frame. A *named tuple* `ArgInfo(args, varargs, keywords, locals)` is returned. *args* is a list of the argument names. *varargs* and *keywords* are the names of the * and ** arguments or None. *locals* is the locals dictionary of the given frame.

注解: This function was inadvertently marked as deprecated in Python 3.5.

`inspect.formatargspec(args[, varargs, varkw, defaults, kwoonlyargs, kwoonlydefaults, annotations[, formatarg, formatvarargs, formatvarkw, formatvalue, formatreturns, formatannotations]])`

Format a pretty argument spec from the values returned by `getfullargspec()`.

The first seven arguments are (*args*, *varargs*, *varkw*, *defaults*, *kwoonlyargs*, *kwoonlydefaults*, *annotations*).

The other six arguments are functions that are called to turn argument names, * argument name, ** argument name, default values, return annotation and individual annotations into strings, respectively.

例如:

```
>>> from inspect import formatargspec, getfullargspec
>>> def f(a: int, b: float):
...     pass
...
>>> formatargspec(*getfullargspec(f))
'(a: int, b: float)'
```

3.5 版后已移除: Use [signature\(\)](#) and [Signature Object](#), which provide a better introspecting API for callables.

`inspect.formatargvalues` (*args*[, *varargs*, *varkw*, *locals*, *formatarg*, *formatvarargs*, *formatvarkw*, *formatvalue*])

Format a pretty argument spec from the four values returned by [getargvalues\(\)](#). The *format** arguments are the corresponding optional formatting functions that are called to turn names and values into strings.

注解: This function was inadvertently marked as deprecated in Python 3.5.

`inspect.getmro` (*cls*)

Return a tuple of class *cls*' s base classes, including *cls*, in method resolution order. No class appears more than once in this tuple. Note that the method resolution order depends on *cls*' s type. Unless a very peculiar user-defined metatype is in use, *cls* will be the first element of the tuple.

`inspect.getcallargs` (*func*, **args*, ***kwargs*)

Bind the *args* and *kwargs* to the argument names of the Python function or method *func*, as if it was called with them. For bound methods, bind also the first argument (typically named *self*) to the associated instance. A dict is returned, mapping the argument names (including the names of the *** and **** arguments, if any) to their values from *args* and *kwargs*. In case of invoking *func* incorrectly, i.e. whenever *func(*args, **kwargs)* would raise an exception because of incompatible signature, an exception of the same type and the same or similar message is raised. For example:

```
>>> from inspect import getcallargs
>>> def f(a, b=1, *pos, **named):
...     pass
...
>>> getcallargs(f, 1, 2, 3) == {'a': 1, 'named': {}, 'b': 2, 'pos': (3,)}
True
>>> getcallargs(f, a=2, x=4) == {'a': 2, 'named': {'x': 4}, 'b': 1, 'pos': ()}
True
>>> getcallargs(f)
Traceback (most recent call last):
...
TypeError: f() missing 1 required positional argument: 'a'
```

3.2 新版功能.

3.5 版后已移除: Use [Signature.bind\(\)](#) and [Signature.bind_partial\(\)](#) instead.

`inspect.getclosurevars` (*func*)

Get the mapping of external name references in a Python function or method *func* to their current values. A [named tuple](#) `ClosureVars(nonlocals, globals, builtins, unbound)` is returned. *nonlocals* maps referenced names to lexical closure variables, *globals* to the function' s module globals and *builtins* to the builtins visible from the function body. *unbound* is the set of names referenced in the function that could not be resolved at all given the current module globals and builtins.

[TypeError](#) is raised if *func* is not a Python function or method.

3.3 新版功能.

`inspect.unwrap(func, *, stop=None)`

Get the object wrapped by *func*. It follows the chain of `__wrapped__` attributes returning the last object in the chain.

stop is an optional callback accepting an object in the wrapper chain as its sole argument that allows the unwrapping to be terminated early if the callback returns a true value. If the callback never returns a true value, the last object in the chain is returned as usual. For example, `signature()` uses this to stop unwrapping if any object in the chain has a `__signature__` attribute defined.

`ValueError` is raised if a cycle is encountered.

3.4 新版功能.

29.12.5 The interpreter stack

When the following functions return “frame records,” each record is a *named tuple* `FrameInfo(frame, filename, lineno, function, code_context, index)`. The tuple contains the frame object, the filename, the line number of the current line, the function name, a list of lines of context from the source code, and the index of the current line within that list.

在 3.5 版更改: Return a named tuple instead of a tuple.

注解: Keeping references to frame objects, as found in the first element of the frame records these functions return, can cause your program to create reference cycles. Once a reference cycle has been created, the lifespan of all objects which can be accessed from the objects which form the cycle can become much longer even if Python’s optional cycle detector is enabled. If such cycles must be created, it is important to ensure they are explicitly broken to avoid the delayed destruction of objects and increased memory consumption which occurs.

Though the cycle detector will catch these, destruction of the frames (and local variables) can be made deterministic by removing the cycle in a `finally` clause. This is also important if the cycle detector was disabled when Python was compiled or using `gc.disable()`. For example:

```
def handle_stackframe_without_leak():
    frame = inspect.currentframe()
    try:
        # do something with the frame
    finally:
        del frame
```

If you want to keep the frame around (for example to print a traceback later), you can also break reference cycles by using the `frame.clear()` method.

The optional *context* argument supported by most of these functions specifies the number of lines of context to return, which are centered around the current line.

`inspect.getframeinfo(frame, context=1)`

Get information about a frame or traceback object. A *named tuple* `Traceback(filename, lineno, function, code_context, index)` is returned.

`inspect.getouterframes(frame, context=1)`

Get a list of frame records for a frame and all outer frames. These frames represent the calls that lead to the creation of *frame*. The first entry in the returned list represents *frame*; the last entry represents the outermost call on *frame*’s stack.

在 3.5 版更改: A list of *named tuples* `FrameInfo(frame, filename, lineno, function, code_context, index)` is returned.

`inspect.getinnerframes(traceback, context=1)`

Get a list of frame records for a traceback's frame and all inner frames. These frames represent calls made as a consequence of *frame*. The first entry in the list represents *traceback*; the last entry represents where the exception was raised.

在 3.5 版更改: A list of *named tuples* `FrameInfo(frame, filename, lineno, function, code_context, index)` is returned.

`inspect.currentframe()`

Return the frame object for the caller's stack frame.

CPython implementation detail: This function relies on Python stack frame support in the interpreter, which isn't guaranteed to exist in all implementations of Python. If running in an implementation without Python stack frame support this function returns `None`.

`inspect.stack(context=1)`

Return a list of frame records for the caller's stack. The first entry in the returned list represents the caller; the last entry represents the outermost call on the stack.

在 3.5 版更改: A list of *named tuples* `FrameInfo(frame, filename, lineno, function, code_context, index)` is returned.

`inspect.trace(context=1)`

Return a list of frame records for the stack between the current frame and the frame in which an exception currently being handled was raised in. The first entry in the list represents the caller; the last entry represents where the exception was raised.

在 3.5 版更改: A list of *named tuples* `FrameInfo(frame, filename, lineno, function, code_context, index)` is returned.

29.12.6 Fetching attributes statically

Both `getattr()` and `hasattr()` can trigger code execution when fetching or checking for the existence of attributes. Descriptors, like properties, will be invoked and `__getattr__()` and `__getattribute__()` may be called.

For cases where you want passive introspection, like documentation tools, this can be inconvenient. `getattr_static()` has the same signature as `getattr()` but avoids executing code when it fetches attributes.

`inspect.getattr_static(obj, attr, default=None)`

Retrieve attributes without triggering dynamic lookup via the descriptor protocol, `__getattr__()` or `__getattribute__()`.

Note: this function may not be able to retrieve all attributes that `getattr` can fetch (like dynamically created attributes) and may find attributes that `getattr` can't (like descriptors that raise `AttributeError`). It can also return descriptors objects instead of instance members.

If the instance `__dict__` is shadowed by another member (for example a property) then this function will be unable to find instance members.

3.2 新版功能.

`getattr_static()` does not resolve descriptors, for example slot descriptors or getset descriptors on objects implemented in C. The descriptor object is returned instead of the underlying attribute.

You can handle these with code like the following. Note that for arbitrary getset descriptors invoking these may trigger code execution:

```
# example code for resolving the builtin descriptor types
class _foo:
    __slots__ = ['foo']
```

(下页继续)

(续上页)

```

slot_descriptor = type(_foo.foo)
getset_descriptor = type(type(open(__file__)).name)
wrapper_descriptor = type(str.__dict__['__add__'])
descriptor_types = (slot_descriptor, getset_descriptor, wrapper_descriptor)

result = getattr_static(some_object, 'foo')
if type(result) in descriptor_types:
    try:
        result = result.__get__()
    except AttributeError:
        # descriptors can raise AttributeError to
        # indicate there is no underlying value
        # in which case the descriptor itself will
        # have to do
        pass

```

29.12.7 Current State of Generators and Coroutines

When implementing coroutine schedulers and for other advanced uses of generators, it is useful to determine whether a generator is currently executing, is waiting to start or resume or execution, or has already terminated. `getgeneratorstate()` allows the current state of a generator to be determined easily.

`inspect.getgeneratorstate(generator)`

Get current state of a generator-iterator.

Possible states are:

- GEN_CREATED: Waiting to start execution.
- GEN_RUNNING: Currently being executed by the interpreter.
- GEN_SUSPENDED: Currently suspended at a yield expression.
- GEN_CLOSED: Execution has completed.

3.2 新版功能.

`inspect.getcoroutinestate(coroutine)`

Get current state of a coroutine object. The function is intended to be used with coroutine objects created by `async def` functions, but will accept any coroutine-like object that has `cr_running` and `cr_frame` attributes.

Possible states are:

- CORO_CREATED: Waiting to start execution.
- CORO_RUNNING: Currently being executed by the interpreter.
- CORO_SUSPENDED: Currently suspended at an await expression.
- CORO_CLOSED: Execution has completed.

3.5 新版功能.

The current internal state of the generator can also be queried. This is mostly useful for testing purposes, to ensure that internal state is being updated as expected:

`inspect.getgeneratorlocals(generator)`

Get the mapping of live local variables in *generator* to their current values. A dictionary is returned that maps from variable names to values. This is the equivalent of calling `locals()` in the body of the generator, and all the same caveats apply.

If *generator* is a *generator* with no currently associated frame, then an empty dictionary is returned. *TypeError* is raised if *generator* is not a Python generator object.

CPython implementation detail: This function relies on the generator exposing a Python stack frame for introspection, which isn't guaranteed to be the case in all implementations of Python. In such cases, this function will always return an empty dictionary.

3.3 新版功能.

`inspect.getcoroutinelocals(coroutine)`

This function is analogous to `getgeneratorlocals()`, but works for coroutine objects created by `async def` functions.

3.5 新版功能.

29.12.8 Code Objects Bit Flags

Python code objects have a `co_flags` attribute, which is a bitmap of the following flags:

`inspect.CO_OPTIMIZED`

The code object is optimized, using fast locals.

`inspect.CO_NEWLOCALS`

If set, a new dict will be created for the frame's `f_locals` when the code object is executed.

`inspect.CO_VARARGS`

The code object has a variable positional parameter (*args-like).

`inspect.CO_VARKEYWORDS`

The code object has a variable keyword parameter (**kwargs-like).

`inspect.CO_NESTED`

The flag is set when the code object is a nested function.

`inspect.CO_GENERATOR`

The flag is set when the code object is a generator function, i.e. a generator object is returned when the code object is executed.

`inspect.CO_NOFREE`

The flag is set if there are no free or cell variables.

`inspect.CO_COROUTINE`

The flag is set when the code object is a coroutine function. When the code object is executed it returns a coroutine object. See [PEP 492](#) for more details.

3.5 新版功能.

`inspect.CO_ITERABLE_COROUTINE`

The flag is used to transform generators into generator-based coroutines. Generator objects with this flag can be used in `await` expression, and can `yield from` coroutine objects. See [PEP 492](#) for more details.

3.5 新版功能.

`inspect.CO_ASYNC_GENERATOR`

The flag is set when the code object is an asynchronous generator function. When the code object is executed it returns an asynchronous generator object. See [PEP 525](#) for more details.

3.6 新版功能.

注解: The flags are specific to CPython, and may not be defined in other Python implementations. Furthermore, the flags are an implementation detail, and can be removed or deprecated in future Python releases. It's recommended to use public APIs from the `inspect` module for any introspection needs.

29.12.9 Command Line Interface

The `inspect` module also provides a basic introspection capability from the command line.

By default, accepts the name of a module and prints the source of that module. A class or function within the module can be printed instead by appended a colon and the qualified name of the target object.

--details

Print information about the specified object rather than the source code

29.13 `site` —Site-specific configuration hook

Source code: [Lib/site.py](#)

This module is automatically imported during initialization. The automatic import can be suppressed using the interpreter's `-S` option.

Importing this module will append site-specific paths to the module search path and add a few builtins, unless `-S` was used. In that case, this module can be safely imported with no automatic modifications to the module search path or additions to the builtins. To explicitly trigger the usual site-specific additions, call the `site.main()` function.

在 3.3 版更改: Importing the module used to trigger paths manipulation even when using `-S`.

It starts by constructing up to four directories from a head and a tail part. For the head part, it uses `sys.prefix` and `sys.exec_prefix`; empty heads are skipped. For the tail part, it uses the empty string and then `lib/site-packages` (on Windows) or `lib/pythonX.Y/site-packages` (on Unix and Macintosh). For each of the distinct head-tail combinations, it sees if it refers to an existing directory, and if so, adds it to `sys.path` and also inspects the newly added path for configuration files.

在 3.5 版更改: Support for the “site-python” directory has been removed.

If a file named “pyvenv.cfg” exists one directory above `sys.executable`, `sys.prefix` and `sys.exec_prefix` are set to that directory and it is also checked for site-packages (`sys.base_prefix` and `sys.base_exec_prefix` will always be the “real” prefixes of the Python installation). If “pyvenv.cfg” (a bootstrap configuration file) contains the key “include-system-site-packages” set to anything other than “false” (case-insensitive), the system-level prefixes will still also be searched for site-packages; otherwise they won't.

A path configuration file is a file whose name has the form `name.pth` and exists in one of the four directories mentioned above; its contents are additional items (one per line) to be added to `sys.path`. Non-existing items are never added to `sys.path`, and no check is made that the item refers to a directory rather than a file. No item is added to `sys.path` more than once. Blank lines and lines beginning with `#` are skipped. Lines starting with `import` (followed by space or tab) are executed.

For example, suppose `sys.prefix` and `sys.exec_prefix` are set to `/usr/local`. The Python X.Y library is then installed in `/usr/local/lib/pythonX.Y`. Suppose this has a subdirectory `/usr/local/lib/pythonX.Y/site-packages` with three subsubdirectories, `foo`, `bar` and `spam`, and two path configuration files, `foo.pth` and `bar.pth`. Assume `foo.pth` contains the following:

```
# foo package configuration

foo
bar
bletch
```

and `bar.pth` contains:

```
# bar package configuration

bar
```

Then the following version-specific directories are added to `sys.path`, in this order:

```
/usr/local/lib/pythonX.Y/site-packages/bar
/usr/local/lib/pythonX.Y/site-packages/foo
```

Note that `bletch` is omitted because it doesn't exist; the `bar` directory precedes the `foo` directory because `bar.pth` comes alphabetically before `foo.pth`; and `spam` is omitted because it is not mentioned in either path configuration file.

After these path manipulations, an attempt is made to import a module named `sitecustomize`, which can perform arbitrary site-specific customizations. It is typically created by a system administrator in the `site-packages` directory. If this import fails with an `ImportError` exception, it is silently ignored. If Python is started without output streams available, as with `pythonw.exe` on Windows (which is used by default to start IDLE), attempted output from `sitecustomize` is ignored. Any exception other than `ImportError` causes a silent and perhaps mysterious failure of the process.

After this, an attempt is made to import a module named `usercustomize`, which can perform arbitrary user-specific customizations, if `ENABLE_USER_SITE` is true. This file is intended to be created in the user `site-packages` directory (see below), which is part of `sys.path` unless disabled by `-s`. An `ImportError` will be silently ignored.

Note that for some non-Unix systems, `sys.prefix` and `sys.exec_prefix` are empty, and the path manipulations are skipped; however the import of `sitecustomize` and `usercustomize` is still attempted.

29.13.1 Readline configuration

On systems that support `readline`, this module will also import and configure the `rlcompleter` module, if Python is started in interactive mode and without the `-S` option. The default behavior is enable tab-completion and to use `~/.python_history` as the history save file. To disable it, delete (or override) the `sys.__interactivehook__` attribute in your `sitecustomize` or `usercustomize` module or your `PYTHONSTARTUP` file.

在 3.4 版更改: Activation of `rlcompleter` and history was made automatic.

29.13.2 Module contents

`site.PREFIXES`

A list of prefixes for `site-packages` directories.

`site.ENABLE_USER_SITE`

Flag showing the status of the user `site-packages` directory. True means that it is enabled and was added to `sys.path`. False means that it was disabled by user request (with `-s` or `PYTHONNOUSERSITE`). None means it was disabled for security reasons (mismatch between user or group id and effective id) or by an administrator.

`site.USER_SITE`

Path to the user `site-packages` for the running Python. Can be None if `getusersitepackages()` hasn't been called yet. Default value is `~/.local/lib/pythonX.Y/site-packages` for UNIX and non-framework Mac OS X builds, `~/Library/Python/X.Y/lib/python/site-packages` for Mac

framework builds, and `%APPDATA%\Python\PythonXY\site-packages` on Windows. This directory is a site directory, which means that `.pth` files in it will be processed.

site.USER_BASE

Path to the base directory for the user site-packages. Can be `None` if `getuserbase()` hasn't been called yet. Default value is `~/local` for UNIX and Mac OS X non-framework builds, `~/Library/Python/X.Y` for Mac framework builds, and `%APPDATA%\Python` for Windows. This value is used by Distutils to compute the installation directories for scripts, data files, Python modules, etc. for the user installation scheme. See also `PYTHONUSERBASE`.

site.main()

Adds all the standard site-specific directories to the module search path. This function is called automatically when this module is imported, unless the Python interpreter was started with the `-S` flag.

在 3.3 版更改: This function used to be called unconditionally.

site.addsitedir(sitedir, known_paths=None)

Add a directory to `sys.path` and process its `.pth` files. Typically used in `sitecustomize` or `usercustomize` (see above).

site.getsitepackages()

Return a list containing all global site-packages directories.

3.2 新版功能.

site.getuserbase()

Return the path of the user base directory, `USER_BASE`. If it is not initialized yet, this function will also set it, respecting `PYTHONUSERBASE`.

3.2 新版功能.

site.getusersitepackages()

Return the path of the user-specific site-packages directory, `USER_SITE`. If it is not initialized yet, this function will also set it, respecting `PYTHONNOUSERSITE` and `USER_BASE`.

3.2 新版功能.

The `site` module also provides a way to get the user directories from the command line:

```
$ python3 -m site --user-site
/home/user/.local/lib/python3.3/site-packages
```

If it is called without arguments, it will print the contents of `sys.path` on the standard output, followed by the value of `USER_BASE` and whether the directory exists, then the same thing for `USER_SITE`, and finally the value of `ENABLE_USER_SITE`.

--user-base

Print the path to the user base directory.

--user-site

Print the path to the user site-packages directory.

If both options are given, user base and user site will be printed (always in this order), separated by `os.pathsep`.

If any option is given, the script will exit with one of these values: 0 if the user site-packages directory is enabled, 1 if it was disabled by the user, 2 if it is disabled for security reasons or by an administrator, and a value greater than 2 if there is an error.

参见:

PEP 370 – 分用户的 site-packages 目录

29.14 `fpectl` —Floating point exception control

注解： The `fpectl` module is not built by default, and its usage is discouraged and may be dangerous except in the hands of experts. See also the section *Limitations and other considerations* on limitations for more details.

Most computers carry out floating point operations in conformance with the so-called IEEE-754 standard. On any real computer, some floating point operations produce results that cannot be expressed as a normal floating point value. For example, try

```
>>> import math
>>> math.exp(1000)
inf
>>> math.exp(1000) / math.exp(1000)
nan
```

(The example above will work on many platforms. DEC Alpha may be one exception.) “Inf” is a special, non-numeric value in IEEE-754 that stands for “infinity”, and “nan” means “not a number.” Note that, other than the non-numeric results, nothing special happened when you asked Python to carry out those calculations. That is in fact the default behaviour prescribed in the IEEE-754 standard, and if it works for you, stop reading now.

In some circumstances, it would be better to raise an exception and stop processing at the point where the faulty operation was attempted. The `fpectl` module is for use in that situation. It provides control over floating point units from several hardware manufacturers, allowing the user to turn on the generation of SIGFPE whenever any of the IEEE-754 exceptions Division by Zero, Overflow, or Invalid Operation occurs. In tandem with a pair of wrapper macros that are inserted into the C code comprising your python system, SIGFPE is trapped and converted into the Python `FloatingPointError` exception.

The `fpectl` module defines the following functions and may raise the given exception:

`fpectl.turnon_sigfpe()`

Turn on the generation of SIGFPE, and set up an appropriate signal handler.

`fpectl.turnoff_sigfpe()`

Reset default handling of floating point exceptions.

exception `fpectl.FloatingPointError`

After `turnon_sigfpe()` has been executed, a floating point operation that raises one of the IEEE-754 exceptions Division by Zero, Overflow, or Invalid operation will in turn raise this standard Python exception.

29.14.1 Example

The following example demonstrates how to start up and test operation of the `fpectl` module.

```
>>> import fpectl
>>> import fpetest
>>> fpectl.turnon_sigfpe()
>>> fpetest.test()
overflow          PASS
FloatingPointError: Overflow

div by 0          PASS
FloatingPointError: Division by zero
[ more output from test elided ]
```

(下页继续)

(续上页)

```
>>> import math
>>> math.exp(1000)
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
FloatingPointError: in math_1
```

29.14.2 Limitations and other considerations

Setting up a given processor to trap IEEE-754 floating point errors currently requires custom code on a per-architecture basis. You may have to modify *fpectl* to control your particular hardware.

Conversion of an IEEE-754 exception to a Python exception requires that the wrapper macros `PyFPE_START_PROTECT` and `PyFPE_END_PROTECT` be inserted into your code in an appropriate fashion. Python itself has been modified to support the *fpectl* module, but many other codes of interest to numerical analysts have not.

The *fpectl* module is not thread-safe.

参见:

Some files in the source distribution may be interesting in learning more about how this module operates. The include file `Include/pyfpe.h` discusses the implementation of this module at some length. `Modules/fpetestmodule.c` gives several examples of use. Many additional examples can be found in `Objects/floatobject.c`.

自定义 Python 解释器

本章中描述的模块允许编写类似于 Python 的交互式解释器的接口。如果你想要一个支持附加一些特殊功能到 Python 语言的 Python 解释器，你应该看看 `code` 模块。（`codeop` 模块是低层级的，用于支持编译可能不完整的 Python 代码块。）

本章描述的完整模块列表如下：

30.1 `code` —Interpreter base classes

源代码 [Lib/code.py](#)

The `code` module provides facilities to implement read-eval-print loops in Python. Two classes and convenience functions are included which can be used to build applications which provide an interactive interpreter prompt.

class `code.InteractiveInterpreter` (*locals=None*)

This class deals with parsing and interpreter state (the user's namespace); it does not deal with input buffering or prompting or input file naming (the filename is always passed in explicitly). The optional *locals* argument specifies the dictionary in which code will be executed; it defaults to a newly created dictionary with key `'__name__'` set to `'__console__'` and key `'__doc__'` set to `None`.

class `code.InteractiveConsole` (*locals=None, filename="<console>"*)

Closely emulate the behavior of the interactive Python interpreter. This class builds on `InteractiveInterpreter` and adds prompting using the familiar `sys.ps1` and `sys.ps2`, and input buffering.

`code.interact` (*banner=None, readfunc=None, local=None, exitmsg=None*)

Convenience function to run a read-eval-print loop. This creates a new instance of `InteractiveConsole` and sets *readfunc* to be used as the `InteractiveConsole.raw_input()` method, if provided. If *local* is provided, it is passed to the `InteractiveConsole` constructor for use as the default namespace for the interpreter loop. The `interact()` method of the instance is then run with *banner* and *exitmsg* passed as the banner and exit message to use, if provided. The console object is discarded after use.

在 3.6 版更改：加入 *exitmsg* 参数。

`code.compile_command(source, filename=<input>, symbol='single')`

This function is useful for programs that want to emulate Python's interpreter main loop (a.k.a. the read-eval-print loop). The tricky part is to determine when the user has entered an incomplete command that can be completed by entering more text (as opposed to a complete command or a syntax error). This function *almost* always makes the same decision as the real interpreter main loop.

source is the source string; *filename* is the optional filename from which source was read, defaulting to '<input>'; and *symbol* is the optional grammar start symbol, which should be either 'single' (the default) or 'eval'.

Returns a code object (the same as `compile(source, filename, symbol)`) if the command is complete and valid; `None` if the command is incomplete; raises `SyntaxError` if the command is complete and contains a syntax error, or raises `OverflowError` or `ValueError` if the command contains an invalid literal.

30.1.1 Interactive Interpreter Objects

`InteractiveInterpreter.runsource(source, filename=<input>, symbol='single')`

Compile and run some source in the interpreter. Arguments are the same as for `compile_command()`; the default for *filename* is '<input>', and for *symbol* is 'single'. One several things can happen:

- The input is incorrect; `compile_command()` raised an exception (`SyntaxError` or `OverflowError`). A syntax traceback will be printed by calling the `showsyntaxerror()` method. `runsource()` returns `False`.
- The input is incomplete, and more input is required; `compile_command()` returned `None`. `runsource()` returns `True`.
- The input is complete; `compile_command()` returned a code object. The code is executed by calling the `runcode()` (which also handles run-time exceptions, except for `SystemExit`). `runsource()` returns `False`.

The return value can be used to decide whether to use `sys.ps1` or `sys.ps2` to prompt the next line.

`InteractiveInterpreter.runcode(code)`

Execute a code object. When an exception occurs, `showtraceback()` is called to display a traceback. All exceptions are caught except `SystemExit`, which is allowed to propagate.

A note about `KeyboardInterrupt`: this exception may occur elsewhere in this code, and may not always be caught. The caller should be prepared to deal with it.

`InteractiveInterpreter.showsyntaxerror(filename=None)`

Display the syntax error that just occurred. This does not display a stack trace because there isn't one for syntax errors. If *filename* is given, it is stuffed into the exception instead of the default filename provided by Python's parser, because it always uses '<string>' when reading from a string. The output is written by the `write()` method.

`InteractiveInterpreter.showtraceback()`

Display the exception that just occurred. We remove the first stack item because it is within the interpreter object implementation. The output is written by the `write()` method.

在 3.5 版更改: The full chained traceback is displayed instead of just the primary traceback.

`InteractiveInterpreter.write(data)`

Write a string to the standard error stream (`sys.stderr`). Derived classes should override this to provide the appropriate output handling as needed.

30.1.2 Interactive Console Objects

The `InteractiveConsole` class is a subclass of `InteractiveInterpreter`, and so offers all the methods of the interpreter objects as well as the following additions.

`InteractiveConsole.interact` (*banner=None, exitmsg=None*)

Closely emulate the interactive Python console. The optional *banner* argument specifies the banner to print before the first interaction; by default it prints a banner similar to the one printed by the standard Python interpreter, followed by the class name of the console object in parentheses (so as not to confuse this with the real interpreter – since it’s so close!).

The optional *exitmsg* argument specifies an exit message printed when exiting. Pass the empty string to suppress the exit message. If *exitmsg* is not given or `None`, a default message is printed.

在 3.4 版更改: To suppress printing any banner, pass an empty string.

在 3.6 版更改: Print an exit message when exiting.

`InteractiveConsole.push` (*line*)

Push a line of source text to the interpreter. The line should not have a trailing newline; it may have internal newlines. The line is appended to a buffer and the interpreter’s `runsource()` method is called with the concatenated contents of the buffer as source. If this indicates that the command was executed or invalid, the buffer is reset; otherwise, the command is incomplete, and the buffer is left as it was after the line was appended. The return value is `True` if more input is required, `False` if the line was dealt with in some way (this is the same as `runsource()`).

`InteractiveConsole.resetbuffer` ()

Remove any unhandled source text from the input buffer.

`InteractiveConsole.raw_input` (*prompt=""*)

Write a prompt and read a line. The returned line does not include the trailing newline. When the user enters the EOF key sequence, `EOFError` is raised. The base implementation reads from `sys.stdin`; a subclass may replace this with a different implementation.

30.2 codeop —Compile Python code

Source code: `Lib/codeop.py`

The `codeop` module provides utilities upon which the Python read-eval-print loop can be emulated, as is done in the `code` module. As a result, you probably don’t want to use the module directly; if you want to include such a loop in your program you probably want to use the `code` module instead.

There are two parts to this job:

1. Being able to tell if a line of input completes a Python statement: in short, telling whether to print ‘>>>’ or ‘...’ next.
2. Remembering which future statements the user has entered, so subsequent input can be compiled with these in effect.

The `codeop` module provides a way of doing each of these things, and a way of doing them both.

To do just the former:

`codeop.compile_command` (*source, filename=<input>, symbol="single"*)

Tries to compile *source*, which should be a string of Python code and return a code object if *source* is valid Python code. In that case, the filename attribute of the code object will be *filename*, which defaults to ‘<input>’. Returns `None` if *source* is not valid Python code, but is a prefix of valid Python code.

If there is a problem with *source*, an exception will be raised. *SyntaxError* is raised if there is invalid Python syntax, and *OverflowError* or *ValueError* if there is an invalid literal.

The *symbol* argument determines whether *source* is compiled as a statement ('single', the default) or as an *expression* ('eval'). Any other value will cause *ValueError* to be raised.

注解: It is possible (but not likely) that the parser stops parsing with a successful outcome before reaching the end of the source; in this case, trailing symbols may be ignored instead of causing an error. For example, a backslash followed by two newlines may be followed by arbitrary garbage. This will be fixed once the API for the parser is better.

class `codeop.Compile`

Instances of this class have `__call__()` methods identical in signature to the built-in function `compile()`, but with the difference that if the instance compiles program text containing a `__future__` statement, the instance ‘remembers’ and compiles all subsequent program texts with the statement in force.

class `codeop.CommandCompiler`

Instances of this class have `__call__()` methods identical in signature to `compile_command()`; the difference is that if the instance compiles program text containing a `__future__` statement, the instance ‘remembers’ and compiles all subsequent program texts with the statement in force.

本章中介绍的模块提供了导入其他 Python 模块和挂钩以自定义导入过程的新方法。
本章描述的完整模块列表如下：

31.1 `zipimport` — Import modules from Zip archives

This module adds the ability to import Python modules (`*.py`, `*.pyc`) and packages from ZIP-format archives. It is usually not needed to use the `zipimport` module explicitly; it is automatically used by the built-in `import` mechanism for `sys.path` items that are paths to ZIP archives.

Typically, `sys.path` is a list of directory names as strings. This module also allows an item of `sys.path` to be a string naming a ZIP file archive. The ZIP archive can contain a subdirectory structure to support package imports, and a path within the archive can be specified to only import from a subdirectory. For example, the path `example.zip/lib/` would only import from the `lib/` subdirectory within the archive.

Any files may be present in the ZIP archive, but only files `.py` and `.pyc` are available for import. ZIP import of dynamic modules (`.pyd`, `.so`) is disallowed. Note that if an archive only contains `.py` files, Python will not attempt to modify the archive by adding the corresponding `.pyc` file, meaning that if a ZIP archive doesn't contain `.pyc` files, importing may be rather slow.

ZIP archives with an archive comment are currently not supported.

参见：

PKZIP Application Note Phil Katz 编写的 ZIP 文件格式文档，此格式和使用的算法的创建者。

PEP 273 - Import Modules from Zip Archives Written by James C. Ahlstrom, who also provided an implementation. Python 2.3 follows the specification in PEP 273, but uses an implementation written by Just van Rossum that uses the import hooks described in PEP 302.

PEP 302 - New Import Hooks The PEP to add the import hooks that help this module work.

This module defines an exception:

exception `zipimport.ZipImportError`

Exception raised by zipimporter objects. It's a subclass of `ImportError`, so it can be caught as `ImportError`, too.

31.1.1 zipimporter Objects

`zipimporter` is the class for importing ZIP files.

class `zipimport.zipimporter` (*archivepath*)

Create a new zipimporter instance. *archivepath* must be a path to a ZIP file, or to a specific path within a ZIP file. For example, an *archivepath* of `foo/bar.zip/lib` will look for modules in the `lib` directory inside the ZIP file `foo/bar.zip` (provided that it exists).

`ZipImportError` is raised if *archivepath* doesn't point to a valid ZIP archive.

find_module (*fullname* [, *path*])

Search for a module specified by *fullname*. *fullname* must be the fully qualified (dotted) module name. It returns the zipimporter instance itself if the module was found, or `None` if it wasn't. The optional *path* argument is ignored—it's there for compatibility with the importer protocol.

get_code (*fullname*)

Return the code object for the specified module. Raise `ZipImportError` if the module couldn't be found.

get_data (*pathname*)

Return the data associated with *pathname*. Raise `OSError` if the file wasn't found.

在 3.3 版更改: `IOError` used to be raised instead of `OSError`.

get_filename (*fullname*)

Return the value `__file__` would be set to if the specified module was imported. Raise `ZipImportError` if the module couldn't be found.

3.1 新版功能.

get_source (*fullname*)

Return the source code for the specified module. Raise `ZipImportError` if the module couldn't be found, return `None` if the archive does contain the module, but has no source for it.

is_package (*fullname*)

Return True if the module specified by *fullname* is a package. Raise `ZipImportError` if the module couldn't be found.

load_module (*fullname*)

Load the module specified by *fullname*. *fullname* must be the fully qualified (dotted) module name. It returns the imported module, or raises `ZipImportError` if it wasn't found.

archive

The file name of the importer's associated ZIP file, without a possible subpath.

prefix

The subpath within the ZIP file where modules are searched. This is the empty string for zipimporter objects which point to the root of the ZIP file.

The *archive* and *prefix* attributes, when combined with a slash, equal the original *archivepath* argument given to the `zipimporter` constructor.

31.1.2 示例

Here is an example that imports a module from a ZIP archive - note that the `zipimport` module is not explicitly used.

```
$ unzip -l example.zip
Archive:  example.zip
  Length      Date    Time    Name
-----
   8467   11-26-02  22:30   jwzthreading.py
-----
   8467                      1 file

$ ./python
Python 2.3 (#1, Aug 1 2003, 19:54:32)
>>> import sys
>>> sys.path.insert(0, 'example.zip') # Add .zip file to front of path
>>> import jwzthreading
>>> jwzthreading.__file__
'example.zip/jwzthreading.py'
```

31.2 pkgutil —Package extension utility

Source code: [Lib/pkgutil.py](#)

This module provides utilities for the import system, in particular package support.

class `pkgutil.ModuleInfo` (*module_finder, name, ispkg*)
 A namedtuple that holds a brief summary of a module's info.

3.6 新版功能.

`pkgutil.extend_path` (*path, name*)

Extend the search path for the modules which comprise a package. Intended use is to place the following code in a package's `__init__.py`:

```
from pkgutil import extend_path
__path__ = extend_path(__path__, __name__)
```

This will add to the package's `__path__` all subdirectories of directories on `sys.path` named after the package. This is useful if one wants to distribute different parts of a single logical package as multiple directories.

It also looks for `*.pkg` files beginning where `*` matches the *name* argument. This feature is similar to `*.pth` files (see the [site](#) module for more information), except that it doesn't special-case lines starting with `import`. A `*.pkg` file is trusted at face value: apart from checking for duplicates, all entries found in a `*.pkg` file are added to the path, regardless of whether they exist on the filesystem. (This is a feature.)

If the input path is not a list (as is the case for frozen packages) it is returned unchanged. The input path is not modified; an extended copy is returned. Items are only appended to the copy at the end.

It is assumed that `sys.path` is a sequence. Items of `sys.path` that are not strings referring to existing directories are ignored. Unicode items on `sys.path` that cause errors when used as filenames may cause this function to raise an exception (in line with `os.path.isdir()` behavior).

class `pkgutil.ImpImporter` (*dirname=None*)

PEP 302 Finder that wraps Python's "classic" import algorithm.

If *dirname* is a string, a **PEP 302** finder is created that searches that directory. If *dirname* is `None`, a **PEP 302** finder is created that searches the current `sys.path`, plus any modules that are frozen or built-in.

Note that `ImpImporter` does not currently support being used by placement on `sys.meta_path`.

3.3 版后已移除: This emulation is no longer needed, as the standard import mechanism is now fully PEP 302 compliant and available in `importlib`.

class `pkgutil.ImpLoader` (*fullname*, *file*, *filename*, *etc*)
Loader that wraps Python's "classic" import algorithm.

3.3 版后已移除: This emulation is no longer needed, as the standard import mechanism is now fully PEP 302 compliant and available in `importlib`.

`pkgutil.find_loader` (*fullname*)
Retrieve a module *loader* for the given *fullname*.

This is a backwards compatibility wrapper around `importlib.util.find_spec()` that converts most failures to `ImportError` and only returns the loader rather than the full `ModuleSpec`.

在 3.3 版更改: Updated to be based directly on `importlib` rather than relying on the package internal PEP 302 import emulation.

在 3.4 版更改: Updated to be based on **PEP 451**

`pkgutil.get_importer` (*path_item*)
Retrieve a *finder* for the given *path_item*.

The returned finder is cached in `sys.path_importer_cache` if it was newly created by a path hook.

The cache (or part of it) can be cleared manually if a rescan of `sys.path_hooks` is necessary.

在 3.3 版更改: Updated to be based directly on `importlib` rather than relying on the package internal PEP 302 import emulation.

`pkgutil.get_loader` (*module_or_name*)
Get a *loader* object for *module_or_name*.

If the module or package is accessible via the normal import mechanism, a wrapper around the relevant part of that machinery is returned. Returns `None` if the module cannot be found or imported. If the named module is not already imported, its containing package (if any) is imported, in order to establish the package `__path__`.

在 3.3 版更改: Updated to be based directly on `importlib` rather than relying on the package internal PEP 302 import emulation.

在 3.4 版更改: Updated to be based on **PEP 451**

`pkgutil.iter_importers` (*fullname=""*)
Yield *finder* objects for the given module name.

If *fullname* contains a `'.'`, the finders will be for the package containing *fullname*, otherwise they will be all registered top level finders (i.e. those on both `sys.meta_path` and `sys.path_hooks`).

If the named module is in a package, that package is imported as a side effect of invoking this function.

If no module name is specified, all top level finders are produced.

在 3.3 版更改: Updated to be based directly on `importlib` rather than relying on the package internal PEP 302 import emulation.

`pkgutil.iter_modules` (*path=None*, *prefix=""*)
Yields *ModuleInfo* for all submodules on *path*, or, if *path* is `None`, all top-level modules on `sys.path`.

path should be either `None` or a list of paths to look for modules in.

prefix is a string to output on the front of every module name on output.

注解: Only works for a *finder* which defines an `iter_modules()` method. This interface is non-standard, so the module also provides implementations for `importlib.machinery.FileFinder` and `zipimport.zipimporter`.

在 3.3 版更改: Updated to be based directly on `importlib` rather than relying on the package internal PEP 302 import emulation.

`pkgutil.walk_packages(path=None, prefix="", onerror=None)`

Yields `ModuleInfo` for all modules recursively on `path`, or, if `path` is `None`, all accessible modules.

`path` should be either `None` or a list of paths to look for modules in.

`prefix` is a string to output on the front of every module name on output.

Note that this function must import all *packages* (not all modules!) on the given `path`, in order to access the `__path__` attribute to find submodules.

`onerror` is a function which gets called with one argument (the name of the package which was being imported) if any exception occurs while trying to import a package. If no `onerror` function is supplied, `ImportErrors` are caught and ignored, while all other exceptions are propagated, terminating the search.

例如:

```
# list all modules python can access
walk_packages()

# list all submodules of ctypes
walk_packages(ctypes.__path__, ctypes.__name__ + '.')

```

注解: Only works for a *finder* which defines an `iter_modules()` method. This interface is non-standard, so the module also provides implementations for `importlib.machinery.FileFinder` and `zipimport.zipimporter`.

在 3.3 版更改: Updated to be based directly on `importlib` rather than relying on the package internal PEP 302 import emulation.

`pkgutil.get_data(package, resource)`

Get a resource from a package.

This is a wrapper for the *loader* `get_data` API. The `package` argument should be the name of a package, in standard module format (`foo.bar`). The `resource` argument should be in the form of a relative filename, using `/` as the path separator. The parent directory name `..` is not allowed, and nor is a rooted name (starting with a `/`).

The function returns a binary string that is the contents of the specified resource.

For packages located in the filesystem, which have already been imported, this is the rough equivalent of:

```
d = os.path.dirname(sys.modules[package].__file__)
data = open(os.path.join(d, resource), 'rb').read()

```

If the package cannot be located or loaded, or it uses a *loader* which does not support `get_data`, then `None` is returned. In particular, the *loader* for *namespace packages* does not support `get_data`.

31.3 modulefinder — 查找脚本使用的模块

源码： `Lib/modulefinder.py`

该模块提供了一个 `ModuleFinder` 类，可用于确定脚本导入的模块集。`modulefinder.py` 也可以作为脚本运行，给出 Python 脚本的文件名作为参数，之后将打印导入模块的报告。

`modulefinder.AddPackagePath(pkg_name, path)`

记录名为 `pkg_name` 的包可以在指定的 `path` 中找到。

`modulefinder.ReplacePackage(oldname, newname)`

允许指定名为 `oldname` 的模块实际上是名为 `newname` 的包。

`class modulefinder.ModuleFinder(path=None, debug=0, excludes=[], replace_paths=[])`

该类提供 `run_script()` 和 `report()` 方法，用于确定脚本导入的模块集。`path` 可以是搜索模块的目录列表；如果没有指定，则使用 `sys.path`。`debug` 设置调试级别；更高的值使类打印调试消息，关于它正在做什么。`excludes` 是要从分析中排除的模块名称列表。`replace_paths` 是将在模块路径中替换的 (`oldpath`, `newpath`) 元组的列表。

`report()`

将报告打印到标准输出，列出脚本导入的模块及其路径，以及缺少或似乎缺失的模块。

`run_script(pathname)`

分析 `pathname` 文件的内容，该文件必须包含 Python 代码。

`modules`

一个将模块名称映射到模块的字典。请参阅 `ModuleFinder` 的示例用法。

31.3.1 ModuleFinder 的示例用法

稍后将分析的脚本 (`bacon.py`)：

```
import re, itertools

try:
    import baconhameggs
except ImportError:
    pass

try:
    import guido.python.ham
except ImportError:
    pass
```

将输出 `bacon.py` 报告的脚本：

```
from modulefinder import ModuleFinder

finder = ModuleFinder()
finder.run_script('bacon.py')

print('Loaded modules:')
for name, mod in finder.modules.items():
    print('%s: ' % name, end='')
    print(', '.join(list(mod.globalnames.keys())[:3]))
```

(下页继续)

(续上页)

```
print('-'*50)
print('Modules not imported:')
print('\n'.join(finder.badmodules.keys()))
```

输出样例（可能因架构而异）：

```
Loaded modules:
_types:
copyreg:  _inverted_registry, _slotnames, __all__
sre_compile:  isstring, _sre, _optimize_unicode
_sre:
sre_constants:  REPEAT_ONE, makedict, AT_END_LINE
sys:
re:  __module__, finditer, _expand
itertools:
__main__:  re, itertools, baconhameggs
sre_parse:  _PATTERNENDERS, SRE_FLAG_UNICODE
array:
types:  __module__, IntType, TypeType
-----
Modules not imported:
guido.python.ham
baconhameggs
```

31.4 runpy — Locating and executing Python modules

Source code: [Lib/runpy.py](#)

The *runpy* module is used to locate and run Python modules without importing them first. Its main use is to implement the `-m` command line switch that allows scripts to be located using the Python module namespace rather than the filesystem.

Note that this is *not* a sandbox module - all code is executed in the current process, and any side effects (such as cached imports of other modules) will remain in place after the functions have returned.

Furthermore, any functions and classes defined by the executed code are not guaranteed to work correctly after a *runpy* function has returned. If that limitation is not acceptable for a given use case, *importlib* is likely to be a more suitable choice than this module.

The *runpy* module provides two functions:

`runpy.run_module(mod_name, init_globals=None, run_name=None, alter_sys=False)`

Execute the code of the specified module and return the resulting module globals dictionary. The module's code is first located using the standard import mechanism (refer to [PEP 302](#) for details) and then executed in a fresh module namespace.

The *mod_name* argument should be an absolute module name. If the module name refers to a package rather than a normal module, then that package is imported and the `__main__` submodule within that package is then executed and the resulting module globals dictionary returned.

The optional dictionary argument *init_globals* may be used to pre-populate the module's globals dictionary before the code is executed. The supplied dictionary will not be modified. If any of the special global variables below are defined in the supplied dictionary, those definitions are overridden by *run_module()*.

The special global variables `__name__`, `__spec__`, `__file__`, `__cached__`, `__loader__` and `__package__` are set in the globals dictionary before the module code is executed (Note that this is a minimal set of variables - other variables may be set implicitly as an interpreter implementation detail).

`__name__` is set to `run_name` if this optional argument is not `None`, to `mod_name + '.__main__'` if the named module is a package and to the `mod_name` argument otherwise.

`__spec__` will be set appropriately for the *actually* imported module (that is, `__spec__.name` will always be `mod_name` or `mod_name + '.__main__'`, never `run_name`).

`__file__`, `__cached__`, `__loader__` and `__package__` are set as normal based on the module spec.

If the argument `alter_sys` is supplied and evaluates to `True`, then `sys.argv[0]` is updated with the value of `__file__` and `sys.modules[__name__]` is updated with a temporary module object for the module being executed. Both `sys.argv[0]` and `sys.modules[__name__]` are restored to their original values before the function returns.

Note that this manipulation of `sys` is not thread-safe. Other threads may see the partially initialised module, as well as the altered list of arguments. It is recommended that the `sys` module be left alone when invoking this function from threaded code.

参见:

The `-m` option offering equivalent functionality from the command line.

在 3.1 版更改: Added ability to execute packages by looking for a `__main__` submodule.

在 3.2 版更改: Added `__cached__` global variable (see [PEP 3147](#)).

在 3.4 版更改: Updated to take advantage of the module spec feature added by [PEP 451](#). This allows `__cached__` to be set correctly for modules run this way, as well as ensuring the real module name is always accessible as `__spec__.name`.

`runpy.run_path(file_path, init_globals=None, run_name=None)`

Execute the code at the named filesystem location and return the resulting module globals dictionary. As with a script name supplied to the CPython command line, the supplied path may refer to a Python source file, a compiled bytecode file or a valid `sys.path` entry containing a `__main__` module (e.g. a zipfile containing a top-level `__main__.py` file).

For a simple script, the specified code is simply executed in a fresh module namespace. For a valid `sys.path` entry (typically a zipfile or directory), the entry is first added to the beginning of `sys.path`. The function then looks for and executes a `__main__` module using the updated path. Note that there is no special protection against invoking an existing `__main__` entry located elsewhere on `sys.path` if there is no such module at the specified location.

The optional dictionary argument `init_globals` may be used to pre-populate the module's globals dictionary before the code is executed. The supplied dictionary will not be modified. If any of the special global variables below are defined in the supplied dictionary, those definitions are overridden by `run_path()`.

The special global variables `__name__`, `__spec__`, `__file__`, `__cached__`, `__loader__` and `__package__` are set in the globals dictionary before the module code is executed (Note that this is a minimal set of variables - other variables may be set implicitly as an interpreter implementation detail).

`__name__` is set to `run_name` if this optional argument is not `None` and to `'<run_path>'` otherwise.

If the supplied path directly references a script file (whether as source or as precompiled byte code), then `__file__` will be set to the supplied path, and `__spec__`, `__cached__`, `__loader__` and `__package__` will all be set to `None`.

If the supplied path is a reference to a valid `sys.path` entry, then `__spec__` will be set appropriately for the imported `__main__` module (that is, `__spec__.name` will always be `__main__`). `__file__`, `__cached__`, `__loader__` and `__package__` will be set as normal based on the module spec.

A number of alterations are also made to the `sys` module. Firstly, `sys.path` may be altered as described above. `sys.argv[0]` is updated with the value of `file_path` and `sys.modules[__name__]` is updated with a temporary module object for the module being executed. All modifications to items in `sys` are reverted before the function returns.

Note that, unlike `run_module()`, the alterations made to `sys` are not optional in this function as these adjustments are essential to allowing the execution of `sys.path` entries. As the thread-safety limitations still apply, use of this function in threaded code should be either serialised with the import lock or delegated to a separate process.

参见:

using-on-interface-options for equivalent functionality on the command line (`python path/to/script`).

3.2 新版功能.

在 3.4 版更改: Updated to take advantage of the module spec feature added by [PEP 451](#). This allows `__cached__` to be set correctly in the case where `__main__` is imported from a valid `sys.path` entry rather than being executed directly.

参见:

[PEP 338](#) – 将模块作为脚本执行 PEP 由 Nick Coghlan 撰写并实现。

[PEP 366](#) – Main module explicit relative imports PEP 由 Nick Coghlan 撰写并实现。

[PEP 451](#) – A ModuleSpec Type for the Import System PEP written and implemented by Eric Snow

using-on-general - CPython command line details

The `importlib.import_module()` function

31.5 importlib —The implementation of import

3.1 新版功能.

源代码 `Lib/importlib/__init__.py`

31.5.1 概述

The purpose of the `importlib` package is two-fold. One is to provide the implementation of the `import` statement (and thus, by extension, the `__import__()` function) in Python source code. This provides an implementation of `import` which is portable to any Python interpreter. This also provides an implementation which is easier to comprehend than one implemented in a programming language other than Python.

第二个目的是实现 `import` 的部分被公开在这个包中, 使得用户更容易创建他们自己的自定义对象 (通常被称为 *importer*) 来参与到导入过程中。

参见:

`import` `import` 语句的语言参考

[包规格说明](#) 包的初始规范。自从编写这个文档开始, 一些语义已经发生改变 (比如基于 `sys.modules` 中 `None` 的重定向)。

`__import__()` 函数 `import` 语句是这个函数的语法糖。

[PEP 235](#) 在忽略大小写的平台上进行导入

[PEP 263](#) 定义 Python 源代码编码

PEP 302 新导入钩子

PEP 328 导入：多行和绝对/相对

PEP 366 主模块显式相对导入

PEP 420 隐式命名空间包

PEP 451 导入系统的一个模块规范类型

PEP 488 消除 PYO 文件

PEP 489 多阶段扩展模块初始化

PEP 3120 使用 UTF-8 作为默认的源编码

PEP 3147 PYC 仓库目录

31.5.2 函数

`importlib.__import__(name, globals=None, locals=None, fromlist=(), level=0)`
内置 `__import__()` 函数的实现。

注解：程序式地导入模块应该使用 `import_module()` 而不是这个函数。

`importlib.import_module(name, package=None)`

导入一个模块。参数 `name` 指定了以绝对或相对导入方式导入什么模块 (比如要么像这样 `pkg.mod` 或者这样 `..mod`)。如果参数 `name` 使用相对导入的方式来指定, 那么那个参数 `package` 必须设置为那个包名, 这个包名作为解析这个包名的锚点 (比如 `import_module('..mod', 'pkg.subpkg')` 将会导入 `pkg.mod`)。

`import_module()` 函数是一个对 `importlib.__import__()` 进行简化的包装器。这意味着该函数的所有主义都来自于 `importlib.__import__()`。这两个函数之间最重要的不同点在于 `import_module()` 返回指定的包或模块 (例如 `pkg.mod`), 而 `__import__()` 返回最高层级的包或模块 (例如 `pkg`)。

如果动态导入一个自从解释器开始执行以来被创建的模块 (即创建了一个 Python 源代码文件), 为了让导入系统知道这个新模块, 可能需要调用 `invalidate_caches()`。

在 3.3 版更改: 父包会被自动导入。

`importlib.find_loader(name, path=None)`

查找一个模块的加载器, 可选择地在指定的 `path` 里面。如果这个模块是在 `sys.modules`, 那么返回 `sys.modules[name].__loader__` (除非这个加载器是 `None` 或者没有被设置, 在这样的情况下, 会引起 `ValueError` 异常)。否则使用 `sys.meta_path` 的一次搜索就结束。如果未发现加载器, 则返回 `None`。

点状的名称没有使得它父包或模块隐式地导入, 因为它需要加载它们并且可能不需要。为了适当地导入一个子模块, 需要导入子模块的所有父包并且使用正确的参数提供给 `path`。

3.3 新版功能.

在 3.4 版更改: 如果没有设置 `__loader__`, 会引起 `ValueError` 异常, 就像属性设置为 `None` 的时候一样。

3.4 版后已移除: 使用 `importlib.util.find_spec()` 来代替。

`importlib.invalidate_caches()`

使查找器存储在 `sys.meta_path` 中的内部缓存无效。如果一个查找器实现了

`invalidate_caches()`，那么它会被调用来执行那个无效过程。如果创建/安装任何模块，同时正在运行的程序是为了保证所有的查找器知道新模块的存在，那么应该调用这个函数。

3.3 新版功能.

`importlib.reload(module)`

重新加载之前导入的 *module*。那个参数必须是一个模块对象，所以它之前必须已经成功导入了。这样做是有用的，如果使用外部编辑器编已经辑过了那个模块的源代码文件并且想在退出 Python 解释器之前试验这个新版本的模块。函数的返回值是那个模块对象（如果重新导入导致一个不同的对象放置在 `sys.modules` 中，那么那个模块对象是有可能不同）。

当执行 `reload()` 的时候：

- Python 模块的代码会被重新编译并且那个模块级的代码被重新执行，通过重新使用一开始加载那个模块的 *loader*，定义一个新的绑定在那个模块字典中的名称的对象集合。扩展模块的“init”函数不会被调用第二次。
- 与 Python 中的所有的其它对象一样，旧的对象只有在它们的引用计数为 0 之后才会被回收。
- 模块命名空间中的名称重新指向任何新的或更改后的对象。
- 其他旧对象的引用（例如那个模块的外部名称）不会被重新绑定到引用的新对象的，并且如果需要，必须在出现的每个命名空间中进行更新。

有一些其他注意事项：

当一个模块被重新加载的时候，它的字典（包含了那个模块的全局变量）会被保留。名称的重新定义会覆盖旧的定义，所以通常来说这不是问题。如果一个新模块没有定义在旧版本模块中定义的名称，则将保留旧版本中的定义。这一特性可用于作为那个模块的优点，如果它维护一个全局表或者对象的缓存——使用 `try` 语句，就可以测试表的存在并且跳过它的初始化，如果有需要的话：

```
try:
    cache
except NameError:
    cache = {}
```

重新加载内置的或者动态加载模块，通常来说不是很有用处。不推荐重新加载” `sys`，`__main__`，`builtins` 和其它关键模块。在很多例子中，扩展模块并不是设计为不止一次的初始化，并且当重新加载时，可能会以任意方式失败。

If a module imports objects from another module using `from ... import ...`, calling `reload()` for the other module does not redefine the objects imported from it — one way around this is to re-execute the `from` statement, another is to use `import` and qualified names (*module.name*) instead.

如果一个模块创建一个类的实例，重新加载定义那个类的模块不影响那些实例的方法定义——它们继续使用旧类中的定义。对于子类来说同样是正确的。

3.4 新版功能.

31.5.3 `importlib.abc` ——关于导入的抽象基类

源代码： [Lib/importlib/abc.py](#)

The `importlib.abc` module contains all of the core abstract base classes used by `import`. Some subclasses of the core abstract base classes are also provided to help in implementing the core ABCs.

ABC 类的层次结构：

```

object
+-- Finder (deprecated)
|   +-- MetaPathFinder
|   +-- PathEntryFinder
+-- Loader
    +-- ResourceLoader -----+
    +-- InspectLoader          |
        +-- ExecutionLoader --+
                                +-- FileLoader
                                +-- SourceLoader

```

class importlib.abc.Finder

代表finder的一个抽象基类

3.3 版后已移除: 使用MetaPathFinder 或PathEntryFinder 来代替。

abstractmethod find_module (fullname, path=None)

为指定的模块查找loader定义的抽象方法。本来是在 PEP 302 指定的, 这个方法是在sys.meta_path 和基于路径的导入子系统中使用。

在 3.4 版更改: 当被调用的时候, 返回 None 而不是引发NotImplementedError。

class importlib.abc.MetaPathFinder

代表meta path finder 的一个抽象基类。为了保持兼容性, 这是Finder 的一个子类。

3.3 新版功能。

find_spec (fullname, path, target=None)

一个为了在指定的模块中查找模块规格而定义的抽象方法。如果这是最高层级导入, path 的值将会是 None。否则, 这是一个查找子包或者模块的方法并且 path 将会是来自父包的 __path__ 的值。如果未能查找到一个规格, 返回 None。当向这个函数传参时, target 是查找器用于做出对返回什么规格更具有根据的猜测的一个模块对象。

3.4 新版功能。

find_module (fullname, path)

一个用于查找指定的模块中loader的遗留方法。如果这是最高层级的导入, path 的值将会是 None。否则, 这是一个查找子包或者模块的方法, 并且 path 的值将会是来自父包的 __path__ 的值。如果未发现加载器, 返回 None。

如果定义了find_spec() 方法, 则提供了向后兼容的功能。

在 3.4 版更改: 当调用这个方法的时候返回 None 而不是引发NotImplementedError。可以使用find_spec() 来提供功能。

3.4 版后已移除: 使用find_spec() 来代替。

invalidate_caches ()

当被调用的时候, 一个可选的方法应该将查找器使用的任何内部缓存进行无效。将在sys.meta_path 上的所有查找器的缓存进行无效的时候, 这个函数被importlib.invalidate_caches() 所使用。

在 3.4 版更改: 当方法被调用的时候, 方法返回是 None 而不是 NotImplemented。

class importlib.abc.PathEntryFinder

path entry finder 的一个抽象基类。尽管这个基类和MetaPathFinder 有一些相似之处, 但是PathEntryFinder 只在由PathFinder 提供的基于路径导入子系统中使用。这个抽象类是Finder 的一个子类, 仅仅是因为兼容性的原因。

3.3 新版功能。

find_spec (fullname, target=None)

一个为了在指定的模块中查找模块规格而定义的抽象方法。查找器只会在赋值给 `path entry` 里面查找那个模块。如果未能找到一个规格，返回 `None`。当向这个函数传参时，`target` 是这个查找器用于做出对于返回什么规格更具有根据的猜测的一个模块对象。

3.4 新版功能。

find_loader (fullname)

一个用于在模块中查找一个 `loader` 的遗留方法。返回一个 `(loader, portion)` 的 2 元组，`portion` 是一个贡献给命名空间包部分的文件系统位置的序列。加载器可能是 `None`，同时正在指定的 `portion` 表示的是贡献给命名空间包的文件系统位置。`portion` 可以使用一个空列表来表示加载器不是命名空间包的一部分。如果 `loader` 是 `None` 并且 `portion` 是一个空列表，那么命名空间包中无加载器或者文件系统位置可查找到（即在那个模块中未能找到任何东西）。

如果定义了 `find_spec()`，则提供了向后兼容的功能。

在 3.4 版更改：返回 `(None, [])` 而不是引发 `NotImplementedError`。当可于提供相应的功能的时候，使用 `find_spec()`。

3.4 版后已移除：使用 `find_spec()` 来代替。

find_module (fullname)

`Finder.find_module` 的具体实现，该方法等价于 `self.find_loader(fullname)[0]`()`。

3.4 版后已移除：使用 `find_spec()` 来代替。

invalidate_caches ()

当被调用的时候，一个可选的方法应该将查找器使用的任何内部缓存进行无效。当将所有缓存的查找器的缓存进行无效的时候，该函数被 `PathFinder.invalidate_caches()` 使用。

class `importlib.abc.Loader`

`loader` 的抽象基类。关于一个加载器的实际定义请查看 [PEP 302](#)。

create_module (spec)

当导入一个模块的时候，一个返回将要使用的那个模块对象的方法。这个方法可能返回 `None`，这暗示着应该发生默认的模式创建语义。”

3.4 新版功能。

在 3.5 版更改：从 Python 3.6 开始，当定义了 `exec_module()` 的时候，这个方法将不会是可选的。

exec_module (module)

当一个模块被导入或重新加载时，一个抽象方法在它自己的命名空间中执行那个模块。当调用 `exec_module()` 的时候，那个模块应该已经被初始化了。当这个方法存在时，必须定义 `create_module()`。

3.4 新版功能。

在 3.6 版更改：`create_module()` 也必须被定义。

load_module (fullname)

用于加载一个模块的传统方法。如果这个模块不能被导入，将引起 `ImportError` 异常，否则返回那个被加载的模块。

如果请求的模块已经存在 `sys.modules`，应该使用并且重新加载那个模块。否则加载器应该是创建一个新的模块并且在任何家过程开始之前将这个新模块插入到 `sys.modules` 中，来阻止递归导入。如果加载器插入了一个模块并且加载失败了，加载器必须从 `sys.modules` 中将这个模块移除。在加载器开始执行之前，已经在 `sys.modules` 中的模块应该被忽略（查看 `importlib.util.module_for_loader()`）。

加载器应该在模块上面设置几个属性。（要知道当重新加载一个模块的时候，那些属性某部分可以改变）：

- `__name__` 模块的名字
- `__file__` 模块数据存储的路径 (不是为了内置的模块而设置)
- `__cached__` 被存储或应该被存储的模块的编译版本的路径 (当这个属性不恰当的时候不设置)。
- `__path__` 指定在一个包中搜索路径的一个字符串列表。这个属性不在模块上面进行设置。
- `__package__` 模块/包的父包。如果这个模块是最上层的, 那么它是一个为空字符串的值。`importlib.util.module_for_loader()` 装饰器可以处理 `__package__` 的细节。
- `__loader__` 用来加载那个模块的加载器。`importlib.util.module_for_loader()` 装饰器可以处理 `__package__` 的细节。

当 `exec_module()` 可用的时候, 那么则提供了向后兼容的功能。

在 3.4 版更改: 当这个方法被调用的时候, 触发 `ImportError` 异常而不是 `NotImplementedError`。当 `exec_module()` 可用的时候, 使用它的功能。

3.4 版后已移除: 加载模块推荐的使用的 API 是 `exec_module()` (和 `create_module()`)。加载器应该实现它而不是 `load_module()`。当 `exec_module()` 被实现的时候, 导入机制关心的是 `load_module()` 所有其他的责任。

`module_repr (module)`

一个遗留方法, 在实现时计算并返回给定模块的 `repr`, 作为字符串。模块类型的默认 `repr()` 将根据需要使用此方法的结果。

3.3 新版功能。

在 3.4 版更改: 是可选的方法而不是一个抽象方法。

3.4 版后已移除: 现在导入机制会自动地关注这个方法。

`class importlib.abc.ResourceLoader`

An abstract base class for a *loader* which implements the optional **PEP 302** protocol for loading arbitrary resources from the storage back-end.

`abstractmethod get_data (path)`

An abstract method to return the bytes for the data located at *path*. Loaders that have a file-like storage back-end that allows storing arbitrary data can implement this abstract method to give direct access to the data stored. `OSError` is to be raised if the *path* cannot be found. The *path* is expected to be constructed using a module's `__file__` attribute or an item from a package's `__path__`.

在 3.4 版更改: Raises `OSError` instead of `NotImplementedError`.

`class importlib.abc.InspectLoader`

An abstract base class for a *loader* which implements the optional **PEP 302** protocol for loaders that inspect modules.

`get_code (fullname)`

Return the code object for a module, or `None` if the module does not have a code object (as would be the case, for example, for a built-in module). Raise an `ImportError` if loader cannot find the requested module.

注解: While the method has a default implementation, it is suggested that it be overridden if possible for performance.

在 3.4 版更改: No longer abstract and a concrete implementation is provided.

`abstractmethod get_source (fullname)`

An abstract method to return the source of a module. It is returned as a text string using *universal newlines*,

translating all recognized line separators into `'\n'` characters. Returns `None` if no source is available (e.g. a built-in module). Raises `ImportError` if the loader cannot find the module specified.

在 3.4 版更改: Raises `ImportError` instead of `NotImplementedError`.

is_package (*fullname*)

An abstract method to return a true value if the module is a package, a false value otherwise. `ImportError` is raised if the *loader* cannot find the module.

在 3.4 版更改: Raises `ImportError` instead of `NotImplementedError`.

static source_to_code (*data*, *path*=`'<string>'`)

Create a code object from Python source.

The *data* argument can be whatever the `compile()` function supports (i.e. string or bytes). The *path* argument should be the “path” to where the source code originated from, which can be an abstract concept (e.g. location in a zip file).

With the subsequent code object one can execute it in a module by running `exec(code, module.__dict__)`.

3.4 新版功能.

在 3.5 版更改: Made the method static.

exec_module (*module*)

Implementation of `Loader.exec_module()`.

3.4 新版功能.

load_module (*fullname*)

Implementation of `Loader.load_module()`.

3.4 版后已移除: use `exec_module()` instead.

class `importlib.abc.ExecutionLoader`

An abstract base class which inherits from `InspectLoader` that, when implemented, helps a module to be executed as a script. The ABC represents an optional **PEP 302** protocol.

abstractmethod `get_filename` (*fullname*)

An abstract method that is to return the value of `__file__` for the specified module. If no path is available, `ImportError` is raised.

If source code is available, then the method should return the path to the source file, regardless of whether a bytecode was used to load the module.

在 3.4 版更改: Raises `ImportError` instead of `NotImplementedError`.

class `importlib.abc.FileLoader` (*fullname*, *path*)

An abstract base class which inherits from `ResourceLoader` and `ExecutionLoader`, providing concrete implementations of `ResourceLoader.get_data()` and `ExecutionLoader.get_filename()`.

The *fullname* argument is a fully resolved name of the module the loader is to handle. The *path* argument is the path to the file for the module.

3.3 新版功能.

name

The name of the module the loader can handle.

path

Path to the file of the module.

load_module (*fullname*)

Calls `super().load_module()`.

3.4 版后已移除: Use `Loader.exec_module()` instead.

abstractmethod `get_filename (fullname)`

Returns *path*.

abstractmethod `get_data (path)`

Reads *path* as a binary file and returns the bytes from it.

class `importlib.abc.SourceLoader`

An abstract base class for implementing source (and optionally bytecode) file loading. The class inherits from both *ResourceLoader* and *ExecutionLoader*, requiring the implementation of:

- *ResourceLoader.get_data()*
- *ExecutionLoader.get_filename()* Should only return the path to the source file; sourceless loading is not supported.

The abstract methods defined by this class are to add optional bytecode file support. Not implementing these optional methods (or causing them to raise *NotImplementedError*) causes the loader to only work with source code. Implementing the methods allows the loader to work with source *and* bytecode files; it does not allow for *sourceless* loading where only bytecode is provided. Bytecode files are an optimization to speed up loading by removing the parsing step of Python's compiler, and so no bytecode-specific API is exposed.

path_stats (path)

Optional abstract method which returns a *dict* containing metadata about the specified path. Supported dictionary keys are:

- 'mtime' (mandatory): an integer or floating-point number representing the modification time of the source code;
- 'size' (optional): the size in bytes of the source code.

Any other keys in the dictionary are ignored, to allow for future extensions. If the path cannot be handled, *OSError* is raised.

3.3 新版功能.

在 3.4 版更改: Raise *OSError* instead of *NotImplementedError*.

path_mtime (path)

Optional abstract method which returns the modification time for the specified path.

3.3 版后已移除: This method is deprecated in favour of *path_stats()*. You don't have to implement it, but it is still available for compatibility purposes. Raise *OSError* if the path cannot be handled.

在 3.4 版更改: Raise *OSError* instead of *NotImplementedError*.

set_data (path, data)

Optional abstract method which writes the specified bytes to a file path. Any intermediate directories which do not exist are to be created automatically.

When writing to the path fails because the path is read-only (*errno.EACCES/PermissionError*), do not propagate the exception.

在 3.4 版更改: No longer raises *NotImplementedError* when called.

get_code (fullname)

Concrete implementation of *InspectLoader.get_code()*.

exec_module (module)

Concrete implementation of *Loader.exec_module()*.

3.4 新版功能.

load_module (*fullname*)

Concrete implementation of *Loader.load_module()*.

3.4 版后已移除: Use *exec_module()* instead.

get_source (*fullname*)

Concrete implementation of *InspectorLoader.get_source()*.

is_package (*fullname*)

Concrete implementation of *InspectorLoader.is_package()*. A module is determined to be a package if its file path (as provided by *ExecutionLoader.get_filename()*) is a file named `__init__` when the file extension is removed **and** the module name itself does not end in `__init__`.

31.5.4 importlib.machinery – Importers and path hooks

Source code: [Lib/importlib/machinery.py](#)

This module contains the various objects that help import find and load modules.

`importlib.machinery.SOURCE_SUFFIXES`

A list of strings representing the recognized file suffixes for source modules.

3.3 新版功能.

`importlib.machinery.DEBUG_BYTECODE_SUFFIXES`

A list of strings representing the file suffixes for non-optimized bytecode modules.

3.3 新版功能.

3.5 版后已移除: Use *BYTECODE_SUFFIXES* instead.

`importlib.machinery.OPTIMIZED_BYTECODE_SUFFIXES`

A list of strings representing the file suffixes for optimized bytecode modules.

3.3 新版功能.

3.5 版后已移除: Use *BYTECODE_SUFFIXES* instead.

`importlib.machinery.BYTECODE_SUFFIXES`

A list of strings representing the recognized file suffixes for bytecode modules (including the leading dot).

3.3 新版功能.

在 3.5 版更改: The value is no longer dependent on `__debug__`.

`importlib.machinery.EXTENSION_SUFFIXES`

A list of strings representing the recognized file suffixes for extension modules.

3.3 新版功能.

`importlib.machinery.all_suffixes()`

Returns a combined list of strings representing all file suffixes for modules recognized by the standard import machinery. This is a helper for code which simply needs to know if a filesystem path potentially refers to a module without needing any details on the kind of module (for example, *inspect.getmodulename()*).

3.3 新版功能.

class `importlib.machinery.BuiltinImporter`

An *importer* for built-in modules. All known built-in modules are listed in *sys.builtin_module_names*. This class implements the *importlib.abc.MetaPathFinder* and *importlib.abc.InspectLoader* ABCs.

Only class methods are defined by this class to alleviate the need for instantiation.

在 3.5 版更改: As part of [PEP 489](#), the builtin importer now implements `Loader.create_module()` and `Loader.exec_module()`

class `importlib.machinery.FrozenImporter`

An *importer* for frozen modules. This class implements the `importlib.abc.MetaPathFinder` and `importlib.abc.InspectLoader` ABCs.

Only class methods are defined by this class to alleviate the need for instantiation.

class `importlib.machinery.WindowsRegistryFinder`

Finder for modules declared in the Windows registry. This class implements the `importlib.abc.MetaPathFinder` ABC.

Only class methods are defined by this class to alleviate the need for instantiation.

3.3 新版功能.

3.6 版后已移除: Use `site` configuration instead. Future versions of Python may not enable this finder by default.

class `importlib.machinery.PathFinder`

A *Finder* for `sys.path` and package `__path__` attributes. This class implements the `importlib.abc.MetaPathFinder` ABC.

Only class methods are defined by this class to alleviate the need for instantiation.

classmethod `find_spec(fullname, path=None, target=None)`

Class method that attempts to find a *spec* for the module specified by *fullname* on `sys.path` or, if defined, on *path*. For each path entry that is searched, `sys.path_importer_cache` is checked. If a non-false object is found then it is used as the *path entry finder* to look for the module being searched for. If no entry is found in `sys.path_importer_cache`, then `sys.path_hooks` is searched for a finder for the path entry and, if found, is stored in `sys.path_importer_cache` along with being queried about the module. If no finder is ever found then `None` is both stored in the cache and returned.

3.4 新版功能.

在 3.5 版更改: If the current working directory – represented by an empty string – is no longer valid then `None` is returned but no value is cached in `sys.path_importer_cache`.

classmethod `find_module(fullname, path=None)`

A legacy wrapper around `find_spec()`.

3.4 版后已移除: 使用 `find_spec()` 来代替.

classmethod `invalidate_caches()`

Calls `importlib.abc.PathEntryFinder.invalidate_caches()` on all finders stored in `sys.path_importer_cache`.

在 3.4 版更改: Calls objects in `sys.path_hooks` with the current working directory for `''` (i.e. the empty string).

class `importlib.machinery.FileFinder(path, *loader_details)`

A concrete implementation of `importlib.abc.PathEntryFinder` which caches results from the file system.

The *path* argument is the directory for which the finder is in charge of searching.

The *loader_details* argument is a variable number of 2-item tuples each containing a loader and a sequence of file suffixes the loader recognizes. The loaders are expected to be callables which accept two arguments of the module's name and the path to the file found.

The finder will cache the directory contents as necessary, making stat calls for each module search to verify the cache is not outdated. Because cache staleness relies upon the granularity of the operating system's state information of

the file system, there is a potential race condition of searching for a module, creating a new file, and then searching for the module the new file represents. If the operations happen fast enough to fit within the granularity of stat calls, then the module search will fail. To prevent this from happening, when you create a module dynamically, make sure to call `importlib.invalidate_caches()`.

3.3 新版功能.

path

The path the finder will search in.

find_spec (*fullname*, *target=None*)

Attempt to find the spec to handle *fullname* within *path*.

3.4 新版功能.

find_loader (*fullname*)

Attempt to find the loader to handle *fullname* within *path*.

invalidate_caches ()

Clear out the internal cache.

classmethod path_hook (**loader_details*)

A class method which returns a closure for use on `sys.path_hooks`. An instance of `FileFinder` is returned by the closure using the *path* argument given to the closure directly and *loader_details* indirectly.

If the argument to the closure is not an existing directory, `ImportError` is raised.

class `importlib.machinery.SourceFileLoader` (*fullname*, *path*)

A concrete implementation of `importlib.abc.SourceLoader` by subclassing `importlib.abc.FileLoader` and providing some concrete implementations of other methods.

3.3 新版功能.

name

The name of the module that this loader will handle.

path

The path to the source file.

is_package (*fullname*)

Return true if *path* appears to be for a package.

path_stats (*path*)

Concrete implementation of `importlib.abc.SourceLoader.path_stats()`.

set_data (*path*, *data*)

Concrete implementation of `importlib.abc.SourceLoader.set_data()`.

load_module (*name=None*)

Concrete implementation of `importlib.abc.Loader.load_module()` where specifying the name of the module to load is optional.

3.6 版后已移除: Use `importlib.abc.Loader.exec_module()` instead.

class `importlib.machinery.SourcelessFileLoader` (*fullname*, *path*)

A concrete implementation of `importlib.abc.FileLoader` which can import bytecode files (i.e. no source code files exist).

Please note that direct use of bytecode files (and thus not source code files) inhibits your modules from being usable by all Python implementations or new versions of Python which change the bytecode format.

3.3 新版功能.

name

The name of the module the loader will handle.

path

The path to the bytecode file.

is_package (*fullname*)

Determines if the module is a package based on *path*.

get_code (*fullname*)

Returns the code object for *name* created from *path*.

get_source (*fullname*)

Returns None as bytecode files have no source when this loader is used.

load_module (*name=None*)

Concrete implementation of `importlib.abc.Loader.load_module()` where specifying the name of the module to load is optional.

3.6 版后已移除: Use `importlib.abc.Loader.exec_module()` instead.

class `importlib.machinery.ExtensionFileLoader` (*fullname, path*)

A concrete implementation of `importlib.abc.ExecutionLoader` for extension modules.

The *fullname* argument specifies the name of the module the loader is to support. The *path* argument is the path to the extension module's file.

3.3 新版功能.

name

Name of the module the loader supports.

path

Path to the extension module.

create_module (*spec*)

Creates the module object from the given specification in accordance with **PEP 489**.

3.5 新版功能.

exec_module (*module*)

Initializes the given module object in accordance with **PEP 489**.

3.5 新版功能.

is_package (*fullname*)

Returns True if the file path points to a package's `__init__` module based on `EXTENSION_SUFFIXES`.

get_code (*fullname*)

Returns None as extension modules lack a code object.

get_source (*fullname*)

Returns None as extension modules do not have source code.

get_filename (*fullname*)

Returns *path*.

3.4 新版功能.

class `importlib.machinery.ModuleSpec` (*name, loader, *, origin=None, loader_state=None, is_package=None*)

A specification for a module's import-system-related state. This is typically exposed as the module's `__spec__` attribute. In the descriptions below, the names in parentheses give the corresponding attribute available directly on the module object. E.g. `module.__spec__.origin == module.__file__`. Note however that

while the *values* are usually equivalent, they can differ since there is no synchronization between the two objects. Thus it is possible to update the module's `__path__` at runtime, and this will not be automatically reflected in `__spec__.submodule_search_locations`.

3.4 新版功能.

name

(`__name__`)

A string for the fully-qualified name of the module.

loader

(`__loader__`)

The loader to use for loading. For namespace packages this should be set to `None`.

origin

(`__file__`)

Name of the place from which the module is loaded, e.g. “builtin” for built-in modules and the filename for modules loaded from source. Normally “origin” should be set, but it may be `None` (the default) which indicates it is unspecified.

submodule_search_locations

(`__path__`)

List of strings for where to find submodules, if a package (`None` otherwise).

loader_state

Container of extra module-specific data for use during loading (or `None`).

cached

(`__cached__`)

String for where the compiled module should be stored (or `None`).

parent

(`__package__`)

(Read-only) Fully-qualified name of the package to which the module belongs as a submodule (or `None`).

has_location

Boolean indicating whether or not the module's “origin” attribute refers to a loadable location.

31.5.5 `importlib.util` – Utility code for importers

Source code: [Lib/importlib/util.py](#)

This module contains the various objects that help in the construction of an *importer*.

`importlib.util.MAGIC_NUMBER`

The bytes which represent the bytecode version number. If you need help with loading/writing bytecode then consider `importlib.abc.SourceLoader`.

3.4 新版功能.

`importlib.util.cache_from_source(path, debug_override=None, *, optimization=None)`

Return the [PEP 3147/PEP 488](#) path to the byte-compiled file associated with the source *path*. For example, if *path* is `/foo/bar/baz.py` the return value would be `/foo/bar/__pycache__/baz.cpython-32.pyc` for Python 3.2. The `cpython-32` string comes from the current magic tag (see `get_tag()`; if `sys.implementation.cache_tag` is not defined then `NotImplementedError` will be raised).

The *optimization* parameter is used to specify the optimization level of the bytecode file. An empty string represents no optimization, so `/foo/bar/baz.py` with an *optimization* of `' '` will result in a bytecode path of `/foo/bar/__pycache__/baz.cpython-32.pyc`. `None` causes the interpreter's optimization level to be used. Any other value's string representation being used, so `/foo/bar/baz.py` with an *optimization* of `2` will lead to the bytecode path of `/foo/bar/__pycache__/baz.cpython-32.opt-2.pyc`. The string representation of *optimization* can only be alphanumeric, else `ValueError` is raised.

The *debug_override* parameter is deprecated and can be used to override the system's value for `__debug__`. A `True` value is the equivalent of setting *optimization* to the empty string. A `False` value is the same as setting *optimization* to `1`. If both *debug_override* and *optimization* are not `None` then `TypeError` is raised.

3.4 新版功能.

在 3.5 版更改: The *optimization* parameter was added and the *debug_override* parameter was deprecated.

在 3.6 版更改: 接受一个类路径对象。

`importlib.util.source_from_cache(path)`

Given the *path* to a [PEP 3147](#) file name, return the associated source code file path. For example, if *path* is `/foo/bar/__pycache__/baz.cpython-32.pyc` the returned path would be `/foo/bar/baz.py`. *path* need not exist, however if it does not conform to [PEP 3147](#) or [PEP 488](#) format, a `ValueError` is raised. If `sys.implementation.cache_tag` is not defined, `NotImplementedError` is raised.

3.4 新版功能.

在 3.6 版更改: 接受一个类路径对象。

`importlib.util.decode_source(source_bytes)`

Decode the given bytes representing source code and return it as a string with universal newlines (as required by `importlib.abc.InspectLoader.get_source()`).

3.4 新版功能.

`importlib.util.resolve_name(name, package)`

Resolve a relative module name to an absolute one.

If **name** has no leading dots, then **name** is simply returned. This allows for usage such as `importlib.util.resolve_name('sys', __package__)` without doing a check to see if the **package** argument is needed.

`ValueError` is raised if **name** is a relative module name but **package** is a false value (e.g. `None` or the empty string). `ValueError` is also raised a relative name would escape its containing package (e.g. requesting `..bacon` from within the `spam` package).

3.3 新版功能.

`importlib.util.find_spec(name, package=None)`

Find the *spec* for a module, optionally relative to the specified **package** name. If the module is in `sys.modules`, then `sys.modules[name].__spec__` is returned (unless the *spec* would be `None` or is not set, in which case `ValueError` is raised). Otherwise a search using `sys.meta_path` is done. `None` is returned if no *spec* is found.

If **name** is for a submodule (contains a dot), the parent module is automatically imported.

name and **package** work the same as for `import_module()`.

3.4 新版功能.

`importlib.util.module_from_spec(spec)`

Create a new module based on `spec` and `spec.loader.create_module`.

If `spec.loader.create_module` does not return `None`, then any pre-existing attributes will not be reset. Also, no `AttributeError` will be raised if triggered while accessing `spec` or setting an attribute on the module.

This function is preferred over using `types.ModuleType` to create a new module as `spec` is used to set as many import-controlled attributes on the module as possible.

3.5 新版功能.

`@importlib.util.module_for_loader`

A *decorator* for `importlib.abc.Loader.load_module()` to handle selecting the proper module object to load with. The decorated method is expected to have a call signature taking two positional arguments (e.g. `load_module(self, module)`) for which the second argument will be the module **object** to be used by the loader. Note that the decorator will not work on static methods because of the assumption of two arguments.

The decorated method will take in the **name** of the module to be loaded as expected for a *loader*. If the module is not found in `sys.modules` then a new one is constructed. Regardless of where the module came from, `__loader__` set to `self` and `__package__` is set based on what `importlib.abc.InspectLoader.is_package()` returns (if available). These attributes are set unconditionally to support reloading.

If an exception is raised by the decorated method and a module was added to `sys.modules`, then the module will be removed to prevent a partially initialized module from being left in `sys.modules`. If the module was already in `sys.modules` then it is left alone.

在 3.3 版更改: `__loader__` and `__package__` are automatically set (when possible).

在 3.4 版更改: Set `__name__`, `__loader__` `__package__` unconditionally to support reloading.

3.4 版后已移除: The import machinery now directly performs all the functionality provided by this function.

`@importlib.util.set_loader`

A *decorator* for `importlib.abc.Loader.load_module()` to set the `__loader__` attribute on the returned module. If the attribute is already set the decorator does nothing. It is assumed that the first positional argument to the wrapped method (i.e. `self`) is what `__loader__` should be set to.

在 3.4 版更改: Set `__loader__` if set to `None`, as if the attribute does not exist.

3.4 版后已移除: The import machinery takes care of this automatically.

`@importlib.util.set_package`

A *decorator* for `importlib.abc.Loader.load_module()` to set the `__package__` attribute on the returned module. If `__package__` is set and has a value other than `None` it will not be changed.

3.4 版后已移除: The import machinery takes care of this automatically.

`importlib.util.spec_from_loader(name, loader, *, origin=None, is_package=None)`

A factory function for creating a `ModuleSpec` instance based on a loader. The parameters have the same meaning as they do for `ModuleSpec`. The function uses available *loader* APIs, such as `InspectLoader.is_package()`, to fill in any missing information on the spec.

3.4 新版功能.

`importlib.util.spec_from_file_location(name, location, *, loader=None, submodule_search_locations=None)`

A factory function for creating a `ModuleSpec` instance based on the path to a file. Missing information will be filled in on the spec by making use of loader APIs and by the implication that the module will be file-based.

3.4 新版功能.

在 3.6 版更改: 接受一个类路径对象。

class `importlib.util.LazyLoader` (*loader*)

A class which postpones the execution of the loader of a module until the module has an attribute accessed.

This class **only** works with loaders that define `exec_module()` as control over what module type is used for the module is required. For those same reasons, the loader's `create_module()` method must return `None` or a type for which its `__class__` attribute can be mutated along with not using `slots`. Finally, modules which substitute the object placed into `sys.modules` will not work as there is no way to properly replace the module references throughout the interpreter safely; `ValueError` is raised if such a substitution is detected.

注解: For projects where startup time is critical, this class allows for potentially minimizing the cost of loading a module if it is never used. For projects where startup time is not essential then use of this class is **heavily** discouraged due to error messages created during loading being postponed and thus occurring out of context.

3.5 新版功能.

在 3.6 版更改: Began calling `create_module()`, removing the compatibility warning for `importlib.machinery.BuiltinImporter` and `importlib.machinery.ExtensionFileLoader`.

classmethod `factory` (*loader*)

A static method which returns a callable that creates a lazy loader. This is meant to be used in situations where the loader is passed by class instead of by instance.

```
suffixes = importlib.machinery.SOURCE_SUFFIXES
loader = importlib.machinery.SourceFileLoader
lazy_loader = importlib.util.LazyLoader.factory(loader)
finder = importlib.machinery.FileFinder(path, (lazy_loader, suffixes))
```

31.5.6 示例

Importing programmatically

To programmatically import a module, use `importlib.import_module()`.

```
import importlib

itertools = importlib.import_module('itertools')
```

Checking if a module can be imported

If you need to find out if a module can be imported without actually doing the import, then you should use `importlib.util.find_spec()`.

```
import importlib.util
import sys

# For illustrative purposes.
name = 'itertools'

spec = importlib.util.find_spec(name)
if spec is None:
    print("can't find the itertools module")
else:
    # If you chose to perform the actual import ...
```

(下页继续)

(续上页)

```

module = importlib.util.module_from_spec(spec)
spec.loader.exec_module(module)
# Adding the module to sys.modules is optional.
sys.modules[name] = module

```

Importing a source file directly

To import a Python source file directly, use the following recipe (Python 3.5 and newer only):

```

import importlib.util
import sys

# For illustrative purposes.
import tokenize
file_path = tokenize.__file__
module_name = tokenize.__name__

spec = importlib.util.spec_from_file_location(module_name, file_path)
module = importlib.util.module_from_spec(spec)
spec.loader.exec_module(module)
# Optional; only necessary if you want to be able to import the module
# by name later.
sys.modules[module_name] = module

```

Setting up an importer

For deep customizations of import, you typically want to implement an *importer*. This means managing both the *finder* and *loader* side of things. For finders there are two flavours to choose from depending on your needs: a *meta path finder* or a *path entry finder*. The former is what you would put on `sys.meta_path` while the latter is what you create using a *path entry hook* on `sys.path_hooks` which works with `sys.path` entries to potentially create a finder. This example will show you how to register your own importers so that import will use them (for creating an importer for yourself, read the documentation for the appropriate classes defined within this package):

```

import importlib.machinery
import sys

# For illustrative purposes only.
SpamMetaPathFinder = importlib.machinery.PathFinder
SpamPathEntryFinder = importlib.machinery.FileFinder
loader_details = (importlib.machinery.SourceFileLoader,
                  importlib.machinery.SOURCE_SUFFIXES)

# Setting up a meta path finder.
# Make sure to put the finder in the proper location in the list in terms of
# priority.
sys.meta_path.append(SpamMetaPathFinder)

# Setting up a path entry finder.
# Make sure to put the path hook in the proper location in the list in terms
# of priority.
sys.path_hooks.append(SpamPathEntryFinder.path_hook(loader_details))

```

Approximating `importlib.import_module()`

Import itself is implemented in Python code, making it possible to expose most of the import machinery through `importlib`. The following helps illustrate the various APIs that `importlib` exposes by providing an approximate implementation of `importlib.import_module()` (Python 3.4 and newer for the `importlib` usage, Python 3.6 and newer for other parts of the code).

```
import importlib.util
import sys

def import_module(name, package=None):
    """An approximate implementation of import."""
    absolute_name = importlib.util.resolve_name(name, package)
    try:
        return sys.modules[absolute_name]
    except KeyError:
        pass

    path = None
    if '.' in absolute_name:
        parent_name, _, child_name = absolute_name.rpartition('.')
        parent_module = import_module(parent_name)
        path = parent_module.__spec__.submodule_search_locations
    for finder in sys.meta_path:
        spec = finder.find_spec(absolute_name, path)
        if spec is not None:
            break
    else:
        msg = f'No module named {absolute_name!r}'
        raise ModuleNotFoundError(msg, name=absolute_name)
    module = importlib.util.module_from_spec(spec)
    spec.loader.exec_module(module)
    sys.modules[absolute_name] = module
    if path is not None:
        setattr(parent_module, child_name, module)
    return module
```

Python 提供了许多模块来帮助使用 Python 语言。这些模块支持标记化、解析、语法分析、字节码反汇编以及各种其他工具。

这些模块包括：

32.1 `parser` —Access Python parse trees

The `parser` module provides an interface to Python's internal parser and byte-code compiler. The primary purpose for this interface is to allow Python code to edit the parse tree of a Python expression and create executable code from this. This is better than trying to parse and modify an arbitrary Python code fragment as a string because parsing is performed in a manner identical to the code forming the application. It is also faster.

注解： From Python 2.5 onward, it's much more convenient to cut in at the Abstract Syntax Tree (AST) generation and compilation stage, using the `ast` module.

There are a few things to note about this module which are important to making use of the data structures created. This is not a tutorial on editing the parse trees for Python code, but some examples of using the `parser` module are presented.

Most importantly, a good understanding of the Python grammar processed by the internal parser is required. For full information on the language syntax, refer to reference-index. The parser itself is created from a grammar specification defined in the file `Grammar/Grammar` in the standard Python distribution. The parse trees stored in the ST objects created by this module are the actual output from the internal parser when created by the `expr()` or `suite()` functions, described below. The ST objects created by `sequence2st()` faithfully simulate those structures. Be aware that the values of the sequences which are considered “correct” will vary from one version of Python to another as the formal grammar for the language is revised. However, transporting code from one Python version to another as source text will always allow correct parse trees to be created in the target version, with the only restriction being that migrating to an older version of the interpreter will not support more recent language constructs. The parse trees are not typically compatible from one version to another, whereas source code has always been forward-compatible.

Each element of the sequences returned by `st2list()` or `st2tuple()` has a simple form. Sequences representing non-terminal elements in the grammar always have a length greater than one. The first element is an integer which identifies a production in the grammar. These integers are given symbolic names in the C header file `Include/graminit.h` and the Python module `symbol`. Each additional element of the sequence represents a component of the production as recognized in the input string: these are always sequences which have the same form as the parent. An important aspect of this structure which should be noted is that keywords used to identify the parent node type, such as the keyword `if` in an `if_stmt`, are included in the node tree without any special treatment. For example, the `if` keyword is represented by the tuple `(1, 'if')`, where `1` is the numeric value associated with all `NAME` tokens, including variable and function names defined by the user. In an alternate form returned when line number information is requested, the same token might be represented as `(1, 'if', 12)`, where the `12` represents the line number at which the terminal symbol was found.

Terminal elements are represented in much the same way, but without any child elements and the addition of the source text which was identified. The example of the `if` keyword above is representative. The various types of terminal symbols are defined in the C header file `Include/token.h` and the Python module `token`.

The ST objects are not required to support the functionality of this module, but are provided for three purposes: to allow an application to amortize the cost of processing complex parse trees, to provide a parse tree representation which conserves memory space when compared to the Python list or tuple representation, and to ease the creation of additional modules in C which manipulate parse trees. A simple “wrapper” class may be created in Python to hide the use of ST objects.

The `parser` module defines functions for a few distinct purposes. The most important purposes are to create ST objects and to convert ST objects to other representations such as parse trees and compiled code objects, but there are also functions which serve to query the type of parse tree represented by an ST object.

参见:

Module `symbol` Useful constants representing internal nodes of the parse tree.

Module `token` Useful constants representing leaf nodes of the parse tree and functions for testing node values.

32.1.1 Creating ST Objects

ST objects may be created from source code or from a parse tree. When creating an ST object from source, different functions are used to create the `'eval'` and `'exec'` forms.

`parser.expr(source)`

The `expr()` function parses the parameter `source` as if it were an input to `compile(source, 'file.py', 'eval')`. If the parse succeeds, an ST object is created to hold the internal parse tree representation, otherwise an appropriate exception is raised.

`parser.suite(source)`

The `suite()` function parses the parameter `source` as if it were an input to `compile(source, 'file.py', 'exec')`. If the parse succeeds, an ST object is created to hold the internal parse tree representation, otherwise an appropriate exception is raised.

`parser.sequence2st(sequence)`

This function accepts a parse tree represented as a sequence and builds an internal representation if possible. If it can validate that the tree conforms to the Python grammar and all nodes are valid node types in the host version of Python, an ST object is created from the internal representation and returned to the caller. If there is a problem creating the internal representation, or if the tree cannot be validated, a `ParserError` exception is raised. An ST object created this way should not be assumed to compile correctly; normal exceptions raised by compilation may still be initiated when the ST object is passed to `compilest()`. This may indicate problems not related to syntax (such as a `MemoryError` exception), but may also be due to constructs such as the result of parsing `del f(0)`, which escapes the Python parser but is checked by the bytecode compiler.

Sequences representing terminal tokens may be represented as either two-element lists of the form `(1, 'name')` or as three-element lists of the form `(1, 'name', 56)`. If the third element is present, it is assumed to be a

valid line number. The line number may be specified for any subset of the terminal symbols in the input tree.

`parser.tuple2st(sequence)`

This is the same function as `sequence2st()`. This entry point is maintained for backward compatibility.

32.1.2 Converting ST Objects

ST objects, regardless of the input used to create them, may be converted to parse trees represented as list- or tuple-trees, or may be compiled into executable code objects. Parse trees may be extracted with or without line numbering information.

`parser.st2list(st, line_info=False, col_info=False)`

This function accepts an ST object from the caller in `st` and returns a Python list representing the equivalent parse tree. The resulting list representation can be used for inspection or the creation of a new parse tree in list form. This function does not fail so long as memory is available to build the list representation. If the parse tree will only be used for inspection, `st2tuple()` should be used instead to reduce memory consumption and fragmentation. When the list representation is required, this function is significantly faster than retrieving a tuple representation and converting that to nested lists.

If `line_info` is true, line number information will be included for all terminal tokens as a third element of the list representing the token. Note that the line number provided specifies the line on which the token *ends*. This information is omitted if the flag is false or omitted.

`parser.st2tuple(st, line_info=False, col_info=False)`

This function accepts an ST object from the caller in `st` and returns a Python tuple representing the equivalent parse tree. Other than returning a tuple instead of a list, this function is identical to `st2list()`.

If `line_info` is true, line number information will be included for all terminal tokens as a third element of the list representing the token. This information is omitted if the flag is false or omitted.

`parser.compilest(st, filename='<syntax-tree>')`

The Python byte compiler can be invoked on an ST object to produce code objects which can be used as part of a call to the built-in `exec()` or `eval()` functions. This function provides the interface to the compiler, passing the internal parse tree from `st` to the parser, using the source file name specified by the `filename` parameter. The default value supplied for `filename` indicates that the source was an ST object.

Compiling an ST object may result in exceptions related to compilation; an example would be a `SyntaxError` caused by the parse tree for `del f(0)`: this statement is considered legal within the formal grammar for Python but is not a legal language construct. The `SyntaxError` raised for this condition is actually generated by the Python byte-compiler normally, which is why it can be raised at this point by the `parser` module. Most causes of compilation failure can be diagnosed programmatically by inspection of the parse tree.

32.1.3 Queries on ST Objects

Two functions are provided which allow an application to determine if an ST was created as an expression or a suite. Neither of these functions can be used to determine if an ST was created from source code via `expr()` or `suite()` or from a parse tree via `sequence2st()`.

`parser.isexpr(st)`

When `st` represents an 'eval' form, this function returns true, otherwise it returns false. This is useful, since code objects normally cannot be queried for this information using existing built-in functions. Note that the code objects created by `compilest()` cannot be queried like this either, and are identical to those created by the built-in `compile()` function.

`parser.issuite(st)`

This function mirrors `isexpr()` in that it reports whether an ST object represents an 'exec' form, commonly

known as a “suite.” It is not safe to assume that this function is equivalent to `not isexpr(st)`, as additional syntactic fragments may be supported in the future.

32.1.4 Exceptions and Error Handling

The parser module defines a single exception, but may also pass other built-in exceptions from other portions of the Python runtime environment. See each function for information about the exceptions it can raise.

exception `parser.ParserError`

Exception raised when a failure occurs within the parser module. This is generally produced for validation failures rather than the built-in `SyntaxError` raised during normal parsing. The exception argument is either a string describing the reason of the failure or a tuple containing a sequence causing the failure from a parse tree passed to `sequence2st()` and an explanatory string. Calls to `sequence2st()` need to be able to handle either type of exception, while calls to other functions in the module will only need to be aware of the simple string values.

Note that the functions `compilest()`, `expr()`, and `suite()` may raise exceptions which are normally raised by the parsing and compilation process. These include the built in exceptions `MemoryError`, `OverflowError`, `SyntaxError`, and `SystemError`. In these cases, these exceptions carry all the meaning normally associated with them. Refer to the descriptions of each function for detailed information.

32.1.5 ST Objects

Ordered and equality comparisons are supported between ST objects. Pickling of ST objects (using the `pickle` module) is also supported.

parser.STType

The type of the objects returned by `expr()`, `suite()` and `sequence2st()`.

ST objects have the following methods:

ST.compile (*filename*='<syntax-tree>')
Same as `compilest(st, filename)`.

ST.isexpr ()
Same as `isexpr(st)`.

ST.issuite ()
Same as `issuite(st)`.

ST.tolist (*line_info*=False, *col_info*=False)
Same as `st2list(st, line_info, col_info)`.

ST.totuple (*line_info*=False, *col_info*=False)
Same as `st2tuple(st, line_info, col_info)`.

32.1.6 Example: Emulation of `compile()`

While many useful operations may take place between parsing and bytecode generation, the simplest operation is to do nothing. For this purpose, using the `parser` module to produce an intermediate data structure is equivalent to the code

```
>>> code = compile('a + 5', 'file.py', 'eval')
>>> a = 5
>>> eval(code)
10
```

The equivalent operation using the `parser` module is somewhat longer, and allows the intermediate internal parse tree to be retained as an ST object:

```
>>> import parser
>>> st = parser.expr('a + 5')
>>> code = st.compile('file.py')
>>> a = 5
>>> eval(code)
10
```

An application which needs both ST and code objects can package this code into readily available functions:

```
import parser

def load_suite(source_string):
    st = parser.suite(source_string)
    return st, st.compile()

def load_expression(source_string):
    st = parser.expr(source_string)
    return st, st.compile()
```

32.2 ast — 抽象语法树

源代码： `Lib/ast.py`

`ast` 模块帮助 Python 程序处理 Python 语法的抽象语法树。抽象语法或许会随着 Python 的更新发布而改变；该模块能够帮助理解当前语法在编程层面的样貌。

抽象语法树可通过将 `ast.PyCF_ONLY_AST` 作为旗标传递给 `compile()` 内置函数来生成，或是使用此模块中提供的 `parse()` 辅助函数。返回结果将是一个对象树，其中的类都继承自 `ast.AST`。抽象语法树可被内置的 `compile()` 函数编译为一个 Python 代码对象。

32.2.1 节点类

class `ast.AST`

这是所有 AST 节点类的基类。实际上，这些节点类派生自 `Parser/Python.asdl` 文件，其中定义的语法树示例如下。它们在 C 语言模块 `_ast` 中定义，并被导出至 `ast` 模块。

抽象语法定义的每个左侧符号（比方说，`ast.stmt` 或者 `ast.expr`）定义了一个类。另外，在抽象语法定义的右侧，对每一个构造器也定义了一个类；这些类继承自树左侧的类。比如，`ast.BinOp` 继承自 `ast.expr`。对于多分支产生式（也就是“和规则”），树右侧的类是抽象的；只有特定构造器结点的实例能被构造。

_fields

每个具体类都有个属性 `_fields`，用来给出所有子节点的名字。

每个具体类的实例对它每个子节点都有一个属性，对应类型如文法中所定义。比如，`ast.BinOp` 的实例有个属性 `left`，类型是 `ast.expr`。

如果这些属性在文法中标记为可选（使用问号），对应值可能会是 `None`。如果这些属性有零或多个（用星号标记），对应值会用 Python 的列表来表示。所有可能的属性必须在用 `compile()` 编译得到 AST 时给出，且是有效的值。

lineno
col_offset

`ast.expr` 和 `ast.stmt` 子类的实例有 `lineno` 和 `col_offset` 属性。`lineno` 是源代码的行数 (从 1 开始, 所以第一行行数是 1), 而 `col_offset` 是该生成节点第一个 token 的 UTF-8 字节偏移量。记录下 UTF-8 偏移量的原因是 `parser` 内部使用 UTF-8。

一个类的构造器 `ast.T` 像下面这样 `parse` 它的参数。

- 如果有位置参数, 它们必须和 `T._fields` 中的元素一样多; 他们会像这些名字的属性一样被赋值。
- 如果有关键字参数, 它们必须被设为和给定值同名的属性。

比方说, 要创建和填充节点 `ast.UnaryOp`, 你得用

```
node = ast.UnaryOp()
node.op = ast.USub()
node.operand = ast.Num()
node.operand.n = 5
node.operand.lineno = 0
node.operand.col_offset = 0
node.lineno = 0
node.col_offset = 0
```

或者更紧凑点

```
node = ast.UnaryOp(ast.USub(), ast.Num(5, lineno=0, col_offset=0),
                  lineno=0, col_offset=0)
```

32.2.2 抽象文法

抽象文法目前定义如下

```
-- ASDL's 7 builtin types are:
-- identifier, int, string, bytes, object, singleton, constant
--
-- singleton: None, True or False
-- constant can be None, whereas None means "no value" for object.

module Python
{
    mod = Module(stmt* body)
        | Interactive(stmt* body)
        | Expression(expr body)

    -- not really an actual node but useful in Jython's typesystem.
    | Suite(stmt* body)

    stmt = FunctionDef(identifier name, arguments args,
                        stmt* body, expr* decorator_list, expr? returns)
        | AsyncFunctionDef(identifier name, arguments args,
                           stmt* body, expr* decorator_list, expr? returns)

    | ClassDef(identifier name,
               expr* bases,
               keyword* keywords,
               stmt* body,
```

(下页继续)

(续上页)

```

    expr* decorator_list)
| Return(expr? value)

| Delete(expr* targets)
| Assign(expr* targets, expr value)
| AugAssign(expr target, operator op, expr value)
-- 'simple' indicates that we annotate simple name without parens
| AnnAssign(expr target, expr annotation, expr? value, int simple)

-- use 'orelse' because else is a keyword in target languages
| For(expr target, expr iter, stmt* body, stmt* orelse)
| AsyncFor(expr target, expr iter, stmt* body, stmt* orelse)
| While(expr test, stmt* body, stmt* orelse)
| If(expr test, stmt* body, stmt* orelse)
| With(withitem* items, stmt* body)
| AsyncWith(withitem* items, stmt* body)

| Raise(expr? exc, expr? cause)
| Try(stmt* body, excepthandler* handlers, stmt* orelse, stmt* finalbody)
| Assert(expr test, expr? msg)

| Import(alias* names)
| ImportFrom(identifier? module, alias* names, int? level)

| Global(identifier* names)
| Nonlocal(identifier* names)
| Expr(expr value)
| Pass | Break | Continue

-- XXX Jython will be different
-- col_offset is the byte offset in the utf8 string the parser uses
attributes (int lineno, int col_offset)

-- BoolOp() can use left & right?
expr = BoolOp(boolop op, expr* values)
| BinOp(expr left, operator op, expr right)
| UnaryOp(unaryop op, expr operand)
| Lambda(arguments args, expr body)
| IfExp(expr test, expr body, expr orelse)
| Dict(expr* keys, expr* values)
| Set(expr* elts)
| ListComp(expr elt, comprehension* generators)
| SetComp(expr elt, comprehension* generators)
| DictComp(expr key, expr value, comprehension* generators)
| GeneratorExp(expr elt, comprehension* generators)
-- the grammar constrains where yield expressions can occur
| Await(expr value)
| Yield(expr? value)
| YieldFrom(expr value)
-- need sequences for compare to distinguish between
-- x < 4 < 3 and (x < 4) < 3
| Compare(expr left, cmpop* ops, expr* comparators)
| Call(expr func, expr* args, keyword* keywords)
| Num(object n) -- a number as a PyObject.
| Str(string s) -- need to specify raw, unicode, etc?
| FormattedValue(expr value, int? conversion, expr? format_spec)

```

(下页继续)

(续上页)

```

    | JoinedStr(expr* values)
    | Bytes(bytes s)
    | NameConstant(singleton value)
    | Ellipsis
    | Constant(constant value)

-- the following expression can appear in assignment context
    | Attribute(expr value, identifier attr, expr_context ctx)
    | Subscript(expr value, slice slice, expr_context ctx)
    | Starred(expr value, expr_context ctx)
    | Name(identifier id, expr_context ctx)
    | List(expr* elts, expr_context ctx)
    | Tuple(expr* elts, expr_context ctx)

    -- col_offset is the byte offset in the utf8 string the parser uses
    attributes (int lineno, int col_offset)

expr_context = Load | Store | Del | AugLoad | AugStore | Param

slice = Slice(expr? lower, expr? upper, expr? step)
    | ExtSlice(slice* dims)
    | Index(expr value)

boolop = And | Or

operator = Add | Sub | Mult | MatMult | Div | Mod | Pow | LShift
    | RShift | BitOr | BitXor | BitAnd | FloorDiv

unaryop = Invert | Not | UAdd | USub

cmpop = Eq | NotEq | Lt | LtE | Gt | GtE | Is | IsNot | In | NotIn

comprehension = (expr target, expr iter, expr* ifs, int is_async)

except_handler = ExceptHandler(expr? type, identifier? name, stmt* body)
    attributes (int lineno, int col_offset)

arguments = (arg* args, arg? vararg, arg* kwonlyargs, expr* kw_defaults,
    arg? kwarg, expr* defaults)

arg = (identifier arg, expr? annotation)
    attributes (int lineno, int col_offset)

-- keyword arguments supplied to call (NULL identifier for **kwargs)
keyword = (identifier? arg, expr value)

-- import name with optional 'as' alias.
alias = (identifier name, identifier? asname)

withitem = (expr context_expr, expr? optional_vars)
}

```

32.2.3 ast 中的辅助函数

除了节点类，`ast` 模块里为遍历抽象语法树定义了这些工具函数和类：

`ast.parse(source, filename='<unknown>', mode='exec')`
把源码解析为 AST 节点。和 `compile(source, filename, mode, ast.PyCF_ONLY_AST)` 等价。

警告： 足够复杂或是巨大的字符串可能导致 Python 解释器的崩溃，因为 Python 的 AST 编译器是有栈深限制的。

`ast.literal_eval(node_or_string)`

Safely evaluate an expression node or a string containing a Python literal or container display. The string or node provided may only consist of the following Python literal structures: strings, bytes, numbers, tuples, lists, dicts, sets, booleans, and `None`.

This can be used for safely evaluating strings containing Python values from untrusted sources without the need to parse the values oneself. It is not capable of evaluating arbitrarily complex expressions, for example involving operators or indexing.

警告： 足够复杂或是巨大的字符串可能导致 Python 解释器的崩溃，因为 Python 的 AST 编译器是有栈深限制的。

在 3.2 版更改: Now allows bytes and set literals.

`ast.get_docstring(node, clean=True)`

Return the docstring of the given *node* (which must be a `FunctionDef`, `ClassDef` or `Module` node), or `None` if it has no docstring. If *clean* is true, clean up the docstring's indentation with `inspect.cleandoc()`.

`ast.fix_missing_locations(node)`

When you compile a node tree with `compile()`, the compiler expects `lineno` and `col_offset` attributes for every node that supports them. This is rather tedious to fill in for generated nodes, so this helper adds these attributes recursively where not already set, by setting them to the values of the parent node. It works recursively starting at *node*.

`ast.increment_lineno(node, n=1)`

Increment the line number of each node in the tree starting at *node* by *n*. This is useful to “move code” to a different location in a file.

`ast.copy_location(new_node, old_node)`

Copy source location (`lineno` and `col_offset`) from *old_node* to *new_node* if possible, and return *new_node*.

`ast.iter_fields(node)`

Yield a tuple of (*fieldname*, *value*) for each field in `node._fields` that is present on *node*.

`ast.iter_child_nodes(node)`

Yield all direct child nodes of *node*, that is, all fields that are nodes and all items of fields that are lists of nodes.

`ast.walk(node)`

Recursively yield all descendant nodes in the tree starting at *node* (including *node* itself), in no specified order. This is useful if you only want to modify nodes in place and don't care about the context.

class `ast.NodeVisitor`

A node visitor base class that walks the abstract syntax tree and calls a visitor function for every node found. This function may return a value which is forwarded by the `visit()` method.

This class is meant to be subclassed, with the subclass adding visitor methods.

visit (*node*)

Visit a node. The default implementation calls the method called `self.visit_classname` where *classname* is the name of the node class, or `generic_visit()` if that method doesn't exist.

generic_visit (*node*)

This visitor calls `visit()` on all children of the node.

Note that child nodes of nodes that have a custom visitor method won't be visited unless the visitor calls `generic_visit()` or visits them itself.

Don't use the `NodeVisitor` if you want to apply changes to nodes during traversal. For this a special visitor exists (`NodeTransformer`) that allows modifications.

class `ast.NodeTransformer`

A `NodeVisitor` subclass that walks the abstract syntax tree and allows modification of nodes.

The `NodeTransformer` will walk the AST and use the return value of the visitor methods to replace or remove the old node. If the return value of the visitor method is `None`, the node will be removed from its location, otherwise it is replaced with the return value. The return value may be the original node in which case no replacement takes place.

Here is an example transformer that rewrites all occurrences of name lookups (`foo`) to `data['foo']`:

```
class RewriteName(NodeTransformer):

    def visit_Name(self, node):
        return copy_location(Subscript(
            value=Name(id='data', ctx=Load()),
            slice=Index(value=Str(s=node.id)),
            ctx=node.ctx
        ), node)
```

Keep in mind that if the node you're operating on has child nodes you must either transform the child nodes yourself or call the `generic_visit()` method for the node first.

For nodes that were part of a collection of statements (that applies to all statement nodes), the visitor may also return a list of nodes rather than just a single node.

Usually you use the transformer like this:

```
node = YourTransformer().visit(node)
```

`ast.dump` (*node*, *annotate_fields=True*, *include_attributes=False*)

Return a formatted dump of the tree in *node*. This is mainly useful for debugging purposes. The returned string will show the names and the values for fields. This makes the code impossible to evaluate, so if evaluation is wanted *annotate_fields* must be set to `False`. Attributes such as line numbers and column offsets are not dumped by default. If this is wanted, *include_attributes* can be set to `True`.

参见:

[Green Tree Snakes](#), an external documentation resource, has good details on working with Python ASTs.

32.3 `symtable` — Access to the compiler's symbol tables

Source code: [Lib/symtable.py](#)

Symbol tables are generated by the compiler from AST just before bytecode is generated. The symbol table is responsible for calculating the scope of every identifier in the code. `symtable` provides an interface to examine these tables.

32.3.1 Generating Symbol Tables

`symtable.symtable (code, filename, compile_type)`

Return the toplevel *SymbolTable* for the Python source *code*. *filename* is the name of the file containing the code. *compile_type* is like the *mode* argument to `compile()`.

32.3.2 Examining Symbol Tables

class `symtable.SymbolTable`

A namespace table for a block. The constructor is not public.

`get_type()`

Return the type of the symbol table. Possible values are 'class', 'module', and 'function'.

`get_id()`

Return the table's identifier.

`get_name()`

Return the table's name. This is the name of the class if the table is for a class, the name of the function if the table is for a function, or 'top' if the table is global (`get_type()` returns 'module').

`get_lineno()`

Return the number of the first line in the block this table represents.

`is_optimized()`

Return True if the locals in this table can be optimized.

`is_nested()`

Return True if the block is a nested class or function.

`has_children()`

Return True if the block has nested namespaces within it. These can be obtained with `get_children()`.

`has_exec()`

Return True if the block uses `exec`.

`get_identifiers()`

Return a list of names of symbols in this table.

`lookup (name)`

Lookup *name* in the table and return a *Symbol* instance.

`get_symbols()`

Return a list of *Symbol* instances for names in the table.

`get_children()`

Return a list of the nested symbol tables.

class `symtable.Function`

A namespace for a function or method. This class inherits *SymbolTable*.

`get_parameters()`

Return a tuple containing names of parameters to this function.

`get_locals()`

Return a tuple containing names of locals in this function.

`get_globals()`

Return a tuple containing names of globals in this function.

get_frees()
Return a tuple containing names of free variables in this function.

class `symtable.Class`
A namespace of a class. This class inherits `SymbolTable`.

get_methods()
Return a tuple containing the names of methods declared in the class.

class `symtable.Symbol`
An entry in a `SymbolTable` corresponding to an identifier in the source. The constructor is not public.

get_name()
Return the symbol's name.

is_referenced()
Return True if the symbol is used in its block.

is_imported()
Return True if the symbol is created from an import statement.

is_parameter()
Return True if the symbol is a parameter.

is_global()
Return True if the symbol is global.

is_declared_global()
Return True if the symbol is declared global with a global statement.

is_local()
Return True if the symbol is local to its block.

is_free()
Return True if the symbol is referenced in its block, but not assigned to.

is_assigned()
Return True if the symbol is assigned to in its block.

is_namespace()
Return True if name binding introduces new namespace.

If the name is used as the target of a function or class statement, this will be true.

例如:

```
>>> table = symtable.symtable("def some_func(): pass", "string", "exec")
>>> table.lookup("some_func").is_namespace()
True
```

Note that a single name can be bound to multiple objects. If the result is True, the name may also be bound to other objects, like an int or list, that does not introduce a new namespace.

get_namespaces()
Return a list of namespaces bound to this name.

get_namespace()
Return the namespace bound to this name. If more than one namespace is bound, `ValueError` is raised.

32.4 `symbol` —与 Python 解析树一起使用的常量

源代码: [Lib/symbol.py](#)

此模块提供用于表示解析树内部节点数值的常量。与大多数 Python 不同，这些常量使用小写字符名称。请参阅 Python 发行版中的 `Grammar/Grammar` 文件来获取该语言语法上下文中对这些名称的定义。这些名称所映射的特定数值可能会在 Python 版本之间更改。

此模块还提供了额外的数据对象：

`symbol.sym_name`

将此模块中定义的常量的数值映射回名称字符串的字典，允许生成更加人类可读的解析树表示。

32.5 `token` —与 Python 解析树一起使用的常量

源码: [Lib/token.py](#)

此模块提供表示解析树（终端令牌）的叶节点的数值的常量。请参阅 Python 发行版中的文件 `Grammar/Grammar`，以获取语言语法上下文中名称的定义。名称映射到的特定数值可能会在 Python 版本之间更改。

该模块还提供从数字代码到名称和一些函数的映射。这些函数镜像了 Python C 头文件中的定义。

`token.tok_name`

将此模块中定义的常量的数值映射回名称字符串的字典，允许生成更加人类可读的解析树表示。

`token.ISTERMINAL(x)`

对终端标记值返回 `true`。

`token.ISNONTERMINAL(x)`

对非终端标记值返回 `true`。

`token.ISEOF(x)`

如果 `x` 是表示输入结束的标记，则返回 `true`。

标记常量是：

`token.ENDMARKER`

`token.NAME`

`token.NUMBER`

`token.STRING`

`token.NEWLINE`

`token.INDENT`

`token.DEDENT`

`token.LPAR`

`token.RPAR`

`token.LSQB`

`token.RSQB`

`token.COLON`

`token.COMMA`

`token.SEMI`

`token.PLUS`

`token.MINUS`

```
token.STAR
token.SLASH
token.VBAR
token.AMPER
token.LESS
token.GREATER
token.EQUAL
token.DOT
token.PERCENT
token.LBRACE
token.RBRACE
token.EQEQUAL
token.NOTEQUAL
token.LESSEQUAL
token.GREATEREQUAL
token.TILDE
token.CIRCUMFLEX
token.LEFTSHIFT
token.RIGHTSHIFT
token.DOUBLESTAR
token.PLUSEQUAL
token.MINEQUAL
token.STAREQUAL
token.SLASHEQUAL
token.PERCENTEQUAL
token.AMPEREQUAL
token.VBAREQUAL
token.CIRCUMFLEXEQUAL
token.LEFTSHIFTEQUAL
token.RIGHTSHIFTEQUAL
token.DOUBLESTAREQUAL
token.DOUBLES�ASH
token.DOUBLES�ASHEQUAL
token.AT
token.ATEQUAL
token.RARROW
token.ELLIPSIS
token.OP
token.AWAIT
token.ASYNC
token.ERRORTOKEN
token.N_TOKENS
token.NT_OFFSET
```

在 3.5 版更改: Added *AWAIT* and *ASYNC* tokens. Starting with Python 3.7, “async” and “await” will be tokenized as *NAME* tokens, and *AWAIT* and *ASYNC* will be removed.

32.6 keyword — 检验 Python 关键字

源码: [Lib/keyword.py](#)

此模块允许 Python 程序确定字符串是否为关键字。

`keyword.iskeyword(s)`

如果 *s* 是 Python 关键字，则返回 `true`。

`keyword.kwlist`

序列包含为解释器定义的所有关键字。如果任何被定义的关键字为仅在 `__future__` 语句生效是特定时间处于活动状态，则也将包含这些关键字。

32.7 tokenize —Tokenizer for Python source

Source code: [Lib/tokenize.py](#)

The `tokenize` module provides a lexical scanner for Python source code, implemented in Python. The scanner in this module returns comments as tokens as well, making it useful for implementing “pretty-printers,” including colorizers for on-screen displays.

To simplify token stream handling, all operator and delimiter tokens and *Ellipsis* are returned using the generic *OP* token type. The exact type can be determined by checking the `exact_type` property on the *named tuple* returned from `tokenize.tokenize()`.

32.7.1 Tokenizing Input

The primary entry point is a *generator*:

`tokenize.tokenize(readline)`

The `tokenize()` generator requires one argument, *readline*, which must be a callable object which provides the same interface as the `io.IOBase.readline()` method of file objects. Each call to the function should return one line of input as bytes.

The generator produces 5-tuples with these members: the token type; the token string; a 2-tuple (*srow*, *scol*) of ints specifying the row and column where the token begins in the source; a 2-tuple (*erow*, *ecol*) of ints specifying the row and column where the token ends in the source; and the line on which the token was found. The line passed (the last tuple item) is the *logical* line; continuation lines are included. The 5 tuple is returned as a *named tuple* with the field names: `type string start end line`.

The returned *named tuple* has an additional property named `exact_type` that contains the exact operator type for `token.OP` tokens. For all other token types `exact_type` equals the *named tuple* `type` field.

在 3.1 版更改: Added support for named tuples.

在 3.3 版更改: Added support for `exact_type`.

`tokenize()` determines the source encoding of the file by looking for a UTF-8 BOM or encoding cookie, according to [PEP 263](#).

All constants from the `token` module are also exported from `tokenize`, as are three additional token type values:

`tokenize.COMMENT`

Token value used to indicate a comment.

`tokenize.NL`

Token value used to indicate a non-terminating newline. The NEWLINE token indicates the end of a logical line of Python code; NL tokens are generated when a logical line of code is continued over multiple physical lines.

`tokenize.ENCODING`

Token value that indicates the encoding used to decode the source bytes into text. The first token returned by `tokenize()` will always be an ENCODING token.

Another function is provided to reverse the tokenization process. This is useful for creating tools that tokenize a script, modify the token stream, and write back the modified script.

`tokenize.untokenize(iterable)`

Converts tokens back into Python source code. The *iterable* must return sequences with at least two elements, the token type and the token string. Any additional sequence elements are ignored.

The reconstructed script is returned as a single string. The result is guaranteed to tokenize back to match the input so that the conversion is lossless and round-trips are assured. The guarantee applies only to the token type and token string as the spacing between tokens (column positions) may change.

It returns bytes, encoded using the `ENCODING` token, which is the first token sequence output by `tokenize()`.

`tokenize()` needs to detect the encoding of source files it tokenizes. The function it uses to do this is available:

`tokenize.detect_encoding(readline)`

The `detect_encoding()` function is used to detect the encoding that should be used to decode a Python source file. It requires one argument, `readline`, in the same way as the `tokenize()` generator.

It will call `readline` a maximum of twice, and return the encoding used (as a string) and a list of any lines (not decoded from bytes) it has read in.

It detects the encoding from the presence of a UTF-8 BOM or an encoding cookie as specified in [PEP 263](#). If both a BOM and a cookie are present, but disagree, a `SyntaxError` will be raised. Note that if the BOM is found, `'utf-8-sig'` will be returned as an encoding.

If no encoding is specified, then the default of `'utf-8'` will be returned.

Use `open()` to open Python source files: it uses `detect_encoding()` to detect the file encoding.

`tokenize.open(filename)`

Open a file in read only mode using the encoding detected by `detect_encoding()`.

3.2 新版功能.

exception `tokenize.TokenError`

Raised when either a docstring or expression that may be split over several lines is not completed anywhere in the file, for example:

```
"""Beginning of
docstring
```

或者:

```
[1,
 2,
 3
```

Note that unclosed single-quoted strings do not cause an error to be raised. They are tokenized as `ERRORTOKEN`, followed by the tokenization of their contents.

32.7.2 Command-Line Usage

3.3 新版功能.

The `tokenize` module can be executed as a script from the command line. It is as simple as:

```
python -m tokenize [-e] [filename.py]
```

The following options are accepted:

-h, --help
show this help message and exit

-e, --exact
display token names using the exact type

If `filename.py` is specified its contents are tokenized to stdout. Otherwise, tokenization is performed on stdin.

32.7.3 示例

Example of a script rewriter that transforms float literals into Decimal objects:

```
from tokenize import tokenize, untokenize, NUMBER, STRING, NAME, OP
from io import BytesIO

def decistmt(s):
    """Substitute Decimals for floats in a string of statements.

    >>> from decimal import Decimal
    >>> s = 'print(+21.3e-5*-.1234/81.7) '
    >>> decistmt(s)
    "print (+Decimal ('21.3e-5')*-Decimal ('.1234')/Decimal ('81.7'))"

    The format of the exponent is inherited from the platform C library.
    Known cases are "e-007" (Windows) and "e-07" (not Windows). Since
    we're only showing 12 digits, and the 13th isn't close to 5, the
    rest of the output should be platform-independent.

    >>> exec(s) #doctest: +ELLIPSIS
    -3.21716034272e-0...7

    Output from calculations with Decimal should be identical across all
    platforms.

    >>> exec(decistmt(s))
    -3.217160342717258261933904529E-7
    """
    result = []
    g = tokenize(BytesIO(s.encode('utf-8')).readline) # tokenize the string
    for toknum, tokval, _, _, _ in g:
        if toknum == NUMBER and '.' in tokval: # replace NUMBER tokens
            result.extend([
                (NAME, 'Decimal'),
                (OP, '('),
                (STRING, repr(tokval)),
                (OP, ')')
            ])
        else:
            result.append((toknum, tokval))
    return untokenize(result).decode('utf-8')
```

Example of tokenizing from the command line. The script:

```
def say_hello():
    print("Hello, World!")

say_hello()
```

will be tokenized to the following output where the first column is the range of the line/column coordinates where the token is found, the second column is the name of the token, and the final column is the value of the token (if any)

```
$ python -m tokenize hello.py
0,0-0,0:      ENCODING      'utf-8'
1,0-1,3:      NAME          'def'
1,4-1,13:     NAME          'say_hello'
1,13-1,14:    OP            '('
1,14-1,15:    OP            ')'
1,15-1,16:    OP            ':'
1,16-1,17:    NEWLINE      '\n'
2,0-2,4:      INDENT       '    '
2,4-2,9:      NAME          'print'
2,9-2,10:     OP            '('
2,10-2,25:    STRING        '"Hello, World!'"
2,25-2,26:    OP            ')'
2,26-2,27:    NEWLINE      '\n'
3,0-3,1:      NL           '\n'
4,0-4,0:      DEDENT       ''
4,0-4,9:      NAME          'say_hello'
4,9-4,10:     OP            '('
4,10-4,11:    OP            ')'
4,11-4,12:    NEWLINE      '\n'
5,0-5,0:      ENDMARKER    ''
```

The exact token type names can be displayed using the `-e` option:

```
$ python -m tokenize -e hello.py
0,0-0,0:      ENCODING      'utf-8'
1,0-1,3:      NAME          'def'
1,4-1,13:     NAME          'say_hello'
1,13-1,14:    LPAR         '('
1,14-1,15:    RPAR         ')'
1,15-1,16:    COLON        ':'
1,16-1,17:    NEWLINE      '\n'
2,0-2,4:      INDENT       '    '
2,4-2,9:      NAME          'print'
2,9-2,10:     LPAR         '('
2,10-2,25:    STRING        '"Hello, World!'"
2,25-2,26:    RPAR         ')'
2,26-2,27:    NEWLINE      '\n'
3,0-3,1:      NL           '\n'
4,0-4,0:      DEDENT       ''
4,0-4,9:      NAME          'say_hello'
4,9-4,10:     LPAR         '('
4,10-4,11:    RPAR         ')'
4,11-4,12:    NEWLINE      '\n'
5,0-5,0:      ENDMARKER    ''
```

32.8 tabnanny — 模糊缩进检测

源代码: [Lib/tabnanny.py](#)

目前, 该模块旨在作为脚本调用。但是可以使用下面描述的 `check()` 函数将其导入 IDE。

注解： 此模块提供的 API 可能会在将来的版本中更改；此类更改可能无法向后兼容。

`tabnanny.check(file_or_dir)`

If *file_or_dir* is a directory and not a symbolic link, then recursively descend the directory tree named by *file_or_dir*, checking all `.py` files along the way. If *file_or_dir* is an ordinary Python source file, it is checked for whitespace related problems. The diagnostic messages are written to standard output using the `print()` function.

`tabnanny.verbose`

Flag indicating whether to print verbose messages. This is incremented by the `-v` option if called as a script.

`tabnanny.filename_only`

Flag indicating whether to print only the filenames of files containing whitespace related problems. This is set to true by the `-q` option if called as a script.

exception `tabnanny.NannyNag`

Raised by `process_tokens()` if detecting an ambiguous indent. Captured and handled in `check()`.

`tabnanny.process_tokens(tokens)`

This function is used by `check()` to process tokens generated by the `tokenize` module.

参见：

模块 `tokenize` 用于 Python 源代码的词法扫描程序。

32.9 pyc1br —Python class browser support

Source code: [Lib/pyc1br.py](#)

The `pyc1br` module can be used to determine some limited information about the classes, methods and top-level functions defined in a module. The information provided is sufficient to implement a traditional three-pane class browser. The information is extracted from the source code rather than by importing the module, so this module is safe to use with untrusted code. This restriction makes it impossible to use this module with modules not implemented in Python, including all standard and optional extension modules.

`pyc1br.readmodule(module, path=None)`

Read a module and return a dictionary mapping class names to class descriptor objects. The parameter *module* should be the name of a module as a string; it may be the name of a module within a package. The *path* parameter should be a sequence, and is used to augment the value of `sys.path`, which is used to locate module source code.

`pyc1br.readmodule_ex(module, path=None)`

Like `readmodule()`, but the returned dictionary, in addition to mapping class names to class descriptor objects, also maps top-level function names to function descriptor objects. Moreover, if the module being read is a package, the key `'__path__'` in the returned dictionary has as its value a list which contains the package search path.

32.9.1 类对象

The `Class` objects used as values in the dictionary returned by `readmodule()` and `readmodule_ex()` provide the following data attributes:

`Class.module`

The name of the module defining the class described by the class descriptor.

`Class.name`

The name of the class.

Class.super

A list of `Class` objects which describe the immediate base classes of the class being described. Classes which are named as superclasses but which are not discoverable by `readmodule()` are listed as a string with the class name instead of as `Class` objects.

Class.methods

A dictionary mapping method names to line numbers.

Class.file

Name of the file containing the `class` statement defining the class.

Class.lineno

The line number of the `class` statement within the file named by *file*.

32.9.2 函数对象

The `Function` objects used as values in the dictionary returned by `readmodule_ex()` provide the following attributes:

Function.module

The name of the module defining the function described by the function descriptor.

Function.name

The name of the function.

Function.file

Name of the file containing the `def` statement defining the function.

Function.lineno

The line number of the `def` statement within the file named by *file*.

32.10 py_compile —Compile Python source files

Source code: [Lib/py_compile.py](#)

The `py_compile` module provides a function to generate a byte-code file from a source file, and another function used when the module source file is invoked as a script.

Though not often needed, this function can be useful when installing modules for shared use, especially if some of the users may not have permission to write the byte-code cache files in the directory containing the source code.

exception py_compile.PyCompileError

Exception raised when an error occurs while attempting to compile the file.

`py_compile.compile(file, cfile=None, dfile=None, doraise=False, optimize=-1)`

Compile a source file to byte-code and write out the byte-code cache file. The source code is loaded from the file named *file*. The byte-code is written to *cfile*, which defaults to the **PEP 3147/PEP 488** path, ending in `.pyc`. For example, if *file* is `/foo/bar/baz.py` *cfile* will default to `/foo/bar/__pycache__/baz.cpython-32.pyc` for Python 3.2. If *dfile* is specified, it is used as the name of the source file in error messages when instead of *file*. If *doraise* is true, a `PyCompileError` is raised when an error is encountered while compiling *file*. If *doraise* is false (the default), an error string is written to `sys.stderr`, but no exception is raised. This function returns the path to byte-compiled file, i.e. whatever *cfile* value was used.

If the path that *cfile* becomes (either explicitly specified or computed) is a symlink or non-regular file, `FileExistsError` will be raised. This is to act as a warning that import will turn those paths into regular

files if it is allowed to write byte-compiled files to those paths. This is a side-effect of import using file renaming to place the final byte-compiled file into place to prevent concurrent file writing issues.

optimize controls the optimization level and is passed to the built-in `compile()` function. The default of `-1` selects the optimization level of the current interpreter.

在 3.2 版更改: Changed default value of *cfile* to be **PEP 3147**-compliant. Previous default was `file + 'c' ('o' if optimization was enabled)`. Also added the *optimize* parameter.

在 3.4 版更改: Changed code to use `importlib` for the byte-code cache file writing. This means file creation/writing semantics now match what `importlib` does, e.g. permissions, write-and-move semantics, etc. Also added the caveat that `FileExistsError` is raised if *cfile* is a symlink or non-regular file.

`py_compile.main(args=None)`

Compile several source files. The files named in *args* (or on the command line, if *args* is `None`) are compiled and the resulting byte-code is cached in the normal manner. This function does not search a directory structure to locate source files; it only compiles files named explicitly. If `'-'` is the only parameter in *args*, the list of files is taken from standard input.

在 3.2 版更改: Added support for `'-'`.

When this module is run as a script, the `main()` is used to compile all the files named on the command line. The exit status is nonzero if one of the files could not be compiled.

参见:

Module `compileall` Utilities to compile all Python source files in a directory tree.

32.11 compileall —Byte-compile Python libraries

Source code: [Lib/compileall.py](#)

This module provides some utility functions to support installing Python libraries. These functions compile Python source files in a directory tree. This module can be used to create the cached byte-code files at library installation time, which makes them available for use even by users who don't have write permission to the library directories.

32.11.1 Command-line use

This module can work as a script (using `python -m compileall`) to compile Python sources.

directory ...

file ...

Positional arguments are files to compile or directories that contain source files, traversed recursively. If no argument is given, behave as if the command line was `-l <directories from sys.path>`.

-l

Do not recurse into subdirectories, only compile source code files directly contained in the named or implied directories.

-f

Force rebuild even if timestamps are up-to-date.

-q

Do not print the list of files compiled. If passed once, error messages will still be printed. If passed twice (`-qq`), all output is suppressed.

- d** *destdir*
Directory prepended to the path to each file being compiled. This will appear in compilation time tracebacks, and is also compiled in to the byte-code file, where it will be used in tracebacks and other messages in cases where the source file does not exist at the time the byte-code file is executed.
- x** *regex*
regex is used to search the full path to each file considered for compilation, and if the regex produces a match, the file is skipped.
- i** *list*
Read the file *list* and add each line that it contains to the list of files and directories to compile. If *list* is `-`, read lines from `stdin`.
- b**
Write the byte-code files to their legacy locations and names, which may overwrite byte-code files created by another version of Python. The default is to write files to their [PEP 3147](#) locations and names, which allows byte-code files from multiple versions of Python to coexist.
- r**
Control the maximum recursion level for subdirectories. If this is given, then `-l` option will not be taken into account. `python -m compileall <directory> -r 0` is equivalent to `python -m compileall <directory> -l`.
- j** *N*
Use *N* workers to compile the files within the given directory. If 0 is used, then the result of `os.cpu_count()` will be used.

在 3.2 版更改: Added the `-i`, `-b` and `-h` options.

在 3.5 版更改: Added the `-j`, `-r`, and `-qq` options. `-q` option was changed to a multilevel value. `-b` will always produce a byte-code file ending in `.pyc`, never `.pyo`.

There is no command-line option to control the optimization level used by the `compile()` function, because the Python interpreter itself already provides the option: `python -O -m compileall`.

32.11.2 Public functions

`compileall.compile_dir` (*dir*, *maxlevels*=10, *ddir*=None, *force*=False, *rx*=None, *quiet*=0, *legacy*=False, *optimize*=-1, *workers*=1)

Recursively descend the directory tree named by *dir*, compiling all `.py` files along the way. Return a true value if all the files compiled successfully, and a false value otherwise.

The *maxlevels* parameter is used to limit the depth of the recursion; it defaults to 10.

If *ddir* is given, it is prepended to the path to each file being compiled for use in compilation time tracebacks, and is also compiled in to the byte-code file, where it will be used in tracebacks and other messages in cases where the source file does not exist at the time the byte-code file is executed.

If *force* is true, modules are re-compiled even if the timestamps are up to date.

If *rx* is given, its search method is called on the complete path to each file considered for compilation, and if it returns a true value, the file is skipped.

If *quiet* is False or 0 (the default), the filenames and other information are printed to standard out. Set to 1, only errors are printed. Set to 2, all output is suppressed.

If *legacy* is true, byte-code files are written to their legacy locations and names, which may overwrite byte-code files created by another version of Python. The default is to write files to their [PEP 3147](#) locations and names, which allows byte-code files from multiple versions of Python to coexist.

optimize specifies the optimization level for the compiler. It is passed to the built-in `compile()` function.

The argument *workers* specifies how many workers are used to compile files in parallel. The default is to not use multiple workers. If the platform can't use multiple workers and *workers* argument is given, then sequential compilation will be used as a fallback. If *workers* is lower than 0, a `ValueError` will be raised.

在 3.2 版更改: Added the *legacy* and *optimize* parameter.

在 3.5 版更改: Added the *workers* parameter.

在 3.5 版更改: *quiet* parameter was changed to a multilevel value.

在 3.5 版更改: The *legacy* parameter only writes out `.pyc` files, not `.pyo` files no matter what the value of *optimize* is.

在 3.6 版更改: 接受一个类路径对象。

```
compileall.compile_file(fullname, ddir=None, force=False, rx=None, quiet=0, legacy=False,
                        optimize=-1)
```

Compile the file with path *fullname*. Return a true value if the file compiled successfully, and a false value otherwise.

If *ddir* is given, it is prepended to the path to the file being compiled for use in compilation time tracebacks, and is also compiled in to the byte-code file, where it will be used in tracebacks and other messages in cases where the source file does not exist at the time the byte-code file is executed.

If *rx* is given, its search method is passed the full path name to the file being compiled, and if it returns a true value, the file is not compiled and `True` is returned.

If *quiet* is `False` or 0 (the default), the filenames and other information are printed to standard out. Set to 1, only errors are printed. Set to 2, all output is suppressed.

If *legacy* is true, byte-code files are written to their legacy locations and names, which may overwrite byte-code files created by another version of Python. The default is to write files to their [PEP 3147](#) locations and names, which allows byte-code files from multiple versions of Python to coexist.

optimize specifies the optimization level for the compiler. It is passed to the built-in `compile()` function.

3.2 新版功能.

在 3.5 版更改: *quiet* parameter was changed to a multilevel value.

在 3.5 版更改: The *legacy* parameter only writes out `.pyc` files, not `.pyo` files no matter what the value of *optimize* is.

```
compileall.compile_path(skip_curdir=True, maxlevels=0, force=False, quiet=0, legacy=False,
                       optimize=-1)
```

Byte-compile all the `.py` files found along `sys.path`. Return a true value if all the files compiled successfully, and a false value otherwise.

If *skip_curdir* is true (the default), the current directory is not included in the search. All other parameters are passed to the `compile_dir()` function. Note that unlike the other compile functions, *maxlevels* defaults to 0.

在 3.2 版更改: Added the *legacy* and *optimize* parameter.

在 3.5 版更改: *quiet* parameter was changed to a multilevel value.

在 3.5 版更改: The *legacy* parameter only writes out `.pyc` files, not `.pyo` files no matter what the value of *optimize* is.

To force a recompile of all the `.py` files in the `Lib/` subdirectory and all its subdirectories:

```
import compileall

compileall.compile_dir('Lib/', force=True)
```

(下页继续)

(续上页)

```
# Perform same compilation, excluding files in .svn directories.
import re
compileall.compile_dir('Lib/', rx=re.compile(r'[/\\][.]svn'), force=True)

# pathlib.Path objects can also be used.
import pathlib
compileall.compile_dir(pathlib.Path('Lib/'), force=True)
```

参见:

模块 `py_compile` Byte-compile a single source file.

32.12 dis — Python 字节码反汇编器

Source code: [Lib/dis.py](#)

The `dis` module supports the analysis of CPython *bytecode* by disassembling it. The CPython bytecode which this module takes as an input is defined in the file `Include/opcode.h` and used by the compiler and the interpreter.

CPython implementation detail: Bytecode is an implementation detail of the CPython interpreter. No guarantees are made that bytecode will not be added, removed, or changed between versions of Python. Use of this module should not be considered to work across Python VMs or Python releases.

在 3.6 版更改: Use 2 bytes for each instruction. Previously the number of bytes varied by instruction.

Example: Given the function `myfunc()`:

```
def myfunc(alist):
    return len(alist)
```

the following command can be used to display the disassembly of `myfunc()`:

```
>>> dis.dis(myfunc)
2          0 LOAD_GLOBAL              0 (len)
          2 LOAD_FAST                0 (alist)
          4 CALL_FUNCTION             1
          6 RETURN_VALUE
```

(The “2” is a line number).

32.12.1 Bytecode analysis

3.4 新版功能.

The bytecode analysis API allows pieces of Python code to be wrapped in a *Bytecode* object that provides easy access to details of the compiled code.

class `dis.Bytecode` (*x*, *, *first_line=None*, *current_offset=None*)

Analyse the bytecode corresponding to a function, generator, method, string of source code, or a code object (as returned by `compile()`).

This is a convenience wrapper around many of the functions listed below, most notably `get_instructions()`, as iterating over a *Bytecode* instance yields the bytecode operations as *Instruction* instances.

If *first_line* is not `None`, it indicates the line number that should be reported for the first source line in the disassembled code. Otherwise, the source line information (if any) is taken directly from the disassembled code object.

If *current_offset* is not `None`, it refers to an instruction offset in the disassembled code. Setting this means `dis()` will display a “current instruction” marker against the specified opcode.

classmethod from_traceback (tb)

Construct a `Bytecode` instance from the given traceback, setting *current_offset* to the instruction responsible for the exception.

codeobj

The compiled code object.

first_line

The first source line of the code object (if available)

dis()

Return a formatted view of the bytecode operations (the same as printed by `dis.dis()`, but returned as a multi-line string).

info()

Return a formatted multi-line string with detailed information about the code object, like `code_info()`.

示例:

```
>>> bytecode = dis.Bytecode(myfunc)
>>> for instr in bytecode:
...     print(instr.opname)
...
LOAD_GLOBAL
LOAD_FAST
CALL_FUNCTION
RETURN_VALUE
```

32.12.2 Analysis functions

The `dis` module also defines the following analysis functions that convert the input directly to the desired output. They can be useful if only a single operation is being performed, so the intermediate analysis object isn't useful:

dis.code_info(x)

Return a formatted multi-line string with detailed code object information for the supplied function, generator, method, source code string or code object.

Note that the exact contents of code info strings are highly implementation dependent and they may change arbitrarily across Python VMs or Python releases.

3.2 新版功能.

dis.show_code(x, *, file=None)

Print detailed code object information for the supplied function, method, source code string or code object to *file* (or `sys.stdout` if *file* is not specified).

This is a convenient shorthand for `print(code_info(x), file=file)`, intended for interactive exploration at the interpreter prompt.

3.2 新版功能.

在 3.4 版更改: Added *file* parameter.

`dis.dis (x=None, *, file=None)`

Disassemble the *x* object. *x* can denote either a module, a class, a method, a function, a generator, a code object, a string of source code or a byte sequence of raw bytecode. For a module, it disassembles all functions. For a class, it disassembles all methods (including class and static methods). For a code object or sequence of raw bytecode, it prints one line per bytecode instruction. Strings are first compiled to code objects with the `compile()` built-in function before being disassembled. If no object is provided, this function disassembles the last traceback.

The disassembly is written as text to the supplied *file* argument if provided and to `sys.stdout` otherwise.

在 3.4 版更改: Added *file* parameter.

`dis.distb (tb=None, *, file=None)`

Disassemble the top-of-stack function of a traceback, using the last traceback if none was passed. The instruction causing the exception is indicated.

The disassembly is written as text to the supplied *file* argument if provided and to `sys.stdout` otherwise.

在 3.4 版更改: Added *file* parameter.

`dis.disassemble (code, lasti=-1, *, file=None)`

`dis.disco (code, lasti=-1, *, file=None)`

Disassemble a code object, indicating the last instruction if *lasti* was provided. The output is divided in the following columns:

1. the line number, for the first instruction of each line
2. the current instruction, indicated as `-->`,
3. a labelled instruction, indicated with `>>`,
4. the address of the instruction,
5. the operation code name,
6. operation parameters, and
7. interpretation of the parameters in parentheses.

The parameter interpretation recognizes local and global variable names, constant values, branch targets, and compare operators.

The disassembly is written as text to the supplied *file* argument if provided and to `sys.stdout` otherwise.

在 3.4 版更改: Added *file* parameter.

`dis.get_instructions (x, *, first_line=None)`

Return an iterator over the instructions in the supplied function, method, source code string or code object.

The iterator generates a series of `Instruction` named tuples giving the details of each operation in the supplied code.

If *first_line* is not `None`, it indicates the line number that should be reported for the first source line in the disassembled code. Otherwise, the source line information (if any) is taken directly from the disassembled code object.

3.4 新版功能.

`dis.findlinestarts (code)`

This generator function uses the `co_firstlineno` and `co_lnotab` attributes of the code object *code* to find the offsets which are starts of lines in the source code. They are generated as (`offset`, `lineno`) pairs. See [Objects/lnotab_notes.txt](#) for the `co_lnotab` format and how to decode it.

在 3.6 版更改: Line numbers can be decreasing. Before, they were always increasing.

`dis.findlabels(code)`
 Detect all offsets in the code object *code* which are jump targets, and return a list of these offsets.

`dis.stack_effect(opcode[, oparg])`
 Compute the stack effect of *opcode* with argument *oparg*.

3.4 新版功能.

32.12.3 Python Bytecode Instructions

The `get_instructions()` function and `Bytecode` class provide details of bytecode instructions as `Instruction` instances:

```
class dis.Instruction
    Details for a bytecode operation

    opcode
        numeric code for operation, corresponding to the opcode values listed below and the bytecode values in the
        Opcode collections.

    opname
        human readable name for operation

    arg
        numeric argument to operation (if any), otherwise None

    argval
        resolved arg value (if known), otherwise same as arg

    argrepr
        human readable description of operation argument

    offset
        start index of operation within bytecode sequence

    starts_line
        line started by this opcode (if any), otherwise None

    is_jump_target
        True if other code jumps to here, otherwise False

    3.4 新版功能.
```

The Python compiler currently generates the following bytecode instructions.

General instructions

NOP
 Do nothing code. Used as a placeholder by the bytecode optimizer.

POP_TOP
 Removes the top-of-stack (TOS) item.

ROT_TWO
 Swaps the two top-most stack items.

ROT_THREE
 Lifts second and third stack item one position up, moves top down to position three.

DUP_TOP
 Duplicates the reference on top of the stack.

3.2 新版功能.

DUP_TOP_TWO

Duplicates the two references on top of the stack, leaving them in the same order.

3.2 新版功能.

Unary operations

Unary operations take the top of the stack, apply the operation, and push the result back on the stack.

UNARY_POSITIVE

Implements `TOS = +TOS`.

UNARY_NEGATIVE

Implements `TOS = -TOS`.

UNARY_NOT

Implements `TOS = not TOS`.

UNARY_INVERT

Implements `TOS = ~TOS`.

GET_ITER

Implements `TOS = iter(TOS)`.

GET_YIELD_FROM_ITER

If `TOS` is a *generator iterator* or *coroutine* object it is left as is. Otherwise, implements `TOS = iter(TOS)`.

3.5 新版功能.

Binary operations

Binary operations remove the top of the stack (`TOS`) and the second top-most stack item (`TOS1`) from the stack. They perform the operation, and put the result back on the stack.

BINARY_POWER

Implements `TOS = TOS1 ** TOS`.

BINARY_MULTIPLY

Implements `TOS = TOS1 * TOS`.

BINARY_MATRIX_MULTIPLY

Implements `TOS = TOS1 @ TOS`.

3.5 新版功能.

BINARY_FLOOR_DIVIDE

Implements `TOS = TOS1 // TOS`.

BINARY_TRUE_DIVIDE

Implements `TOS = TOS1 / TOS`.

BINARY_MODULO

Implements `TOS = TOS1 % TOS`.

BINARY_ADD

Implements `TOS = TOS1 + TOS`.

BINARY_SUBTRACT

Implements `TOS = TOS1 - TOS`.

BINARY_SUBSCR

Implements `TOS = TOS1[TOS]`.

BINARY_LSHIFT

Implements `TOS = TOS1 << TOS`.

BINARY_RSHIFT

Implements $TOS = TOS1 \gg TOS$.

BINARY_AND

Implements $TOS = TOS1 \& TOS$.

BINARY_XOR

Implements $TOS = TOS1 \wedge TOS$.

BINARY_OR

Implements $TOS = TOS1 | TOS$.

In-place operations

In-place operations are like binary operations, in that they remove TOS and TOS1, and push the result back on the stack, but the operation is done in-place when TOS1 supports it, and the resulting TOS may be (but does not have to be) the original TOS1.

INPLACE_POWER

Implements in-place $TOS = TOS1 ** TOS$.

INPLACE_MULTIPLY

Implements in-place $TOS = TOS1 * TOS$.

INPLACE_MATRIX_MULTIPLY

Implements in-place $TOS = TOS1 @ TOS$.

3.5 新版功能.

INPLACE_FLOOR_DIVIDE

Implements in-place $TOS = TOS1 // TOS$.

INPLACE_TRUE_DIVIDE

Implements in-place $TOS = TOS1 / TOS$.

INPLACE_MODULO

Implements in-place $TOS = TOS1 \% TOS$.

INPLACE_ADD

Implements in-place $TOS = TOS1 + TOS$.

INPLACE_SUBTRACT

Implements in-place $TOS = TOS1 - TOS$.

INPLACE_LSHIFT

Implements in-place $TOS = TOS1 \ll TOS$.

INPLACE_RSHIFT

Implements in-place $TOS = TOS1 \gg TOS$.

INPLACE_AND

Implements in-place $TOS = TOS1 \& TOS$.

INPLACE_XOR

Implements in-place $TOS = TOS1 \wedge TOS$.

INPLACE_OR

Implements in-place $TOS = TOS1 | TOS$.

STORE_SUBSCR

Implements $TOS1[TOS] = TOS2$.

DELETE_SUBSCR

Implements $\text{del } TOS1[TOS]$.

Coroutine opcodes

GET_AWAITABLE

Implements `TOS = get_awaitable(TOS)`, where `get_awaitable(o)` returns `o` if `o` is a coroutine object or a generator object with the `CO_ITERABLE_COROUTINE` flag, or resolves `o.__await__`.

3.5 新版功能.

GET_AITER

Implements `TOS = get_awaitable(TOS.__aiter__())`. See `GET_AWAITABLE` for details about `get_awaitable`.

3.5 新版功能.

GET_ANEXT

Implements `PUSH(get_awaitable(TOS.__anext__()))`. See `GET_AWAITABLE` for details about `get_awaitable`.

3.5 新版功能.

BEFORE_ASYNC_WITH

Resolves `__aenter__` and `__aexit__` from the object on top of the stack. Pushes `__aexit__` and result of `__aenter__()` to the stack.

3.5 新版功能.

SETUP_ASYNC_WITH

Creates a new frame object.

3.5 新版功能.

Miscellaneous opcodes

PRINT_EXPR

Implements the expression statement for the interactive mode. `TOS` is removed from the stack and printed. In non-interactive mode, an expression statement is terminated with `POP_TOP`.

BREAK_LOOP

Terminates a loop due to a `break` statement.

CONTINUE_LOOP (*target*)

Continues a loop due to a `continue` statement. *target* is the address to jump to (which should be a `FOR_ITER` instruction).

SET_ADD (*i*)

Calls `set.add(TOS1[-i], TOS)`. Used to implement set comprehensions.

LIST_APPEND (*i*)

Calls `list.append(TOS[-i], TOS)`. Used to implement list comprehensions.

MAP_ADD (*i*)

Calls `dict.setdefault(TOS1[-i], TOS, TOS1)`. Used to implement dict comprehensions.

3.1 新版功能.

For all of the `SET_ADD`, `LIST_APPEND` and `MAP_ADD` instructions, while the added value or key/value pair is popped off, the container object remains on the stack so that it is available for further iterations of the loop.

RETURN_VALUE

Returns with `TOS` to the caller of the function.

YIELD_VALUE

Pops `TOS` and yields it from a *generator*.

YIELD_FROM

Pops TOS and delegates to it as a subiterator from a *generator*.

3.3 新版功能.

SETUP_ANNOTATIONS

Checks whether `__annotations__` is defined in `locals()`, if not it is set up to an empty dict. This opcode is only emitted if a class or module body contains *variable annotations* statically.

3.6 新版功能.

IMPORT_STAR

Loads all symbols not starting with `'_'` directly from the module TOS to the local namespace. The module is popped after loading all names. This opcode implements `from module import *`.

POP_BLOCK

Removes one block from the block stack. Per frame, there is a stack of blocks, denoting nested loops, try statements, and such.

POP_EXCEPT

Removes one block from the block stack. The popped block must be an exception handler block, as implicitly created when entering an except handler. In addition to popping extraneous values from the frame stack, the last three popped values are used to restore the exception state.

END_FINALLY

Terminates a *finally* clause. The interpreter recalls whether the exception has to be re-raised, or whether the function returns, and continues with the outer-next block.

LOAD_BUILD_CLASS

Pushes `builtins.__build_class__()` onto the stack. It is later called by *CALL_FUNCTION* to construct a class.

SETUP_WITH (*delta*)

This opcode performs several operations before a *with* block starts. First, it loads `__exit__()` from the context manager and pushes it onto the stack for later use by *WITH_CLEANUP*. Then, `__enter__()` is called, and a finally block pointing to *delta* is pushed. Finally, the result of calling the enter method is pushed onto the stack. The next opcode will either ignore it (*POP_TOP*), or store it in (a) variable(s) (*STORE_FAST*, *STORE_NAME*, or *UNPACK_SEQUENCE*).

3.2 新版功能.

WITH_CLEANUP_START

Cleans up the stack when a *with* statement block exits. TOS is the context manager's `__exit__()` bound method. Below TOS are 1–3 values indicating how/why the finally clause was entered:

- `SECOND = None`
- `(SECOND, THIRD) = (WHY_{RETURN, CONTINUE}), retval`
- `SECOND = WHY_*`; no `retval` below it
- `(SECOND, THIRD, FOURTH) = exc_info()`

In the last case, `TOS(SECOND, THIRD, FOURTH)` is called, otherwise `TOS(None, None, None)`. Pushes `SECOND` and result of the call to the stack.

WITH_CLEANUP_FINISH

Pops exception type and result of `'exit'` function call from the stack.

If the stack represents an exception, *and* the function call returns a `'true'` value, this information is “zapped” and replaced with a single `WHY_SILENCED` to prevent *END_FINALLY* from re-raising the exception. (But non-local *gotos* will still be resumed.)

All of the following opcodes use their arguments.

STORE_NAME (*namei*)

Implements `name = TOS`. *namei* is the index of *name* in the attribute `co_names` of the code object. The compiler tries to use *STORE_FAST* or *STORE_GLOBAL* if possible.

DELETE_NAME (*namei*)

Implements `del name`, where *namei* is the index into `co_names` attribute of the code object.

UNPACK_SEQUENCE (*count*)

Unpacks TOS into *count* individual values, which are put onto the stack right-to-left.

UNPACK_EX (*counts*)

Implements assignment with a starred target: Unpacks an iterable in TOS into individual values, where the total number of values can be smaller than the number of items in the iterable: one of the new values will be a list of all leftover items.

The low byte of *counts* is the number of values before the list value, the high byte of *counts* the number of values after it. The resulting values are put onto the stack right-to-left.

STORE_ATTR (*namei*)

Implements `TOS.name = TOS1`, where *namei* is the index of *name* in `co_names`.

DELETE_ATTR (*namei*)

Implements `del TOS.name`, using *namei* as index into `co_names`.

STORE_GLOBAL (*namei*)

Works as *STORE_NAME*, but stores the name as a global.

DELETE_GLOBAL (*namei*)

Works as *DELETE_NAME*, but deletes a global name.

LOAD_CONST (*consti*)

Pushes `co_consts[consti]` onto the stack.

LOAD_NAME (*namei*)

Pushes the value associated with `co_names[namei]` onto the stack.

BUILD_TUPLE (*count*)

Creates a tuple consuming *count* items from the stack, and pushes the resulting tuple onto the stack.

BUILD_LIST (*count*)

Works as *BUILD_TUPLE*, but creates a list.

BUILD_SET (*count*)

Works as *BUILD_TUPLE*, but creates a set.

BUILD_MAP (*count*)

Pushes a new dictionary object onto the stack. Pops $2 * count$ items so that the dictionary holds *count* entries: `{..., TOS3: TOS2, TOS1: TOS}`.

在 3.5 版更改: The dictionary is created from stack items instead of creating an empty dictionary pre-sized to hold *count* items.

BUILD_CONST_KEY_MAP (*count*)

The version of *BUILD_MAP* specialized for constant keys. *count* values are consumed from the stack. The top element on the stack contains a tuple of keys.

3.6 新版功能.

BUILD_STRING (*count*)

Concatenates *count* strings from the stack and pushes the resulting string onto the stack.

3.6 新版功能.

BUILD_TUPLE_UNPACK (*count*)

Pops *count* iterables from the stack, joins them in a single tuple, and pushes the result. Implements iterable unpacking in tuple displays `(*x, *y, *z)`.

3.5 新版功能.

BUILD_TUPLE_UNPACK_WITH_CALL (*count*)

This is similar to [BUILD_TUPLE_UNPACK](#), but is used for `f(*x, *y, *z)` call syntax. The stack item at position *count* + 1 should be the corresponding callable *f*.

3.6 新版功能.

BUILD_LIST_UNPACK (*count*)

This is similar to [BUILD_TUPLE_UNPACK](#), but pushes a list instead of tuple. Implements iterable unpacking in list displays `[*x, *y, *z]`.

3.5 新版功能.

BUILD_SET_UNPACK (*count*)

This is similar to [BUILD_TUPLE_UNPACK](#), but pushes a set instead of tuple. Implements iterable unpacking in set displays `{*x, *y, *z}`.

3.5 新版功能.

BUILD_MAP_UNPACK (*count*)

Pops *count* mappings from the stack, merges them into a single dictionary, and pushes the result. Implements dictionary unpacking in dictionary displays `{**x, **y, **z}`.

3.5 新版功能.

BUILD_MAP_UNPACK_WITH_CALL (*count*)

This is similar to [BUILD_MAP_UNPACK](#), but is used for `f(**x, **y, **z)` call syntax. The stack item at position *count* + 2 should be the corresponding callable *f*.

3.5 新版功能.

在 3.6 版更改: The position of the callable is determined by adding 2 to the opcode argument instead of encoding it in the second byte of the argument.

LOAD_ATTR (*namei*)

Replaces TOS with `getattr(TOS, co_names[namei])`.

COMPARE_OP (*opname*)

Performs a Boolean operation. The operation name can be found in `cmp_op[opname]`.

IMPORT_NAME (*namei*)

Imports the module `co_names[namei]`. TOS and TOS1 are popped and provide the *fromlist* and *level* arguments of `__import__()`. The module object is pushed onto the stack. The current namespace is not affected: for a proper import statement, a subsequent [STORE_FAST](#) instruction modifies the namespace.

IMPORT_FROM (*namei*)

Loads the attribute `co_names[namei]` from the module found in TOS. The resulting object is pushed onto the stack, to be subsequently stored by a [STORE_FAST](#) instruction.

JUMP_FORWARD (*delta*)

Increments bytecode counter by *delta*.

POP_JUMP_IF_TRUE (*target*)

If TOS is true, sets the bytecode counter to *target*. TOS is popped.

3.1 新版功能.

POP_JUMP_IF_FALSE (*target*)

If TOS is false, sets the bytecode counter to *target*. TOS is popped.

3.1 新版功能.

JUMP_IF_TRUE_OR_POP (*target*)

If TOS is true, sets the bytecode counter to *target* and leaves TOS on the stack. Otherwise (TOS is false), TOS is popped.

3.1 新版功能.

JUMP_IF_FALSE_OR_POP (*target*)

If TOS is false, sets the bytecode counter to *target* and leaves TOS on the stack. Otherwise (TOS is true), TOS is popped.

3.1 新版功能.

JUMP_ABSOLUTE (*target*)

Set bytecode counter to *target*.

FOR_ITER (*delta*)

TOS is an *iterator*. Call its `__next__()` method. If this yields a new value, push it on the stack (leaving the iterator below it). If the iterator indicates it is exhausted TOS is popped, and the byte code counter is incremented by *delta*.

LOAD_GLOBAL (*namei*)

Loads the global named `co_names[namei]` onto the stack.

SETUP_LOOP (*delta*)

Pushes a block for a loop onto the block stack. The block spans from the current instruction with a size of *delta* bytes.

SETUP_EXCEPT (*delta*)

Pushes a try block from a try-except clause onto the block stack. *delta* points to the first except block.

SETUP_FINALLY (*delta*)

Pushes a try block from a try-except clause onto the block stack. *delta* points to the finally block.

LOAD_FAST (*var_num*)

Pushes a reference to the local `co_varnames[var_num]` onto the stack.

STORE_FAST (*var_num*)

Stores TOS into the local `co_varnames[var_num]`.

DELETE_FAST (*var_num*)

Deletes local `co_varnames[var_num]`.

STORE_ANNOTATION (*namei*)

Stores TOS as `locals()['__annotations__'][co_names[namei]] = TOS`.

3.6 新版功能.

LOAD_CLOSURE (*i*)

Pushes a reference to the cell contained in slot *i* of the cell and free variable storage. The name of the variable is `co_cellvars[i]` if *i* is less than the length of `co_cellvars`. Otherwise it is `co_freevars[i - len(co_cellvars)]`.

LOAD_DEREF (*i*)

Loads the cell contained in slot *i* of the cell and free variable storage. Pushes a reference to the object the cell contains on the stack.

LOAD_CLASSDEREF (*i*)

Much like `LOAD_DEREF` but first checks the locals dictionary before consulting the cell. This is used for loading free variables in class bodies.

3.4 新版功能.

STORE_DEREF (*i*)

Stores TOS into the cell contained in slot *i* of the cell and free variable storage.

DELETE_DEREF (*i*)

Empties the cell contained in slot *i* of the cell and free variable storage. Used by the `del` statement.

3.2 新版功能.

RAISE_VARARGS (*argc*)

Raises an exception. *argc* indicates the number of arguments to the raise statement, ranging from 0 to 3. The handler will find the traceback as TOS2, the parameter as TOS1, and the exception as TOS.

CALL_FUNCTION (*argc*)

Calls a callable object with positional arguments. *argc* indicates the number of positional arguments. The top of the stack contains positional arguments, with the right-most argument on top. Below the arguments is a callable object to call. **CALL_FUNCTION** pops all arguments and the callable object off the stack, calls the callable object with those arguments, and pushes the return value returned by the callable object.

在 3.6 版更改: This opcode is used only for calls with positional arguments.

CALL_FUNCTION_KW (*argc*)

Calls a callable object with positional (if any) and keyword arguments. *argc* indicates the total number of positional and keyword arguments. The top element on the stack contains a tuple of keyword argument names. Below that are keyword arguments in the order corresponding to the tuple. Below that are positional arguments, with the right-most parameter on top. Below the arguments is a callable object to call. **CALL_FUNCTION_KW** pops all arguments and the callable object off the stack, calls the callable object with those arguments, and pushes the return value returned by the callable object.

在 3.6 版更改: Keyword arguments are packed in a tuple instead of a dictionary, *argc* indicates the total number of arguments.

CALL_FUNCTION_EX (*flags*)

Calls a callable object with variable set of positional and keyword arguments. If the lowest bit of *flags* is set, the top of the stack contains a mapping object containing additional keyword arguments. Below that is an iterable object containing positional arguments and a callable object to call. [BUILD_MAP_UNPACK_WITH_CALL](#) and [BUILD_TUPLE_UNPACK_WITH_CALL](#) can be used for merging multiple mapping objects and iterables containing arguments. Before the callable is called, the mapping object and iterable object are each “unpacked” and their contents passed in as keyword and positional arguments respectively. **CALL_FUNCTION_EX** pops all arguments and the callable object off the stack, calls the callable object with those arguments, and pushes the return value returned by the callable object.

3.6 新版功能.

MAKE_FUNCTION (*argc*)

Pushes a new function object on the stack. From bottom to top, the consumed stack must consist of values if the argument carries a specified flag value

- 0x01 a tuple of default values for positional-only and positional-or-keyword parameters in positional order
- 0x02 a dictionary of keyword-only parameters’ default values
- 0x04 an annotation dictionary
- 0x08 a tuple containing cells for free variables, making a closure
- the code associated with the function (at TOS1)
- the *qualified name* of the function (at TOS)

BUILD_SLICE (*argc*)

Pushes a slice object on the stack. *argc* must be 2 or 3. If it is 2, `slice(TOS1, TOS)` is pushed; if it is 3, `slice(TOS2, TOS1, TOS)` is pushed. See the [slice\(\)](#) built-in function for more information.

EXTENDED_ARG (*ext*)

Prefixes any opcode which has an argument too big to fit into the default two bytes. *ext* holds two additional bytes which, taken together with the subsequent opcode's argument, comprise a four-byte argument, *ext* being the two most-significant bytes.

FORMAT_VALUE (*flags*)

Used for implementing formatted literal strings (f-strings). Pops an optional *fmt_spec* from the stack, then a required *value*. *flags* is interpreted as follows:

- (flags & 0x03) == 0x00: *value* is formatted as-is.
- (flags & 0x03) == 0x01: call `str()` on *value* before formatting it.
- (flags & 0x03) == 0x02: call `repr()` on *value* before formatting it.
- (flags & 0x03) == 0x03: call `ascii()` on *value* before formatting it.
- (flags & 0x04) == 0x04: pop *fmt_spec* from the stack and use it, else use an empty *fmt_spec*.

Formatting is performed using `PyObject_Format()`. The result is pushed on the stack.

3.6 新版功能.

HAVE_ARGUMENT

This is not really an opcode. It identifies the dividing line between opcodes which don't use their argument and those that do (< HAVE_ARGUMENT and >= HAVE_ARGUMENT, respectively).

在 3.6 版更改: Now every instruction has an argument, but opcodes < HAVE_ARGUMENT ignore it. Before, only opcodes >= HAVE_ARGUMENT had an argument.

32.12.4 Opcode collections

These collections are provided for automatic introspection of bytecode instructions:

dis.opname

Sequence of operation names, indexable using the bytecode.

dis.opmap

Dictionary mapping operation names to bytecodes.

dis.cmp_op

Sequence of all compare operation names.

dis.hasconst

Sequence of bytecodes that access a constant.

dis.hasfree

Sequence of bytecodes that access a free variable (note that 'free' in this context refers to names in the current scope that are referenced by inner scopes or names in outer scopes that are referenced from this scope. It does *not* include references to global or builtin scopes).

dis.hasname

Sequence of bytecodes that access an attribute by name.

dis.hasjrel

Sequence of bytecodes that have a relative jump target.

dis.hasjabs

Sequence of bytecodes that have an absolute jump target.

dis.haslocal

Sequence of bytecodes that access a local variable.

`dis.hascompare`

Sequence of bytecodes of Boolean operations.

32.13 `pickletools` —Tools for pickle developers

Source code: [Lib/pickletools.py](#)

This module contains various constants relating to the intimate details of the `pickle` module, some lengthy comments about the implementation, and a few useful functions for analyzing pickled data. The contents of this module are useful for Python core developers who are working on the `pickle`; ordinary users of the `pickle` module probably won't find the `pickletools` module relevant.

32.13.1 Command line usage

3.2 新版功能.

When invoked from the command line, `python -m pickletools` will disassemble the contents of one or more pickle files. Note that if you want to see the Python object stored in the pickle rather than the details of pickle format, you may want to use `-m pickle` instead. However, when the pickle file that you want to examine comes from an untrusted source, `-m pickletools` is a safer option because it does not execute pickle bytecode.

For example, with a tuple `(1, 2)` pickled in file `x.pickle`:

```
$ python -m pickle x.pickle
(1, 2)

$ python -m pickletools x.pickle
0: \x80 PROTO      3
2: K    BININT1    1
4: K    BININT1    2
6: \x86 TUPLE2
7: q    BINPUT     0
9: .    STOP
highest protocol among opcodes = 2
```

Command line options

-a, --annotate

Annotate each line with a short opcode description.

-o, --output=<file>

Name of a file where the output should be written.

-l, --indentlevel=<num>

The number of blanks by which to indent a new MARK level.

-m, --memo

When multiple objects are disassembled, preserve memo between disassemblies.

-p, --preamble=<preamble>

When more than one pickle file are specified, print given preamble before each disassembly.

32.13.2 Programmatic Interface

`pickletools.dis` (*pickle*, *out=None*, *memo=None*, *indentlevel=4*, *annotate=0*)

Outputs a symbolic disassembly of the pickle to the file-like object *out*, defaulting to `sys.stdout`. *pickle* can be a string or a file-like object. *memo* can be a Python dictionary that will be used as the pickle's memo; it can be used to perform disassemblies across multiple pickles created by the same pickler. Successive levels, indicated by MARK opcodes in the stream, are indented by *indentlevel* spaces. If a nonzero value is given to *annotate*, each opcode in the output is annotated with a short description. The value of *annotate* is used as a hint for the column where annotation should start.

3.2 新版功能: The *annotate* argument.

`pickletools.genops` (*pickle*)

Provides an *iterator* over all of the opcodes in a pickle, returning a sequence of (*opcode*, *arg*, *pos*) triples. *opcode* is an instance of an `OpcodeInfo` class; *arg* is the decoded value, as a Python object, of the opcode's argument; *pos* is the position at which this opcode is located. *pickle* can be a string or a file-like object.

`pickletools.optimize` (*picklestring*)

Returns a new equivalent pickle string after eliminating unused PUT opcodes. The optimized pickle is shorter, takes less transmission time, requires less storage space, and unpickles more efficiently.

本章中介绍的模块提供了所有 Python 版本中提供的各种杂项服务。这是一个概述：

33.1 `formatter` —Generic output formatting

3.4 版后已移除: Due to lack of usage, the `formatter` module has been deprecated.

This module supports two interface definitions, each with multiple implementations: The *formatter* interface, and the *writer* interface which is required by the *formatter* interface.

Formatter objects transform an abstract flow of formatting events into specific output events on writer objects. Formatters manage several stack structures to allow various properties of a writer object to be changed and restored; writers need not be able to handle relative changes nor any sort of “change back” operation. Specific writer properties which may be controlled via formatter objects are horizontal alignment, font, and left margin indentations. A mechanism is provided which supports providing arbitrary, non-exclusive style settings to a writer as well. Additional interfaces facilitate formatting events which are not reversible, such as paragraph separation.

Writer objects encapsulate device interfaces. Abstract devices, such as file formats, are supported as well as physical devices. The provided implementations all work with abstract devices. The interface makes available mechanisms for setting the properties which formatter objects manage and inserting data into the output.

33.1.1 The Formatter Interface

Interfaces to create formatters are dependent on the specific formatter class being instantiated. The interfaces described below are the required interfaces which all formatters must support once initialized.

One data element is defined at the module level:

`formatter.AS_IS`

Value which can be used in the font specification passed to the `push_font()` method described below, or as the new value to any other `push_property()` method. Pushing the `AS_IS` value allows the corresponding `pop_property()` method to be called without having to track whether the property was changed.

The following attributes are defined for formatter instance objects:

`formatter.writer`

The writer instance with which the formatter interacts.

`formatter.end_paragraph (blanklines)`

Close any open paragraphs and insert at least *blanklines* before the next paragraph.

`formatter.add_line_break ()`

Add a hard line break if one does not already exist. This does not break the logical paragraph.

`formatter.add_hor_rule (*args, **kw)`

Insert a horizontal rule in the output. A hard break is inserted if there is data in the current paragraph, but the logical paragraph is not broken. The arguments and keywords are passed on to the writer's `send_line_break()` method.

`formatter.add_flow_data (data)`

Provide data which should be formatted with collapsed whitespace. Whitespace from preceding and successive calls to `add_flow_data()` is considered as well when the whitespace collapse is performed. The data which is passed to this method is expected to be word-wrapped by the output device. Note that any word-wrapping still must be performed by the writer object due to the need to rely on device and font information.

`formatter.add_literal_data (data)`

Provide data which should be passed to the writer unchanged. Whitespace, including newline and tab characters, are considered legal in the value of *data*.

`formatter.add_label_data (format, counter)`

Insert a label which should be placed to the left of the current left margin. This should be used for constructing bulleted or numbered lists. If the *format* value is a string, it is interpreted as a format specification for *counter*, which should be an integer. The result of this formatting becomes the value of the label; if *format* is not a string it is used as the label value directly. The label value is passed as the only argument to the writer's `send_label_data()` method. Interpretation of non-string label values is dependent on the associated writer.

Format specifications are strings which, in combination with a counter value, are used to compute label values. Each character in the format string is copied to the label value, with some characters recognized to indicate a transform on the counter value. Specifically, the character '1' represents the counter value formatter as an Arabic number, the characters 'A' and 'a' represent alphabetic representations of the counter value in upper and lower case, respectively, and 'I' and 'i' represent the counter value in Roman numerals, in upper and lower case. Note that the alphabetic and roman transforms require that the counter value be greater than zero.

`formatter.flush_softspace ()`

Send any pending whitespace buffered from a previous call to `add_flow_data()` to the associated writer object. This should be called before any direct manipulation of the writer object.

`formatter.push_alignment (align)`

Push a new alignment setting onto the alignment stack. This may be `AS_IS` if no change is desired. If the alignment value is changed from the previous setting, the writer's `new_alignment()` method is called with the *align* value.

`formatter.pop_alignment ()`

Restore the previous alignment.

`formatter.push_font ((size, italic, bold, teletype))`

Change some or all font properties of the writer object. Properties which are not set to `AS_IS` are set to the values passed in while others are maintained at their current settings. The writer's `new_font()` method is called with the fully resolved font specification.

`formatter.pop_font ()`

Restore the previous font.

`formatter.push_margin(margin)`

Increase the number of left margin indentations by one, associating the logical tag *margin* with the new indentation. The initial margin level is 0. Changed values of the logical tag must be true values; false values other than `AS_IS` are not sufficient to change the margin.

`formatter.pop_margin()`

Restore the previous margin.

`formatter.push_style(*styles)`

Push any number of arbitrary style specifications. All styles are pushed onto the styles stack in order. A tuple representing the entire stack, including `AS_IS` values, is passed to the writer's `new_styles()` method.

`formatter.pop_style(n=1)`

Pop the last *n* style specifications passed to `push_style()`. A tuple representing the revised stack, including `AS_IS` values, is passed to the writer's `new_styles()` method.

`formatter.set_spacing(spacing)`

Set the spacing style for the writer.

`formatter.assert_line_data(flag=1)`

Inform the formatter that data has been added to the current paragraph out-of-band. This should be used when the writer has been manipulated directly. The optional *flag* argument can be set to false if the writer manipulations produced a hard line break at the end of the output.

33.1.2 Formatter Implementations

Two implementations of formatter objects are provided by this module. Most applications may use one of these classes without modification or subclassing.

class `formatter.NullFormatter(writer=None)`

A formatter which does nothing. If *writer* is omitted, a `NullWriter` instance is created. No methods of the writer are called by `NullFormatter` instances. Implementations should inherit from this class if implementing a writer interface but don't need to inherit any implementation.

class `formatter.AbstractFormatter(writer)`

The standard formatter. This implementation has demonstrated wide applicability to many writers, and may be used directly in most circumstances. It has been used to implement a full-featured World Wide Web browser.

33.1.3 The Writer Interface

Interfaces to create writers are dependent on the specific writer class being instantiated. The interfaces described below are the required interfaces which all writers must support once initialized. Note that while most applications can use the `AbstractFormatter` class as a formatter, the writer must typically be provided by the application.

`writer.flush()`

Flush any buffered output or device control events.

`writer.new_alignment(align)`

Set the alignment style. The *align* value can be any object, but by convention is a string or None, where None indicates that the writer's "preferred" alignment should be used. Conventional *align* values are 'left', 'center', 'right', and 'justify'.

`writer.new_font(font)`

Set the font style. The value of *font* will be None, indicating that the device's default font should be used, or a tuple of the form (size, italic, bold, teletype). Size will be a string indicating the size of font that should be used; specific strings and their interpretation must be defined by the application. The *italic*, *bold*, and *teletype* values are Boolean values specifying which of those font attributes should be used.

`writer.new_margin(margin, level)`

Set the margin level to the integer *level* and the logical tag to *margin*. Interpretation of the logical tag is at the writer's discretion; the only restriction on the value of the logical tag is that it not be a false value for non-zero values of *level*.

`writer.new_spacing(spacing)`

Set the spacing style to *spacing*.

`writer.new_styles(styles)`

Set additional styles. The *styles* value is a tuple of arbitrary values; the value `AS_IS` should be ignored. The *styles* tuple may be interpreted either as a set or as a stack depending on the requirements of the application and writer implementation.

`writer.send_line_break()`

Break the current line.

`writer.send_paragraph(blankline)`

Produce a paragraph separation of at least *blankline* blank lines, or the equivalent. The *blankline* value will be an integer. Note that the implementation will receive a call to `send_line_break()` before this call if a line break is needed; this method should not include ending the last line of the paragraph. It is only responsible for vertical spacing between paragraphs.

`writer.send_hor_rule(*args, **kw)`

Display a horizontal rule on the output device. The arguments to this method are entirely application- and writer-specific, and should be interpreted with care. The method implementation may assume that a line break has already been issued via `send_line_break()`.

`writer.send_flow_data(data)`

Output character data which may be word-wrapped and re-flowed as needed. Within any sequence of calls to this method, the writer may assume that spans of multiple whitespace characters have been collapsed to single space characters.

`writer.send_literal_data(data)`

Output character data which has already been formatted for display. Generally, this should be interpreted to mean that line breaks indicated by newline characters should be preserved and no new line breaks should be introduced. The data may contain embedded newline and tab characters, unlike data provided to the `send_formatted_data()` interface.

`writer.send_label_data(data)`

Set *data* to the left of the current left margin, if possible. The value of *data* is not restricted; treatment of non-string values is entirely application- and writer-dependent. This method will only be called at the beginning of a line.

33.1.4 Writer Implementations

Three implementations of the writer object interface are provided as examples by this module. Most applications will need to derive new writer classes from the `NullWriter` class.

class `formatter.NullWriter`

A writer which only provides the interface definition; no actions are taken on any methods. This should be the base class for all writers which do not need to inherit any implementation methods.

class `formatter.AbstractWriter`

A writer which can be used in debugging formatters, but not much else. Each method simply announces itself by printing its name and arguments on standard output.

class `formatter.DumbWriter` (*file=None*, *maxcol=72*)

Simple writer class which writes output on the *file object* passed in as *file* or, if *file* is omitted, on standard output.

The output is simply word-wrapped to the number of columns specified by *maxcol*. This class is suitable for reflowing a sequence of paragraphs.

本章节叙述的模块只在 Windows 平台上可用。

34.1 msilib —Read and write Microsoft Installer files

Source code: [Lib/msilib/__init__.py](#)

The *msilib* supports the creation of Microsoft Installer (.msi) files. Because these files often contain an embedded “cabinet” file (.cab), it also exposes an API to create CAB files. Support for reading .cab files is currently not implemented; read support for the .msi database is possible.

This package aims to provide complete access to all tables in an .msi file, therefore, it is a fairly low-level API. Two primary applications of this package are the *distutils* command `bdist_msi`, and the creation of Python installer package itself (although that currently uses a different version of *msilib*).

The package contents can be roughly split into four parts: low-level CAB routines, low-level MSI routines, higher-level MSI routines, and standard table structures.

msilib.FCICreate (*cabname, files*)

Create a new CAB file named *cabname*. *files* must be a list of tuples, each containing the name of the file on disk, and the name of the file inside the CAB file.

The files are added to the CAB file in the order they appear in the list. All files are added into a single CAB file, using the MSZIP compression algorithm.

Callbacks to Python for the various steps of MSI creation are currently not exposed.

msilib.UuidCreate ()

Return the string representation of a new unique identifier. This wraps the Windows API functions `UuidCreate()` and `UuidToString()`.

msilib.OpenDatabase (*path, persist*)

Return a new database object by calling `MsiOpenDatabase`. *path* is the file name of the MSI file; *persist* can be one of the constants `MSIDBOPEN_CREATEDIRECT`, `MSIDBOPEN_CREATE`, `MSIDBOPEN_DIRECT`,

MSIDBOPEN_READONLY, or MSIDBOPEN_TRANSACT, and may include the flag MSIDBOPEN_PATCHFILE. See the Microsoft documentation for the meaning of these flags; depending on the flags, an existing database is opened, or a new one created.

`msilib.CreateRecord(count)`

Return a new record object by calling `MSICreateRecord()`. *count* is the number of fields of the record.

`msilib.init_database(name, schema, ProductName, ProductCode, ProductVersion, Manufacturer)`

Create and return a new database *name*, initialize it with *schema*, and set the properties *ProductName*, *ProductCode*, *ProductVersion*, and *Manufacturer*.

schema must be a module object containing `tables` and `_Validation_records` attributes; typically, `msilib.schema` should be used.

The database will contain just the schema and the validation records when this function returns.

`msilib.add_data(database, table, records)`

Add all *records* to the table named *table* in *database*.

The *table* argument must be one of the predefined tables in the MSI schema, e.g. 'Feature', 'File', 'Component', 'Dialog', 'Control', etc.

records should be a list of tuples, each one containing all fields of a record according to the schema of the table. For optional fields, `None` can be passed.

Field values can be ints, strings, or instances of the Binary class.

class `msilib.Binary(filename)`

Represents entries in the Binary table; inserting such an object using `add_data()` reads the file named *filename* into the table.

`msilib.add_tables(database, module)`

Add all table content from *module* to *database*. *module* must contain an attribute *tables* listing all tables for which content should be added, and one attribute per table that has the actual content.

This is typically used to install the sequence tables.

`msilib.add_stream(database, name, path)`

Add the file *path* into the `_Stream` table of *database*, with the stream name *name*.

`msilib.gen_uuid()`

Return a new UUID, in the format that MSI typically requires (i.e. in curly braces, and with all hexdigits in upper-case).

参见:

[FCICreate UuidCreate UuidToString](#)

34.1.1 Database Objects

`Database.OpenView(sql)`

Return a view object, by calling `MSIDatabaseOpenView()`. *sql* is the SQL statement to execute.

`Database.Commit()`

Commit the changes pending in the current transaction, by calling `MSIDatabaseCommit()`.

`Database.GetSummaryInformation(count)`

Return a new summary information object, by calling `MsiGetSummaryInformation()`. *count* is the maximum number of updated values.

参见:

[MSIDatabaseOpenView](#) [MSIDatabaseCommit](#) [MSIGetSummaryInformation](#)

34.1.2 View Objects

View.Execute (*params*)

Execute the SQL query of the view, through `MsiViewExecute()`. If *params* is not `None`, it is a record describing actual values of the parameter tokens in the query.

View.GetColumnInfo (*kind*)

Return a record describing the columns of the view, through calling `MsiViewGetColumnInfo()`. *kind* can be either `MSICOLINFO_NAMES` or `MSICOLINFO_TYPES`.

View.Fetch ()

Return a result record of the query, through calling `MsiViewFetch()`.

View.Modify (*kind*, *data*)

Modify the view, by calling `MsiViewModify()`. *kind* can be one of `MSIMODIFY_SEEK`, `MSIMODIFY_REFRESH`, `MSIMODIFY_INSERT`, `MSIMODIFY_UPDATE`, `MSIMODIFY_ASSIGN`, `MSIMODIFY_REPLACE`, `MSIMODIFY_MERGE`, `MSIMODIFY_DELETE`, `MSIMODIFY_INSERT_TEMPORARY`, `MSIMODIFY_VALIDATE`, `MSIMODIFY_VALIDATE_NEW`, `MSIMODIFY_VALIDATE_FIELD`, or `MSIMODIFY_VALIDATE_DELETE`.

data must be a record describing the new data.

View.Close ()

Close the view, through `MsiViewClose()`.

参见:

[MsiViewExecute](#) [MsiViewGetColumnInfo](#) [MsiViewFetch](#) [MsiViewModify](#) [MsiViewClose](#)

34.1.3 Summary Information Objects

SummaryInformation.GetProperty (*field*)

Return a property of the summary, through `MsiSummaryInfoGetProperty()`. *field* is the name of the property, and can be one of the constants `PID_CODEPAGE`, `PID_TITLE`, `PID_SUBJECT`, `PID_AUTHOR`, `PID_KEYWORDS`, `PID_COMMENTS`, `PID_TEMPLATE`, `PID_LASTAUTHOR`, `PID_REVNUMBER`, `PID_LASTPRINTED`, `PID_CREATE_DTM`, `PID_LASTSAVE_DTM`, `PID_PAGECOUNT`, `PID_WORDCOUNT`, `PID_CHARCOUNT`, `PID_APPNAME`, or `PID_SECURITY`.

SummaryInformation.GetPropertyCount ()

Return the number of summary properties, through `MsiSummaryInfoGetPropertyCount()`.

SummaryInformation.SetProperty (*field*, *value*)

Set a property through `MsiSummaryInfoSetProperty()`. *field* can have the same values as in [GetProperty\(\)](#), *value* is the new value of the property. Possible value types are integer and string.

SummaryInformation.Persist ()

Write the modified properties to the summary information stream, using `MsiSummaryInfoPersist()`.

参见:

[MsiSummaryInfoGetProperty](#) [MsiSummaryInfoGetPropertyCount](#) [MsiSummaryInfoSetProperty](#) [MsiSummaryInfoPersist](#)

34.1.4 Record Objects

Record.GetFieldCount()

Return the number of fields of the record, through `MsiRecordGetFieldCount()`.

Record.GetInteger(*field*)

Return the value of *field* as an integer where possible. *field* must be an integer.

Record.GetString(*field*)

Return the value of *field* as a string where possible. *field* must be an integer.

Record.SetString(*field*, *value*)

Set *field* to *value* through `MsiRecordSetString()`. *field* must be an integer; *value* a string.

Record.SetStream(*field*, *value*)

Set *field* to the contents of the file named *value*, through `MsiRecordSetStream()`. *field* must be an integer; *value* a string.

Record.SetInteger(*field*, *value*)

Set *field* to *value* through `MsiRecordSetInteger()`. Both *field* and *value* must be an integer.

Record.ClearData()

Set all fields of the record to 0, through `MsiRecordClearData()`.

参见:

[MsiRecordGetFieldCount](#) [MsiRecordSetString](#) [MsiRecordSetStream](#) [MsiRecordSetInteger](#) [MsiRecordClearData](#)

34.1.5 Errors

All wrappers around MSI functions raise `MSIError`; the string inside the exception will contain more detail.

34.1.6 CAB Objects

class msilib.CAB(*name*)

The class `CAB` represents a CAB file. During MSI construction, files will be added simultaneously to the `Files` table, and to a CAB file. Then, when all files have been added, the CAB file can be written, then added to the MSI file.

name is the name of the CAB file in the MSI file.

append(*full*, *file*, *logical*)

Add the file with the pathname *full* to the CAB file, under the name *logical*. If there is already a file named *logical*, a new file name is created.

Return the index of the file in the CAB file, and the new name of the file inside the CAB file.

commit(*database*)

Generate a CAB file, add it as a stream to the MSI file, put it into the `Media` table, and remove the generated file from the disk.

34.1.7 Directory Objects

class msilib.Directory(*database*, *cab*, *basedir*, *physical*, *logical*, *default*[, *componentflags*])

Create a new directory in the Directory table. There is a current component at each point in time for the directory, which is either explicitly created through `start_component()`, or implicitly when files are added for the first time. Files are added into the current component, and into the cab file. To create a directory, a base directory

object needs to be specified (can be `None`), the path to the physical directory, and a logical directory name. *default* specifies the DefaultDir slot in the directory table. *componentflags* specifies the default flags that new components get.

start_component (*component=None, feature=None, flags=None, keyfile=None, uuid=None*)

Add an entry to the Component table, and make this component the current component for this directory. If no component name is given, the directory name is used. If no *feature* is given, the current feature is used. If no *flags* are given, the directory's default flags are used. If no *keyfile* is given, the KeyPath is left null in the Component table.

add_file (*file, src=None, version=None, language=None*)

Add a file to the current component of the directory, starting a new one if there is no current component. By default, the file name in the source and the file table will be identical. If the *src* file is specified, it is interpreted relative to the current directory. Optionally, a *version* and a *language* can be specified for the entry in the File table.

glob (*pattern, exclude=None*)

Add a list of files to the current component as specified in the glob pattern. Individual files can be excluded in the *exclude* list.

remove_pyc ()

Remove .pyc files on uninstall.

参见:

[Directory Table](#) [File Table](#) [Component Table](#) [FeatureComponents Table](#)

34.1.8 相关特性

class `msilib.Feature` (*db, id, title, desc, display, level=1, parent=None, directory=None, attributes=0*)

Add a new record to the Feature table, using the values *id*, *parent.id*, *title*, *desc*, *display*, *level*, *directory*, and *attributes*. The resulting feature object can be passed to the `start_component()` method of `Directory`.

set_current ()

Make this feature the current feature of `msilib`. New components are automatically added to the default feature, unless a feature is explicitly specified.

参见:

[Feature Table](#)

34.1.9 GUI classes

`msilib` provides several classes that wrap the GUI tables in an MSI database. However, no standard user interface is provided; use `bdist_msi` to create MSI files with a user-interface for installing Python packages.

class `msilib.Control` (*dlg, name*)

Base class of the dialog controls. *dlg* is the dialog object the control belongs to, and *name* is the control's name.

event (*event, argument, condition=1, ordering=None*)

Make an entry into the ControlEvent table for this control.

mapping (*event, attribute*)

Make an entry into the EventMapping table for this control.

condition (*action, condition*)

Make an entry into the ControlCondition table for this control.

class `msilib.RadioButtonGroup` (*dlg, name, property*)

Create a radio button control named *name*. *property* is the installer property that gets set when a radio button is selected.

add (*name, x, y, width, height, text, value=None*)

Add a radio button named *name* to the group, at the coordinates *x, y, width, height*, and with the label *text*. If *value* is `None`, it defaults to *name*.

class `msilib.Dialog` (*db, name, x, y, w, h, attr, title, first, default, cancel*)

Return a new `Dialog` object. An entry in the `Dialog` table is made, with the specified coordinates, dialog attributes, title, name of the first, default, and cancel controls.

control (*name, type, x, y, width, height, attributes, property, text, control_next, help*)

Return a new `Control` object. An entry in the `Control` table is made with the specified parameters.

This is a generic method; for specific types, specialized methods are provided.

text (*name, x, y, width, height, attributes, text*)

Add and return a `Text` control.

bitmap (*name, x, y, width, height, text*)

Add and return a `Bitmap` control.

line (*name, x, y, width, height*)

Add and return a `Line` control.

pushbutton (*name, x, y, width, height, attributes, text, next_control*)

Add and return a `PushButton` control.

radiogroup (*name, x, y, width, height, attributes, property, text, next_control*)

Add and return a `RadioButtonGroup` control.

checkbox (*name, x, y, width, height, attributes, property, text, next_control*)

Add and return a `CheckBox` control.

参见:

[Dialog Table Control Table Control Types ControlCondition Table ControlEvent Table EventMapping Table RadioButton Table](#)

34.1.10 Precomputed tables

`msilib` provides a few subpackages that contain only schema and table definitions. Currently, these definitions are based on MSI version 2.0.

`msilib.schema`

This is the standard MSI schema for MSI 2.0, with the *tables* variable providing a list of table definitions, and *_Validation_records* providing the data for MSI validation.

`msilib.sequence`

This module contains table contents for the standard sequence tables: *AdminExecuteSequence*, *AdminUISequence*, *AdvtExecuteSequence*, *InstallExecuteSequence*, and *InstallUISequence*.

`msilib.text`

This module contains definitions for the *UIText* and *ActionText* tables, for the standard installer actions.

34.2 msvcrt —Useful routines from the MS VC++ runtime

These functions provide access to some useful capabilities on Windows platforms. Some higher-level modules use these functions to build the Windows implementations of their services. For example, the `getpass` module uses this in the implementation of the `getpass()` function.

Further documentation on these functions can be found in the Platform API documentation.

The module implements both the normal and wide char variants of the console I/O api. The normal API deals only with ASCII characters and is of limited use for internationalized applications. The wide char API should be used where ever possible.

在 3.3 版更改: Operations in this module now raise `OSError` where `IOError` was raised.

34.2.1 File Operations

`msvcrt.locking(fd, mode, nbytes)`

Lock part of a file based on file descriptor `fd` from the C runtime. Raises `OSError` on failure. The locked region of the file extends from the current file position for `nbytes` bytes, and may continue beyond the end of the file. `mode` must be one of the `LK_*` constants listed below. Multiple regions in a file may be locked at the same time, but may not overlap. Adjacent regions are not merged; they must be unlocked individually.

`msvcrt.LK_LOCK`

`msvcrt.LK_RLCK`

Locks the specified bytes. If the bytes cannot be locked, the program immediately tries again after 1 second. If, after 10 attempts, the bytes cannot be locked, `OSError` is raised.

`msvcrt.LK_NBLCK`

`msvcrt.LK_NBLCK`

Locks the specified bytes. If the bytes cannot be locked, `OSError` is raised.

`msvcrt.LK_UNLCK`

Unlocks the specified bytes, which must have been previously locked.

`msvcrt.setmode(fd, flags)`

Set the line-end translation mode for the file descriptor `fd`. To set it to text mode, `flags` should be `os.O_TEXT`; for binary, it should be `os.O_BINARY`.

`msvcrt.open_osfhandle(handle, flags)`

Create a C runtime file descriptor from the file handle `handle`. The `flags` parameter should be a bitwise OR of `os.O_APPEND`, `os.O_RDONLY`, and `os.O_TEXT`. The returned file descriptor may be used as a parameter to `os.fdopen()` to create a file object.

`msvcrt.get_osfhandle(fd)`

Return the file handle for the file descriptor `fd`. Raises `OSError` if `fd` is not recognized.

34.2.2 Console I/O

`msvcrt.kbhit()`

Return true if a keypress is waiting to be read.

`msvcrt.getch()`

Read a keypress and return the resulting character as a byte string. Nothing is echoed to the console. This call will block if a keypress is not already available, but will not wait for Enter to be pressed. If the pressed key was a special function key, this will return `'\000'` or `'\xe0'`; the next call will return the keycode. The Control-C keypress cannot be read with this function.

`msvcrt.getwch()`

Wide char variant of `getch()`, returning a Unicode value.

`msvcrt.getche()`

Similar to `getch()`, but the keypress will be echoed if it represents a printable character.

`msvcrt.getwche()`

Wide char variant of `getche()`, returning a Unicode value.

`msvcrt.putch(char)`

Print the byte string `char` to the console without buffering.

`msvcrt.putwch(unicode_char)`

Wide char variant of `putch()`, accepting a Unicode value.

`msvcrt.ungetch(char)`

Cause the byte string `char` to be “pushed back” into the console buffer; it will be the next character read by `getch()` or `getche()`.

`msvcrt.ungetwch(unicode_char)`

Wide char variant of `ungetch()`, accepting a Unicode value.

34.2.3 Other Functions

`msvcrt.heapmin()`

Force the `malloc()` heap to clean itself up and return unused blocks to the operating system. On failure, this raises `OSError`.

34.3 winreg — Windows 注册表访问

这些函数将 Windows 注册表 API 暴露给 Python。为了确保即便程序员忽略了显式关闭句柄，该句柄依然能够正确关闭，它使用了一个 *handle* 对象 而不是整数来作为注册表句柄。

在 3.3 版更改：该模块中的几个函数被用于引发 `WindowsError`，该异常现在是 `OSError` 的别名。

34.3.1 函数

该模块提供了下列函数：

`winreg.CloseKey(hkey)`

关闭之前打开的注册表键。参数 `hkey` 指之前打开的键。

注解：如果没有使用该方法关闭 `hkey` (或者通过 `hkey.Close()`)，在对象 `hkey` 被 Python 销毁时会将其关闭。

`winreg.ConnectRegistry(computer_name, key)`

建立到另一台计算机上的预定义注册表句柄的连接，并返回一个 *handle* 对象。

`computer_name` 是远程计算机的名称，以 `r"\\computername"` 的形式。如果是 `None`，将会使用本地计算机。

`key` 是所连接到的预定义句柄。

返回值是所开打键的句柄。如果函数失败，则引发一个 `OSError` 异常。

在 3.3 版更改：参考上文。

`winreg.CreateKey(key, sub_key)`

创建或打开特定的键，返回一个 *handle* 对象。

key 为某个已经打开的键，或者预定义的 *HKEY_** 常量 之一。

sub_key 是用于命名该方法所打开或创建的键的字符串。

如果 *key* 是预定义键之一，*sub_key* 可能会是 `None`。该情况下，返回的句柄就是传入函数的句柄。

如果键已经存在，则该函数打开已经存在的该键。

返回值是所开打键的句柄。如果函数失败，则引发一个 *OSError* 异常。

在 3.3 版更改: 参考上文。

`winreg.CreateKeyEx(key, sub_key, reserved=0, access=KEY_WRITE)`

创建或打开特定的键，返回一个 *handle* 对象。

key 为某个已经打开的键，或者预定义的 *HKEY_** 常量 之一。

sub_key 是用于命名该方法所打开或创建的键的字符串。

reserved 是一个保留的证书，必须为零。默认值为零。

access 为一个整数，用于给键的预期安全访问指定访问掩码。默认值为 *KEY_WRITE*。参阅 *Access Rights* 了解其它允许值。

如果 *key* 是预定义键之一，*sub_key* 可能会是 `None`。该情况下，返回的句柄就是传入函数的句柄。

如果键已经存在，则该函数打开已经存在的该键。

返回值是所开打键的句柄。如果函数失败，则引发一个 *OSError* 异常。

3.2 新版功能.

在 3.3 版更改: 参考上文。

`winreg.DeleteKey(key, sub_key)`

删除指定的键。

key 为某个已经打开的键，或者预定义的 *HKEY_** 常量 之一。

sub_key 这个字符串必须是由 *key* 参数所指定键的一个子项。该值项不可以是 `None`，同时键也不可以有子项。

该方法不能删除带有子项的键。

如果方法成功，则整个键，包括其所有值项都会被移除。如果方法失败，则引发一个 *OSError* 异常。

在 3.3 版更改: 参考上文。

`winreg.DeleteKeyEx(key, sub_key, access=KEY_WOW64_64KEY, reserved=0)`

删除指定的键。

注解: 函数 *DeleteKeyEx()* 通过 *RegDeleteKeyEx* 这个 Windows API 函数实现，该函数为 Windows 的 64 位版本专属。参阅 *RegDeleteKeyEx* 文档。

key 为某个已经打开的键，或者预定义的 *HKEY_** 常量 之一。

sub_key 这个字符串必须是由 *key* 参数所指定键的一个子项。该值项不可以是 `None`，同时键也不可以有子项。

reserved 是一个保留的证书，必须为零。默认值为零。

access 为一个整数，用于给键的预期安全访问指定访问掩码。默认值为常量 *_WOW64_64KEY*。参阅 *Access Rights* 了解其它允许值。

该方法不能删除带有子项的键。

如果方法成功，则整个键，包括其所有值项都会被移除。如果方法失败，则引发一个 `OSError` 异常。
在不支持的 Windows 版本之上，将会引发 `NotImplementedError` 异常。

3.2 新版功能.

在 3.3 版更改: 参考上文。

`winreg.DeleteValue(key, value)`

从某个注册键中删除一个命名值项。

`key` 为某个已经打开的键，或者预定义的 `HKEY_*` 常量 之一。

`value` 为标识所要删除值项的字符串。

`winreg.EnumKey(key, index)`

列举某个已经打开注册表键的子项，并返回一个字符串。

`key` 为某个已经打开的键，或者预定义的 `HKEY_*` 常量 之一。

`index` 为一个整数，用于标识所获取键的索引。

每次调用该函数都会获取一个子项的名字。通常它会被反复调用，直到引发 `OSError` 异常，这说明已经没有更多的可用值了。

在 3.3 版更改: 参考上文。

`winreg.EnumValue(key, index)`

列举某个已经打开注册表键的值项，并返回一个元组。

`key` 为某个已经打开的键，或者预定义的 `HKEY_*` 常量 之一。

`index` 为一个整数，用于标识要获取值项的索引。

每次调用该函数都会获取一个子项的名字。通常它会被反复调用，直到引发 `OSError` 异常，这说明已经没有更多的可用值了。

结果为 3 元素的元组。

索引	意义
0	用于标识值项名称的字符串。
1	保存值项数据的对象，其类型取决于背后的注册表类型。
2	标识值项数据类型的整数。(请查阅 <code>SetValueEx()</code> 文档中的表格)

在 3.3 版更改: 参考上文。

`winreg.ExpandEnvironmentStrings(str)`

Expands environment variable placeholders `%NAME%` in strings like `REG_EXPAND_SZ`:

```
>>> ExpandEnvironmentStrings('%windir%')
'C:\\Windows'
```

`winreg.FlushKey(key)`

将某个键的所有属性写入注册表。

`key` 为某个已经打开的键，或者预定义的 `HKEY_*` 常量 之一。

It is not necessary to call `FlushKey()` to change a key. Registry changes are flushed to disk by the registry using its lazy flusher. Registry changes are also flushed to disk at system shutdown. Unlike `CloseKey()`, the `FlushKey()` method returns only when all the data has been written to the registry. An application should only call `FlushKey()` if it requires absolute certainty that registry changes are on disk.

注解: If you don't know whether a `FlushKey()` call is required, it probably isn't.

`winreg.LoadKey(key, sub_key, file_name)`
Creates a subkey under the specified key and stores registration information from a specified file into that subkey.
`key` is a handle returned by `ConnectRegistry()` or one of the constants `HKEY_USERS` or `HKEY_LOCAL_MACHINE`.
`sub_key` is a string that identifies the subkey to load.
`file_name` is the name of the file to load registry data from. This file must have been created with the `SaveKey()` function. Under the file allocation table (FAT) file system, the filename may not have an extension.
A call to `LoadKey()` fails if the calling process does not have the `SE_RESTORE_PRIVILEGE` privilege. Note that privileges are different from permissions – see the [RegLoadKey documentation](#) for more details.
If `key` is a handle returned by `ConnectRegistry()`, then the path specified in `file_name` is relative to the remote computer.

`winreg.OpenKey(key, sub_key, reserved=0, access=KEY_READ)`
`winreg.OpenKeyEx(key, sub_key, reserved=0, access=KEY_READ)`
Opens the specified key, returning a *handle object*.
`key` 为某个已经打开的键，或者预定义的 `HKEY_*` 常量 之一。
`sub_key` is a string that identifies the sub_key to open.
`reserved` is a reserved integer, and must be zero. The default is zero.
`access` is an integer that specifies an access mask that describes the desired security access for the key. Default is `KEY_READ`. See [Access Rights](#) for other allowed values.
The result is a new handle to the specified key.
If the function fails, `OSError` is raised.
在 3.2 版更改: Allow the use of named arguments.
在 3.3 版更改: 参考上文。

`winreg.QueryInfoKey(key)`
Returns information about a key, as a tuple.
`key` 为某个已经打开的键，或者预定义的 `HKEY_*` 常量 之一。
结果为 3 元素的元组。

索引	意义
0	An integer giving the number of sub keys this key has.
1	An integer giving the number of values this key has.
2	An integer giving when the key was last modified (if available) as 100' s of nanoseconds since Jan 1, 1601.

`winreg.QueryValue(key, sub_key)`
Retrieves the unnamed value for a key, as a string.
`key` 为某个已经打开的键，或者预定义的 `HKEY_*` 常量 之一。
`sub_key` is a string that holds the name of the subkey with which the value is associated. If this parameter is `None` or empty, the function retrieves the value set by the `SetValue()` method for the key identified by `key`.

Values in the registry have name, type, and data components. This method retrieves the data for a key's first value that has a NULL name. But the underlying API call doesn't return the type, so always use `QueryValueEx()` if possible.

`winreg.QueryValueEx(key, value_name)`

Retrieves the type and data for a specified value name associated with an open registry key.

`key` 为某个已经打开的键，或者预定义的 `HKEY_*` 常量之一。

`value_name` is a string indicating the value to query.

The result is a tuple of 2 items:

索引	意义
0	The value of the registry item.
1	An integer giving the registry type for this value (see table in docs for <code>SetValueEx()</code>)

`winreg.SaveKey(key, file_name)`

Saves the specified key, and all its subkeys to the specified file.

`key` 为某个已经打开的键，或者预定义的 `HKEY_*` 常量之一。

`file_name` is the name of the file to save registry data to. This file cannot already exist. If this filename includes an extension, it cannot be used on file allocation table (FAT) file systems by the `LoadKey()` method.

If `key` represents a key on a remote computer, the path described by `file_name` is relative to the remote computer. The caller of this method must possess the `SeBackupPrivilege` security privilege. Note that privileges are different than permissions – see the [Conflicts Between User Rights and Permissions](#) documentation for more details.

This function passes NULL for `security_attributes` to the API.

`winreg.SetValue(key, sub_key, type, value)`

Associates a value with a specified key.

`key` 为某个已经打开的键，或者预定义的 `HKEY_*` 常量之一。

`sub_key` is a string that names the subkey with which the value is associated.

`type` is an integer that specifies the type of the data. Currently this must be `REG_SZ`, meaning only strings are supported. Use the `SetValueEx()` function for support for other data types.

`value` is a string that specifies the new value.

If the key specified by the `sub_key` parameter does not exist, the `SetValue` function creates it.

Value lengths are limited by available memory. Long values (more than 2048 bytes) should be stored as files with the filenames stored in the configuration registry. This helps the registry perform efficiently.

The key identified by the `key` parameter must have been opened with `KEY_SET_VALUE` access.

`winreg.SetValueEx(key, value_name, reserved, type, value)`

Stores data in the value field of an open registry key.

`key` 为某个已经打开的键，或者预定义的 `HKEY_*` 常量之一。

`value_name` is a string that names the subkey with which the value is associated.

`reserved` can be anything – zero is always passed to the API.

`type` is an integer that specifies the type of the data. See [Value Types](#) for the available types.

`value` is a string that specifies the new value.

This method can also set additional value and type information for the specified key. The key identified by the `key` parameter must have been opened with `KEY_SET_VALUE` access.

To open the key, use the `CreateKey()` or `OpenKey()` methods.

Value lengths are limited by available memory. Long values (more than 2048 bytes) should be stored as files with the filenames stored in the configuration registry. This helps the registry perform efficiently.

`winreg.DisableReflectionKey(key)`

Disables registry reflection for 32-bit processes running on a 64-bit operating system.

`key` 为某个已经打开的键，或者预定义的 `HKEY_*` 常量 之一。

Will generally raise `NotImplemented` if executed on a 32-bit operating system.

If the key is not on the reflection list, the function succeeds but has no effect. Disabling reflection for a key does not affect reflection of any subkeys.

`winreg.EnableReflectionKey(key)`

Restores registry reflection for the specified disabled key.

`key` 为某个已经打开的键，或者预定义的 `HKEY_*` 常量 之一。

Will generally raise `NotImplemented` if executed on a 32-bit operating system.

Restoring reflection for a key does not affect reflection of any subkeys.

`winreg.QueryReflectionKey(key)`

Determines the reflection state for the specified key.

`key` 为某个已经打开的键，或者预定义的 `HKEY_*` 常量 之一。

Returns `True` if reflection is disabled.

Will generally raise `NotImplemented` if executed on a 32-bit operating system.

34.3.2 常数

The following constants are defined for use in many `_winreg` functions.

HKEY_* Constants

`winreg.HKEY_CLASSES_ROOT`

Registry entries subordinate to this key define types (or classes) of documents and the properties associated with those types. Shell and COM applications use the information stored under this key.

`winreg.HKEY_CURRENT_USER`

Registry entries subordinate to this key define the preferences of the current user. These preferences include the settings of environment variables, data about program groups, colors, printers, network connections, and application preferences.

`winreg.HKEY_LOCAL_MACHINE`

Registry entries subordinate to this key define the physical state of the computer, including data about the bus type, system memory, and installed hardware and software.

`winreg.HKEY_USERS`

Registry entries subordinate to this key define the default user configuration for new users on the local computer and the user configuration for the current user.

`winreg.HKEY_PERFORMANCE_DATA`

Registry entries subordinate to this key allow you to access performance data. The data is not actually stored in the registry; the registry functions cause the system to collect the data from its source.

`winreg.HKEY_CURRENT_CONFIG`

Contains information about the current hardware profile of the local computer system.

`winreg.HKEY_DYN_DATA`

This key is not used in versions of Windows after 98.

Access Rights

For more information, see [Registry Key Security and Access](#).

`winreg.KEY_ALL_ACCESS`

Combines the `STANDARD_RIGHTS_REQUIRED`, `KEY_QUERY_VALUE`, `KEY_SET_VALUE`, `KEY_CREATE_SUB_KEY`, `KEY_ENUMERATE_SUB_KEYS`, `KEY_NOTIFY`, and `KEY_CREATE_LINK` access rights.

`winreg.KEY_WRITE`

Combines the `STANDARD_RIGHTS_WRITE`, `KEY_SET_VALUE`, and `KEY_CREATE_SUB_KEY` access rights.

`winreg.KEY_READ`

Combines the `STANDARD_RIGHTS_READ`, `KEY_QUERY_VALUE`, `KEY_ENUMERATE_SUB_KEYS`, and `KEY_NOTIFY` values.

`winreg.KEY_EXECUTE`

Equivalent to `KEY_READ`.

`winreg.KEY_QUERY_VALUE`

Required to query the values of a registry key.

`winreg.KEY_SET_VALUE`

Required to create, delete, or set a registry value.

`winreg.KEY_CREATE_SUB_KEY`

Required to create a subkey of a registry key.

`winreg.KEY_ENUMERATE_SUB_KEYS`

Required to enumerate the subkeys of a registry key.

`winreg.KEY_NOTIFY`

Required to request change notifications for a registry key or for subkeys of a registry key.

`winreg.KEY_CREATE_LINK`

Reserved for system use.

64-bit Specific

For more information, see [Accessing an Alternate Registry View](#).

`winreg.KEY_WOW64_64KEY`

Indicates that an application on 64-bit Windows should operate on the 64-bit registry view.

`winreg.KEY_WOW64_32KEY`

Indicates that an application on 64-bit Windows should operate on the 32-bit registry view.

Value Types

For more information, see [Registry Value Types](#).

`winreg.REG_BINARY`
Binary data in any form.

`winreg.REG_DWORD`
32-bit number.

`winreg.REG_DWORD_LITTLE_ENDIAN`
A 32-bit number in little-endian format. Equivalent to `REG_DWORD`.

`winreg.REG_DWORD_BIG_ENDIAN`
A 32-bit number in big-endian format.

`winreg.REG_EXPAND_SZ`
Null-terminated string containing references to environment variables (%PATH%).

`winreg.REG_LINK`
A Unicode symbolic link.

`winreg.REG_MULTI_SZ`
A sequence of null-terminated strings, terminated by two null characters. (Python handles this termination automatically.)

`winreg.REG_NONE`
No defined value type.

`winreg.REG_QWORD`
A 64-bit number.
3.6 新版功能.

`winreg.REG_QWORD_LITTLE_ENDIAN`
A 64-bit number in little-endian format. Equivalent to `REG_QWORD`.
3.6 新版功能.

`winreg.REG_RESOURCE_LIST`
A device-driver resource list.

`winreg.REG_FULL_RESOURCE_DESCRIPTOR`
A hardware setting.

`winreg.REG_RESOURCE_REQUIREMENTS_LIST`
A hardware resource list.

`winreg.REG_SZ`
A null-terminated string.

34.3.3 Registry Handle Objects

This object wraps a Windows HKEY object, automatically closing it when the object is destroyed. To guarantee cleanup, you can call either the `Close()` method on the object, or the `CloseKey()` function.

All registry functions in this module return one of these objects.

All registry functions in this module which accept a handle object also accept an integer, however, use of the handle object is encouraged.

Handle objects provide semantics for `__bool__()` – thus

```
if handle:
    print("Yes")
```

will print `Yes` if the handle is currently valid (has not been closed or detached).

The object also support comparison semantics, so handle objects will compare true if they both reference the same underlying Windows handle value.

Handle objects can be converted to an integer (e.g., using the built-in `int()` function), in which case the underlying Windows handle value is returned. You can also use the `Detach()` method to return the integer handle, and also disconnect the Windows handle from the handle object.

`PyHKEY.Close()`

Closes the underlying Windows handle.

If the handle is already closed, no error is raised.

`PyHKEY.Detach()`

Detaches the Windows handle from the handle object.

The result is an integer that holds the value of the handle before it is detached. If the handle is already detached or closed, this will return zero.

After calling this function, the handle is effectively invalidated, but the handle is not closed. You would call this function when you need the underlying Win32 handle to exist beyond the lifetime of the handle object.

`PyHKEY.__enter__()`

`PyHKEY.__exit__(*exc_info)`

The `HKEY` object implements `__enter__()` and `__exit__()` and thus supports the context protocol for the `with` statement:

```
with OpenKey(HKEY_LOCAL_MACHINE, "foo") as key:
    ... # work with key
```

will automatically close `key` when control leaves the `with` block.

34.4 winsound — Sound-playing interface for Windows

The `winsound` module provides access to the basic sound-playing machinery provided by Windows platforms. It includes functions and several constants.

`winsound.Beep(frequency, duration)`

Beep the PC's speaker. The *frequency* parameter specifies frequency, in hertz, of the sound, and must be in the range 37 through 32,767. The *duration* parameter specifies the number of milliseconds the sound should last. If the system is not able to beep the speaker, `RuntimeError` is raised.

`winsound.PlaySound(sound, flags)`

Call the underlying `PlaySound()` function from the Platform API. The *sound* parameter may be a filename, a system sound alias, audio data as a *bytes-like object*, or `None`. Its interpretation depends on the value of *flags*, which can be a bitwise ORed combination of the constants described below. If the *sound* parameter is `None`, any currently playing waveform sound is stopped. If the system indicates an error, `RuntimeError` is raised.

`winsound.MessageBeep(type=MB_OK)`

Call the underlying `MessageBeep()` function from the Platform API. This plays a sound as specified in the registry. The *type* argument specifies which sound to play; possible values are `-1`, `MB_ICONASTERISK`, `MB_ICONEXCLAMATION`, `MB_ICONHAND`, `MB_ICONQUESTION`, and `MB_OK`, all described below. The value `-1` produces a "simple beep"; this is the final fallback if a sound cannot be played otherwise. If the system indicates an error, `RuntimeError` is raised.

winsound.SND_FILENAME

The *sound* parameter is the name of a WAV file. Do not use with *SND_ALIAS*.

winsound.SND_ALIAS

The *sound* parameter is a sound association name from the registry. If the registry contains no such name, play the system default sound unless *SND_NODEFAULT* is also specified. If no default sound is registered, raise *RuntimeError*. Do not use with *SND_FILENAME*.

All Win32 systems support at least the following; most systems support many more:

<i>PlaySound()</i> name	Corresponding Control Panel Sound name
'SystemAsterisk'	Asterisk
'SystemExclamation'	Exclamation
'SystemExit'	Exit Windows
'SystemHand'	Critical Stop
'SystemQuestion'	Question

例如:

```
import winsound
# Play Windows exit sound.
winsound.PlaySound("SystemExit", winsound.SND_ALIAS)

# Probably play Windows default sound, if any is registered (because
# "" probably isn't the registered name of any sound).
winsound.PlaySound("", winsound.SND_ALIAS)
```

winsound.SND_LOOP

Play the sound repeatedly. The *SND_ASYNC* flag must also be used to avoid blocking. Cannot be used with *SND_MEMORY*.

winsound.SND_MEMORY

The *sound* parameter to *PlaySound()* is a memory image of a WAV file, as a *bytes-like object*.

注解: This module does not support playing from a memory image asynchronously, so a combination of this flag and *SND_ASYNC* will raise *RuntimeError*.

winsound.SND_PURGE

Stop playing all instances of the specified sound.

注解: This flag is not supported on modern Windows platforms.

winsound.SND_ASYNC

Return immediately, allowing sounds to play asynchronously.

winsound.SND_NODEFAULT

If the specified sound cannot be found, do not play the system default sound.

winsound.SND_NOSTOP

Do not interrupt sounds currently playing.

winsound.SND_NOWAIT

Return immediately if the sound driver is busy.

注解: This flag is not supported on modern Windows platforms.

`winsound.MB_ICONASTERISK`

Play the `SystemDefault` sound.

`winsound.MB_ICONEXCLAMATION`

Play the `SystemExclamation` sound.

`winsound.MB_ICONHAND`

Play the `SystemHand` sound.

`winsound.MB_ICONQUESTION`

Play the `SystemQuestion` sound.

`winsound.MB_OK`

Play the `SystemDefault` sound.

本章描述的模块提供了 Unix 操作系统独有特性的接口，在某些情况下也适用于它的某些或许多衍生版。以下为模块概览：

35.1 `posix` —The most common POSIX system calls

This module provides access to operating system functionality that is standardized by the C Standard and the POSIX standard (a thinly disguised Unix interface).

Do not import this module directly. Instead, import the module `os`, which provides a *portable* version of this interface. On Unix, the `os` module provides a superset of the `posix` interface. On non-Unix operating systems the `posix` module is not available, but a subset is always available through the `os` interface. Once `os` is imported, there is *no* performance penalty in using it instead of `posix`. In addition, `os` provides some additional functionality, such as automatically calling `putenv()` when an entry in `os.environ` is changed.

Errors are reported as exceptions; the usual exceptions are given for type errors, while errors reported by the system calls raise `OSError`.

35.1.1 Large File Support

Several operating systems (including AIX, HP-UX, Irix and Solaris) provide support for files that are larger than 2 GiB from a C programming model where `int` and `long` are 32-bit values. This is typically accomplished by defining the relevant size and offset types as 64-bit values. Such files are sometimes referred to as *large files*.

Large file support is enabled in Python when the size of an `off_t` is larger than a `long` and the `long long` type is available and is at least as large as an `off_t`. It may be necessary to configure and compile Python with certain compiler flags to enable this mode. For example, it is enabled by default with recent versions of Irix, but with Solaris 2.6 and 2.7 you need to do something like:

```
CFLAGS="`getconf LFS_CFLAGS`" OPT="-g -O2 $CFLAGS" \
./configure
```

On large-file-capable Linux systems, this might work:

```
CFLAGS='-D_LARGEFILE64_SOURCE -D_FILE_OFFSET_BITS=64' OPT="-g -O2 $CFLAGS" \
./configure
```

35.1.2 Notable Module Contents

In addition to many functions described in the *os* module documentation, *posix* defines the following data item:

posix.envIRON
A dictionary representing the string environment at the time the interpreter was started. Keys and values are bytes on Unix and str on Windows. For example, `environ[b'HOME']` (`environ['HOME']` on Windows) is the pathname of your home directory, equivalent to `getenv("HOME")` in C.
Modifying this dictionary does not affect the string environment passed on by *execv()*, *popen()* or *system()*; if you need to change the environment, pass *environ* to *execve()* or add variable assignments and export statements to the command string for *system()* or *popen()*.

在 3.2 版更改: On Unix, keys and values are bytes.

注解: The *os* module provides an alternate implementation of *environ* which updates the environment on modification. Note also that updating *os.envIRON* will render this dictionary obsolete. Use of the *os* module version of this is recommended over direct access to the *posix* module.

35.2 pwd — 用户密码数据库

此模块可以访问 Unix 用户账户名及密码数据库，在所有 Unix 版本上均可使用。
密码数据库中的条目以元组对象返回，属性对应 *passwd* 中的结构（属性如下所示，可参考 *<pwd.h>*）:

索引	属性	意义
0	<i>pw_name</i>	登录名
1	<i>pw_passwd</i>	密码，可能已经加密
2	<i>pw_uid</i>	用户 ID 数值
3	<i>pw_gid</i>	组 ID 数值
4	<i>pw_gecos</i>	用户名或备注
5	<i>pw_dir</i>	用户主目录
6	<i>pw_shell</i>	用户的命令解释器

其中 *uid* 和 *gid* 是整数，其他是字符串，如果找不到对应的项目，抛出 *KeyError* 异常。

注解: 传统的 Unix 系统中，*pw_passwd* 的值通常使用 DES 导出的算法加密（参阅 *crypt* 模块）。不过现在的 unix 系统使用 影子密码系统。在这些 unix 上，*pw_passwd* 只包含星号（*'*'*）或字母（*'x'*），而加密的

密码存储在文件 `/etc/shadow` 中，此文件不是全局可读的。在 `pw_passwd` 中是否包含有用信息是系统相关的。如果可以访问到加密的密码，就需要使用 `spwd` 模块了。

本模块定义如下内容：

`pwd.getpwuid(uid)`
给定用户的数值 ID，返回密码数据库的对应项目。

`pwd.getpwnam(name)`
给定用户名，返回密码数据库的对应项目。

`pwd.getpwall()`
返回密码数据库中所有项目的列表，顺序不是固定的。

参见：

模块 `grp` 针对用户组数据库的接口，与本模块类似。

模块 `spwd` 针对影子密码数据库的接口，与本模块类似。

35.3 spwd —The shadow password database

This module provides access to the Unix shadow password database. It is available on various Unix versions.

You must have enough privileges to access the shadow password database (this usually means you have to be root).

Shadow password database entries are reported as a tuple-like object, whose attributes correspond to the members of the `spwd` structure (Attribute field below, see `<shadow.h>`):

索引	属性	意义
0	<code>sp_namp</code>	登录名
1	<code>sp_pwdp</code>	Encrypted password
2	<code>sp_lstchg</code>	Date of last change
3	<code>sp_min</code>	Minimal number of days between changes
4	<code>sp_max</code>	Maximum number of days between changes
5	<code>sp_warn</code>	Number of days before password expires to warn user about it
6	<code>sp_inact</code>	Number of days after password expires until account is disabled
7	<code>sp_expire</code>	Number of days since 1970-01-01 when account expires
8	<code>sp_flag</code>	Reserved

The `sp_namp` and `sp_pwdp` items are strings, all others are integers. `KeyError` is raised if the entry asked for cannot be found.

定义了以下函数：

`spwd.getspnam(name)`
Return the shadow password database entry for the given user name.

在 3.6 版更改: Raises a `PermissionError` instead of `KeyError` if the user doesn't have privileges.

`spwd.getspall()`
Return a list of all available shadow password database entries, in arbitrary order.

参见：

模块 `grp` 针对用户组数据库的接口，与本模块类似。

Module `pwd` An interface to the normal password database, similar to this.

35.4 `grp` —The group database

This module provides access to the Unix group database. It is available on all Unix versions.

Group database entries are reported as a tuple-like object, whose attributes correspond to the members of the `group` structure (Attribute field below, see `<pwd.h>`):

索引	属性	意义
0	<code>gr_name</code>	the name of the group
1	<code>gr_passwd</code>	the (encrypted) group password; often empty
2	<code>gr_gid</code>	the numerical group ID
3	<code>gr_mem</code>	all the group member's user names

The `gid` is an integer, name and password are strings, and the member list is a list of strings. (Note that most users are not explicitly listed as members of the group they are in according to the password database. Check both databases to get complete membership information. Also note that a `gr_name` that starts with a `+` or `-` is likely to be a YP/NIS reference and may not be accessible via `getgrnam()` or `getgrgid()`.)

本模块定义如下内容：

`grp.getgrgid(gid)`

Return the group database entry for the given numeric group ID. `KeyError` is raised if the entry asked for cannot be found.

3.6 版后已移除: Since Python 3.6 the support of non-integer arguments like floats or strings in `getgrgid()` is deprecated.

`grp.getgrnam(name)`

Return the group database entry for the given group name. `KeyError` is raised if the entry asked for cannot be found.

`grp.getgrall()`

Return a list of all available group entries, in arbitrary order.

参见：

Module `pwd` An interface to the user database, similar to this.

模块 `spwd` 针对影子密码数据库的接口，与本模块类似。

35.5 `crypt` —Function to check Unix passwords

Source code: [Lib/crypt.py](#)

This module implements an interface to the `crypt(3)` routine, which is a one-way hash function based upon a modified DES algorithm; see the Unix man page for further details. Possible uses include storing hashed passwords so you can check passwords without storing the actual password, or attempting to crack Unix passwords with a dictionary.

Notice that the behavior of this module depends on the actual implementation of the `crypt(3)` routine in the running system. Therefore, any extensions available on the current implementation will also be available on this module.

35.5.1 Hashing Methods

3.3 新版功能.

The `crypt` module defines the list of hashing methods (not all methods are available on all platforms):

`crypt.METHOD_SHA512`

A Modular Crypt Format method with 16 character salt and 86 character hash. This is the strongest method.

`crypt.METHOD_SHA256`

Another Modular Crypt Format method with 16 character salt and 43 character hash.

`crypt.METHOD_MD5`

Another Modular Crypt Format method with 8 character salt and 22 character hash.

`crypt.METHOD_CRYPT`

The traditional method with a 2 character salt and 13 characters of hash. This is the weakest method.

35.5.2 Module Attributes

3.3 新版功能.

`crypt.methods`

A list of available password hashing algorithms, as `crypt.METHOD_*` objects. This list is sorted from strongest to weakest.

35.5.3 模块函数

The `crypt` module defines the following functions:

`crypt.crypt(word, salt=None)`

word will usually be a user's password as typed at a prompt or in a graphical interface. The optional *salt* is either a string as returned from `mk salt()`, one of the `crypt.METHOD_*` values (though not all may be available on all platforms), or a full encrypted password including salt, as returned by this function. If *salt* is not provided, the strongest method will be used (as returned by `methods()`).

Checking a password is usually done by passing the plain-text password as *word* and the full results of a previous `crypt()` call, which should be the same as the results of this call.

salt (either a random 2 or 16 character string, possibly prefixed with `$digit$` to indicate the method) which will be used to perturb the encryption algorithm. The characters in *salt* must be in the set `[./a-zA-Z0-9]`, with the exception of Modular Crypt Format which prefixes a `$digit$`.

Returns the hashed password as a string, which will be composed of characters from the same alphabet as the salt.

Since a few `crypt(3)` extensions allow different values, with different sizes in the *salt*, it is recommended to use the full crypted password as salt when checking for a password.

在 3.3 版更改: Accept `crypt.METHOD_*` values in addition to strings for *salt*.

`crypt.mk salt(method=None)`

Return a randomly generated salt of the specified method. If no *method* is given, the strongest method available as returned by `methods()` is used.

The return value is a string either of 2 characters in length for `crypt.METHOD_CRYPT`, or 19 characters starting with `$digit$` and 16 random characters from the set `[./a-zA-Z0-9]`, suitable for passing as the *salt* argument to `crypt()`.

3.3 新版功能.

35.5.4 示例

A simple example illustrating typical use (a constant-time comparison operation is needed to limit exposure to timing attacks. `hmac.compare_digest()` is suitable for this purpose):

```
import pwd
import crypt
import getpass
from hmac import compare_digest as compare_hash

def login():
    username = input('Python login: ')
    cryptpasswd = pwd.getpwnam(username)[1]
    if cryptpasswd:
        if cryptpasswd == 'x' or cryptpasswd == '*':
            raise ValueError('no support for shadow passwords')
        cleartext = getpass.getpass()
        return compare_hash(crypt.crypt(cleartext, cryptpasswd), cryptpasswd)
    else:
        return True
```

To generate a hash of a password using the strongest available method and check it against the original:

```
import crypt
from hmac import compare_digest as compare_hash

hashed = crypt.crypt(plaintext)
if not compare_hash(hashed, crypt.crypt(plaintext, hashed)):
    raise ValueError("hashed version doesn't validate against original")
```

35.6 termios —POSIX style tty control

This module provides an interface to the POSIX calls for tty I/O control. For a complete description of these calls, see *termios(3)* Unix manual page. It is only available for those Unix versions that support POSIX *termios* style tty I/O control configured during installation.

All functions in this module take a file descriptor *fd* as their first argument. This can be an integer file descriptor, such as returned by `sys.stdin.fileno()`, or a *file object*, such as `sys.stdin` itself.

This module also defines all the constants needed to work with the functions provided here; these have the same name as their counterparts in C. Please refer to your system documentation for more information on using these terminal control interfaces.

这个模块定义了以下函数:

`termios.tcgetattr(fd)`

Return a list containing the tty attributes for file descriptor *fd*, as follows: [*iflag*, *oflag*, *cflag*, *lflag*, *ispeed*, *ospeed*, *cc*] where *cc* is a list of the tty special characters (each a string of length 1, except the items with indices *VMIN* and *VTIME*, which are integers when these fields are defined). The interpretation of the flags and the speeds as well as the indexing in the *cc* array must be done using the symbolic constants defined in the *termios* module.

`termios.tcsetattr(fd, when, attributes)`

Set the tty attributes for file descriptor *fd* from the *attributes*, which is a list like the one returned by *tcgetattr()*.

The *when* argument determines when the attributes are changed: `TCSANOW` to change immediately, `TCSADRAIN` to change after transmitting all queued output, or `TCSAFLUSH` to change after transmitting all queued output and discarding all queued input.

`termios.tcsendbreak(fd, duration)`

Send a break on file descriptor *fd*. A zero *duration* sends a break for 0.25–0.5 seconds; a nonzero *duration* has a system dependent meaning.

`termios.tcdrain(fd)`

Wait until all output written to file descriptor *fd* has been transmitted.

`termios.tcflush(fd, queue)`

Discard queued data on file descriptor *fd*. The *queue* selector specifies which queue: `TCIFLUSH` for the input queue, `TCOFLUSH` for the output queue, or `TCIOFLUSH` for both queues.

`termios.tcflow(fd, action)`

Suspend or resume input or output on file descriptor *fd*. The *action* argument can be `TCOOFF` to suspend output, `TCOON` to restart output, `TCIOFF` to suspend input, or `TCION` to restart input.

参见:

Module `tty` Convenience functions for common terminal control operations.

35.6.1 示例

Here's a function that prompts for a password with echoing turned off. Note the technique using a separate `tcgetattr()` call and a `try ... finally` statement to ensure that the old `tty` attributes are restored exactly no matter what happens:

```
def getpass(prompt="Password: "):
    import termios, sys
    fd = sys.stdin.fileno()
    old = termios.tcgetattr(fd)
    new = termios.tcgetattr(fd)
    new[3] = new[3] & ~termios.ECHO          # lflags
    try:
        termios.tcsetattr(fd, termios.TCSADRAIN, new)
        passwd = input(prompt)
    finally:
        termios.tcsetattr(fd, termios.TCSADRAIN, old)
    return passwd
```

35.7 tty —终端控制功能

Source code: [Lib/tty.py](#)

`tty` 模块定义了将 `tty` 放入 `cbreak` 和 `raw` 模式的函数。

因为它需要 `termios` 模块，所以只能在 Unix 上运行。

`tty` 模块定义了以下函数：

`tty.setraw(fd, when=termios.TCSAFLUSH)`

将文件描述符 *fd* 的模式更改为 `raw`。如果 *when* 被省略，则默认为 `termios.TCSAFLUSH`，并传递给 `termios.tcsetattr()`。

`tty.setcbreak(fd, when=termios.TCSAFLUSH)`

将文件描述符 *fd* 的模式更改为 `cbreak`。如果 *when* 被省略，则默认为 `termios.TCSAFLUSH`，并传递给 `termios.tcsetattr()`。

参见：

模块 `termios` 低级终端控制接口。

35.8 pty —Pseudo-terminal utilities

Source code: [Lib/pty.py](#)

The `pty` module defines operations for handling the pseudo-terminal concept: starting another process and being able to write to and read from its controlling terminal programmatically.

Because pseudo-terminal handling is highly platform dependent, there is code to do it only for Linux. (The Linux code is supposed to work on other platforms, but hasn't been tested yet.)

The `pty` module defines the following functions:

`pty.fork()`

Fork. Connect the child's controlling terminal to a pseudo-terminal. Return value is `(pid, fd)`. Note that the child gets `pid 0`, and the `fd` is *invalid*. The parent's return value is the `pid` of the child, and `fd` is a file descriptor connected to the child's controlling terminal (and also to the child's standard input and output).

`pty.openpty()`

Open a new pseudo-terminal pair, using `os.openpty()` if possible, or emulation code for generic Unix systems. Return a pair of file descriptors `(master, slave)`, for the master and the slave end, respectively.

`pty.spawn(argv[, master_read[, stdin_read]])`

Spawn a process, and connect its controlling terminal with the current process's standard io. This is often used to baffle programs which insist on reading from the controlling terminal.

The functions `master_read` and `stdin_read` should be functions which read from a file descriptor. The defaults try to read 1024 bytes each time they are called.

在 3.4 版更改: `spawn()` now returns the status value from `os.waitpid()` on the child process.

35.8.1 示例

The following program acts like the Unix command `script(1)`, using a pseudo-terminal to record all input and output of a terminal session in a “typescript”.

```
import argparse
import os
import pty
import sys
import time

parser = argparse.ArgumentParser()
parser.add_argument('-a', dest='append', action='store_true')
parser.add_argument('-p', dest='use_python', action='store_true')
parser.add_argument('filename', nargs='?', default='typescript')
options = parser.parse_args()
```

(下页继续)

(续上页)

```

shell = sys.executable if options.use_python else os.environ.get('SHELL', 'sh')
filename = options.filename
mode = 'ab' if options.append else 'wb'

with open(filename, mode) as script:
    def read(fd):
        data = os.read(fd, 1024)
        script.write(data)
        return data

    print('Script started, file is', filename)
    script.write(('Script started on %s\n' % time.asctime()).encode())

    pty.spawn(shell, read)

    script.write(('Script done on %s\n' % time.asctime()).encode())
    print('Script done, file is', filename)

```

35.9 fcntl —The fcntl and ioctl system calls

This module performs file control and I/O control on file descriptors. It is an interface to the `fcntl()` and `ioctl()` Unix routines. For a complete description of these calls, see *fcntl(2)* and *ioctl(2)* Unix manual pages.

All functions in this module take a file descriptor *fd* as their first argument. This can be an integer file descriptor, such as returned by `sys.stdin.fileno()`, or an *io.IOBase* object, such as `sys.stdin` itself, which provides a *fileno()* that returns a genuine file descriptor.

在 3.3 版更改: Operations in this module used to raise an *IOError* where they now raise an *OSError*.

这个模块定义了以下函数:

`fcntl.fcntl(fd, cmd, arg=0)`

Perform the operation *cmd* on file descriptor *fd* (file objects providing a *fileno()* method are accepted as well). The values used for *cmd* are operating system dependent, and are available as constants in the *fcntl* module, using the same names as used in the relevant C header files. The argument *arg* can either be an integer value, or a *bytes* object. With an integer value, the return value of this function is the integer return value of the C `fcntl()` call. When the argument is bytes it represents a binary structure, e.g. created by `struct.pack()`. The binary data is copied to a buffer whose address is passed to the C `fcntl()` call. The return value after a successful call is the contents of the buffer, converted to a *bytes* object. The length of the returned object will be the same as the length of the *arg* argument. This is limited to 1024 bytes. If the information returned in the buffer by the operating system is larger than 1024 bytes, this is most likely to result in a segmentation violation or a more subtle data corruption.

If the `fcntl()` fails, an *OSError* is raised.

`fcntl.ioctl(fd, request, arg=0, mutate_flag=True)`

This function is identical to the *fcntl()* function, except that the argument handling is even more complicated.

The *request* parameter is limited to values that can fit in 32-bits. Additional constants of interest for use as the *request* argument can be found in the *termios* module, under the same names as used in the relevant C header files.

The parameter *arg* can be one of an integer, an object supporting the read-only buffer interface (like *bytes*) or an object supporting the read-write buffer interface (like *bytearray*).

In all but the last case, behaviour is as for the `fcntl()` function.

If a mutable buffer is passed, then the behaviour is determined by the value of the `mutate_flag` parameter.

If it is false, the buffer's mutability is ignored and behaviour is as for a read-only buffer, except that the 1024 byte limit mentioned above is avoided – so long as the buffer you pass is at least as long as what the operating system wants to put there, things should work.

If `mutate_flag` is true (the default), then the buffer is (in effect) passed to the underlying `ioctl()` system call, the latter's return code is passed back to the calling Python, and the buffer's new contents reflect the action of the `ioctl()`. This is a slight simplification, because if the supplied buffer is less than 1024 bytes long it is first copied into a static buffer 1024 bytes long which is then passed to `ioctl()` and copied back into the supplied buffer.

If the `ioctl()` fails, an `OSError` exception is raised.

举个例子：

```
>>> import array, fcntl, struct, termios, os
>>> os.getpgrp()
13341
>>> struct.unpack('h', fcntl.ioctl(0, termios.TIOCGPRG, "  ")[0])
13341
>>> buf = array.array('h', [0])
>>> fcntl.ioctl(0, termios.TIOCGPRG, buf, 1)
0
>>> buf
array('h', [13341])
```

`fcntl.flock(fd, operation)`

Perform the lock operation *operation* on file descriptor *fd* (file objects providing a `fileno()` method are accepted as well). See the Unix manual `flock(2)` for details. (On some systems, this function is emulated using `fcntl()`.)

If the `flock()` fails, an `OSError` exception is raised.

`fcntl.lockf(fd, cmd, len=0, start=0, whence=0)`

This is essentially a wrapper around the `fcntl()` locking calls. *fd* is the file descriptor of the file to lock or unlock, and *cmd* is one of the following values:

- `LOCK_UN` – unlock
- `LOCK_SH` – acquire a shared lock
- `LOCK_EX` – acquire an exclusive lock

When *cmd* is `LOCK_SH` or `LOCK_EX`, it can also be bitwise ORed with `LOCK_NB` to avoid blocking on lock acquisition. If `LOCK_NB` is used and the lock cannot be acquired, an `OSError` will be raised and the exception will have an `errno` attribute set to `EACCES` or `EAGAIN` (depending on the operating system; for portability, check for both values). On at least some systems, `LOCK_EX` can only be used if the file descriptor refers to a file opened for writing.

len is the number of bytes to lock, *start* is the byte offset at which the lock starts, relative to *whence*, and *whence* is as with `io.IOBase.seek()`, specifically:

- 0 – relative to the start of the file (`os.SEEK_SET`)
- 1 – relative to the current buffer position (`os.SEEK_CUR`)
- 2 – relative to the end of the file (`os.SEEK_END`)

The default for *start* is 0, which means to start at the beginning of the file. The default for *len* is 0 which means to lock to the end of the file. The default for *whence* is also 0.

Examples (all on a SVR4 compliant system):

```
import struct, fcntl, os

f = open(...)
rv = fcntl.fcntl(f, fcntl.F_SETFL, os.O_NDELAY)

lockdata = struct.pack('hhllhh', fcntl.F_WRLCK, 0, 0, 0, 0, 0)
rv = fcntl.fcntl(f, fcntl.F_SETLKW, lockdata)
```

Note that in the first example the return value variable *rv* will hold an integer value; in the second example it will hold a *bytes* object. The structure lay-out for the *lockdata* variable is system dependent —therefore using the *flock()* call may be better.

参见:

模块 *os* If the locking flags *O_SHLOCK* and *O_EXLOCK* are present in the *os* module (on BSD only), the *os.open()* function provides an alternative to the *lockf()* and *flock()* functions.

35.10 pipes —Interface to shell pipelines

源代码: [Lib/pipes.py](#)

The *pipes* module defines a class to abstract the concept of a *pipeline* —a sequence of converters from one file to another.

Because the module uses */bin/sh* command lines, a POSIX or compatible shell for *os.system()* and *os.popen()* is required.

The *pipes* module defines the following class:

class *pipes.Template*
An abstraction of a pipeline.

示例:

```
>>> import pipes
>>> t = pipes.Template()
>>> t.append('tr a-z A-Z', '--')
>>> f = t.open('pipefile', 'w')
>>> f.write('hello world')
>>> f.close()
>>> open('pipefile').read()
'HELLO WORLD'
```

35.10.1 Template Objects

Template objects following methods:

Template.reset()
Restore a pipeline template to its initial state.

Template.clone()
Return a new, equivalent, pipeline template.

`Template.debug(flag)`

If *flag* is true, turn debugging on. Otherwise, turn debugging off. When debugging is on, commands to be executed are printed, and the shell is given `set -x` command to be more verbose.

`Template.append(cmd, kind)`

Append a new action at the end. The *cmd* variable must be a valid bourne shell command. The *kind* variable consists of two letters.

The first letter can be either of `'-'` (which means the command reads its standard input), `'f'` (which means the commands reads a given file on the command line) or `'.'` (which means the commands reads no input, and hence must be first.)

Similarly, the second letter can be either of `'-'` (which means the command writes to standard output), `'f'` (which means the command writes a file on the command line) or `'.'` (which means the command does not write anything, and hence must be last.)

`Template.prepend(cmd, kind)`

Add a new action at the beginning. See [append\(\)](#) for explanations of the arguments.

`Template.open(file, mode)`

Return a file-like object, open to *file*, but read from or written to by the pipeline. Note that only one of `'r'`, `'w'` may be given.

`Template.copy(infile, outfile)`

Copy *infile* to *outfile* through the pipe.

35.11 resource —Resource usage information

This module provides basic mechanisms for measuring and controlling system resources utilized by a program.

Symbolic constants are used to specify particular system resources and to request usage information about either the current process or its children.

An [OSError](#) is raised on syscall failure.

exception `resource.error`

A deprecated alias of [OSError](#).

在 3.3 版更改: Following [PEP 3151](#), this class was made an alias of [OSError](#).

35.11.1 Resource Limits

Resources usage can be limited using the [setrlimit\(\)](#) function described below. Each resource is controlled by a pair of limits: a soft limit and a hard limit. The soft limit is the current limit, and may be lowered or raised by a process over time. The soft limit can never exceed the hard limit. The hard limit can be lowered to any value greater than the soft limit, but not raised. (Only processes with the effective UID of the super-user can raise a hard limit.)

The specific resources that can be limited are system dependent. They are described in the [getrlimit\(2\)](#) man page. The resources listed below are supported when the underlying operating system supports them; resources which cannot be checked or controlled by the operating system are not defined in this module for those platforms.

`resource.RLIM_INFINITY`

Constant used to represent the limit for an unlimited resource.

`resource.getrlimit(resource)`

Returns a tuple (soft, hard) with the current soft and hard limits of *resource*. Raises `ValueError` if an invalid resource is specified, or `error` if the underlying system call fails unexpectedly.

`resource.setrlimit(resource, limits)`

Sets new limits of consumption of *resource*. The *limits* argument must be a tuple (soft, hard) of two integers describing the new limits. A value of `RLIM_INFINITY` can be used to request a limit that is unlimited.

Raises `ValueError` if an invalid resource is specified, if the new soft limit exceeds the hard limit, or if a process tries to raise its hard limit. Specifying a limit of `RLIM_INFINITY` when the hard or system limit for that resource is not unlimited will result in a `ValueError`. A process with the effective UID of super-user can request any valid limit value, including unlimited, but `ValueError` will still be raised if the requested limit exceeds the system imposed limit.

`setrlimit` may also raise `error` if the underlying system call fails.

`resource.prlimit(pid, resource[, limits])`

Combines `setrlimit()` and `getrlimit()` in one function and supports to get and set the resources limits of an arbitrary process. If *pid* is 0, then the call applies to the current process. *resource* and *limits* have the same meaning as in `setrlimit()`, except that *limits* is optional.

When *limits* is not given the function returns the *resource* limit of the process *pid*. When *limits* is given the *resource* limit of the process is set and the former resource limit is returned.

Raises `ProcessLookupError` when *pid* can't be found and `PermissionError` when the user doesn't have `CAP_SYS_RESOURCE` for the process.

Availability: Linux 2.6.36 or later with glibc 2.13 or later

3.4 新版功能.

These symbols define resources whose consumption can be controlled using the `setrlimit()` and `getrlimit()` functions described below. The values of these symbols are exactly the constants used by C programs.

The Unix man page for `getrlimit(2)` lists the available resources. Note that not all systems use the same symbol or same value to denote the same resource. This module does not attempt to mask platform differences —symbols not defined for a platform will not be available from this module on that platform.

`resource.RLIMIT_CORE`

The maximum size (in bytes) of a core file that the current process can create. This may result in the creation of a partial core file if a larger core would be required to contain the entire process image.

`resource.RLIMIT_CPU`

The maximum amount of processor time (in seconds) that a process can use. If this limit is exceeded, a `SIGXCPU` signal is sent to the process. (See the `signal` module documentation for information about how to catch this signal and do something useful, e.g. flush open files to disk.)

`resource.RLIMIT_FSIZE`

The maximum size of a file which the process may create.

`resource.RLIMIT_DATA`

The maximum size (in bytes) of the process's heap.

`resource.RLIMIT_STACK`

The maximum size (in bytes) of the call stack for the current process. This only affects the stack of the main thread in a multi-threaded process.

`resource.RLIMIT_RSS`

The maximum resident set size that should be made available to the process.

`resource.RLIMIT_NPROC`

The maximum number of processes the current process may create.

`resource.RLIMIT_NOFILE`

The maximum number of open file descriptors for the current process.

`resource.RLIMIT_OFILE`

The BSD name for `RLIMIT_NOFILE`.

`resource.RLIMIT_MEMLOCK`

The maximum address space which may be locked in memory.

`resource.RLIMIT_VMEM`

The largest area of mapped memory which the process may occupy.

`resource.RLIMIT_AS`

The maximum area (in bytes) of address space which may be taken by the process.

`resource.RLIMIT_MSGQUEUE`

The number of bytes that can be allocated for POSIX message queues.

Availability: Linux 2.6.8 or later.

3.4 新版功能.

`resource.RLIMIT_NICE`

The ceiling for the process' s nice level (calculated as `20 - rlim_cur`).

Availability: Linux 2.6.12 or later.

3.4 新版功能.

`resource.RLIMIT_RTPRIO`

The ceiling of the real-time priority.

Availability: Linux 2.6.12 or later.

3.4 新版功能.

`resource.RLIMIT_RTTIME`

The time limit (in microseconds) on CPU time that a process can spend under real-time scheduling without making a blocking syscall.

Availability: Linux 2.6.25 or later.

3.4 新版功能.

`resource.RLIMIT_SIGPENDING`

The number of signals which the process may queue.

Availability: Linux 2.6.8 or later.

3.4 新版功能.

`resource.RLIMIT_SBSIZE`

The maximum size (in bytes) of socket buffer usage for this user. This limits the amount of network memory, and hence the amount of mbufs, that this user may hold at any time.

Availability: FreeBSD 9 or later.

3.4 新版功能.

`resource.RLIMIT_SWAP`

The maximum size (in bytes) of the swap space that may be reserved or used by all of this user id' s processes. This limit is enforced only if bit 1 of the `vm.overcommit` sysctl is set. Please see *tuning(7)* for a complete description of this sysctl.

Availability: FreeBSD 9 or later.

3.4 新版功能.

`resource.RLIMIT_NPTS`

The maximum number of pseudo-terminals created by this user id.

Availability: FreeBSD 9 or later.

3.4 新版功能.

35.11.2 Resource Usage

These functions are used to retrieve resource usage information:

`resource.getrusage(who)`

This function returns an object that describes the resources consumed by either the current process or its children, as specified by the *who* parameter. The *who* parameter should be specified using one of the `RUSAGE_*` constants described below.

The fields of the return value each describe how a particular system resource has been used, e.g. amount of time spent running in user mode or number of times the process was swapped out of main memory. Some values are dependent on the clock tick interval, e.g. the amount of memory the process is using.

For backward compatibility, the return value is also accessible as a tuple of 16 elements.

The fields `ru_utime` and `ru_stime` of the return value are floating point values representing the amount of time spent executing in user mode and the amount of time spent executing in system mode, respectively. The remaining values are integers. Consult the `getrusage(2)` man page for detailed information about these values. A brief summary is presented here:

索引	Field	Resource
0	<code>ru_utime</code>	time in user mode (float)
1	<code>ru_stime</code>	time in system mode (float)
2	<code>ru_maxrss</code>	maximum resident set size
3	<code>ru_ixrss</code>	shared memory size
4	<code>ru_idrss</code>	unshared memory size
5	<code>ru_isrss</code>	unshared stack size
6	<code>ru_minflt</code>	page faults not requiring I/O
7	<code>ru_majflt</code>	page faults requiring I/O
8	<code>ru_nswap</code>	number of swap outs
9	<code>ru_inblock</code>	block input operations
10	<code>ru_oublock</code>	block output operations
11	<code>ru_msgsnd</code>	messages sent
12	<code>ru_msgrcv</code>	messages received
13	<code>ru_nsignals</code>	signals received
14	<code>ru_nvcsw</code>	voluntary context switches
15	<code>ru_nivcsw</code>	involuntary context switches

This function will raise a `ValueError` if an invalid *who* parameter is specified. It may also raise `error` exception in unusual circumstances.

`resource.getpagesize()`

Returns the number of bytes in a system page. (This need not be the same as the hardware page size.)

The following `RUSAGE_*` symbols are passed to the `getrusage()` function to specify which processes information should be provided for.

`resource.RUSAGE_SELF`

Pass to `getrusage()` to request resources consumed by the calling process, which is the sum of resources used by all threads in the process.

`resource.RUSAGE_CHILDREN`

Pass to `getrusage()` to request resources consumed by child processes of the calling process which have been terminated and waited for.

`resource.RUSAGE_BOTH`

Pass to `getrusage()` to request resources consumed by both the current process and child processes. May not be available on all systems.

`resource.RUSAGE_THREAD`

Pass to `getrusage()` to request resources consumed by the current thread. May not be available on all systems.

3.2 新版功能.

35.12 nis —Interface to Sun' s NIS (Yellow Pages)

The `nis` module gives a thin wrapper around the NIS library, useful for central administration of several hosts.

Because NIS exists only on Unix systems, this module is only available for Unix.

The `nis` module defines the following functions:

`nis.match(key, mapname, domain=default_domain)`

Return the match for `key` in map `mapname`, or raise an error (`nis.error`) if there is none. Both should be strings, `key` is 8-bit clean. Return value is an arbitrary array of bytes (may contain NULL and other joys).

Note that `mapname` is first checked if it is an alias to another name.

The `domain` argument allows overriding the NIS domain used for the lookup. If unspecified, lookup is in the default NIS domain.

`nis.cat(mapname, domain=default_domain)`

Return a dictionary mapping `key` to `value` such that `match(key, mapname) == value`. Note that both keys and values of the dictionary are arbitrary arrays of bytes.

Note that `mapname` is first checked if it is an alias to another name.

The `domain` argument allows overriding the NIS domain used for the lookup. If unspecified, lookup is in the default NIS domain.

`nis.maps(domain=default_domain)`

Return a list of all valid maps.

The `domain` argument allows overriding the NIS domain used for the lookup. If unspecified, lookup is in the default NIS domain.

`nis.get_default_domain()`

Return the system default NIS domain.

The `nis` module defines the following exception:

exception `nis.error`

An error raised when a NIS function returns an error code.

35.13 Unix syslog 库例程

This module provides an interface to the Unix `syslog` library routines. Refer to the Unix manual pages for a detailed description of the `syslog` facility.

This module wraps the system `syslog` family of routines. A pure Python library that can speak to a syslog server is available in the `logging.handlers` module as `SysLogHandler`.

这个模块定义了以下函数：

`syslog.syslog(message)`

`syslog.syslog(priority, message)`

Send the string *message* to the system logger. A trailing newline is added if necessary. Each message is tagged with a priority composed of a *facility* and a *level*. The optional *priority* argument, which defaults to `LOG_INFO`, determines the message priority. If the facility is not encoded in *priority* using logical-or (`LOG_INFO | LOG_USER`), the value given in the `openlog()` call is used.

If `openlog()` has not been called prior to the call to `syslog()`, `openlog()` will be called with no arguments.

`syslog.openlog([ident[, logoption[, facility]]])`

Logging options of subsequent `syslog()` calls can be set by calling `openlog()`. `syslog()` will call `openlog()` with no arguments if the log is not currently open.

The optional *ident* keyword argument is a string which is prepended to every message, and defaults to `sys.argv[0]` with leading path components stripped. The optional *logoption* keyword argument (default is 0) is a bit field – see below for possible values to combine. The optional *facility* keyword argument (default is `LOG_USER`) sets the default facility for messages which do not have a facility explicitly encoded.

在 3.2 版更改：In previous versions, keyword arguments were not allowed, and *ident* was required. The default for *ident* was dependent on the system libraries, and often was `python` instead of the name of the python program file.

`syslog.closelog()`

Reset the syslog module values and call the system library `closelog()`.

This causes the module to behave as it does when initially imported. For example, `openlog()` will be called on the first `syslog()` call (if `openlog()` hasn't already been called), and *ident* and other `openlog()` parameters are reset to defaults.

`syslog.setlogmask(maskpri)`

Set the priority mask to *maskpri* and return the previous mask value. Calls to `syslog()` with a priority level not set in *maskpri* are ignored. The default is to log all priorities. The function `LOG_MASK(pri)` calculates the mask for the individual priority *pri*. The function `LOG_UPTO(pri)` calculates the mask for all priorities up to and including *pri*.

The module defines the following constants:

Priority levels (high to low): `LOG_EMERG`, `LOG_ALERT`, `LOG_CRIT`, `LOG_ERR`, `LOG_WARNING`, `LOG_NOTICE`, `LOG_INFO`, `LOG_DEBUG`.

Facilities: `LOG_KERN`, `LOG_USER`, `LOG_MAIL`, `LOG_DAEMON`, `LOG_AUTH`, `LOG_LPR`, `LOG_NEWS`, `LOG_UUCP`, `LOG_CRON`, `LOG_SYSLOG`, `LOG_LOCAL0` to `LOG_LOCAL7`, and, if defined in `<syslog.h>`, `LOG_AUTHPRIV`.

Log options: `LOG_PID`, `LOG_CONS`, `LOG_NDELAY`, and, if defined in `<syslog.h>`, `LOG_ODELAY`, `LOG_NOWAIT`, and `LOG_PERROR`.

35.13.1 示例

Simple example

A simple set of examples:

```
import syslog

syslog.syslog('Processing started')
if error:
    syslog.syslog(syslog.LOG_ERR, 'Processing started')
```

An example of setting some log options, these would include the process ID in logged messages, and write the messages to the destination facility used for mail logging:

```
syslog.openlog(logoption=syslog.LOG_PID, facility=syslog.LOG_MAIL)
syslog.syslog('E-mail processing initiated...')
```

被取代的模块

本章中描述的模块均已弃用，仅保留用于向后兼容。它们已经被其他模块所取代。

36.1 `optparse` — Parser for command line options

Source code: [Lib/optparse.py](#)

3.2 版后已移除: The `optparse` module is deprecated and will not be developed further; development will continue with the `argparse` module.

`optparse` is a more convenient, flexible, and powerful library for parsing command-line options than the old `getopt` module. `optparse` uses a more declarative style of command-line parsing: you create an instance of `OptionParser`, populate it with options, and parse the command line. `optparse` allows users to specify options in the conventional GNU/POSIX syntax, and additionally generates usage and help messages for you.

Here's an example of using `optparse` in a simple script:

```
from optparse import OptionParser
...
parser = OptionParser()
parser.add_option("-f", "--file", dest="filename",
                  help="write report to FILE", metavar="FILE")
parser.add_option("-q", "--quiet",
                  action="store_false", dest="verbose", default=True,
                  help="don't print status messages to stdout")

(options, args) = parser.parse_args()
```

With these few lines of code, users of your script can now do the “usual thing” on the command-line, for example:

```
<yourscript> --file=outfile -q
```

As it parses the command line, `optparse` sets attributes of the `options` object returned by `parse_args()` based on user-supplied command-line values. When `parse_args()` returns from parsing this command line, `options.filename` will be "outfile" and `options.verbose` will be `False`. `optparse` supports both long and short options, allows short options to be merged together, and allows options to be associated with their arguments in a variety of ways. Thus, the following command lines are all equivalent to the above example:

```
<yourscript> -f outfile --quiet
<yourscript> --quiet --file outfile
<yourscript> -q -foutfile
<yourscript> -qfoutfile
```

Additionally, users can run one of

```
<yourscript> -h
<yourscript> --help
```

and `optparse` will print out a brief summary of your script's options:

```
Usage: <yourscript> [options]

Options:
  -h, --help            show this help message and exit
  -f FILE, --file=FILE  write report to FILE
  -q, --quiet           don't print status messages to stdout
```

where the value of `yourscript` is determined at runtime (normally from `sys.argv[0]`).

36.1.1 Background

`optparse` was explicitly designed to encourage the creation of programs with straightforward, conventional command-line interfaces. To that end, it supports only the most common command-line syntax and semantics conventionally used under Unix. If you are unfamiliar with these conventions, read this section to acquaint yourself with them.

Terminology

参数 a string entered on the command-line, and passed by the shell to `execl()` or `execv()`. In Python, arguments are elements of `sys.argv[1:]` (`sys.argv[0]` is the name of the program being executed). Unix shells also use the term "word".

It is occasionally desirable to substitute an argument list other than `sys.argv[1:]`, so you should read "argument" as "an element of `sys.argv[1:]`, or of some other list provided as a substitute for `sys.argv[1:]`".

option an argument used to supply extra information to guide or customize the execution of a program. There are many different syntaxes for options; the traditional Unix syntax is a hyphen ("-") followed by a single letter, e.g. `-x` or `-F`. Also, traditional Unix syntax allows multiple options to be merged into a single argument, e.g. `-x -F` is equivalent to `-xF`. The GNU project introduced `--` followed by a series of hyphen-separated words, e.g. `--file` or `--dry-run`. These are the only two option syntaxes provided by `optparse`.

Some other option syntaxes that the world has seen include:

- a hyphen followed by a few letters, e.g. `-pf` (this is *not* the same as multiple options merged into a single argument)
- a hyphen followed by a whole word, e.g. `-file` (this is technically equivalent to the previous syntax, but they aren't usually seen in the same program)

- a plus sign followed by a single letter, or a few letters, or a word, e.g. `+f`, `+rgb`
- a slash followed by a letter, or a few letters, or a word, e.g. `/f`, `/file`

These option syntaxes are not supported by `optparse`, and they never will be. This is deliberate: the first three are non-standard on any environment, and the last only makes sense if you’re exclusively targeting VMS, MS-DOS, and/or Windows.

option argument an argument that follows an option, is closely associated with that option, and is consumed from the argument list when that option is. With `optparse`, option arguments may either be in a separate argument from their option:

```
-f foo
--file foo
```

or included in the same argument:

```
-ffoo
--file=foo
```

Typically, a given option either takes an argument or it doesn’t. Lots of people want an “optional option arguments” feature, meaning that some options will take an argument if they see it, and won’t if they don’t. This is somewhat controversial, because it makes parsing ambiguous: if `-a` takes an optional argument and `-b` is another option entirely, how do we interpret `-ab`? Because of this ambiguity, `optparse` does not support this feature.

positional 参数 something leftover in the argument list after options have been parsed, i.e. after options and their arguments have been parsed and removed from the argument list.

required option an option that must be supplied on the command-line; note that the phrase “required option” is self-contradictory in English. `optparse` doesn’t prevent you from implementing required options, but doesn’t give you much help at it either.

For example, consider this hypothetical command-line:

```
prog -v --report report.txt foo bar
```

`-v` and `--report` are both options. Assuming that `--report` takes one argument, `report.txt` is an option argument. `foo` and `bar` are positional arguments.

What are options for?

Options are used to provide extra information to tune or customize the execution of a program. In case it wasn’t clear, options are usually *optional*. A program should be able to run just fine with no options whatsoever. (Pick a random program from the Unix or GNU toolsets. Can it run without any options at all and still make sense? The main exceptions are `find`, `tar`, and `dd`—all of which are mutant oddballs that have been rightly criticized for their non-standard syntax and confusing interfaces.)

Lots of people want their programs to have “required options”. Think about it. If it’s required, then it’s *not optional*! If there is a piece of information that your program absolutely requires in order to run successfully, that’s what positional arguments are for.

As an example of good command-line interface design, consider the humble `cp` utility, for copying files. It doesn’t make much sense to try to copy files without supplying a destination and at least one source. Hence, `cp` fails if you run it with no arguments. However, it has a flexible, useful syntax that does not require any options at all:

```
cp SOURCE DEST
cp SOURCE ... DEST-DIR
```

You can get pretty far with just that. Most `cp` implementations provide a bunch of options to tweak exactly how the files are copied: you can preserve mode and modification time, avoid following symlinks, ask before clobbering existing files, etc. But none of this distracts from the core mission of `cp`, which is to copy either one file to another, or several files to another directory.

What are positional arguments for?

Positional arguments are for those pieces of information that your program absolutely, positively requires to run.

A good user interface should have as few absolute requirements as possible. If your program requires 17 distinct pieces of information in order to run successfully, it doesn't matter *how* you get that information from the user—most people will give up and walk away before they successfully run the program. This applies whether the user interface is a command-line, a configuration file, or a GUI: if you make that many demands on your users, most of them will simply give up.

In short, try to minimize the amount of information that users are absolutely required to supply—use sensible defaults whenever possible. Of course, you also want to make your programs reasonably flexible. That's what options are for. Again, it doesn't matter if they are entries in a config file, widgets in the “Preferences” dialog of a GUI, or command-line options—the more options you implement, the more flexible your program is, and the more complicated its implementation becomes. Too much flexibility has drawbacks as well, of course; too many options can overwhelm users and make your code much harder to maintain.

36.1.2 教程

While `optparse` is quite flexible and powerful, it's also straightforward to use in most cases. This section covers the code patterns that are common to any `optparse`-based program.

First, you need to import the `OptionParser` class; then, early in the main program, create an `OptionParser` instance:

```
from optparse import OptionParser
...
parser = OptionParser()
```

Then you can start defining options. The basic syntax is:

```
parser.add_option(opt_str, ...,
                  attr=value, ...)
```

Each option has one or more option strings, such as `-f` or `--file`, and several option attributes that tell `optparse` what to expect and what to do when it encounters that option on the command line.

Typically, each option will have one short option string and one long option string, e.g.:

```
parser.add_option("-f", "--file", ...)
```

You're free to define as many short option strings and as many long option strings as you like (including zero), as long as there is at least one option string overall.

The option strings passed to `OptionParser.add_option()` are effectively labels for the option defined by that call. For brevity, we will frequently refer to *encountering an option* on the command line; in reality, `optparse` encounters *option strings* and looks up options from them.

Once all of your options are defined, instruct `optparse` to parse your program's command line:

```
(options, args) = parser.parse_args()
```

(If you like, you can pass a custom argument list to `parse_args()`, but that's rarely necessary: by default it uses `sys.argv[1:]`.)

`parse_args()` returns two values:

- `options`, an object containing values for all of your options—e.g. if `--file` takes a single string argument, then `options.file` will be the filename supplied by the user, or `None` if the user did not supply that option
- `args`, the list of positional arguments leftover after parsing options

This tutorial section only covers the four most important option attributes: `action`, `type`, `dest` (destination), and `help`. Of these, `action` is the most fundamental.

Understanding option actions

Actions tell `optparse` what to do when it encounters an option on the command line. There is a fixed set of actions hard-coded into `optparse`; adding new actions is an advanced topic covered in section [Extending optparse](#). Most actions tell `optparse` to store a value in some variable—for example, take a string from the command line and store it in an attribute of `options`.

If you don't specify an option action, `optparse` defaults to `store`.

The store action

The most common option action is `store`, which tells `optparse` to take the next argument (or the remainder of the current argument), ensure that it is of the correct type, and store it to your chosen destination.

例如:

```
parser.add_option("-f", "--file",
                 action="store", type="string", dest="filename")
```

Now let's make up a fake command line and ask `optparse` to parse it:

```
args = ["-f", "foo.txt"]
(options, args) = parser.parse_args(args)
```

When `optparse` sees the option string `-f`, it consumes the next argument, `foo.txt`, and stores it in `options.filename`. So, after this call to `parse_args()`, `options.filename` is `"foo.txt"`.

Some other option types supported by `optparse` are `int` and `float`. Here's an option that expects an integer argument:

```
parser.add_option("-n", type="int", dest="num")
```

Note that this option has no long option string, which is perfectly acceptable. Also, there's no explicit action, since the default is `store`.

Let's parse another fake command-line. This time, we'll jam the option argument right up against the option: since `-n42` (one argument) is equivalent to `-n 42` (two arguments), the code

```
(options, args) = parser.parse_args(["-n42"])
print(options.num)
```

will print 42.

If you don't specify a type, `optparse` assumes `string`. Combined with the fact that the default action is `store`, that means our first example can be a lot shorter:

```
parser.add_option("-f", "--file", dest="filename")
```

If you don't supply a destination, *optparse* figures out a sensible default from the option strings: if the first long option string is `--foo-bar`, then the default destination is `foo_bar`. If there are no long option strings, *optparse* looks at the first short option string: the default destination for `-f` is `f`.

optparse also includes the built-in `complex` type. Adding types is covered in section [Extending *optparse*](#).

Handling boolean (flag) options

Flag options—set a variable to true or false when a particular option is seen—are quite common. *optparse* supports them with two separate actions, `store_true` and `store_false`. For example, you might have a `verbose` flag that is turned on with `-v` and off with `-q`:

```
parser.add_option("-v", action="store_true", dest="verbose")
parser.add_option("-q", action="store_false", dest="verbose")
```

Here we have two different options with the same destination, which is perfectly OK. (It just means you have to be a bit careful when setting default values—see below.)

When *optparse* encounters `-v` on the command line, it sets `options.verbose` to `True`; when it encounters `-q`, `options.verbose` is set to `False`.

Other actions

Some other actions supported by *optparse* are:

- "store_const"** store a constant value
- "append"** append this option's argument to a list
- "count"** increment a counter by one
- "callback"** call a specified function

These are covered in section [Reference Guide](#), Reference Guide and section [Option Callbacks](#).

Default values

All of the above examples involve setting some variable (the “destination”) when certain command-line options are seen. What happens if those options are never seen? Since we didn't supply any defaults, they are all set to `None`. This is usually fine, but sometimes you want more control. *optparse* lets you supply a default value for each destination, which is assigned before the command line is parsed.

First, consider the verbose/quiet example. If we want *optparse* to set `verbose` to `True` unless `-q` is seen, then we can do this:

```
parser.add_option("-v", action="store_true", dest="verbose", default=True)
parser.add_option("-q", action="store_false", dest="verbose")
```

Since default values apply to the *destination* rather than to any particular option, and these two options happen to have the same destination, this is exactly equivalent:

```
parser.add_option("-v", action="store_true", dest="verbose")
parser.add_option("-q", action="store_false", dest="verbose", default=True)
```

Consider this:

```
parser.add_option("-v", action="store_true", dest="verbose", default=False)
parser.add_option("-q", action="store_false", dest="verbose", default=True)
```

Again, the default value for `verbose` will be `True`: the last default value supplied for any particular destination is the one that counts.

A clearer way to specify default values is the `set_defaults()` method of `OptionParser`, which you can call at any time before calling `parse_args()`:

```
parser.set_defaults(verbose=True)
parser.add_option(...)
(options, args) = parser.parse_args()
```

As before, the last value specified for a given option destination is the one that counts. For clarity, try to use one method or the other of setting default values, not both.

Generating help

`optparse`'s ability to generate help and usage text automatically is useful for creating user-friendly command-line interfaces. All you have to do is supply a `help` value for each option, and optionally a short usage message for your whole program. Here's an `OptionParser` populated with user-friendly (documented) options:

```
usage = "usage: %prog [options] arg1 arg2"
parser = OptionParser(usage=usage)
parser.add_option("-v", "--verbose",
                  action="store_true", dest="verbose", default=True,
                  help="make lots of noise [default]")
parser.add_option("-q", "--quiet",
                  action="store_false", dest="verbose",
                  help="be vewwy quiet (I'm hunting wabbits)")
parser.add_option("-f", "--filename",
                  metavar="FILE", help="write output to FILE")
parser.add_option("-m", "--mode",
                  default="intermediate",
                  help="interaction mode: novice, intermediate, "
                       "or expert [default: %default]")
```

If `optparse` encounters either `-h` or `--help` on the command-line, or if you just call `parser.print_help()`, it prints the following to standard output:

```
Usage: <yourscript> [options] arg1 arg2

Options:
  -h, --help            show this help message and exit
  -v, --verbose          make lots of noise [default]
  -q, --quiet            be vewwy quiet (I'm hunting wabbits)
  -f FILE, --filename=FILE
                        write output to FILE
  -m MODE, --mode=MODE  interaction mode: novice, intermediate, or
                        expert [default: intermediate]
```

(If the help output is triggered by a help option, `optparse` exits after printing the help text.)

There's a lot going on here to help `optparse` generate the best possible help message:

- the script defines its own usage message:

```
usage = "usage: %prog [options] arg1 arg2"
```

`optparse` expands `%prog` in the usage string to the name of the current program, i.e. `os.path.basename(sys.argv[0])`. The expanded string is then printed before the detailed option help.

If you don't supply a usage string, `optparse` uses a bland but sensible default: `"Usage: %prog [options]"`, which is fine if your script doesn't take any positional arguments.

- every option defines a help string, and doesn't worry about line-wrapping—`optparse` takes care of wrapping lines and making the help output look good.
- options that take a value indicate this fact in their automatically-generated help message, e.g. for the “mode” option:

```
-m MODE, --mode=MODE
```

Here, “MODE” is called the meta-variable: it stands for the argument that the user is expected to supply to `-m/--mode`. By default, `optparse` converts the destination variable name to uppercase and uses that for the meta-variable. Sometimes, that's not what you want—for example, the `--filename` option explicitly sets `metavar="FILE"`, resulting in this automatically-generated option description:

```
-f FILE, --filename=FILE
```

This is important for more than just saving space, though: the manually written help text uses the meta-variable `FILE` to clue the user in that there's a connection between the semi-formal syntax `-f FILE` and the informal semantic description “write output to FILE”. This is a simple but effective way to make your help text a lot clearer and more useful for end users.

- options that have a default value can include `%default` in the help string—`optparse` will replace it with `str()` of the option's default value. If an option has no default value (or the default value is `None`), `%default` expands to `none`.

Grouping Options

When dealing with many options, it is convenient to group these options for better help output. An `OptionParser` can contain several option groups, each of which can contain several options.

An option group is obtained using the class `OptionGroup`:

```
class optparse.OptionGroup (parser, title, description=None)
    where
```

- `parser` is the `OptionParser` instance the group will be inserted in to
- `title` is the group title
- `description`, optional, is a long description of the group

`OptionGroup` inherits from `OptionContainer` (like `OptionParser`) and so the `add_option()` method can be used to add an option to the group.

Once all the options are declared, using the `OptionParser` method `add_option_group()` the group is added to the previously defined parser.

Continuing with the parser defined in the previous section, adding an `OptionGroup` to a parser is easy:

```
group = OptionGroup(parser, "Dangerous Options",
                    "Caution: use these options at your own risk.  "
                    "It is believed that some of them bite.")
```

(下页继续)

(续上页)

```
group.add_option("-g", action="store_true", help="Group option.")
parser.add_option_group(group)
```

This would result in the following help output:

```
Usage: <yourscript> [options] arg1 arg2

Options:
  -h, --help            show this help message and exit
  -v, --verbose          make lots of noise [default]
  -q, --quiet           be vewwy quiet (I'm hunting wabbits)
  -f FILE, --filename=FILE
                        write output to FILE
  -m MODE, --mode=MODE  interaction mode: novice, intermediate, or
                        expert [default: intermediate]

Dangerous Options:
  Caution: use these options at your own risk.  It is believed that some
  of them bite.

  -g                    Group option.
```

A bit more complete example might involve using more than one group: still extending the previous example:

```
group = OptionGroup(parser, "Dangerous Options",
                    "Caution: use these options at your own risk.  "
                    "It is believed that some of them bite.")
group.add_option("-g", action="store_true", help="Group option.")
parser.add_option_group(group)

group = OptionGroup(parser, "Debug Options")
group.add_option("-d", "--debug", action="store_true",
                help="Print debug information")
group.add_option("-s", "--sql", action="store_true",
                help="Print all SQL statements executed")
group.add_option("-e", action="store_true", help="Print every action done")
parser.add_option_group(group)
```

that results in the following output:

```
Usage: <yourscript> [options] arg1 arg2

Options:
  -h, --help            show this help message and exit
  -v, --verbose          make lots of noise [default]
  -q, --quiet           be vewwy quiet (I'm hunting wabbits)
  -f FILE, --filename=FILE
                        write output to FILE
  -m MODE, --mode=MODE  interaction mode: novice, intermediate, or expert
                        [default: intermediate]

Dangerous Options:
  Caution: use these options at your own risk.  It is believed that some
  of them bite.

  -g                    Group option.
```

(下页继续)

(续上页)

```

Debug Options:
  -d, --debug          Print debug information
  -s, --sql            Print all SQL statements executed
  -e                  Print every action done

```

Another interesting method, in particular when working programmatically with option groups is:

`OptionParser.get_option_group(opt_str)`

Return the *OptionGroup* to which the short or long option string *opt_str* (e.g. `'-o'` or `'--option'`) belongs. If there's no such *OptionGroup*, return `None`.

Printing a version string

Similar to the brief usage string, *optparse* can also print a version string for your program. You have to supply the string as the *version* argument to *OptionParser*:

```
parser = OptionParser(usage="%prog [-f] [-q]", version="%prog 1.0")
```

`%prog` is expanded just like it is in usage. Apart from that, *version* can contain anything you like. When you supply it, *optparse* automatically adds a `--version` option to your parser. If it encounters this option on the command line, it expands your version string (by replacing `%prog`), prints it to stdout, and exits.

For example, if your script is called `/usr/bin/foo`:

```

$ /usr/bin/foo --version
foo 1.0

```

The following two methods can be used to print and get the version string:

`OptionParser.print_version(file=None)`

Print the version message for the current program (`self.version`) to *file* (default stdout). As with *print_usage()*, any occurrence of `%prog` in `self.version` is replaced with the name of the current program. Does nothing if `self.version` is empty or undefined.

`OptionParser.get_version()`

Same as *print_version()* but returns the version string instead of printing it.

How optparse handles errors

There are two broad classes of errors that *optparse* has to worry about: programmer errors and user errors. Programmer errors are usually erroneous calls to *OptionParser.add_option()*, e.g. invalid option strings, unknown option attributes, missing option attributes, etc. These are dealt with in the usual way: raise an exception (either *OptionParser.OptionError* or *TypeError*) and let the program crash.

Handling user errors is much more important, since they are guaranteed to happen no matter how stable your code is. *optparse* can automatically detect some user errors, such as bad option arguments (passing `-n 4x` where `-n` takes an integer argument), missing arguments (`-n` at the end of the command line, where `-n` takes an argument of any type). Also, you can call *OptionParser.error()* to signal an application-defined error condition:

```

(options, args) = parser.parse_args()
...
if options.a and options.b:
    parser.error("options -a and -b are mutually exclusive")

```

In either case, *optparse* handles the error the same way: it prints the program's usage message and an error message to standard error and exits with error status 2.

Consider the first example above, where the user passes 4x to an option that takes an integer:

```
$ /usr/bin/foo -n 4x
Usage: foo [options]

foo: error: option -n: invalid integer value: '4x'
```

Or, where the user fails to pass a value at all:

```
$ /usr/bin/foo -n
Usage: foo [options]

foo: error: -n option requires an argument
```

optparse-generated error messages take care always to mention the option involved in the error; be sure to do the same when calling `OptionParser.error()` from your application code.

If *optparse*'s default error-handling behaviour does not suit your needs, you'll need to subclass `OptionParser` and override its `exit()` and/or `error()` methods.

Putting it all together

Here's what *optparse*-based scripts usually look like:

```
from optparse import OptionParser
...
def main():
    usage = "usage: %prog [options] arg"
    parser = OptionParser(usage)
    parser.add_option("-f", "--file", dest="filename",
                      help="read data from FILENAME")
    parser.add_option("-v", "--verbose",
                      action="store_true", dest="verbose")
    parser.add_option("-q", "--quiet",
                      action="store_false", dest="verbose")
    ...
    (options, args) = parser.parse_args()
    if len(args) != 1:
        parser.error("incorrect number of arguments")
    if options.verbose:
        print("reading %s..." % options.filename)
    ...

if __name__ == "__main__":
    main()
```

36.1.3 Reference Guide

Creating the parser

The first step in using *optparse* is to create an `OptionParser` instance.

class `optparse.OptionParser(...)`

The `OptionParser` constructor has no required arguments, but a number of optional keyword arguments. You should always pass them as keyword arguments, i.e. do not rely on the order in which the arguments are declared.

usage (default: `"%prog [options]"`) The usage summary to print when your program is run incorrectly or with a help option. When `optparse` prints the usage string, it expands `%prog` to `os.path.basename(sys.argv[0])` (or to `prog` if you passed that keyword argument). To suppress a usage message, pass the special value `optparse.SUPPRESS_USAGE`.

option_list (default: `[]`) A list of `Option` objects to populate the parser with. The options in `option_list` are added after any options in `standard_option_list` (a class attribute that may be set by `OptionParser` subclasses), but before any version or help options. Deprecated; use `add_option()` after creating the parser instead.

option_class (default: `optparse.Option`) Class to use when adding options to the parser in `add_option()`.

version (default: `None`) A version string to print when the user supplies a version option. If you supply a true value for `version`, `optparse` automatically adds a version option with the single option string `--version`. The substring `%prog` is expanded the same as for `usage`.

conflict_handler (default: `"error"`) Specifies what to do when options with conflicting option strings are added to the parser; see section *Conflicts between options*.

description (default: `None`) A paragraph of text giving a brief overview of your program. `optparse` reformats this paragraph to fit the current terminal width and prints it when the user requests help (after `usage`, but before the list of options).

formatter (default: a new `IndentedHelpFormatter`) An instance of `optparse.HelpFormatter` that will be used for printing help text. `optparse` provides two concrete classes for this purpose: `IndentedHelpFormatter` and `TitledHelpFormatter`.

add_help_option (default: `True`) If true, `optparse` will add a help option (with option strings `-h` and `--help`) to the parser.

prog The string to use when expanding `%prog` in `usage` and `version` instead of `os.path.basename(sys.argv[0])`.

epilog (default: `None`) A paragraph of help text to print after the option help.

Populating the parser

There are several ways to populate the parser with options. The preferred way is by using `OptionParser.add_option()`, as shown in section *教程*. `add_option()` can be called in one of two ways:

- pass it an `Option` instance (as returned by `make_option()`)
- pass it any combination of positional and keyword arguments that are acceptable to `make_option()` (i.e., to the `Option` constructor), and it will create the `Option` instance for you

The other alternative is to pass a list of pre-constructed `Option` instances to the `OptionParser` constructor, as in:

```
option_list = [
    make_option("-f", "--filename",
                action="store", type="string", dest="filename"),
    make_option("-q", "--quiet",
                action="store_false", dest="verbose"),
]
parser = OptionParser(option_list=option_list)
```

(`make_option()` is a factory function for creating `Option` instances; currently it is an alias for the `Option` constructor. A future version of `optparse` may split `Option` into several classes, and `make_option()` will pick the right class to instantiate. Do not instantiate `Option` directly.)

Defining options

Each `Option` instance represents a set of synonymous command-line option strings, e.g. `-f` and `--file`. You can specify any number of short or long option strings, but you must specify at least one overall option string.

The canonical way to create an `Option` instance is with the `add_option()` method of `OptionParser`.

```
OptionParser.add_option(option)
OptionParser.add_option(*opt_str, attr=value, ...)
```

To define an option with only a short option string:

```
parser.add_option("-f", attr=value, ...)
```

And to define an option with only a long option string:

```
parser.add_option("--foo", attr=value, ...)
```

The keyword arguments define attributes of the new `Option` object. The most important option attribute is `action`, and it largely determines which other attributes are relevant or required. If you pass irrelevant option attributes, or fail to pass required ones, `optparse` raises an `OptionError` exception explaining your mistake.

An option's `action` determines what `optparse` does when it encounters this option on the command-line. The standard option actions hard-coded into `optparse` are:

"store" store this option's argument (default)

"store_const" store a constant value

"store_true" store a true value

"store_false" store a false value

"append" append this option's argument to a list

"append_const" append a constant value to a list

"count" increment a counter by one

"callback" call a specified function

"help" print a usage message including all options and the documentation for them

(If you don't supply an action, the default is `"store"`. For this action, you may also supply `type` and `dest` option attributes; see [Standard option actions](#).)

As you can see, most actions involve storing or updating a value somewhere. `optparse` always creates a special object for this, conventionally called `options` (it happens to be an instance of `optparse.Values`). Option arguments (and various other values) are stored as attributes of this object, according to the `dest` (destination) option attribute.

For example, when you call

```
parser.parse_args()
```

one of the first things `optparse` does is create the `options` object:

```
options = Values()
```

If one of the options in this parser is defined with

```
parser.add_option("-f", "--file", action="store", type="string", dest="filename")
```

and the command-line being parsed includes any of the following:

```
-ffoo
-f foo
--file=foo
--file foo
```

then *optparse*, on seeing this option, will do the equivalent of

```
options.filename = "foo"
```

The *type* and *dest* option attributes are almost as important as *action*, but *action* is the only one that makes sense for *all* options.

Option attributes

The following option attributes may be passed as keyword arguments to *OptionParser.add_option()*. If you pass an option attribute that is not relevant to a particular option, or fail to pass a required option attribute, *optparse* raises *OptionError*.

Option.action
(default: "store")

Determines *optparse*'s behaviour when this option is seen on the command line; the available options are documented [here](#).

Option.type
(default: "string")

The argument type expected by this option (e.g., "string" or "int"); the available option types are documented [here](#).

Option.dest
(default: derived from option strings)

If the option's action implies writing or modifying a value somewhere, this tells *optparse* where to write it: *dest* names an attribute of the *options* object that *optparse* builds as it parses the command line.

Option.default
The value to use for this option's destination if the option is not seen on the command line. See also *OptionParser.set_defaults()*.

Option.nargs
(default: 1)

How many arguments of type *type* should be consumed when this option is seen. If > 1, *optparse* will store a tuple of values to *dest*.

Option.const
For actions that store a constant value, the constant value to store.

Option.choices
For options of type "choice", the list of strings the user may choose from.

Option.callback
For options with action "callback", the callable to call when this option is seen. See section [Option Callbacks](#) for detail on the arguments passed to the callable.

`Option.callback_args`

`Option.callback_kwargs`

Additional positional and keyword arguments to pass to `callback` after the four standard callback arguments.

`Option.help`

Help text to print for this option when listing all available options after the user supplies a [help](#) option (such as `--help`). If no help text is supplied, the option will be listed without help text. To hide this option, use the special value `optparse.SUPPRESS_HELP`.

`Option.metavar`

(default: derived from option strings)

Stand-in for the option argument(s) to use when printing help text. See section [教程](#) for an example.

Standard option actions

The various option actions all have slightly different requirements and effects. Most actions have several relevant option attributes which you may specify to guide `optparse`'s behaviour; a few have required attributes, which you must specify for any option using that action.

- "store" [relevant: `type`, `dest`, `nargs`, `choices`]

The option must be followed by an argument, which is converted to a value according to `type` and stored in `dest`. If `nargs > 1`, multiple arguments will be consumed from the command line; all will be converted according to `type` and stored to `dest` as a tuple. See the [Standard option types](#) section.

If `choices` is supplied (a list or tuple of strings), the type defaults to "choice".

If `type` is not supplied, it defaults to "string".

If `dest` is not supplied, `optparse` derives a destination from the first long option string (e.g., `--foo-bar` implies `foo_bar`). If there are no long option strings, `optparse` derives a destination from the first short option string (e.g., `-f` implies `f`).

示例:

```
parser.add_option("-f")
parser.add_option("-p", type="float", nargs=3, dest="point")
```

As it parses the command line

```
-f foo.txt -p 1 -3.5 4 -fbar.txt
```

`optparse` will set

```
options.f = "foo.txt"
options.point = (1.0, -3.5, 4.0)
options.f = "bar.txt"
```

- "store_const" [required: `const`; relevant: `dest`]

The value `const` is stored in `dest`.

示例:

```
parser.add_option("-q", "--quiet",
                  action="store_const", const=0, dest="verbose")
parser.add_option("-v", "--verbose",
                  action="store_const", const=1, dest="verbose")
```

(下页继续)

(续上页)

```
parser.add_option("--noisy",
                  action="store_const", const=2, dest="verbose")
```

If `--noisy` is seen, `optparse` will set

```
options.verbose = 2
```

- "store_true" [relevant: *dest*]

A special case of "store_const" that stores a true value to *dest*.

- "store_false" [relevant: *dest*]

Like "store_true", but stores a false value.

示例:

```
parser.add_option("--clobber", action="store_true", dest="clobber")
parser.add_option("--no-clobber", action="store_false", dest="clobber")
```

- "append" [relevant: *type*, *dest*, *nargs*, *choices*]

The option must be followed by an argument, which is appended to the list in *dest*. If no default value for *dest* is supplied, an empty list is automatically created when `optparse` first encounters this option on the command-line. If *nargs* > 1, multiple arguments are consumed, and a tuple of length *nargs* is appended to *dest*.

The defaults for *type* and *dest* are the same as for the "store" action.

示例:

```
parser.add_option("-t", "--tracks", action="append", type="int")
```

If `-t 3` is seen on the command-line, `optparse` does the equivalent of:

```
options.tracks = []
options.tracks.append(int("3"))
```

If, a little later on, `--tracks=4` is seen, it does:

```
options.tracks.append(int("4"))
```

The append action calls the append method on the current value of the option. This means that any default value specified must have an append method. It also means that if the default value is non-empty, the default elements will be present in the parsed value for the option, with any values from the command line appended after those default values:

```
>>> parser.add_option("--files", action="append", default=['~/mypkg/defaults'])
>>> opts, args = parser.parse_args(['--files', 'overrides.mypkg'])
>>> opts.files
['~/mypkg/defaults', 'overrides.mypkg']
```

- "append_const" [required: *const*; relevant: *dest*]

Like "store_const", but the value *const* is appended to *dest*; as with "append", *dest* defaults to None, and an empty list is automatically created the first time the option is encountered.

- "count" [relevant: *dest*]

Increment the integer stored at *dest*. If no default value is supplied, *dest* is set to zero before being incremented the first time.

示例:

```
parser.add_option("-v", action="count", dest="verbosity")
```

The first time `-v` is seen on the command line, *optparse* does the equivalent of:

```
options.verbosity = 0
options.verbosity += 1
```

Every subsequent occurrence of `-v` results in

```
options.verbosity += 1
```

- "callback" [required: *callback*; relevant: *type*, *nargs*, *callback_args*, *callback_kwargs*]

Call the function specified by *callback*, which is called as

```
func(option, opt_str, value, parser, *args, **kwargs)
```

See section *Option Callbacks* for more detail.

- "help"

Prints a complete help message for all the options in the current option parser. The help message is constructed from the usage string passed to *OptionParser*'s constructor and the *help* string passed to every option.

If no *help* string is supplied for an option, it will still be listed in the help message. To omit an option entirely, use the special value *optparse.SUPPRESS_HELP*.

optparse automatically adds a *help* option to all *OptionParsers*, so you do not normally need to create one.

示例:

```
from optparse import OptionParser, SUPPRESS_HELP

# usually, a help option is added automatically, but that can
# be suppressed using the add_help_option argument
parser = OptionParser(add_help_option=False)

parser.add_option("-h", "--help", action="help")
parser.add_option("-v", action="store_true", dest="verbose",
                  help="Be moderately verbose")
parser.add_option("--file", dest="filename",
                  help="Input file to read data from")
parser.add_option("--secret", help=SUPPRESS_HELP)
```

If *optparse* sees either `-h` or `--help` on the command line, it will print something like the following help message to stdout (assuming `sys.argv[0]` is `"foo.py"`):

```
Usage: foo.py [options]

Options:
  -h, --help            Show this help message and exit
  -v                    Be moderately verbose
  --file=FILENAME       Input file to read data from
```

After printing the help message, *optparse* terminates your process with `sys.exit(0)`.

- "version"

Prints the version number supplied to the `OptionParser` to stdout and exits. The version number is actually formatted and printed by the `print_version()` method of `OptionParser`. Generally only relevant if the `version` argument is supplied to the `OptionParser` constructor. As with `help` options, you will rarely create `version` options, since `optparse` automatically adds them when needed.

Standard option types

`optparse` has five built-in option types: "string", "int", "choice", "float" and "complex". If you need to add new option types, see section [Extending optparse](#).

Arguments to string options are not checked or converted in any way: the text on the command line is stored in the destination (or passed to the callback) as-is.

Integer arguments (type "int") are parsed as follows:

- if the number starts with 0x, it is parsed as a hexadecimal number
- if the number starts with 0, it is parsed as an octal number
- if the number starts with 0b, it is parsed as a binary number
- otherwise, the number is parsed as a decimal number

The conversion is done by calling `int()` with the appropriate base (2, 8, 10, or 16). If this fails, so will `optparse`, although with a more useful error message.

"float" and "complex" option arguments are converted directly with `float()` and `complex()`, with similar error-handling.

"choice" options are a subtype of "string" options. The `choices` option attribute (a sequence of strings) defines the set of allowed option arguments. `optparse.check_choice()` compares user-supplied option arguments against this master list and raises `OptionValueError` if an invalid string is given.

解析参数

The whole point of creating and populating an `OptionParser` is to call its `parse_args()` method:

```
(options, args) = parser.parse_args(args=None, values=None)
```

where the input parameters are

args the list of arguments to process (default: `sys.argv[1:]`)

values an `optparse.Values` object to store option arguments in (default: a new instance of `Values`) – if you give an existing object, the option defaults will not be initialized on it

and the return values are

options the same object that was passed in as `values`, or the `optparse.Values` instance created by `optparse`

args the leftover positional arguments after all options have been processed

The most common usage is to supply neither keyword argument. If you supply `values`, it will be modified with repeated `setattr()` calls (roughly one for every option argument stored to an option destination) and returned by `parse_args()`.

If `parse_args()` encounters any errors in the argument list, it calls the `OptionParser`'s `error()` method with an appropriate end-user error message. This ultimately terminates your process with an exit status of 2 (the traditional Unix exit status for command-line errors).

Querying and manipulating your option parser

The default behavior of the option parser can be customized slightly, and you can also poke around your option parser and see what's there. `OptionParser` provides several methods to help you out:

`OptionParser.disable_interspersed_args()`

Set parsing to stop on the first non-option. For example, if `-a` and `-b` are both simple options that take no arguments, `optparse` normally accepts this syntax:

```
prog -a arg1 -b arg2
```

and treats it as equivalent to

```
prog -a -b arg1 arg2
```

To disable this feature, call `disable_interspersed_args()`. This restores traditional Unix syntax, where option parsing stops with the first non-option argument.

Use this if you have a command processor which runs another command which has options of its own and you want to make sure these options don't get confused. For example, each command might have a different set of options.

`OptionParser.enable_interspersed_args()`

Set parsing to not stop on the first non-option, allowing interspersing switches with command arguments. This is the default behavior.

`OptionParser.get_option(opt_str)`

Returns the `Option` instance with the option string `opt_str`, or `None` if no options have that option string.

`OptionParser.has_option(opt_str)`

Return true if the `OptionParser` has an option with option string `opt_str` (e.g., `-q` or `--verbose`).

`OptionParser.remove_option(opt_str)`

If the `OptionParser` has an option corresponding to `opt_str`, that option is removed. If that option provided any other option strings, all of those option strings become invalid. If `opt_str` does not occur in any option belonging to this `OptionParser`, raises `ValueError`.

Conflicts between options

If you're not careful, it's easy to define options with conflicting option strings:

```
parser.add_option("-n", "--dry-run", ...)
...
parser.add_option("-n", "--noisy", ...)
```

(This is particularly true if you've defined your own `OptionParser` subclass with some standard options.)

Every time you add an option, `optparse` checks for conflicts with existing options. If it finds any, it invokes the current conflict-handling mechanism. You can set the conflict-handling mechanism either in the constructor:

```
parser = OptionParser(..., conflict_handler=handler)
```

or with a separate call:

```
parser.set_conflict_handler(handler)
```

The available conflict handlers are:

"error" (default) assume option conflicts are a programming error and raise `OptionConflictError`

"resolve" resolve option conflicts intelligently (see below)

As an example, let's define an `OptionParser` that resolves conflicts intelligently and add conflicting options to it:

```
parser = OptionParser(conflict_handler="resolve")
parser.add_option("-n", "--dry-run", ..., help="do no harm")
parser.add_option("-n", "--noisy", ..., help="be noisy")
```

At this point, `optparse` detects that a previously-added option is already using the `-n` option string. Since `conflict_handler` is `"resolve"`, it resolves the situation by removing `-n` from the earlier option's list of option strings. Now `--dry-run` is the only way for the user to activate that option. If the user asks for help, the help message will reflect that:

```
Options:
  --dry-run      do no harm
  ...
  -n, --noisy    be noisy
```

It's possible to whittle away the option strings for a previously-added option until there are none left, and the user has no way of invoking that option from the command-line. In that case, `optparse` removes that option completely, so it doesn't show up in help text or anywhere else. Carrying on with our existing `OptionParser`:

```
parser.add_option("--dry-run", ..., help="new dry-run option")
```

At this point, the original `-n/--dry-run` option is no longer accessible, so `optparse` removes it, leaving this help text:

```
Options:
  ...
  -n, --noisy    be noisy
  --dry-run      new dry-run option
```

Cleanup

`OptionParser` instances have several cyclic references. This should not be a problem for Python's garbage collector, but you may wish to break the cyclic references explicitly by calling `destroy()` on your `OptionParser` once you are done with it. This is particularly useful in long-running applications where large object graphs are reachable from your `OptionParser`.

Other methods

`OptionParser` supports several other public methods:

`OptionParser.set_usage(usage)`

Set the usage string according to the rules described above for the `usage` constructor keyword argument. Passing `None` sets the default usage string; use `optparse.SUPPRESS_USAGE` to suppress a usage message.

`OptionParser.print_usage(file=None)`

Print the usage message for the current program (`self.usage`) to *file* (default `stdout`). Any occurrence of the string `%prog` in `self.usage` is replaced with the name of the current program. Does nothing if `self.usage` is empty or not defined.

`OptionParser.get_usage()`

Same as `print_usage()` but returns the usage string instead of printing it.

`OptionParser.set_defaults(dest=value, ...)`

Set default values for several option destinations at once. Using `set_defaults()` is the preferred way to set default values for options, since multiple options can share the same destination. For example, if several “mode” options all set the same destination, any one of them can set the default, and the last one wins:

```
parser.add_option("--advanced", action="store_const",
                  dest="mode", const="advanced",
                  default="novice")    # overridden below
parser.add_option("--novice", action="store_const",
                  dest="mode", const="novice",
                  default="advanced")  # overrides above setting
```

To avoid this confusion, use `set_defaults()`:

```
parser.set_defaults(mode="advanced")
parser.add_option("--advanced", action="store_const",
                  dest="mode", const="advanced")
parser.add_option("--novice", action="store_const",
                  dest="mode", const="novice")
```

36.1.4 Option Callbacks

When `optparse`’s built-in actions and types aren’t quite enough for your needs, you have two choices: extend `optparse` or define a callback option. Extending `optparse` is more general, but overkill for a lot of simple cases. Quite often a simple callback is all you need.

There are two steps to defining a callback option:

- define the option itself using the “callback” action
- write the callback; this is a function (or method) that takes at least four arguments, as described below

Defining a callback option

As always, the easiest way to define a callback option is by using the `OptionParser.add_option()` method. Apart from `action`, the only option attribute you must specify is `callback`, the function to call:

```
parser.add_option("-c", action="callback", callback=my_callback)
```

`callback` is a function (or other callable object), so you must have already defined `my_callback()` when you create this callback option. In this simple case, `optparse` doesn’t even know if `-c` takes any arguments, which usually means that the option takes no arguments—the mere presence of `-c` on the command-line is all it needs to know. In some circumstances, though, you might want your callback to consume an arbitrary number of command-line arguments. This is where writing callbacks gets tricky; it’s covered later in this section.

`optparse` always passes four particular arguments to your callback, and it will only pass additional arguments if you specify them via `callback_args` and `callback_kwargs`. Thus, the minimal callback function signature is:

```
def my_callback(option, opt, value, parser):
```

The four arguments to a callback are described below.

There are several other option attributes that you can supply when you define a callback option:

type has its usual meaning: as with the “store” or “append” actions, it instructs `optparse` to consume one argument and convert it to `type`. Rather than storing the converted value(s) anywhere, though, `optparse` passes it to your callback function.

nargs also has its usual meaning: if it is supplied and > 1 , *optparse* will consume *nargs* arguments, each of which must be convertible to *type*. It then passes a tuple of converted values to your callback.

callback_args a tuple of extra positional arguments to pass to the callback

callback_kwargs a dictionary of extra keyword arguments to pass to the callback

How callbacks are called

All callbacks are called as follows:

```
func(option, opt_str, value, parser, *args, **kwargs)
```

where

option is the Option instance that's calling the callback

opt_str is the option string seen on the command-line that's triggering the callback. (If an abbreviated long option was used, *opt_str* will be the full, canonical option string—e.g. if the user puts `--foo` on the command-line as an abbreviation for `--foobar`, then *opt_str* will be `"--foobar"`.)

value is the argument to this option seen on the command-line. *optparse* will only expect an argument if *type* is set; the type of *value* will be the type implied by the option's type. If *type* for this option is `None` (no argument expected), then *value* will be `None`. If *nargs* > 1 , *value* will be a tuple of values of the appropriate type.

parser is the OptionParser instance driving the whole thing, mainly useful because you can access some other interesting data through its instance attributes:

parser.largs the current list of leftover arguments, ie. arguments that have been consumed but are neither options nor option arguments. Feel free to modify *parser.largs*, e.g. by adding more arguments to it. (This list will become *args*, the second return value of *parse_args()*.)

parser.rargs the current list of remaining arguments, ie. with *opt_str* and *value* (if applicable) removed, and only the arguments following them still there. Feel free to modify *parser.rargs*, e.g. by consuming more arguments.

parser.values the object where option values are by default stored (an instance of *optparse.OptionValues*). This lets callbacks use the same mechanism as the rest of *optparse* for storing option values; you don't need to mess around with globals or closures. You can also access or modify the value(s) of any options already encountered on the command-line.

args is a tuple of arbitrary positional arguments supplied via the *callback_args* option attribute.

kwargs is a dictionary of arbitrary keyword arguments supplied via *callback_kwargs*.

Raising errors in a callback

The callback function should raise *OptionValueError* if there are any problems with the option or its argument(s). *optparse* catches this and terminates the program, printing the error message you supply to *stderr*. Your message should be clear, concise, accurate, and mention the option at fault. Otherwise, the user will have a hard time figuring out what they did wrong.

Callback example 1: trivial callback

Here's an example of a callback option that takes no arguments, and simply records that the option was seen:

```
def record_foo_seen(option, opt_str, value, parser):
    parser.values.saw_foo = True

parser.add_option("--foo", action="callback", callback=record_foo_seen)
```

Of course, you could do that with the "store_true" action.

Callback example 2: check option order

Here's a slightly more interesting example: record the fact that -a is seen, but blow up if it comes after -b in the command-line.

```
def check_order(option, opt_str, value, parser):
    if parser.values.b:
        raise OptionValueError("can't use -a after -b")
    parser.values.a = 1
...
parser.add_option("-a", action="callback", callback=check_order)
parser.add_option("-b", action="store_true", dest="b")
```

Callback example 3: check option order (generalized)

If you want to re-use this callback for several similar options (set a flag, but blow up if -b has already been seen), it needs a bit of work: the error message and the flag that it sets must be generalized.

```
def check_order(option, opt_str, value, parser):
    if parser.values.b:
        raise OptionValueError("can't use %s after -b" % opt_str)
    setattr(parser.values, option.dest, 1)
...
parser.add_option("-a", action="callback", callback=check_order, dest='a')
parser.add_option("-b", action="store_true", dest="b")
parser.add_option("-c", action="callback", callback=check_order, dest='c')
```

Callback example 4: check arbitrary condition

Of course, you could put any condition in there—you're not limited to checking the values of already-defined options. For example, if you have options that should not be called when the moon is full, all you have to do is this:

```
def check_moon(option, opt_str, value, parser):
    if is_moon_full():
        raise OptionValueError("%s option invalid when moon is full"
                                % opt_str)
    setattr(parser.values, option.dest, 1)
...
parser.add_option("--foo",
                  action="callback", callback=check_moon, dest="foo")
```

(The definition of `is_moon_full()` is left as an exercise for the reader.)

Callback example 5: fixed arguments

Things get slightly more interesting when you define callback options that take a fixed number of arguments. Specifying that a callback option takes arguments is similar to defining a "store" or "append" option: if you define `type`, then the option takes one argument that must be convertible to that type; if you further define `nargs`, then the option takes `nargs` arguments.

Here's an example that just emulates the standard "store" action:

```
def store_value(option, opt_str, value, parser):
    setattr(parser.values, option.dest, value)
...
parser.add_option("--foo",
                  action="callback", callback=store_value,
                  type="int", nargs=3, dest="foo")
```

Note that `optparse` takes care of consuming 3 arguments and converting them to integers for you; all you have to do is store them. (Or whatever; obviously you don't need a callback for this example.)

Callback example 6: variable arguments

Things get hairy when you want an option to take a variable number of arguments. For this case, you must write a callback, as `optparse` doesn't provide any built-in capabilities for it. And you have to deal with certain intricacies of conventional Unix command-line parsing that `optparse` normally handles for you. In particular, callbacks should implement the conventional rules for bare `--` and `-` arguments:

- either `--` or `-` can be option arguments
- bare `--` (if not the argument to some option): halt command-line processing and discard the `--`
- bare `-` (if not the argument to some option): halt command-line processing but keep the `-` (append it to `parser.largs`)

If you want an option that takes a variable number of arguments, there are several subtle, tricky issues to worry about. The exact implementation you choose will be based on which trade-offs you're willing to make for your application (which is why `optparse` doesn't support this sort of thing directly).

Nevertheless, here's a stab at a callback for an option with variable arguments:

```
def vararg_callback(option, opt_str, value, parser):
    assert value is None
    value = []

    def floatable(str):
        try:
            float(str)
            return True
        except ValueError:
            return False

    for arg in parser.rargs:
        # stop on --foo like options
        if arg[:2] == "--" and len(arg) > 2:
            break
        # stop on -a, but not on -3 or -3.0
        if arg[:1] == "-" and len(arg) > 1 and not floatable(arg):
            break
        value.append(arg)
```

(下页继续)

(续上页)

```

del parser.rargs[:len(value)]
setattr(parser.values, option.dest, value)

...
parser.add_option("-c", "--callback", dest="vararg_attr",
                  action="callback", callback=vararg_callback)

```

36.1.5 Extending optparse

Since the two major controlling factors in how *optparse* interprets command-line options are the action and type of each option, the most likely direction of extension is to add new actions and new types.

Adding new types

To add new types, you need to define your own subclass of *optparse*'s *Option* class. This class has a couple of attributes that define *optparse*'s types: *TYPES* and *TYPE_CHECKER*.

Option.*TYPES*

A tuple of type names; in your subclass, simply define a new tuple *TYPES* that builds on the standard one.

Option.*TYPE_CHECKER*

A dictionary mapping type names to type-checking functions. A type-checking function has the following signature:

```
def check_mytype(option, opt, value)
```

where *option* is an *Option* instance, *opt* is an option string (e.g., -f), and *value* is the string from the command line that must be checked and converted to your desired type. *check_mytype()* should return an object of the hypothetical type *mytype*. The value returned by a type-checking function will wind up in the *OptionValues* instance returned by *OptionParser.parse_args()*, or be passed to a callback as the *value* parameter.

Your type-checking function should raise *OptionValueError* if it encounters any problems. *OptionValueError* takes a single string argument, which is passed as-is to *OptionParser*'s *error()* method, which in turn prepends the program name and the string "error:" and prints everything to *stderr* before terminating the process.

Here's a silly example that demonstrates adding a "complex" option type to parse Python-style complex numbers on the command line. (This is even sillier than it used to be, because *optparse* 1.3 added built-in support for complex numbers, but never mind.)

First, the necessary imports:

```

from copy import copy
from optparse import Option, OptionValueError

```

You need to define your type-checker first, since it's referred to later (in the *TYPE_CHECKER* class attribute of your *Option* subclass):

```

def check_complex(option, opt, value):
    try:
        return complex(value)
    except ValueError:
        raise OptionValueError(
            "option %s: invalid complex value: %r" % (opt, value))

```

Finally, the Option subclass:

```
class MyOption (Option):
    TYPES = Option.TYPES + ("complex",)
    TYPE_CHECKER = copy(Option.TYPE_CHECKER)
    TYPE_CHECKER["complex"] = check_complex
```

(If we didn't make a `copy()` of `Option.TYPE_CHECKER`, we would end up modifying the `TYPE_CHECKER` attribute of `optparse`'s Option class. This being Python, nothing stops you from doing that except good manners and common sense.)

That's it! Now you can write a script that uses the new option type just like any other `optparse`-based script, except you have to instruct your OptionParser to use MyOption instead of Option:

```
parser = OptionParser(option_class=MyOption)
parser.add_option("-c", type="complex")
```

Alternately, you can build your own option list and pass it to OptionParser; if you don't use `add_option()` in the above way, you don't need to tell OptionParser which option class to use:

```
option_list = [MyOption("-c", action="store", type="complex", dest="c")]
parser = OptionParser(option_list=option_list)
```

Adding new actions

Adding new actions is a bit trickier, because you have to understand that `optparse` has a couple of classifications for actions:

“store” actions actions that result in `optparse` storing a value to an attribute of the current OptionValues instance; these options require a `dest` attribute to be supplied to the Option constructor.

“typed” actions actions that take a value from the command line and expect it to be of a certain type; or rather, a string that can be converted to a certain type. These options require a `type` attribute to the Option constructor.

These are overlapping sets: some default “store” actions are "store", "store_const", "append", and "count", while the default “typed” actions are "store", "append", and "callback".

When you add an action, you need to categorize it by listing it in at least one of the following class attributes of Option (all are lists of strings):

`Option.ACTIONS`

All actions must be listed in ACTIONS.

`Option.STORE_ACTIONS`

“store” actions are additionally listed here.

`Option.TYPED_ACTIONS`

“typed” actions are additionally listed here.

`Option.ALWAYS_TYPED_ACTIONS`

Actions that always take a type (i.e. whose options always take a value) are additionally listed here. The only effect of this is that `optparse` assigns the default type, "string", to options with no explicit type whose action is listed in `ALWAYS_TYPED_ACTIONS`.

In order to actually implement your new action, you must override Option's `take_action()` method and add a case that recognizes your action.

For example, let's add an "extend" action. This is similar to the standard "append" action, but instead of taking a single value from the command-line and appending it to an existing list, "extend" will take multiple values in a single

comma-delimited string, and extend an existing list with them. That is, if `--names` is an "extend" option of type "string", the command line

```
--names=foo,bar --names blah --names ding,dong
```

would result in a list

```
["foo", "bar", "blah", "ding", "dong"]
```

Again we define a subclass of `Option`:

```
class MyOption(Option):

    ACTIONS = Option.ACTIONS + ("extend",)
    STORE_ACTIONS = Option.STORE_ACTIONS + ("extend",)
    TYPED_ACTIONS = Option.TYPED_ACTIONS + ("extend",)
    ALWAYS_TYPED_ACTIONS = Option.ALWAYS_TYPED_ACTIONS + ("extend",)

    def take_action(self, action, dest, opt, value, values, parser):
        if action == "extend":
            lvalue = value.split(",")
            values.ensure_value(dest, []).extend(lvalue)
        else:
            Option.take_action(
                self, action, dest, opt, value, values, parser)
```

Features of note:

- "extend" both expects a value on the command-line and stores that value somewhere, so it goes in both `STORE_ACTIONS` and `TYPED_ACTIONS`.
- to ensure that `optparse` assigns the default type of "string" to "extend" actions, we put the "extend" action in `ALWAYS_TYPED_ACTIONS` as well.
- `MyOption.take_action()` implements just this one new action, and passes control back to `Option.take_action()` for the standard `optparse` actions.
- `values` is an instance of the `optparse_parser.Values` class, which provides the very useful `ensure_value()` method. `ensure_value()` is essentially `getattr()` with a safety valve; it is called as

```
values.ensure_value(attr, value)
```

If the `attr` attribute of `values` doesn't exist or is `None`, then `ensure_value()` first sets it to `value`, and then returns `value`. This is very handy for actions like "extend", "append", and "count", all of which accumulate data in a variable and expect that variable to be of a certain type (a list for the first two, an integer for the latter). Using `ensure_value()` means that scripts using your action don't have to worry about setting a default value for the option destinations in question; they can just leave the default as `None` and `ensure_value()` will take care of getting it right when it's needed.

36.2 `imp` —Access the import internals

Source code: [Lib/imp.py](#)

3.4 版后已移除: The `imp` package is pending deprecation in favor of `importlib`.

This module provides an interface to the mechanisms used to implement the `import` statement. It defines the following constants and functions:

`imp.get_magic()`

Return the magic string value used to recognize byte-compiled code files (`.pyc` files). (This value may be different for each Python version.)

3.4 版后已移除: Use `importlib.util.MAGIC_NUMBER` instead.

`imp.get_suffixes()`

Return a list of 3-element tuples, each describing a particular type of module. Each triple has the form `(suffix, mode, type)`, where `suffix` is a string to be appended to the module name to form the filename to search for, `mode` is the mode string to pass to the built-in `open()` function to open the file (this can be `'r'` for text files or `'rb'` for binary files), and `type` is the file type, which has one of the values `PY_SOURCE`, `PY_COMPILED`, or `C_EXTENSION`, described below.

3.3 版后已移除: Use the constants defined on `importlib.machinery` instead.

`imp.find_module(name[, path])`

Try to find the module `name`. If `path` is omitted or `None`, the list of directory names given by `sys.path` is searched, but first a few special places are searched: the function tries to find a built-in module with the given name (`C_BUILTIN`), then a frozen module (`PY_FROZEN`), and on some systems some other places are looked in as well (on Windows, it looks in the registry which may point to a specific file).

Otherwise, `path` must be a list of directory names; each directory is searched for files with any of the suffixes returned by `get_suffixes()` above. Invalid names in the list are silently ignored (but all list items must be strings).

If search is successful, the return value is a 3-element tuple `(file, pathname, description)`:

`file` is an open *file object* positioned at the beginning, `pathname` is the pathname of the file found, and `description` is a 3-element tuple as contained in the list returned by `get_suffixes()` describing the kind of module found.

If the module does not live in a file, the returned `file` is `None`, `pathname` is the empty string, and the `description` tuple contains empty strings for its suffix and mode; the module type is indicated as given in parentheses above. If the search is unsuccessful, `ImportError` is raised. Other exceptions indicate problems with the arguments or environment.

If the module is a package, `file` is `None`, `pathname` is the package path and the last item in the `description` tuple is `PKG_DIRECTORY`.

This function does not handle hierarchical module names (names containing dots). In order to find `P.M`, that is, submodule `M` of package `P`, use `find_module()` and `load_module()` to find and load package `P`, and then use `find_module()` with the `path` argument set to `P.__path__`. When `P` itself has a dotted name, apply this recipe recursively.

3.3 版后已移除: Use `importlib.util.find_spec()` instead unless Python 3.3 compatibility is required, in which case use `importlib.find_loader()`. For example usage of the former case, see the 示例 section of the `importlib` documentation.

`imp.load_module(name, file, pathname, description)`

Load a module that was previously found by `find_module()` (or by an otherwise conducted search yielding compatible results). This function does more than importing the module: if the module was already imported, it will reload the module! The `name` argument indicates the full module name (including the package name, if this is a submodule of a package). The `file` argument is an open file, and `pathname` is the corresponding file name; these can be `None` and `' '`, respectively, when the module is a package or not being loaded from a file. The `description` argument is a tuple, as would be returned by `get_suffixes()`, describing what kind of module must be loaded.

If the load is successful, the return value is the module object; otherwise, an exception (usually `ImportError`) is raised.

Important: the caller is responsible for closing the *file* argument, if it was not `None`, even when an exception is raised. This is best done using a `try ... finally` statement.

3.3 版后已移除: If previously used in conjunction with `imp.find_module()` then consider using `importlib.import_module()`, otherwise use the loader returned by the replacement you chose for `imp.find_module()`. If you called `imp.load_module()` and related functions directly with file path arguments then use a combination of `importlib.util.spec_from_file_location()` and `importlib.util.module_from_spec()`. See the 示例 section of the `importlib` documentation for details of the various approaches.

`imp.new_module(name)`

Return a new empty module object called *name*. This object is *not* inserted in `sys.modules`.

3.4 版后已移除: Use `importlib.util.module_from_spec()` instead.

`imp.reload(module)`

Reload a previously imported *module*. The argument must be a module object, so it must have been successfully imported before. This is useful if you have edited the module source file using an external editor and want to try out the new version without leaving the Python interpreter. The return value is the module object (the same as the *module* argument).

When `reload(module)` is executed:

- Python modules' code is recompiled and the module-level code reexecuted, defining a new set of objects which are bound to names in the module's dictionary. The `init` function of extension modules is not called a second time.
- 与 Python 中的所有的其它对象一样，旧的对象只有在引用计数为 0 之后才会被回收。
- 模块命名空间中的名称重新指向任何新的或更改后的对象。
- 其他旧对象的引用（例如那个模块的外部名称）不会被重新绑定到引用的新对象的，并且如果有需要，必须在出现的每个命名空间中进行更新。

有一些其他注意事项：

当一个模块被重新加载的时候，它的字典（包含了那个模块的全局变量）会被保留。名称的重新定义会覆盖旧的定义，所以通常来说这不是问题。如果一个新模块没有定义在旧版本模块中定义的名称，则将保留旧版本中的定义。这一特性可用于作为那个模块的优点，如果它维护一个全局表或者对象的缓存——使用 `try` 语句，就可以测试表的存在并且跳过它的初始化，如果有需要的话：

```
try:
    cache
except NameError:
    cache = {}
```

It is legal though generally not very useful to reload built-in or dynamically loaded modules, except for `sys`, `__main__` and `builtins`. In many cases, however, extension modules are not designed to be initialized more than once, and may fail in arbitrary ways when reloaded.

If a module imports objects from another module using `from ... import ...`, calling `reload()` for the other module does not redefine the objects imported from it—one way around this is to re-execute the `from` statement, another is to use `import` and qualified names (`module.*name*`) instead.

如果一个模块创建一个类的实例，重新加载定义那个类的模块不影响这个实例的方法定义——它们继续使用旧类的定义。对于子类，也是一样的。

在 3.3 版更改: Relies on both `__name__` and `__loader__` being defined on the module being reloaded instead of just `__name__`.

3.4 版后已移除: Use `importlib.reload()` instead.

The following functions are conveniences for handling [PEP 3147](#) byte-compiled file paths.

3.2 新版功能.

`imp.cache_from_source(path, debug_override=None)`

Return the [PEP 3147](#) path to the byte-compiled file associated with the source *path*. For example, if *path* is `/foo/bar/baz.py` the return value would be `/foo/bar/__pycache__/baz.cpython-32.pyc` for Python 3.2. The `cpython-32` string comes from the current magic tag (see `get_tag()`; if `sys.implementation.cache_tag` is not defined then `NotImplementedError` will be raised). By passing in `True` or `False` for *debug_override* you can override the system's value for `__debug__`, leading to optimized bytecode.

path need not exist.

在 3.3 版更改: If `sys.implementation.cache_tag` is `None`, then `NotImplementedError` is raised.

3.4 版后已移除: Use `importlib.util.cache_from_source()` instead.

在 3.5 版更改: The *debug_override* parameter no longer creates a `.pyo` file.

`imp.source_from_cache(path)`

Given the *path* to a [PEP 3147](#) file name, return the associated source code file path. For example, if *path* is `/foo/bar/__pycache__/baz.cpython-32.pyc` the returned path would be `/foo/bar/baz.py`. *path* need not exist, however if it does not conform to [PEP 3147](#) format, a `ValueError` is raised. If `sys.implementation.cache_tag` is not defined, `NotImplementedError` is raised.

在 3.3 版更改: Raise `NotImplementedError` when `sys.implementation.cache_tag` is not defined.

3.4 版后已移除: Use `importlib.util.source_from_cache()` instead.

`imp.get_tag()`

Return the [PEP 3147](#) magic tag string matching this version of Python's magic number, as returned by `get_magic()`.

3.4 版后已移除: Use `sys.implementation.cache_tag` directly starting in Python 3.3.

The following functions help interact with the import system's internal locking mechanism. Locking semantics of imports are an implementation detail which may vary from release to release. However, Python ensures that circular imports work without any deadlocks.

`imp.lock_held()`

Return `True` if the global import lock is currently held, else `False`. On platforms without threads, always return `False`.

On platforms with threads, a thread executing an import first holds a global import lock, then sets up a per-module lock for the rest of the import. This blocks other threads from importing the same module until the original import completes, preventing other threads from seeing incomplete module objects constructed by the original thread. An exception is made for circular imports, which by construction have to expose an incomplete module object at some point.

在 3.3 版更改: The locking scheme has changed to per-module locks for the most part. A global import lock is kept for some critical tasks, such as initializing the per-module locks.

3.4 版后已移除.

`imp.acquire_lock()`

Acquire the interpreter's global import lock for the current thread. This lock should be used by import hooks to ensure thread-safety when importing modules.

Once a thread has acquired the import lock, the same thread may acquire it again without blocking; the thread must release it once for each time it has acquired it.

On platforms without threads, this function does nothing.

在 3.3 版更改: The locking scheme has changed to per-module locks for the most part. A global import lock is kept for some critical tasks, such as initializing the per-module locks.

3.4 版后已移除.

`imp.release_lock()`

Release the interpreter's global import lock. On platforms without threads, this function does nothing.

在 3.3 版更改: The locking scheme has changed to per-module locks for the most part. A global import lock is kept for some critical tasks, such as initializing the per-module locks.

3.4 版后已移除.

The following constants with integer values, defined in this module, are used to indicate the search result of `find_module()`.

`imp.PY_SOURCE`

The module was found as a source file.

3.3 版后已移除.

`imp.PY_COMPILED`

The module was found as a compiled code object file.

3.3 版后已移除.

`imp.C_EXTENSION`

The module was found as dynamically loadable shared library.

3.3 版后已移除.

`imp.PKG_DIRECTORY`

The module was found as a package directory.

3.3 版后已移除.

`imp.C_BUILTIN`

The module was found as a built-in module.

3.3 版后已移除.

`imp.PY_FROZEN`

The module was found as a frozen module.

3.3 版后已移除.

class `imp.NullImporter` (*path_string*)

The *NullImporter* type is a [PEP 302](#) import hook that handles non-directory path strings by failing to find any modules. Calling this type with an existing directory or empty string raises *ImportError*. Otherwise, a *NullImporter* instance is returned.

Instances have only one method:

find_module (*fullname* [, *path*])

This method always returns *None*, indicating that the requested module could not be found.

在 3.3 版更改: *None* is inserted into `sys.path_importer_cache` instead of an instance of *NullImporter*.

3.4 版后已移除: Insert *None* into `sys.path_importer_cache` instead.

36.2.1 示例

The following function emulates what was the standard `import` statement up to Python 1.4 (no hierarchical module names). (This *implementation* wouldn't work in that version, since `find_module()` has been extended and `load_module()` has been added in 1.4.)

```
import imp
import sys

def __import__(name, globals=None, locals=None, fromlist=None):
    # Fast path: see if the module has already been imported.
    try:
        return sys.modules[name]
    except KeyError:
        pass

    # If any of the following calls raises an exception,
    # there's a problem we can't handle -- let the caller handle it.

    fp, pathname, description = imp.find_module(name)

    try:
        return imp.load_module(name, fp, pathname, description)
    finally:
        # Since we may exit via an exception, close fp explicitly.
        if fp:
            fp.close()
```

未创建文档的模块

以下是目前未创建文档模块的速览列表，但其它它们应创建文档。欢迎为他们提供文档！（通过电子邮件发送到 docs@python.org）

本章的想法和原内容取自 Fredrik Lundh 的帖子；本章的具体内容已经大幅修改。

37.1 平台特定模块

这些模块用于实现 `os.path` 模块，除此之外没有文档。几乎没有必要创建这些文档。

ntpath — 在 Win32 和 Win64 平台上实现 `os.path`。

posixpath — 在 POSIX 上实现 `os.path`。

术语对照表

>>> 交互式终端中默认的 Python 提示符。往往会显示于能以交互方式在解释器里执行的样例代码之前。

... The default Python prompt of the interactive shell when entering code for an indented code block, when within a pair of matching left and right delimiters (parentheses, square brackets, curly braces or triple quotes), or after specifying a decorator.

2to3 一个将 Python 2.x 代码转换为 Python 3.x 代码的工具，能够处理大部分通过解析源码并遍历解析树可检测到的不兼容问题。

2to3 包含在标准库中，模块名为 `lib2to3`；并提供一个独立入口点 `Tools/scripts/2to3`。参见 [2to3](#) - 自动将 Python 2 代码转为 Python 3 代码。

abstract base class – 抽象基类 抽象基类简称 ABC，是对 [duck-typing](#) 的补充，它提供了一种定义接口的新方式，相比之下其他技巧例如 `hasattr()` 显得过于笨拙或有微妙错误（例如使用 魔术方法）。ABC 引入了虚拟子类，这种类并非继承自其他类，但却仍能被 `isinstance()` 和 `issubclass()` 所认可；详见 [abc](#) 模块文档。Python 自带许多内置的 ABC 用于实现数据结构（在 `collections.abc` 模块中）、数字（在 `numbers` 模块中）、流（在 `io` 模块中）、导入查找器和加载器（在 `importlib.abc` 模块中）。你可以使用 `abc` 模块来创建自己的 ABC。

annotation – 标注 关联到某个变量、类属性、函数形参或返回值的标签，被约定作为 [type hint](#) 来使用。

局部变量的标注在运行时不可访问，但全局变量、类属性和函数的标注会分别存放模块、类和函数的 `__annotations__` 特殊属性中。

参见 [variable annotation](#)、[function annotation](#)、[PEP 484](#) 和 [PEP 526](#)，对此功能均有介绍。

argument – 参数 在调用函数时传给 [function](#)（或 [method](#)）的值。参数分为两种：

- 关键字参数: 在函数调用中前面带有标识符（例如 `name=`）或者作为包含在前面带有 `**` 的字典里的值传入。举例来说，3 和 5 在以下对 `complex()` 的调用中均属于关键字参数：

```
complex(real=3, imag=5)
complex(**{'real': 3, 'imag': 5})
```

- 位置参数: 不属于关键字参数的参数。位置参数可出现于参数列表的开头以及/或者作为前面带有 `*` 的 [iterable](#) 里的元素被传入。举例来说，3 和 5 在以下调用中均属于位置参数：

```
complex(3, 5)
complex(*(3, 5))
```

参数会被赋值给函数体中对应的局部变量。有关赋值规则参见 [calls](#) 一节。根据语法，任何表达式都可用来表示一个参数；最终算出的值会被赋给对应的局部变量。

另参见 [parameter](#) 术语表条目，常见问题中 参数与形参的区别以及 [PEP 362](#)。

asynchronous context manager – 异步上下文管理器 此种对象通过定义 `__aenter__()` 和 `__aexit__()` 方法来对 `async with` 语句中的环境进行控制。由 [PEP 492](#) 引入。

asynchronous generator – 异步生成器 返回值为 [asynchronous generator iterator](#) 的函数。它与使用 `async def` 定义的协程函数很相似，不同之处在于它包含 `yield` 表达式以产生一系列可在 `async for` 循环中使用的值。

此术语通常是指异步生成器函数，但在某些情况下则可能是指 异步生成器迭代器。如果需要清楚表达具体含义，请使用全称以避免歧义。

一个异步生成器函数可能包含 `await` 表达式或者 `async for` 以及 `async with` 语句。

asynchronous generator iterator – 异步生成器迭代器 [asynchronous generator](#) 函数所创建的对象。

此对象属于 [asynchronous iterator](#)，当使用 `__anext__()` 方法调用时会返回一个可等待对象来执行异步生成器函数的代码直到下一个 `yield` 表达式。

每个 `yield` 会临时暂停处理，记住当前位置执行状态（包括局部变量和挂起的 `try` 语句）。当该 异步生成器迭代器与其他 `__anext__()` 返回的可等待对象有效恢复时，它会从离开位置继续执行。参见 [PEP 492](#) 和 [PEP 525](#)。

asynchronous iterable – 异步可迭代对象 可在 `async for` 语句中被使用的对象。必须通过它的 `__aiter__()` 方法返回一个 [asynchronous iterator](#)。由 [PEP 492](#) 引入。

asynchronous iterator – 异步迭代器 实现了 `__aiter__()` 和 `__anext__()` 方法的对象。`__anext__` 必须返回一个 [awaitable](#) 对象。`async for` 会处理异步迭代器的 `__anext__()` 方法所返回的可等待对象，直到其引发一个 [StopAsyncIteration](#) 异常。由 [PEP 492](#) 引入。

attribute – 属性 关联到一个对象的值，可以使用点号表达式通过其名称来引用。例如，如果一个对象 `o` 具有一个属性 `a`，就可以用 `o.a` 来引用它。

awaitable – 可等待对象 能在 `await` 表达式中使用的对象。可以是 [coroutine](#) 或是具有 `__await__()` 方法的对象。参见 [PEP 492](#)。

BDFL Benevolent Dictator For Life, a.k.a. [Guido van Rossum](#), Python’ s creator.

binary file – 二进制文件 [file object](#) 能够读写字节类对象。二进制文件的例子包括以二进制模式（`'rb'`, `'wb'` 或 `'rb+'`）打开的文件、`sys.stdin.buffer`、`sys.stdout.buffer` 以及 [io.BytesIO](#) 和 [gzip.GzipFile](#) 的实例。

另请参见 [text file](#) 了解能够读写 `str` 对象的文件对象。

bytes-like object – 字节类对象 支持 `bufferobjects` 并且能导出 [C-contiguous](#) 缓冲的对象。这包括所有 [bytes](#)、[bytearray](#) 和 `array.array` 对象，以及许多普通 [memoryview](#) 对象。字节类对象可在多种二进制数据操作中使用；这些操作包括压缩、保存为二进制文件以及通过套接字发送等。

某些操作需要可变的二进制数据。这种对象在文档中常被称为“可读写字节类对象”。可变缓冲对象的例子包括 [bytearray](#) 以及 [bytearray](#) 的 [memoryview](#)。其他操作要求二进制数据存放于不可变对象（“只读字节类对象”）；这种对象的例子包括 [bytes](#) 以及 [bytes](#) 对象的 [memoryview](#)。

bytecode – 字节码 Python 源代码会被编译为字节码，即 CPython 解释器中表示 Python 程序的内部代码。字节码还会缓存在 `.pyc` 文件中，这样第二次执行同一文件时速度更快（可以免去将源码重新编译为字节码）。这种“中间语言”运行在根据字节码执行相应机器码的 [virtual machine](#) 之上。请注意不同 Python 虚拟机上的字节码不一定通用，也不一定能在不同 Python 版本上兼容。

字节码指令列表可以在 *dis* 模块 的文档中查看。

class – 类 用来创建用户定义对象的模板。类定义通常包含对该类的实例进行操作的方法定义。

class variable – 类变量 在类中定义的变量，并且仅限在类的层级上修改 (而不是在类的实例中修改)。

coercion – 强制类型转换 The implicit conversion of an instance of one type to another during an operation which involves two arguments of the same type. For example, `int(3.15)` converts the floating point number to the integer 3, but in `3+4.5`, each argument is of a different type (one int, one float), and both must be converted to the same type before they can be added or it will raise a `TypeError`. Without coercion, all arguments of even compatible types would have to be normalized to the same value by the programmer, e.g., `float(3)+4.5` rather than just `3+4.5`.

complex number – 复数 对普通实数系统的扩展，其中所有数字都被表示为一个实部和一个虚部的和。虚数是虚数单位 (-1 的平方根) 的实倍数，通常在数学中写为 i ，在工程学中写为 j 。Python 内置了对复数的支持，采用工程学标记方式；虚部带有一个 j 后缀，例如 `3+1j`。如果需要 *math* 模块内对象的对应复数版本，请使用 *cmath*，复数的使用是一个比较高级的数学特性。如果你感觉没有必要，忽略它们也几乎不会有任何问题。

context manager – 上下文管理器 在 `with` 语句中使用，通过定义 `__enter__()` 和 `__exit__()` 方法来控制环境状态的对象。参见 [PEP 343](#)。

contiguous – 连续 一个缓冲如果是 *C* 连续或 *Fortran* 连续就会被认为是连续的。零维缓冲是 *C* 和 *Fortran* 连续的。在一维数组中，所有条目必须在内存中彼此相邻地排列，采用从零开始的递增索引顺序。在多维 *C*-连续数组中，当按内存地址排列时用最后一个索引访问条目时速度最快。但是在 *Fortran* 连续数组中则是用第一个索引最快。

coroutine – 协程 协程是子例程的更一般形式。子例程可以在某一点进入并在另一点退出。协程则可以在许多不同的点上进入、退出和恢复。它们可通过 `async def` 语句来实现。参见 [PEP 492](#)。

coroutine function – 协程函数 返回一个 *coroutine* 对象的函数。协程函数可通过 `async def` 语句来定义，并可能包含 `await`、`async for` 和 `async with` 关键字。这些特性是由 [PEP 492](#) 引入的。

CPython Python 编程语言的规范实现，在 [python.org](#) 上发布。” CPython” 一词用于在必要时将此实现与其他实现例如 Jython 或 IronPython 相区别。

decorator – 装饰器 返回值为另一个函数的函数，通常使用 `@wrapper` 语法形式来进行函数变换。装饰器的常见例子包括 `classmethod()` 和 `staticmethod()`。

装饰器语法只是一种语法糖，以下两个函数定义在语义上完全等价：

```
def f(...):
    ...
f = staticmethod(f)

@staticmethod
def f(...):
    ...
```

同样的概念也适用于类，但通常较少这样使用。有关装饰器的详情可参见 [函数定义](#) 和 [类定义](#) 的文档。

descriptor – 描述器 任何定义了 `__get__()`、`__set__()` 或 `__delete__()` 方法的对象。当一个类属性为描述器时，它的特殊绑定行为就会在属性查找时被触发。通常情况下，使用 `a.b` 来获取、设置或删除一个属性时会在 `a` 的类字典中查找名称为 `b` 的对象，但如果 `b` 是一个描述器，则会调用对应的描述器方法。理解描述器的概念是更深层次理解 Python 的关键，因为这是许多重要特性的基础，包括函数、方法、属性、类方法、静态方法以及对超类的引用等等。

有关描述符的方法的详情可参看 [descriptors](#)。

dictionary – 字典 一个关联数组，其中的任意键都映射到相应的值。键可以是任何具有 `__hash__()` 和 `__eq__()` 方法的对象。在 Perl 语言中称为 `hash`。

dictionary view – 字典视图 从 `dict.keys()`, `dict.values()` 和 `dict.items()` 返回的对象被称为字典视图。它们提供了字典条目的一个动态视图，这意味着当字典改变时，视图也会相应改变。要将字典视图强制转换为真正的列表，可使用 `list(dictview)`。参见字典视图对象。

docstring – 文档字符串 作为类、函数或模块之内的第一个表达式出现的字符串字面值。它在代码执行时会被忽略，但会被解释器识别并放入所在类、函数或模块的 `__doc__` 属性中。由于它可用于代码内省，因此是对象存放文档的规范位置。

duck-typing – 鸭子类型 指一种编程风格，它并不依靠查找对象类型来确定其是否具有正确的接口，而是直接调用或使用其方法或属性（“看起来像鸭子，叫起来也像鸭子，那么肯定就是鸭子。”）由于强调接口而非特定类型，设计良好的代码可通过允许多态替代来提升灵活性。鸭子类型避免使用 `type()` 或 `isinstance()` 检测。（但要注意鸭子类型可以使用抽象基类作为补充。）而往往会采用 `hasattr()` 检测或是 *EAFP* 编程。

EAFP “求原谅比求许可更容易”的英文缩写。这种 Python 常用代码编写风格会假定所需的键或属性存在，并在假定错误时捕获异常。这种简洁快速风格的特点就是大量运用 `try` 和 `except` 语句。于其相对的则是所谓 *LBYL* 风格，常见于 C 等许多其他语言。

expression – 表达式 A piece of syntax which can be evaluated to some value. In other words, an expression is an accumulation of expression elements like literals, names, attribute access, operators or function calls which all return a value. In contrast to many other languages, not all language constructs are expressions. There are also *statements* which cannot be used as expressions, such as `if`. Assignments are also statements, not expressions.

extension module – 扩展模块 以 C 或 C++ 编写的模块，使用 Python 的 C API 来与语言核心以及用户代码进行交互。

f-string – f-字符串 带有 'f' 或 'F' 前缀的字符串字面值通常被称为“f-字符串”即 格式化字符串字面值的简写。参见 [PEP 498](#)。

file object – 文件对象 对外提供面向文件 API 以使用下层资源的对象（带有 `read()` 或 `write()` 这样的方法）。根据其创建方式的不同，文件对象可以处理对真实磁盘文件，对其他类型存储，或是对通讯设备的访问（例如标准输入/输出、内存缓冲区、套接字、管道等等）。文件对象也被称为 文件类对象或 流。

实际上共有三种类别的文件对象：原始二进制文件，缓冲二进制文件以及文本文件。它们的接口定义均在 `io` 模块中。创建文件对象的规范方式是使用 `open()` 函数。

file-like object – 文件类对象 *file object* 的同义词。

finder – 查找器 一种会尝试查找被导入模块的 *loader* 的对象。

从 Python 3.3 起存在两种类型的查找器：元路径查找器 配合 `sys.meta_path` 使用，以及 *path entry finders* 配合 `sys.path_hooks` 使用。

更多详情可参见 [PEP 302](#), [PEP 420](#) 和 [PEP 451](#)。

floor division – 向下取整除法 向下舍入到最接近的整数的数学除法。向下取整除法的运算符是 `//`。例如，表达式 `11 // 4` 的计算结果是 2，而与之相反的是浮点数的真正除法返回 2.75。注意 `(-11) // 4` 会返回 -3 因为这是 -2.75 向下舍入得到的结果。见 [PEP 238](#)。

function – 函数 可以向调用者返回某个值的一组语句。还可以向其传入零个或多个参数并在函数体执行中被使用。另见 *parameter*, *method* 和 *function* 等节。

function annotation – 函数标注 即针对函数形参或返回值的 *annotation*。

函数标注通常用于类型提示：例如以下函数预期接受两个 `int` 参数并预期返回一个 `int` 值：

```
def sum_two_numbers(a: int, b: int) -> int:
    return a + b
```

函数标注语法的详解见 *function* 一节。

请参看 *variable annotation* 和 [PEP 484](#) 对此功能的描述。

__future__ 一种伪模块，可被程序员用来启用与当前解释器不兼容的新语言特性。

通过导入 `__future__` 模块并对其中的变量求值，你可以查看新特性何时首次加入语言以及何时成为默认：

```
>>> import __future__
>>> __future__.division
_Feature((2, 2, 0, 'alpha', 2), (3, 0, 0, 'alpha', 0), 8192)
```

garbage collection – 垃圾回收 释放不再被使用的内存空间的过程。Python 是通过引用计数和一个能够检测和打破循环引用的循环垃圾回收器来执行垃圾回收的。可以使用 `gc` 模块来控制垃圾回收器。

generator – 生成器 返回一个 *generator iterator* 的函数。它看起来很像普通函数，不同点在于其包含 `yield` 表达式以便产生一系列值供给 `for`-循环使用或是通过 `next()` 函数逐一获取。

通常是指生成器函数，但在某些情况下也可能是指生成器迭代器。如果需要清楚表达具体含义，请使用全称以避免歧义。

generator iterator – 生成器迭代器 *generator* 函数所创建的对象。

每个 `yield` 会临时暂停处理，记住当前位置执行状态（包括局部变量和挂起的 `try` 语句）。当该生成器迭代器恢复时，它会从离开位置继续执行（这与每次调用都从新开始的普通函数差别很大）。

generator expression – 生成器表达式 An expression that returns an iterator. It looks like a normal expression followed by a `for` expression defining a loop variable, range, and an optional `if` expression. The combined expression generates values for an enclosing function:

```
>>> sum(i*i for i in range(10))           # sum of squares 0, 1, 4, ... 81
285
```

generic function – 泛型函数 为不同的类型实现相同操作的多个函数所组成的函数。在调用时会由调度算法来确定应该使用哪个实现。

另请参见 *single dispatch* 术语表条目、`functools.singledispatch()` 装饰器以及 **PEP 443**。

GIL 参见 *global interpreter lock*。

global interpreter lock – 全局解释器锁 CPython 解释器所采用的一种机制，它确保同一时刻只有一个线程在执行 Python *bytecode*。此机制通过设置对象模型（包括 *dict* 等重要内置类型）针对并发访问的隐式安全简化了 CPython 实现。给整个解释器加锁使得解释器多线程运行更方便，其代价则是牺牲了在多处理器上的并行性。

不过，某些标准库或第三方库的扩展模块被设计为在执行计算密集型任务如压缩或哈希时释放 GIL。此外，在执行 I/O 操作时也总是会释放 GIL。

创建一个（以更精细粒度来锁定共享数据的）“自由线程”解释器的努力从未获得成功，因为这会牺牲在普通单处理器情况下的性能。据信克服这种性能问题的措施将导致实现变得更复杂，从而更难以维护。

hashable – 可哈希 一个对象的哈希值如果在其生命周期内绝不改变，就被称为可哈希（它需要具有 `__hash__()` 方法），并可以同其他对象进行比较（它需要具有 `__eq__()` 方法）。可哈希对象必须具有相同的哈希值比较结果才会相同。

可哈希性使得对象能够作为字典键或集合成员使用，因为这些数据结构要在内部使用哈希值。

All of Python's immutable built-in objects are hashable; mutable containers (such as lists or dictionaries) are not. Objects which are instances of user-defined classes are hashable by default. They all compare unequal (except with themselves), and their hash value is derived from their `id()`.

IDLE Python 的 IDE，“集成开发与学习环境”的英文缩写。是 Python 标准发行版附带的基本编程器和解释器环境。

immutable – 不可变 具有固定值的对象。不可变对象包括数字、字符串和元组。这样的对象不能被改变。如果必须存储一个不同的值，则必须创建新的对象。它们在需要常量哈希值的地方起着重要作用，例如作为字典中的键。

import path – 导入路径 由多个位置（或路径条目）组成的列表，会被模块的 *path based finder* 用来查找导入目标。在导入时，此位置列表通常来自 `sys.path`，但对次级包来说也可能来自上级包的 `__path__` 属性。

importing – 导入 令一个模块中的 Python 代码能为另一个模块中的 Python 代码所使用的过程。

importer – 导入器 查找并加载模块的对象；此对象既属于 *finder* 又属于 *loader*。

interactive – 交互 Python 带有一个交互式解释器，即你可以在解释器提示符后输入语句和表达式，立即执行并查看其结果。只需不带参数地启动 `python` 命令（也可以在你的计算机开始菜单中选择相应菜单项）。在测试新想法或检验模块和包的时候用这种方式会非常方便（请记得使用 `help(x)`）。

interpreted – 解释型 Python 一是种解释型语言，与之相对的是编译型语言，虽然两者的区别由于字节码编译器的存在而会有所模糊。这意味着源文件可以直接运行而不必显式地创建可执行文件再运行。解释型语言通常具有比编译型语言更短的开发/调试周期，但是其程序往往运行得更慢。参见 *interactive*。

interpreter shutdown – 解释器关闭 当被要求关闭时，Python 解释器将进入一个特殊运行阶段并逐步释放所有已分配资源，例如模块和各种关键内部结构等。它还会多次调用垃圾回收器。这会触发用户定义析构器或弱引用回调中的代码执行。在关闭阶段执行的代码可能会遇到各种异常，因为其所依赖的资源已不再有效（常见的例子有库模块或警告机制等）。

解释器需要关闭的主要原因有 `__main__` 模块或所运行的脚本已完成执行。

iterable – 可迭代对象 能够逐一返回其成员项的对象。可迭代对象的例子包括所有序列类型（例如 *list*、*str* 和 *tuple*）以及某些非序列类型例如 *dict*、文件对象以及定义了 `__iter__()` 方法或是实现了 *Sequence* 语义的 `__getitem__()` 方法的任意自定义类对象。

可迭代对象被可用于 `for` 循环以及许多其他需要一个序列的地方（*zip()*、*map()* ...）。当一个可迭代对象作为参数传给内置函数 *iter()* 时，它会返回该对象的迭代器。这种迭代器适用于对值集合的一次性遍历。在使用可迭代对象时，你通常不需要调用 *iter()* 或者自己处理迭代器对象。`for` 语句会为你自动处理那些操作，创建一个临时的未命名变量用来在循环期间保存迭代器。参见 *iterator*、*sequence* 以及 *generator*。

iterator – 迭代器 用来表示一连串数据流的对象。重复调用迭代器的 `__next__()` 方法（或将其传给内置函数 *next()*）将逐个返回流中的项。当没有数据可用时则将引发 *StopIteration* 异常。到这时迭代器对象中的数据项已耗尽，继续调用其 `__next__()` 方法只会再次引发 *StopIteration* 异常。迭代器必须具有 `__iter__()` 方法用来返回该迭代器对象自身，因此迭代器必定也是可迭代对象，可被用于其他可迭代对象适用的大部分场合。一个显著的例外是那些会多次重复访问迭代项的代码。容器对象（例如 *list*）在你每次向其传入 *iter()* 函数或是在 `for` 循环中使用它时都会产生一个新的迭代器。如果在此情况下你尝试用迭代器则会返回在之前迭代过程中被耗尽的同一迭代器对象，使其看起来就像是一个空容器。

更多信息可查看 *迭代器类型*。

key function – 键函数 键函数或称整理函数，是能够返回用于排序或排位的值的可调用对象。例如，*locale.strxfrm()* 可用于生成一个符合特定区域排序约定的排序键。

Python 中有许多工具都允许用键函数来控制元素的排位或分组方式。其中包括 *min()*、*max()*、*sorted()*、*list.sort()*、*heapq.merge()*、*heapq.nsmallest()*、*heapq.nlargest()* 以及 *itertools.groupby()*。

要创建一个键函数有多种方式。例如，*str.lower()* 方法可以用作忽略大小写排序的键函数。另外，键函数也可通过 `lambda` 表达式来创建，例如 `lambda r: (r[0], r[2])`。还有 *operator* 模块提供了三个键函数构造器：*attrgetter()*、*itemgetter()* 和 *methodcaller()*。请查看 *如何排序* 一节以获取创建和使用键函数的示例。

keyword argument – 关键字参数 参见 *argument*。

lambda 由一个单独 *expression* 构成的匿名内联函数，表达式会在调用时被求值。创建 lambda 函数的句法为 `lambda [parameters]: expression`

LBYL “先查看后跳跃”的英文缩写。这种代码编写风格会在进行调用或查找之前显式地检查前提条件。此风格与 *EAFP* 方式恰成对比，其特点是大量使用 `if` 语句。

在多线程环境中，LBYL 方式会导致“查看”和“跳跃”之间发生条件竞争风险。例如，以下代码 `if key in mapping: return mapping[key]` 可能由于在检查操作之后其他线程从 *mapping* 中移除了 *key* 而出错。这种问题可通过加锁或使用 *EAFP* 方式来解决。

list – 列表 Python 内置的一种 *sequence*。虽然名为列表，但更类似于其他语言中的数组而非链接列表，因为访问元素的时间复杂度为 $O(1)$ 。

list comprehension – 列表推导式 处理一个序列中的所有或部分元素并返回结果列表的一种紧凑写法。`result = ['{:04x}'.format(x) for x in range(256) if x % 2 == 0]` 将生成一个 0 到 255 范围内的十六进制偶数对应字符串 (0x..) 的列表。其中 `if` 子句是可选的，如果省略则 `range(256)` 中的所有元素都会被处理。

loader – 加载器 负责加载模块的对象。它必须定义名为 `load_module()` 的方法。加载器通常由一个 *finder* 返回。详情参见 [PEP 302](#)，对于 *abstract base class* 可参见 `importlib.abc.Loader`。

mapping – 映射 一种支持任意键查找并实现了 *Mapping* 或 *MutableMapping* 抽象基类 中所规定方法的容器对象。此类对象的例子包括 `dict`, `collections.defaultdict`, `collections.OrderedDict` 以及 `collections.Counter`。

meta path finder – 元路径查找器 `sys.meta_path` 的搜索所返回的 *finder*。元路径查找器与 *path entry finders* 存在关联但并不相同。

请查看 `importlib.abc.MetaPathFinder` 了解元路径查找器所实现的方法。

metaclass – 元类 一种用于创建类的类。类定义包含类名、类字典和基类列表。元类负责接受上述三个参数并创建相应的类。大部分面向对象的编程语言都会提供一个默认实现。Python 的特别之处在于可以创建自定义元类。大部分用户永远不需要这个工具，但当需要出现时，元类可提供强大而优雅的解决方案。它们已被用于记录属性访问日志、添加线程安全性、跟踪对象创建、实现单例，以及其他许多任务。

更多详情参见 `metaclasses`。

method – 方法 在类内部定义的函数。如果作为该类的实例的一个属性来调用，方法将会获取实例对象作为其第一个 *argument* (通常命名为 `self`)。参见 *function* 和 *nested scope*。

method resolution order – 方法解析顺序 方法解析顺序就是在查找成员时搜索全部基类所用的先后顺序。请查看 [Python 2.3 方法解析顺序](#) 了解自 2.3 版起 Python 解析器所用相关算法的详情。

module – 模块 此对象是 Python 代码的一种组织单位。各模块具有独立的命名空间，可包含任意 Python 对象。模块可通过 *importing* 操作被加载到 Python 中。

另见 *package*。

module spec – 模块规格 一个命名空间，其中包含用于加载模块的相关导入信息。是 `importlib.machinery.ModuleSpec` 的实例。

MRO 参见 *method resolution order*。

mutable – 可变 可变对象可以在其 `id()` 保持固定的情况下改变其取值。另请参见 *immutable*。

named tuple – 具名元组 任何类似元组的类，其中的可索引元素也能使用名称属性来访问。(例如，`time.localtime()` 会返回一个类似元组的对象，其中的 `year` 既可以通过索引访问如 `t[0]` 也可以通过名称属性访问如 `t.tm_year`)。

具名元组可以是一个内置类型例如 `time.struct_time`，也可以通过正规的类定义来创建。一个完备的具名元组还可以通过工厂函数 `collections.namedtuple()` 来创建。后面这种方式会自动提供一些额外特性，例如 `Employee(name='jones', title='programmer')` 这样的自包含文档表示形式。

namespace – 命名空间 命名空间是存放变量的场所。命名空间有局部、全局和内置的，还有对象中的嵌套命名空间（在方法之内）。命名空间通过防止命名冲突来支持模块化。例如，函数 `builtins.open` 与 `os.open()` 可通过各自的命名空间来区分。命名空间还通过明确哪个模块实现那个函数来帮助提高可读性和可维护性。例如，`random.seed()` 或 `itertools.islice()` 这种写法明确了这些函数是由 `random` 与 `itertools` 模块分别实现的。

namespace package – 命名空间包 PEP 420 所引入的一种仅被用作子包的容器的 `package`，命名空间包可以没有实体表示物，其描述方式与 `regular package` 不同，因为它们没有 `__init__.py` 文件。

另可参见 `module`。

nested scope – 嵌套作用域 在一个定义范围内引用变量的能力。例如，在另一函数之内定义的函数可以引用前者的变量。请注意嵌套作用域默认只对引用有效而对赋值无效。局部变量的读写都受限於最内层作用域。类似的，全局变量的读写则作用于全局命名空间。通过 `nonlocal` 关键字可允许写入外层作用域。

new-style class – 新式类 对于目前已被应用于所有类对象的类形式的旧称谓。在早先的 Python 版本中，只有新式类能够使用 Python 新增的更灵活特性，例如 `__slots__`、描述符、特征属性、`__getattr__()`、类方法和静态方法等。

object – 对象 任何具有状态（属性或值）以及预定义行为（方法）的数据。`object` 也是任何 `new-style class` 的最顶层基类名。

package – 包 一种可包含子模块或递归地包含子包的 Python `module`。从技术上说，包是带有 `__path__` 属性的 Python 模块。

另参见 `regular package` 和 `namespace package`。

parameter – 形参 `function`（或方法）定义中的命名实体，它指定函数可以接受的一个 `argument`（或在某些情况下，多个实参）。有五种形参：

- *positional-or-keyword*：位置或关键字，指定一个可以作为位置参数传入也可以作为关键字参数传入的实参。这是默认的形参类型，例如下面的 `foo` 和 `bar`：

```
def func(foo, bar=None): ...
```

- *positional-only*：仅限位置，指定一个只能按位置传入的参数。Python 中没有定义仅限位置形参的语法。但是一些内置函数有仅限位置形参（比如 `abs()`）。
- *keyword-only*：仅限关键字，指定一个只能通过关键字传入的参数。仅限关键字形参可通过在函数定义的形参列表中包含单个可变位置形参或者在多个可变位置形参之前放一个 `*` 来定义，例如下面的 `kw_only1` 和 `kw_only2`：

```
def func(arg, *, kw_only1, kw_only2): ...
```

- *var-positional*：可变位置，指定可以提供由一个任意数量的位置参数构成的序列（附加在其他形参已接受的位置参数之后）。这种形参可通过在形参名称前加缀 `*` 来定义，例如下面的 `args`：

```
def func(*args, **kwargs): ...
```

- *var-keyword*：可变关键字，指定可以提供任意数量的关键字参数（附加在其他形参已接受的关键字参数之后）。这种形参可通过在形参名称前加缀 `**` 来定义，例如上面的 `kwargs`。

形参可以同时指定可选和必选参数，也可以为某些可选参数指定默认值。

另参见 `argument` 术语表条目、参数与形参的区别中的常见问题、`inspect.Parameter` 类、`function` 一节以及 PEP 362。

path entry – 路径入口 `import path` 中的一个单独位置，会被 `path based finder` 用来查找要导入的模块。

path entry finder – 路径入口查找器 任一可调用对象使用 `sys.path_hooks`（即 `path entry hook`）返回的 `finder`，此种对象能通过 `path entry` 来定位模块。

请参看 `importlib.abc.PathEntryFinder` 以了解路径入口查找器所实现的各个方法。

path entry hook – 路径入口钩子 一种可调用对象，在知道如何查找特定 *path entry* 中的模块的情况下能够使用 `sys.path_hook` 列表返回一个 *path entry finder*。

path based finder – 基于路径的查找器 默认的一种元路径查找器，可在一个 *import path* 中查找模块。

path-like object – 路径类对象 代表一个文件系统路径的对象。类路径对象可以是一个表示路径的 *str* 或者 *bytes* 对象，还可以是一个实现了 `os.PathLike` 协议的对象。一个支持 `os.PathLike` 协议的对象可通过调用 `os.fspath()` 函数转换为 *str* 或者 *bytes* 类型的文件系统路径；`os.fsdecode()` 和 `os.fsencode()` 可被分别用来确保获得 *str* 或 *bytes* 类型的结果。此对象是由 **PEP 519** 引入的。

PEP “Python 增强提议”的英文缩写。一个 PEP 就是一份设计文档，用来向 Python 社区提供信息，或描述一个 Python 的新增特性及其进度或环境。PEP 应当提供精确的技术规格和所提议特性的原理说明。

PEP 应被作为提出主要新特性建议、收集社区对特定问题反馈以及为必须加入 Python 的设计决策编写文档的首选机制。PEP 的作者有责任在社区内部建立共识，并应将不同意见也记入文档。

参见 **PEP 1**。

portion – 部分 构成一个命名空间包的单个目录内文件集合（也可能存放于一个 zip 文件内），具体定义见 **PEP 420**。

positional argument – 位置参数 参见 *argument*。

provisional API – 暂定 API 暂定 API 是指被有意排除在标准库的向后兼容性保证之外的应用编程接口。虽然此类接口通常不会再有重大改变，但只要其被标记为暂定，就可能在核心开发者确定有必要的情况下进行向后不兼容的更改（甚至包括移除该接口）。此种更改并不会随意进行 – 仅在 API 被加入之前未考虑到的严重基础性缺陷被发现时才可能会这样做。

即便是对暂定 API 来说，向后不兼容的更改也会被视为“最后的解决方案”——任何问题被确认时都会尽可能先尝试找到一种向后兼容的解决方案。

这种处理过程允许标准库持续不断地演进，不至于被有问题的长期性设计缺陷所困。详情见 **PEP 411**。

provisional package – 暂定包 参见 *provisional API*。

Python 3000 Python 3.x 发布路线的昵称（这个名字在版本 3 的发布还遥遥无期的时候就已出现了）。有时也被缩写为“Py3k”。

Pythonic 指一个思路或一段代码紧密遵循了 Python 语言最常用的风格和理念，而不是使用其他语言中通用的概念来实现代码。例如，Python 的常用风格是使用 `for` 语句循环来遍历一个可迭代对象中的所有元素。许多其他语言没有这样的结构，因此不熟悉 Python 的人有时会选择使用一个数字计数器：

```
for i in range(len(food)):
    print(food[i])
```

而相应的更简洁更 Pythonic 的方法是这样的：

```
for piece in food:
    print(piece)
```

qualified name – 限定名称 一个以点号分隔的名称，显示从模块的全局作用域到该模块中定义的某个类、函数或方法的“路径”，相关定义见 **PEP 3155**。对于最高层级的函数和类，限定名称与对象名称一致：

```
>>> class C:
...     class D:
...         def meth(self):
...             pass
...
>>> C.__qualname__
'C'
```

(下页继续)

(续上页)

```
>>> C.D.__qualname__
'C.D'
>>> C.D.meth.__qualname__
'C.D.meth'
```

当被用于引用模块时，完整限定名称意为标示该模块的以点号分隔的整个路径，其中包含其所有的父包，例如 `email.mime.text`：

```
>>> import email.mime.text
>>> email.mime.text.__name__
'email.mime.text'
```

reference count – 引用计数 对特定对象的引用的数量。当一个对象的引用计数降为零时，所分配资源将被释放。引用计数对 Python 代码来说通常是不可见的，但它是 CPython 实现的一个关键元素。`sys` 模块定义了一个 `getrefcount()` 函数，程序员可调用它来返回特定对象的引用计数。

regular package – 正规包 传统型的 *package*，例如包含有一个 `__init__.py` 文件的目录。

另参见 *namespace package*。

__slots__ 一种写在类内部的声明，通过预先声明实例属性等对象并移除实例字典来节省内存。虽然这种技巧很流行，但想要用好却并不容易，最好是只保留在少数情况下采用，例如极耗内存的应用程序，并且其中包含大量实例。

sequence – 序列 一种 *iterable*，它支持通过 `__getitem__()` 特殊方法来使用整数索引进行高效的元素访问，并定义了一个返回序列长度的 `__len__()` 方法。内置的序列类型有 *list*、*str*、*tuple* 和 *bytes*。注意虽然 *dict* 也支持 `__getitem__()` 和 `__len__()`，但它被认为属于映射而非序列，因为它查找时使用任意的 *immutable* 键而非整数。

`collections.abc.Sequence` 抽象基类定义了一个更丰富的接口，它超越了 `__getitem__()` 和 `__len__()`，添加了 `count()`、`index()`、`__contains__()` 和 `__reversed__()`。可以使用 `register()` 显式注册实现此扩展接口的类型。

single dispatch – 单分派 一种 *generic function* 分派形式，其实现是基于单个参数的类型来选择的。

slice – 切片 通常只包含了特定 *sequence* 的一部分的对象。切片是通过使用下标标记来创建的，在 `[]` 中给出几个以冒号分隔的数字，例如 `variable_name[1:3:5]`。方括号（下标）标记在内部使用 *slice* 对象。

special method – 特殊方法 一种由 Python 隐式调用的方法，用来对某个类型执行特定操作例如相加等等。这种方法的名称的首尾都为双下划线。特殊方法的文档参见 *specialnames*。

statement – 语句 语句是程序段（一个代码“块”）的组成单位。一条语句可以是一个 *expression* 或某个带有关键字的结构，例如 `if`、`while` 或 `for`。

struct sequence – 结构序列 具有命名元素的元组。结构序列所暴露的接口类似于 *named tuple*，其元素既可通过索引也可作为属性来访问。不过，它们没有任何具名元组的方法，例如 `_make()` 或 `_asdict()`。结构序列的例子包括 `sys.float_info` 以及 `os.stat()` 的返回值。

text encoding – 文本编码 用于将 Unicode 字符串编码为字节串的编码器。

text file – 文本文件 一种能够读写 *str* 对象的 *file object*。通常一个文本文件实际是访问一个面向字节的数据流并自动处理 *text encoding*。文本文件的例子包括以文本模式（`'r'` 或 `'w'`）打开的文件、`sys.stdin`、`sys.stdout` 以及 `io.StringIO` 的实例。

另请参看 *binary file* 了解能够读写字节类对象的文件对象。

triple-quoted string – 三引号字符串 首尾各带三个连续双引号（`"""`）或者单引号（`'`）的字符串。它们在功能上与首尾各用一个引号标注的字符串没有什么不同，但是有多种用处。它们允许你在字符串内包含未经转义的单引号和双引号，并且可以跨越多行而无需使用连接符，在编写文档字符串时特别好用。

type – 类型 类型决定一个 Python 对象属于什么种类；每个对象都具有一种类型。要知道对象的类型，可以访问它的 `__class__` 属性，或是通过 `type(obj)` 来获取。

type alias – 类型别名 一个类型的同义词，创建方式是把类型赋值给特定的标识符。

类型别名的作用是简化类型提示。例如：

```
from typing import List, Tuple

def remove_gray_shades(
    colors: List[Tuple[int, int, int]]) -> List[Tuple[int, int, int]]:
    pass
```

可以这样提高可读性：

```
from typing import List, Tuple

Color = Tuple[int, int, int]

def remove_gray_shades(colors: List[Color]) -> List[Color]:
    pass
```

参见 *typing* 和 **PEP 484**，其中有对此功能的详细描述。

type hint – 类型提示 *annotation* 为变量、类属性、函数的形参或返回值指定预期的类型。

类型提示属于可选项，Python 不要求提供，但其可对静态类型分析工具起作用，并可协助 IDE 实现代码补全与重构。

全局变量、类属性和函数的类型提示可以使用 `typing.get_type_hints()` 来访问，但局部变量则不可以。

参见 *typing* 和 **PEP 484**，其中有对此功能的详细描述。

universal newlines – 通用换行 一种解读文本流的方式，将以下所有符号都识别为行结束标志：Unix 的行结束约定 `'\n'`、Windows 的约定 `'\r\n'` 以及旧版 Macintosh 的约定 `'\r'`。参见 **PEP 278** 和 **PEP 3116** 和 `bytes.splitlines()` 了解更多用法说明。

variable annotation – 变量标注 对变量或类属性的 *annotation*。

在标注变量或类属性时，还可选择为其赋值：

```
class C:
    field: 'annotation'
```

变量标注通常被用作类型提示：例如以下变量预期接受 `int` 类型的值：

```
count: int = 0
```

变量标注语法的详细解释见 *annassign* 一节。

请参看 *function annotation*、**PEP 484** 和 **PEP 526**，其中对此功能有详细描述。

virtual environment – 虚拟环境 一种采用协作式隔离的运行时环境，允许 Python 用户和应用程序在安装和升级 Python 分发包时不会干扰到同一系统上运行的其他 Python 应用程序的行为。

另参见 *venv*。

virtual machine – 虚拟机 一台完全通过软件定义的计算机。Python 虚拟机可执行字节码编译器所生成的 *bytecode*。

Zen of Python – Python 之禅 列出 Python 设计的原则与哲学，有助于理解与使用这种语言。查看其具体内容可在交互模式提示符中输入 `“import this”`。

文档说明

这些文档生成自 [reStructuredText](#) 原文档，由 [Sphinx](#)（一个专门为 Python 文档写的文档生成器）创建。

本文档和它所用工具链的开发完全是由志愿者完成的，这和 Python 本身一样。如果您想参与进来，请阅读 [reporting-bugs](#) 了解如何参与。我们随时欢迎新的志愿者！

非常感谢：

- Fred L. Drake, Jr., 创造了用于早期 Python 文档的工具链，以及撰写了非常多的文档；
- [Docutils](#) 软件包 项目，创建了 [reStructuredText](#) 文本格式和 [Docutils](#) 软件套件；
- Fredrik Lundh, [Sphinx](#) 从他的 [Alternative Python Reference](#) 项目中获得了很多好的想法。

B.1 Python 文档贡献者

有很多对 Python 语言，Python 标准库和 Python 文档有贡献的人，随 Python 源代码发布的 [Misc/ACKS](#) 文件列出了部分贡献者。

有了 Python 社区的输入和贡献，Python 才有了如此出色的文档 - 谢谢你们！

历史和许可证

C.1 软件历史

Python 作为 ABC 语言的继承者，于 20 世纪 90 年代初由荷兰的 Guido van Rossum 在荷兰数学和计算机科学研究学会（CWI，见 <https://www.cwi.nl/>）创建。Guido 仍然是 Python 的主要作者，尽管它包含了许多其他人的贡献。

1995 年，Guido 在弗吉尼亚州的国家创新研究公司（CNRI，见 <https://www.cnri.reston.va.us/>）继续他的 Python 工作，在那里他发布了该软件的几个版本。

2000 年 5 月，Guido 和 Python 核心开发团队转移到 BeOpen.com，组建了 BeOpen PythonLabs 团队。同年 10 月，PythonLabs 团队转向 Digital Creations（现为 Zope Corporation；见 <http://www.zope.com/>）。2001 年，Python 软件基金会（PSF，请参阅 <https://www.python.org/psf/>）成立，这是一个专门为拥有与 Python 相关的知识产权而创建的非营利组织。Zope Corporation 是 PSF 的赞助成员。

所有的 Python 版本都是开源的（有关开源的定义参阅 <https://opensource.org/>）。历史上，绝大多数 Python 版本是 GPL 兼容的；下表总结了各个版本情况。

发布版本	源自	年	所有者	GPL 兼容？
0.9.0 至 1.2	n/a	1991-1995	CWI	yes
1.3 至 1.5.2	1.2	1995-1999	CNRI	yes
1.6	1.5.2	2000	CNRI	no
2.0	1.6	2000	BeOpen.com	no
1.6.1	1.6	2001	CNRI	no
2.1	2.0+1.6.1	2001	PSF	no
2.0.1	2.0+1.6.1	2001	PSF	yes
2.1.1	2.1+2.0.1	2001	PSF	yes
2.1.2	2.1.1	2002	PSF	yes
2.1.3	2.1.2	2002	PSF	yes
2.2 及更高版本	2.1.1	2001 至今	PSF	yes

注解： GPL 兼容并不意味着我们在 GPL 下发布 Python。与 GPL 不同的是，所有 Python 许可都允许您分发修

改后的版本，而不需要对所做的更改进行开源。GPL 兼容的许可证使得可以将 Python 与在 GPL 下发布的其他软件组合在一起；其他则不行。

感谢许多在 Guido 指导下工作的外部志愿者，使这些发布成为可能。

C.2 访问 Python 或以其他方式使用 Python 的条款和条件

C.2.1 Python 3.7.3 的 PSF 许可协议

1. This LICENSE AGREEMENT is between the Python Software Foundation ("PSF"),
→and
the Individual or Organization ("Licensee") accessing and otherwise using
→Python
3.6.8 software in source or binary form and its associated documentation.
2. Subject to the terms and conditions of this License Agreement, PSF hereby
grants Licensee a nonexclusive, royalty-free, world-wide license to
→reproduce,
analyze, test, perform and/or display publicly, prepare derivative works,
distribute, and otherwise use Python 3.6.8 alone or in any derivative
version, provided, however, that PSF's License Agreement and PSF's notice
→of
copyright, i.e., "Copyright © 2001-2019 Python Software Foundation; All
→Rights
Reserved" are retained in Python 3.6.8 alone or in any derivative version
prepared by Licensee.
3. In the event Licensee prepares a derivative work that is based on or
incorporates Python 3.6.8 or any part thereof, and wants to make the
derivative work available to others as provided herein, then Licensee
→hereby
agrees to include in any such work a brief summary of the changes made to
→Python
3.6.8.
4. PSF is making Python 3.6.8 available to Licensee on an "AS IS" basis.
PSF MAKES NO REPRESENTATIONS OR WARRANTIES, EXPRESS OR IMPLIED. BY WAY OF
EXAMPLE, BUT NOT LIMITATION, PSF MAKES NO AND DISCLAIMS ANY REPRESENTATION
→OR
WARRANTY OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE OR THAT
→THE
USE OF PYTHON 3.6.8 WILL NOT INFRINGE ANY THIRD PARTY RIGHTS.
5. PSF SHALL NOT BE LIABLE TO LICENSEE OR ANY OTHER USERS OF PYTHON 3.6.8
FOR ANY INCIDENTAL, SPECIAL, OR CONSEQUENTIAL DAMAGES OR LOSS AS A RESULT
→OF
MODIFYING, DISTRIBUTING, OR OTHERWISE USING PYTHON 3.6.8, OR ANY DERIVATIVE
THEREOF, EVEN IF ADVISED OF THE POSSIBILITY THEREOF.
6. This License Agreement will automatically terminate upon a material breach
→of

its terms and conditions.

7. Nothing in this License Agreement shall be deemed to create any relationship of agency, partnership, or joint venture between PSF and Licensee. This License Agreement does not grant permission to use PSF trademarks or trade name in a trademark sense to endorse or promote products or services of Licensee, or any third party.
8. By copying, installing or otherwise using Python 3.6.8, Licensee agrees to be bound by the terms and conditions of this License Agreement.

C.2.2 Python 2.0 的 BeOpen.com 许可协议

Beopen Python 开源许可协议第 1 版

1. This LICENSE AGREEMENT is between BeOpen.com ("BeOpen"), having an office at 160 Saratoga Avenue, Santa Clara, CA 95051, and the Individual or Organization ("Licensee") accessing and otherwise using this software in source or binary form and its associated documentation ("the Software").
2. Subject to the terms and conditions of this BeOpen Python License Agreement, BeOpen hereby grants Licensee a non-exclusive, royalty-free, world-wide license to reproduce, analyze, test, perform and/or display publicly, prepare derivative works, distribute, and otherwise use the Software alone or in any derivative version, provided, however, that the BeOpen Python License is retained in the Software, alone or in any derivative version prepared by Licensee.
3. BeOpen is making the Software available to Licensee on an "AS IS" basis. BEOPEN MAKES NO REPRESENTATIONS OR WARRANTIES, EXPRESS OR IMPLIED. BY WAY OF EXAMPLE, BUT NOT LIMITATION, BEOPEN MAKES NO AND DISCLAIMS ANY REPRESENTATION OR WARRANTY OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE OR THAT THE USE OF THE SOFTWARE WILL NOT INFRINGE ANY THIRD PARTY RIGHTS.
4. BEOPEN SHALL NOT BE LIABLE TO LICENSEE OR ANY OTHER USERS OF THE SOFTWARE FOR ANY INCIDENTAL, SPECIAL, OR CONSEQUENTIAL DAMAGES OR LOSS AS A RESULT OF USING, MODIFYING OR DISTRIBUTING THE SOFTWARE, OR ANY DERIVATIVE THEREOF, EVEN IF ADVISED OF THE POSSIBILITY THEREOF.
5. This License Agreement will automatically terminate upon a material breach of its terms and conditions.
6. This License Agreement shall be governed by and interpreted in all respects by the law of the State of California, excluding conflict of law provisions. Nothing in this License Agreement shall be deemed to create any relationship of agency, partnership, or joint venture between BeOpen and Licensee. This License Agreement does not grant permission to use BeOpen trademarks or trade names in a trademark sense to endorse or promote products or services of Licensee, or any third party. As an exception, the "BeOpen Python" logos available at <http://www.pythonlabs.com/logos.html> may be used according to the permissions granted on that web page.

(下页继续)

(续上页)

7. By copying, installing or otherwise using the software, Licensee agrees to be bound by the terms and conditions of this License Agreement.

C.2.3 Python 1.6.1 的 CNRI 许可协议

1. This LICENSE AGREEMENT is between the Corporation for National Research Initiatives, having an office at 1895 Preston White Drive, Reston, VA 20191 ("CNRI"), and the Individual or Organization ("Licensee") accessing and otherwise using Python 1.6.1 software in source or binary form and its associated documentation.
2. Subject to the terms and conditions of this License Agreement, CNRI hereby grants Licensee a nonexclusive, royalty-free, world-wide license to reproduce, analyze, test, perform and/or display publicly, prepare derivative works, distribute, and otherwise use Python 1.6.1 alone or in any derivative version, provided, however, that CNRI's License Agreement and CNRI's notice of copyright, i.e., "Copyright © 1995-2001 Corporation for National Research Initiatives; All Rights Reserved" are retained in Python 1.6.1 alone or in any derivative version prepared by Licensee. Alternately, in lieu of CNRI's License Agreement, Licensee may substitute the following text (omitting the quotes): "Python 1.6.1 is made available subject to the terms and conditions in CNRI's License Agreement. This Agreement together with Python 1.6.1 may be located on the Internet using the following unique, persistent identifier (known as a handle): 1895.22/1013. This Agreement may also be obtained from a proxy server on the Internet using the following URL: <http://hdl.handle.net/1895.22/1013>."
3. In the event Licensee prepares a derivative work that is based on or incorporates Python 1.6.1 or any part thereof, and wants to make the derivative work available to others as provided herein, then Licensee hereby agrees to include in any such work a brief summary of the changes made to Python 1.6.1.
4. CNRI is making Python 1.6.1 available to Licensee on an "AS IS" basis. CNRI MAKES NO REPRESENTATIONS OR WARRANTIES, EXPRESS OR IMPLIED. BY WAY OF EXAMPLE, BUT NOT LIMITATION, CNRI MAKES NO AND DISCLAIMS ANY REPRESENTATION OR WARRANTY OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE OR THAT THE USE OF PYTHON 1.6.1 WILL NOT INFRINGE ANY THIRD PARTY RIGHTS.
5. CNRI SHALL NOT BE LIABLE TO LICENSEE OR ANY OTHER USERS OF PYTHON 1.6.1 FOR ANY INCIDENTAL, SPECIAL, OR CONSEQUENTIAL DAMAGES OR LOSS AS A RESULT OF MODIFYING, DISTRIBUTING, OR OTHERWISE USING PYTHON 1.6.1, OR ANY DERIVATIVE THEREOF, EVEN IF ADVISED OF THE POSSIBILITY THEREOF.
6. This License Agreement will automatically terminate upon a material breach of its terms and conditions.
7. This License Agreement shall be governed by the federal intellectual property law of the United States, including without limitation the federal copyright law, and, to the extent such U.S. federal law does not apply, by the law of the Commonwealth of Virginia, excluding Virginia's conflict of law provisions. Notwithstanding the foregoing, with regard to derivative works based on Python 1.6.1 that incorporate non-separable material that was previously distributed under the GNU General Public License (GPL), the law of the Commonwealth of Virginia shall govern this License Agreement only as to issues arising under or with respect to Paragraphs 4, 5, and 7 of this License Agreement. Nothing in

(下页继续)

(续上页)

this License Agreement shall be deemed to create any relationship of agency, partnership, or joint venture between CNRI and Licensee. This License Agreement does not grant permission to use CNRI trademarks or trade name in a trademark sense to endorse or promote products or services of Licensee, or any third party.

8. By clicking on the "ACCEPT" button where indicated, or by copying, installing or otherwise using Python 1.6.1, Licensee agrees to be bound by the terms and conditions of this License Agreement.

C.2.4 Python 0.9.0 至 1.2 的 CWI 许可协议

Copyright © 1991 - 1995, Stichting Mathematisch Centrum Amsterdam, The Netherlands. All rights reserved.

Permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of Stichting Mathematisch Centrum or CWI not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission.

STICHTING MATHEMATISCH CENTRUM DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS, IN NO EVENT SHALL STICHTING MATHEMATISCH CENTRUM BE LIABLE FOR ANY SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

C.3 集成软件的许可和认可

This section is an incomplete, but growing list of licenses and acknowledgements for third-party software incorporated in the Python distribution.

C.3.1 Mersenne Twister

The `_random` module includes code based on a download from <http://www.math.sci.hiroshima-u.ac.jp/~m-mat/MT/MT2002/emt19937ar.html>. The following are the verbatim comments from the original code:

A C-program for MT19937, with initialization improved 2002/1/26.
Coded by Takuji Nishimura and Makoto Matsumoto.

Before using, initialize the state by using `init_genrand(seed)`
or `init_by_array(init_key, key_length)`.

Copyright (C) 1997 - 2002, Makoto Matsumoto and Takuji Nishimura,
All rights reserved.

(下页继续)

(续上页)

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. The names of its contributors may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Any feedback is very welcome.

<http://www.math.sci.hiroshima-u.ac.jp/~m-mat/MT/emt.html>

email: m-mat @ math.sci.hiroshima-u.ac.jp (remove space)

C.3.2 套接字

The `socket` module uses the functions, `getaddrinfo()`, and `getnameinfo()`, which are coded in separate source files from the WIDE Project, <http://www.wide.ad.jp/>.

Copyright (C) 1995, 1996, 1997, and 1998 WIDE Project.
All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. Neither the name of the project nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE PROJECT AND CONTRIBUTORS ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE

(下页继续)

(续上页)

ARE DISCLAIMED. IN NO EVENT SHALL THE PROJECT OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

C.3.3 Floating point exception control

The source for the `fpectl` module includes the following notice:

```

-----
/                               Copyright (c) 1996.                               \
|                               The Regents of the University of California.          |
|                               All rights reserved.                                |
|                                                                                 |
|  Permission to use, copy, modify, and distribute this software for               |
|  any purpose without fee is hereby granted, provided that this en-               |
|  tire notice is included in all copies of any software which is or               |
|  includes a copy or modification of this software and in all                     |
|  copies of the supporting documentation for such software.                       |
|                                                                                 |
|  This work was produced at the University of California, Lawrence                  |
|  Livermore National Laboratory under contract no. W-7405-ENG-48                   |
|  between the U.S. Department of Energy and The Regents of the                    |
|  University of California for the operation of UC LLNL.                          |
|                                                                                 |
|                               DISCLAIMER                                           |
|                                                                                 |
|  This software was prepared as an account of work sponsored by an                 |
|  agency of the United States Government. Neither the United States                |
|  Government nor the University of California nor any of their em-                 |
|  ployees, makes any warranty, express or implied, or assumes any                 |
|  liability or responsibility for the accuracy, completeness, or                   |
|  usefulness of any information, apparatus, product, or process                    |
|  disclosed, or represents that its use would not infringe                        |
|  privately-owned rights. Reference herein to any specific commer-                 |
|  cial products, process, or service by trade name, trademark,                     |
|  manufacturer, or otherwise, does not necessarily constitute or                   |
|  imply its endorsement, recommendation, or favoring by the United                |
|  States Government or the University of California. The views and                 |
|  opinions of authors expressed herein do not necessarily state or                 |
|  reflect those of the United States Government or the University                   |
|  of California, and shall not be used for advertising or product                 |
|  endorsement purposes.                                                            |
\                                                                                 /
-----

```

C.3.4 Asynchronous socket services

The `asynchat` and `asyncore` modules contain the following notice:

Copyright 1996 by Sam Rushing

All Rights Reserved

Permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of Sam Rushing not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission.

SAM RUSHING DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS, IN NO EVENT SHALL SAM RUSHING BE LIABLE FOR ANY SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

C.3.5 Cookie management

The `http.cookies` module contains the following notice:

Copyright 2000 by Timothy O'Malley <timo@alum.mit.edu>

All Rights Reserved

Permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of Timothy O'Malley not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission.

Timothy O'Malley DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS, IN NO EVENT SHALL Timothy O'Malley BE LIABLE FOR ANY SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

C.3.6 Execution tracing

The `trace` module contains the following notice:

portions copyright 2001, Autonomous Zones Industries, Inc., all rights...
err... reserved and offered to the public under the terms of the
Python 2.2 license.

(下页继续)

(续上页)

Author: Zooko O'Whielacronx
<http://zooko.com/>
<mailto:zooko@zooko.com>

Copyright 2000, Mojam Media, Inc., all rights reserved.
 Author: Skip Montanaro

Copyright 1999, Bioreason, Inc., all rights reserved.
 Author: Andrew Dalke

Copyright 1995-1997, Automatrix, Inc., all rights reserved.
 Author: Skip Montanaro

Copyright 1991-1995, Stichting Mathematisch Centrum, all rights reserved.

Permission to use, copy, modify, and distribute this Python software and its associated documentation for any purpose without fee is hereby granted, provided that the above copyright notice appears in all copies, and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of neither Automatrix, Bioreason or Mojam Media be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission.

C.3.7 UUencode and UUdecode functions

The `uu` module contains the following notice:

Copyright 1994 by Lance Ellinghouse
 Cathedral City, California Republic, United States of America.
 All Rights Reserved

Permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of Lance Ellinghouse not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission.

LANCE ELLINGHOUSE DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS, IN NO EVENT SHALL LANCE ELLINGHOUSE CENTRUM BE LIABLE FOR ANY SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

Modified by Jack Jansen, CWI, July 1995:

- Use `binascii` module to do the actual line-by-line conversion between `ascii` and binary. This results in a 1000-fold speedup. The C version is still 5 times faster, though.
- Arguments more compliant with Python standard

C.3.8 XML Remote Procedure Calls

The `xmlrpc.client` module contains the following notice:

```
The XML-RPC client interface is

Copyright (c) 1999-2002 by Secret Labs AB
Copyright (c) 1999-2002 by Fredrik Lundh

By obtaining, using, and/or copying this software and/or its
associated documentation, you agree that you have read, understood,
and will comply with the following terms and conditions:

Permission to use, copy, modify, and distribute this software and
its associated documentation for any purpose and without fee is
hereby granted, provided that the above copyright notice appears in
all copies, and that both that copyright notice and this permission
notice appear in supporting documentation, and that the name of
Secret Labs AB or the author not be used in advertising or publicity
pertaining to distribution of the software without specific, written
prior permission.

SECRET LABS AB AND THE AUTHOR DISCLAIMS ALL WARRANTIES WITH REGARD
TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANT-
ABILITY AND FITNESS.  IN NO EVENT SHALL SECRET LABS AB OR THE AUTHOR
BE LIABLE FOR ANY SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES OR ANY
DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS,
WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS
ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE
OF THIS SOFTWARE.
```

C.3.9 test_epoll

The `test_epoll` module contains the following notice:

```
Copyright (c) 2001-2006 Twisted Matrix Laboratories.

Permission is hereby granted, free of charge, to any person obtaining
a copy of this software and associated documentation files (the
"Software"), to deal in the Software without restriction, including
without limitation the rights to use, copy, modify, merge, publish,
distribute, sublicense, and/or sell copies of the Software, and to
permit persons to whom the Software is furnished to do so, subject to
the following conditions:

The above copyright notice and this permission notice shall be
included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND,
EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF
MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND
NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE
LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION
OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION
WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.
```

C.3.10 Select kqueue

The `select` module contains the following notice for the kqueue interface:

Copyright (c) 2000 Doug White, 2006 James Knight, 2007 Christian Heimes
All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR AND CONTRIBUTORS ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

C.3.11 SipHash24

The file `Python/pyhash.c` contains Marek Majkowski's implementation of Dan Bernstein's SipHash24 algorithm. The contains the following note:

<MIT License>

Copyright (c) 2013 Marek Majkowski <marek@popcount.org>

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

</MIT License>

Original location:

<https://github.com/majek/csiphash/>

Solution inspired by code from:

Samuel Neves ([supercop/crypto_auth/siphash24/little](https://github.com/sneves/siphash24/little))

djb ([supercop/crypto_auth/siphash24/little2](https://github.com/djb/siphash24/little2))

Jean-Philippe Aumasson (<https://131002.net/siphash/siphash24.c>)

C.3.12 strtod and dtoa

The file `Python/dtoa.c`, which supplies C functions `dtoa` and `strtod` for conversion of C doubles to and from strings, is derived from the file of the same name by David M. Gay, currently available from <http://www.netlib.org/fp/>. The original file, as retrieved on March 16, 2009, contains the following copyright and licensing notice:

```

/*****
 *
 * The author of this software is David M. Gay.
 *
 * Copyright (c) 1991, 2000, 2001 by Lucent Technologies.
 *
 * Permission to use, copy, modify, and distribute this software for any
 * purpose without fee is hereby granted, provided that this entire notice
 * is included in all copies of any software which is or includes a copy
 * or modification of this software and in all copies of the supporting
 * documentation for such software.
 *
 * THIS SOFTWARE IS BEING PROVIDED "AS IS", WITHOUT ANY EXPRESS OR IMPLIED
 * WARRANTY. IN PARTICULAR, NEITHER THE AUTHOR NOR LUCENT MAKES ANY
 * REPRESENTATION OR WARRANTY OF ANY KIND CONCERNING THE MERCHANTABILITY
 * OF THIS SOFTWARE OR ITS FITNESS FOR ANY PARTICULAR PURPOSE.
 *
 *****/

```

C.3.13 OpenSSL

The modules `hashlib`, `posix`, `ssl`, `crypt` use the OpenSSL library for added performance if made available by the operating system. Additionally, the Windows and Mac OS X installers for Python may include a copy of the OpenSSL libraries, so we include a copy of the OpenSSL license here:

```

LICENSE ISSUES
=====

The OpenSSL toolkit stays under a dual license, i.e. both the conditions of
the OpenSSL License and the original SSLeay license apply to the toolkit.
See below for the actual license texts. Actually both licenses are BSD-style
Open Source licenses. In case of any license issues related to OpenSSL
please contact openssl-core@openssl.org.

OpenSSL License
-----

/* =====
 * Copyright (c) 1998-2008 The OpenSSL Project. All rights reserved.
 *
 * Redistribution and use in source and binary forms, with or without
 * modification, are permitted provided that the following conditions
 * are met:
 *
 * 1. Redistributions of source code must retain the above copyright
 * notice, this list of conditions and the following disclaimer.
 *
 * 2. Redistributions in binary form must reproduce the above copyright
 * notice, this list of conditions and the following disclaimer in
 * the documentation and/or other materials provided with the

```

(下页继续)

(续上页)

```

*   distribution.
*
* 3. All advertising materials mentioning features or use of this
*   software must display the following acknowledgment:
*   "This product includes software developed by the OpenSSL Project
*   for use in the OpenSSL Toolkit. (http://www.openssl.org/)"
*
* 4. The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used to
*   endorse or promote products derived from this software without
*   prior written permission. For written permission, please contact
*   openssl-core@openssl.org.
*
* 5. Products derived from this software may not be called "OpenSSL"
*   nor may "OpenSSL" appear in their names without prior written
*   permission of the OpenSSL Project.
*
* 6. Redistributions of any form whatsoever must retain the following
*   acknowledgment:
*   "This product includes software developed by the OpenSSL Project
*   for use in the OpenSSL Toolkit (http://www.openssl.org/)"
*
* THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT ``AS IS'' AND ANY
* EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
* IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR
* PURPOSE ARE DISCLAIMED.  IN NO EVENT SHALL THE OpenSSL PROJECT OR
* ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL,
* SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT
* NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES;
* LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
* HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
* STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
* ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
* OF THE POSSIBILITY OF SUCH DAMAGE.
* =====
*
* This product includes cryptographic software written by Eric Young
* (eay@cryptsoft.com).  This product includes software written by Tim
* Hudson (tjh@cryptsoft.com).
*
*/

```

Original SSLeay License

```

/* Copyright (C) 1995-1998 Eric Young (eay@cryptsoft.com)
 * All rights reserved.
 *
 * This package is an SSL implementation written
 * by Eric Young (eay@cryptsoft.com).
 * The implementation was written so as to conform with Netscapes SSL.
 *
 * This library is free for commercial and non-commercial use as long as
 * the following conditions are aheared to.  The following conditions
 * apply to all code found in this distribution, be it the RC4, RSA,
 * lhash, DES, etc., code; not just the SSL code.  The SSL documentation
 * included with this distribution is covered by the same copyright terms

```

(下页继续)

(续上页)

```

* except that the holder is Tim Hudson (tjh@cryptsoft.com).
*
* Copyright remains Eric Young's, and as such any Copyright notices in
* the code are not to be removed.
* If this package is used in a product, Eric Young should be given attribution
* as the author of the parts of the library used.
* This can be in the form of a textual message at program startup or
* in documentation (online or textual) provided with the package.
*
* Redistribution and use in source and binary forms, with or without
* modification, are permitted provided that the following conditions
* are met:
* 1. Redistributions of source code must retain the copyright
*   notice, this list of conditions and the following disclaimer.
* 2. Redistributions in binary form must reproduce the above copyright
*   notice, this list of conditions and the following disclaimer in the
*   documentation and/or other materials provided with the distribution.
* 3. All advertising materials mentioning features or use of this software
*   must display the following acknowledgement:
*   "This product includes cryptographic software written by
*    Eric Young (eay@cryptsoft.com)"
*   The word 'cryptographic' can be left out if the rouines from the library
*   being used are not cryptographic related :-).
* 4. If you include any Windows specific code (or a derivative thereof) from
*   the apps directory (application code) you must include an acknowledgement:
*   "This product includes software written by Tim Hudson (tjh@cryptsoft.com)"
*
* THIS SOFTWARE IS PROVIDED BY ERIC YOUNG ``AS IS'' AND
* ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
* IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE
* ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE
* FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL
* DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS
* OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
* HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT
* LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY
* OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF
* SUCH DAMAGE.
*
* The licence and distribution terms for any publically available version or
* derivative of this code cannot be changed. i.e. this code cannot simply be
* copied and put under another distribution licence
* [including the GNU Public Licence.]
*/

```

C.3.14 expat

The pyexpat extension is built using an included copy of the expat sources unless the build is configured `--with-system-expat`:

```

Copyright (c) 1998, 1999, 2000 Thai Open Source Software Center Ltd
and Clark Cooper

Permission is hereby granted, free of charge, to any person obtaining
a copy of this software and associated documentation files (the

```

(下页继续)

(续上页)

"Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

C.3.15 libffi

The `_ctypes` extension is built using an included copy of the libffi sources unless the build is configured `--with-system-libffi`:

Copyright (c) 1996-2008 Red Hat, Inc and others.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the ``Software''), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED ``AS IS'', WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

C.3.16 zlib

The `zlib` extension is built using an included copy of the zlib sources if the zlib version found on the system is too old to be used for the build:

Copyright (C) 1995-2011 Jean-loup Gailly and Mark Adler

This software is provided 'as-is', without any express or implied warranty. In no event will the authors be held liable for any damages

(下页继续)

(续上页)

arising from the use of this software.

Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions:

1. The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required.
2. Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.
3. This notice may not be removed or altered from any source distribution.

Jean-loup Gailly
jloup@gzip.org

Mark Adler
madler@alumni.caltech.edu

C.3.17 cfuhash

The implementation of the hash table used by the *tracemalloc* is based on the cfuhash project:

Copyright (c) 2005 Don Owens
All rights reserved.

This code is released under the BSD license:

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- * Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- * Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- * Neither the name of the author nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED

(下页继续)

(续上页)

OF THE POSSIBILITY OF SUCH DAMAGE.

C.3.18 libmpdec

The `_decimal` module is built using an included copy of the libmpdec library unless the build is configured `--with-system-libmpdec`:

Copyright (c) 2008-2016 Stefan Krah. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

版权

Python 和这份文档的版权说明：

Copyright © 2001-2019 Python Software Foundation. All rights reserved.

Copyright © 2000 BeOpen.com. All rights reserved.

Copyright © 1995-2000 Corporation for National Research Initiatives. All rights reserved.

Copyright © 1991-1995 Stichting Mathematisch Centrum. All rights reserved.

更完整的版权信息，请参考[历史和许可证](#)。

Bibliography

- [Frie09] Friedl, Jeffrey. Mastering Regular Expressions. 第三版, O’ Reilly Media, 2009. 第三版不再使用 Python, 但第一版提供了编写正则表达式的良好细节。
- [C99] ISO/IEC 9899:1999. “Programming languages – C.” A public draft of this standard is available at <http://www.open-std.org/jtc1/sc22/wg14/www/docs/n1256.pdf>.

—
__future__, 1591
__main__, 1563
_dummy_thread, 787
_thread, 786

a

abc, 1579
aifc, 1230
argparse, 573
array, 220
ast, 1649
asynchat, 917
asyncio, 852
asyncore, 912
atexit, 1584
audioop, 1227

b

base64, 1017
bdb, 1487
binascii, 1020
binhex, 1019
bisect, 218
builtins, 1562
bz2, 435

c

calendar, 192
cgi, 1087
cgitb, 1094
chunk, 1238
cmath, 269
cmd, 1296
code, 1615
codecs, 145
codeop, 1617
collections, 195
collections.abc, 210

colorsys, 1239
compileall, 1665
concurrent.futures, 761
configparser, 468
contextlib, 1568
copy, 234
copyreg, 400
cProfile, 1502
crypt (*Unix*), 1710
csv, 461
ctypes, 675
curses (*Unix*), 644
curses.ascii, 662
curses.panel, 665
curses.textpad, 661

d

datetime, 163
dbm, 404
dbm.dumb, 408
dbm.gnu (*Unix*), 406
dbm.ndbm (*Unix*), 407
decimal, 272
difflib, 117
dis, 1668
distutils, 1525
doctest, 1368
dummy_threading, 785

e

email, 929
email.charset, 980
email.contentmanager, 957
email.encoders, 982
email.errors, 951
email.generator, 942
email.header, 977
email.headerregistry, 952
email.iterators, 985
email.message, 930

email.mime, 975
email.parser, 938
email.policy, 944
email.utils, 983
encodings.idna, 160
encodings.mbc, 161
encodings.utf_8_sig, 161
ensurepip, 1526
enum, 242
errno, 669

f

faulthandler, 1491
fcntl (*Unix*), 1715
filecmp, 369
fileinput, 361
fnmatch, 376
formatter, 1683
fpectl (*Unix*), 1613
fractions, 298
ftplib, 1139
functools, 327

g

gc, 1592
getopt, 603
getpass, 643
gettext, 1247
glob, 375
grp (*Unix*), 1710
gzip, 432

h

hashlib, 493
heapq, 214
hmac, 503
html, 1025
html.entities, 1030
html.parser, 1026
http, 1131
http.client, 1133
http.cookiejar, 1192
http.cookies, 1189
http.server, 1184

i

imaplib, 1147
imghdr, 1240
imp, 1751
importlib, 1627
importlib.abc, 1629
importlib.machinery, 1635
importlib.util, 1639
inspect, 1595

io, 554
ipaddress, 1214
itertools, 313

j

json, 986
json.tool, 994

k

keyword, 1658

l

lib2to3, 1477
linecache, 377
locale, 1255
logging, 605
logging.config, 621
logging.handlers, 631
lzma, 438

m

macpath, 385
mailbox, 996
mailcap, 995
marshal, 403
math, 264
mimetypes, 1014
mmap, 924
modulefinder, 1624
msilib (*Windows*), 1689
msvcrt (*Windows*), 1694
multiprocessing, 720
multiprocessing.connection, 748
multiprocessing.dummy, 752
multiprocessing.managers, 739
multiprocessing.pool, 745
multiprocessing.sharedctypes, 737

n

netrc, 485
nis (*Unix*), 1722
nntplib, 1153
numbers, 261

o

operator, 333
optparse, 1725
os, 509
os.path, 357
ossaudiodev (*Linux, FreeBSD*), 1241

p

parser, 1645

pathlib, 341
 pdb, 1494
 pickle, 387
 pickletools, 1681
 pipes (*Unix*), 1717
 pkgutil, 1621
 platform, 666
 plistlib, 489
 poplib, 1144
 posix (*Unix*), 1707
 pprint, 235
 profile, 1502
 pstats, 1503
 pty (*Linux*), 1714
 pwd (*Unix*), 1708
 py_compile, 1664
 pycldr, 1663
 pydoc, 1366

q

queue, 783
 quopri, 1022

r

random, 300
 re, 99
 readline (*Unix*), 133
 reprlib, 240
 resource (*Unix*), 1718
 rlcompleter, 138
 runpy, 1625

s

sched, 782
 secrets, 505
 select, 841
 selectors, 848
 shelve, 401
 shlex, 1301
 shutil, 378
 signal, 919
 site, 1610
 smtpd, 1167
 smtplib, 1160
 sndhdr, 1240
 socket, 789
 socketserver, 1176
 spwd (*Unix*), 1709
 sqlite3, 408
 ssl, 810
 stat, 364
 statistics, 306
 string, 89
 stringprep, 132

struct, 139
 subprocess, 767
 sunau, 1232
 symbol, 1657
 symtable, 1654
 sys, 1543
 sysconfig, 1558
 syslog (*Unix*), 1723

t

tabnanny, 1662
 tarfile, 451
 telnetlib, 1170
 tempfile, 371
 termios (*Unix*), 1712
 test, 1477
 test.support, 1480
 textwrap, 127
 threading, 709
 time, 566
 timeit, 1507
 tkinter, 1307
 tkinter.scrolledtext (*Tk*), 1340
 tkinter.tix, 1335
 tkinter.ttk, 1318
 token, 1657
 tokenize, 1659
 trace, 1512
 traceback, 1585
 tracemalloc, 1515
 tty (*Unix*), 1713
 turtle, 1263
 turtledemo, 1294
 types, 230
 typing, 1351

u

unicodedata, 130
 unittest, 1391
 unittest.mock, 1418
 urllib, 1104
 urllib.error, 1129
 urllib.parse, 1122
 urllib.request, 1104
 urllib.response, 1122
 urllib.robotparser, 1130
 uu, 1022
 uuid, 1173

v

venv, 1527

w

warnings, 1563

wave, 1235
weakref, 223
webbrowser, 1085
winreg (*Windows*), 1696
winsound (*Windows*), 1704
wsgiref, 1095
wsgiref.handlers, 1100
wsgiref.headers, 1097
wsgiref.simple_server, 1098
wsgiref.util, 1095
wsgiref.validate, 1099

X

xdrlib, 486
xml, 1030
xml.dom, 1047
xml.dom.minidom, 1057
xml.dom.pulldom, 1061
xml.etree.ElementTree, 1032
xml.parsers.expat, 1075
xml.parsers.expat.errors, 1082
xml.parsers.expat.model, 1081
xml.sax, 1063
xml.sax.handler, 1065
xml.sax.saxutils, 1070
xml.sax.xmlreader, 1071
xmlrpc.client, 1201
xmlrpc.server, 1209

Z

zipapp, 1535
zipfile, 443
zipimport, 1619
zlib, 429

非字母

- ??
 - in regular expressions, 100
- ..
 - in pathnames, 552
- ..., 1759
 - ellipsis literal, 25, 77
 - in doctests, 1375
 - interpreter prompt, 1372, 1554
 - placeholder, 130, 235, 240
- . (dot)
 - in glob-style wildcards, 375
 - in pathnames, 552
 - in printf-style formatting, 48, 60
 - in regular expressions, 99
 - in string formatting, 91
 - in Tkinter, 1310
- ! (exclamation)
 - in a command interpreter, 1296
 - in curses module, 664
 - in glob-style wildcards, 375, 376
 - in string formatting, 91
 - in struct format strings, 140
- (minus)
 - binary operator, 29
 - in doctests, 1377
 - in glob-style wildcards, 375, 376
 - in printf-style formatting, 48, 61
 - in regular expressions, 100
 - in string formatting, 93
 - unary operator, 29
- ! (pdb command), 1499
- ? (question mark)
 - in a command interpreter, 1296
 - in argparse module, 585
 - in AST grammar, 1649
 - in glob-style wildcards, 375, 376
 - in regular expressions, 100
 - in SQL statements, 417
 - in struct format strings, 142, 143
 - replacement character, 148
- # (hash)
 - comment, 1610
 - in doctests, 1377
 - in printf-style formatting, 48, 61
 - in regular expressions, 105
 - in string formatting, 93
- \$ (dollar)
 - environment variables expansion, 358
 - in regular expressions, 99
 - in template strings, 97
 - interpolation in configuration files, 472
- % (percent)
 - datetime format, 189, 569, 570
 - environment variables expansion (Windows), 358, 1698
 - interpolation in configuration files, 471
 - printf-style formatting, 48, 60
 - 运算符, 29
- & (ampersand)
 - 运算符, 30
- (?
 - in regular expressions, 101
- (?!
 - in regular expressions, 101
- (?#
 - in regular expressions, 101
- () (parentheses)
 - in printf-style formatting, 48, 60
 - in regular expressions, 100
- (?:
 - in regular expressions, 101
- (<?!
 - in regular expressions, 102
- (<=
 - in regular expressions, 101
- (<P<
 -

in regular expressions, 101
 (?P=
 in regular expressions, 101
 *?
 in regular expressions, 100
 * (*asterisk*)
 in argparse module, 586
 in AST grammar, 1649
 in glob-style wildcards, 375, 376
 in printf-style formatting, 48, 60
 in regular expressions, 100
 运算符, 29
 **
 in glob-style wildcards, 375
 运算符, 29
 +?
 in regular expressions, 100
 + (*plus*)
 binary operator, 29
 in argparse module, 586
 in doctests, 1377
 in printf-style formatting, 48, 61
 in regular expressions, 100
 in string formatting, 93
 unary operator, 29
 , (*comma*)
 in string formatting, 93
 / (*slash*)
 in pathnames, 552
 运算符, 29
 //
 运算符, 29
 2to3, 1759
 : (*colon*)
 in SQL statements, 417
 in string formatting, 91
 path separator (*POSIX*), 552
 ; (*semicolon*), 552
 < (*less*)
 in string formatting, 93
 in struct format strings, 140
 运算符, 28
 <<
 运算符, 30
 <=
 运算符, 28
 <BLANKLINE>, 1375
 !=
 运算符, 28
 = (*equals*)
 in string formatting, 93
 in struct format strings, 140
 ==
 运算符, 28
 > (*greater*)
 in string formatting, 93
 in struct format strings, 140
 运算符, 28
 >=
 运算符, 28
 >>
 运算符, 30
 >>>
 interpreter prompt, 1372, 1554
 @ (*at*)
 in struct format strings, 140
 [] (*square brackets*)
 in glob-style wildcards, 375, 376
 in regular expressions, 100
 in string formatting, 91
 \ (*backslash*)
 escape sequence, 148
 in pathnames (*Windows*), 552
 in regular expressions, 100, 102
 \\
 in regular expressions, 103
 \A
 in regular expressions, 102
 \a
 in regular expressions, 103
 \B
 in regular expressions, 102
 \b
 in regular expressions, 102, 103
 \D
 in regular expressions, 102
 \d
 in regular expressions, 102
 \f
 in regular expressions, 103
 \g
 in regular expressions, 107
 \N
 escape sequence, 148
 in regular expressions, 103
 \n
 in regular expressions, 103
 \r
 in regular expressions, 103
 \S
 in regular expressions, 103
 \s
 in regular expressions, 102
 \t
 in regular expressions, 103
 >>>, 1759
 \U
 escape sequence, 148

in regular expressions, 103
 \u
 escape sequence, 148
 in regular expressions, 103
 \v
 in regular expressions, 103
 \W
 in regular expressions, 103
 \w
 in regular expressions, 103
 \x
 escape sequence, 148
 in regular expressions, 103
 \Z
 in regular expressions, 103
 ^ (caret)
 in curses module, 664
 in regular expressions, 99, 100
 in string formatting, 93
 marker, 1374, 1586
 运算符, 30
 _ (underscore)
 gettext, 1248
 in string formatting, 93
 __abs__() (在 operator 模块中), 334
 __add__() (在 operator 模块中), 334
 __and__() (在 operator 模块中), 334
 __bases__ (class 属性), 77
 __bytes__() (email.message.EmailMessage 方法), 931
 __bytes__() (email.message.Message 方法), 968
 __call__() (email.headerregistry.HeaderRegistry 方法), 956
 __call__() (weakref.finalize 方法), 226
 __callback__ (weakref.ref 属性), 224
 __cause__ (traceback.TracebackException 属性), 1587
 __ceil__() (fractions.Fraction 方法), 300
 __class__ (instance 属性), 77
 __class__ (unittest.mock.Mock 属性), 1428
 __code__ (function object attribute), 76
 __concat__() (在 operator 模块中), 336
 __contains__() (email.message.EmailMessage 方法), 932
 __contains__() (email.message.Message 方法), 970
 __contains__() (mailbox.Mailbox 方法), 999
 __contains__() (在 operator 模块中), 336
 __context__ (traceback.TracebackException 属性), 1587
 __copy__() (copy protocol), 234
 __debug__ (设置变量), 25
 __deepcopy__() (copy protocol), 234
 __del__() (io.IOBase 方法), 558
 __delitem__() (email.message.EmailMessage 方法), 932
 __delitem__() (email.message.Message 方法), 970
 __delitem__() (mailbox.Mailbox 方法), 997
 __delitem__() (mailbox.MH 方法), 1003
 __delitem__() (在 operator 模块中), 336
 __dict__ (object 属性), 77
 __dir__() (unittest.mock.Mock 方法), 1424
 __displayhook__() (在 sys 模块中), 1545
 __doc__ (types.ModuleType 属性), 232
 __enter__() (contextmanager 方法), 74
 __enter__() (winreg.PyHKEY 方法), 1704
 __eq__() (email.charset.Charset 方法), 981
 __eq__() (email.header.Header 方法), 979
 __eq__() (instance method), 28
 __eq__() (memoryview 方法), 64
 __eq__() (在 operator 模块中), 334
 __excepthook__() (在 sys 模块中), 1545
 __exit__() (contextmanager 方法), 75
 __exit__() (winreg.PyHKEY 方法), 1704
 __floor__() (fractions.Fraction 方法), 300
 __floordiv__() (在 operator 模块中), 334
 __format__, 11
 __format__() (datetime.date 方法), 170
 __format__() (datetime.datetime 方法), 177
 __format__() (datetime.time 方法), 181
 __fspath__() (os.PathLike 方法), 511
 __future__, 1763
 __future__ (模块), 1591
 __ge__() (instance method), 28
 __ge__() (在 operator 模块中), 334
 __getitem__() (email.headerregistry.HeaderRegistry 方法), 956
 __getitem__() (email.message.EmailMessage 方法), 932
 __getitem__() (email.message.Message 方法), 970
 __getitem__() (mailbox.Mailbox 方法), 998
 __getitem__() (re.match 方法), 110
 __getitem__() (在 operator 模块中), 336
 __getnewargs__() (object 方法), 393
 __getnewargs_ex__() (object 方法), 393
 __getstate__() (copy protocol), 397
 __getstate__() (object 方法), 393
 __gt__() (instance method), 28
 __gt__() (在 operator 模块中), 334
 __iadd__() (在 operator 模块中), 339
 __iand__() (在 operator 模块中), 339
 __iconcat__() (在 operator 模块中), 339
 __ifloordiv__() (在 operator 模块中), 339
 __ilshift__() (在 operator 模块中), 339
 __imatmul__() (在 operator 模块中), 339
 __imod__() (在 operator 模块中), 339
 __import__() (设置函数), 22
 __import__() (在 importlib 模块中), 1628
 __imul__() (在 operator 模块中), 339
 __index__() (在 operator 模块中), 335
 __init__() (difflib.HtmlDiff 方法), 117

- `__init__()` (*logging.Handler* 方法), 610
- `__interactivehook__()` (在 *sys* 模块中), 1551
- `__inv__()` (在 *operator* 模块中), 335
- `__invert__()` (在 *operator* 模块中), 335
- `__ior__()` (在 *operator* 模块中), 340
- `__ipow__()` (在 *operator* 模块中), 340
- `__irshift__()` (在 *operator* 模块中), 340
- `__isub__()` (在 *operator* 模块中), 340
- `__iter__()` (*container* 方法), 34
- `__iter__()` (*iterator* 方法), 34
- `__iter__()` (*mailbox.Mailbox* 方法), 998
- `__iter__()` (*unittest.TestSuite* 方法), 1409
- `__itruediv__()` (在 *operator* 模块中), 340
- `__ixor__()` (在 *operator* 模块中), 340
- `__le__()` (*instance method*), 28
- `__le__()` (在 *operator* 模块中), 334
- `__len__()` (*email.message.EmailMessage* 方法), 932
- `__len__()` (*email.message.Message* 方法), 970
- `__len__()` (*mailbox.Mailbox* 方法), 999
- `__loader__` (*types.ModuleType* 属性), 232
- `__lshift__()` (在 *operator* 模块中), 335
- `__lt__()` (*instance method*), 28
- `__lt__()` (在 *operator* 模块中), 334
- `__main__`
 - 模块, 1625, 1626
- `__main__` (模块), 1563
- `__matmul__()` (在 *operator* 模块中), 335
- `__missing__()`, 72
- `__missing__()` (*collections.defaultdict* 方法), 203
- `__mod__()` (在 *operator* 模块中), 335
- `__mro__` (*class* 属性), 78
- `__mul__()` (在 *operator* 模块中), 335
- `__name__` (*definition* 属性), 77
- `__name__` (*types.ModuleType* 属性), 232
- `__ne__()` (*email.charset.Charset* 方法), 981
- `__ne__()` (*email.header.Header* 方法), 979
- `__ne__()` (*instance method*), 28
- `__ne__()` (在 *operator* 模块中), 334
- `__neg__()` (在 *operator* 模块中), 335
- `__next__()` (*csv.csvreader* 方法), 466
- `__next__()` (*iterator* 方法), 34
- `__not__()` (在 *operator* 模块中), 334
- `__or__()` (在 *operator* 模块中), 335
- `__package__` (*types.ModuleType* 属性), 232
- `__pos__()` (在 *operator* 模块中), 335
- `__pow__()` (在 *operator* 模块中), 335
- `__qualname__` (*definition* 属性), 77
- `__reduce__()` (*object* 方法), 394
- `__reduce_ex__()` (*object* 方法), 394
- `__repr__()` (*multiprocessing.managers.BaseProxy* 方法), 745
- `__repr__()` (*netrc.netrc* 方法), 485
- `__round__()` (*fractions.Fraction* 方法), 300
- `__rshift__()` (在 *operator* 模块中), 335
- `__setitem__()` (*email.message.EmailMessage* 方法), 932
- `__setitem__()` (*email.message.Message* 方法), 970
- `__setitem__()` (*mailbox.Mailbox* 方法), 997
- `__setitem__()` (*mailbox.Maildir* 方法), 1000
- `__setitem__()` (在 *operator* 模块中), 336
- `__setstate__()` (*copy protocol*), 397
- `__setstate__()` (*object* 方法), 393
- `__slots__`, 1768
- `__stderr__()` (在 *sys* 模块中), 1557
- `__stdin__()` (在 *sys* 模块中), 1557
- `__stdout__()` (在 *sys* 模块中), 1557
- `__str__()` (*datetime.date* 方法), 169
- `__str__()` (*datetime.datetime* 方法), 177
- `__str__()` (*datetime.time* 方法), 181
- `__str__()` (*email.charset.Charset* 方法), 981
- `__str__()` (*email.header.Header* 方法), 979
- `__str__()` (*email.headerregistry.Address* 方法), 957
- `__str__()` (*email.headerregistry.Group* 方法), 957
- `__str__()` (*email.message.EmailMessage* 方法), 931
- `__str__()` (*email.message.Message* 方法), 967
- `__str__()` (*multiprocessing.managers.BaseProxy* 方法), 745
- `__sub__()` (在 *operator* 模块中), 335
- `__subclasses__()` (*class* 方法), 78
- `__subclasshook__()` (*abc.ABCMeta* 方法), 1580
- `__suppress_context__` (*traceback.TracebackException* 属性), 1587
- `__truediv__()` (在 *operator* 模块中), 335
- `__xor__()` (在 *operator* 模块中), 335
- `__anonymous__` (*ctypes.Structure* 属性), 706
- `__asdict__()` (*collections.somenamedtuple* 方法), 205
- `__b_base__` (*ctypes._CData* 属性), 703
- `__b_needsfree__` (*ctypes._CData* 属性), 703
- `__callmethod__()` (*multiprocessing.managers.BaseProxy* 方法), 744
- `__CData` (*ctypes* 中的类), 702
- `__clear_type_cache__()` (在 *sys* 模块中), 1544
- `__current_frames__()` (在 *sys* 模块中), 1544
- `__debugmallocstats__()` (在 *sys* 模块中), 1544
- `__dummy_thread` (模块), 787
- `__enablelegacywindowsfsencoding__()` (在 *sys* 模块中), 1556
- `__exit__()` (在 *os* 模块中), 543
- `__fields` (*ast.AST* 属性), 1649
- `__fields` (*collections.somenamedtuple* 属性), 206
- `__fields__` (*ctypes.Structure* 属性), 706
- `__flush__()` (*wsgiref.handlers.BaseHandler* 方法), 1101
- `__FuncPtr` (*ctypes* 中的类), 697
- `__get_child_mock__()` (*unittest.mock.Mock* 方法), 1424
- `__getframe__()` (在 *sys* 模块中), 1549
- `__getvalue__()` (*multiprocessing.managers.BaseProxy* 方法), 745

- `_handle` (*ctypes.PyDLL* 属性), 696
- `_length_` (*ctypes.Array* 属性), 707
- `_locale`
 - 模块, 1255
- `_make()` (*collections.somenamedtuple* 类方法), 205
- `_makeResult()` (*unittest.TextTestRunner* 方法), 1414
- `_name` (*ctypes.PyDLL* 属性), 696
- `_objects` (*ctypes._CData* 属性), 703
- `_pack_` (*ctypes.Structure* 属性), 706
- `_parse()` (*gettext.NullTranslations* 方法), 1250
- `_Pointer` (*ctypes* 中的类), 707
- `_replace()` (*collections.somenamedtuple* 方法), 206
- `_setroot()` (*xml.etree.ElementTree.ElementTree* 方法), 1043
- `_SimpleCData` (*ctypes* 中的类), 703
- `_source` (*collections.somenamedtuple* 属性), 206
- `_structure()` (在 *email.iterators* 模块中), 985
- `_thread` (模块), 786
- `_type_` (*ctypes._Pointer* 属性), 707
- `_type_` (*ctypes.Array* 属性), 707
- `_write()` (*wsgiref.handlers.BaseHandler* 方法), 1101
- `_xoptions()` (在 *sys* 模块中), 1558
- `{}` (*curly brackets*)
 - in regular expressions, 100
 - in string formatting, 91
- `|` (*vertical bar*)
 - in regular expressions, 100
 - 运算符, 30
- `~` (*tilde*)
 - home directory expansion, 358
 - 运算符, 30
- 环境变量
 - AUDIODEV, 1241
 - BROWSER, 1085, 1086
 - COLS, 650
 - COLUMNS, 650
 - COMSPEC, 547, 771
 - HOME, 358
 - HOMEDRIVE, 358
 - HOMEPATH, 358
 - http_proxy, 1105, 1118
 - IDLESTARTUP, 1346
 - KDEDIR, 1087
 - LANG, 1247, 1249, 1256, 1258
 - LANGUAGE, 1247, 1249
 - LC_ALL, 1247, 1249
 - LC_MESSAGES, 1247, 1249
 - LINES, 646, 650
 - LNAME, 644
 - LOGNAME, 512, 644
 - MIXERDEV, 1242
 - no_proxy, 1107
 - PAGER, 1367
 - PATH, 542, 546, 552, 1085, 1092, 1094
 - POSIXLY_CORRECT, 604
 - PYTHON_DOM, 1048
 - PYTHONASYNCIODEBUG, 862, 907
 - PYTHONDOCS, 1367
 - PYTHONDONTWRITEBYTECODE, 1545
 - PYTHONFAULTHANDLER, 1492
 - PYTHONIOENCODING, 1557
 - PYTHONLEGACYWINDOWSFSENCODING, 1556
 - PYTHONNOUSERSITE, 1611, 1612
 - PYTHONPATH, 1092, 1552
 - PYTHONSTARTUP, 136, 1346, 1551, 1611
 - PYTHONTRACEMALLOC, 1515, 1520
 - PYTHONUSERBASE, 1612
 - SSL_CERT_FILE, 841
 - SSL_CERT_PATH, 841
 - SystemRoot, 772
 - TEMP, 373
 - TERM, 649
 - TIX_LIBRARY, 1336
 - TMP, 373
 - TMPDIR, 373
 - TZ, 571, 572
 - USER, 644
 - USERNAME, 512, 644
 - USERPROFILE, 358
- 语句
 - assert, 80
 - del, 37, 71
 - except, 79
 - if, 27
 - import, 22, 1610, 1751
 - raise, 79
 - try, 79
 - while, 27
- 运算符
 - % (*percent*), 29
 - & (*ampersand*), 30
 - * (*asterisk*), 29
 - ** , 29
 - / (*slash*), 29
 - // , 29
 - < (*less*), 28
 - <<, 30
 - <=, 28
 - !=, 28
 - ==, 28
 - > (*greater*), 28
 - >=, 28
 - >>, 30
 - ^ (*caret*), 30
 - | (*vertical bar*), 30
 - ~ (*tilde*), 30
 - and, 27, 28
 - in, 28, 35

is, 28
 is not, 28
 not, 28
 not in, 28, 35
 or, 27, 28

A

A() (在 *re* 模块中), 104

-a, --annotate

 pickletools 命令行选项, 1681

a2b_base64() (在 *binascii* 模块中), 1020

a2b_hex() (在 *binascii* 模块中), 1021

a2b_hqx() (在 *binascii* 模块中), 1021

a2b_qp() (在 *binascii* 模块中), 1021

a2b_uu() (在 *binascii* 模块中), 1020

a85decode() (在 *base64* 模块中), 1018

a85encode() (在 *base64* 模块中), 1018

ABC (*abc* 中的类), 1579

abc (模块), 1579

ABCMeta (*abc* 中的类), 1580

abiflags() (在 *sys* 模块中), 1543

abort() (*asyncio.DatagramTransport* 方法), 881

abort() (*asyncio.WriteTransport* 方法), 880

abort() (*ftplib.FTP* 方法), 1142

abort() (*threading.Barrier* 方法), 719

abort() (在 *os* 模块中), 542

above() (*curses.panel.Panel* 方法), 665

abs() (*decimal.Context* 方法), 285

abs() (置函数), 5

abs() (在 *operator* 模块中), 334

abspath() (在 *os.path* 模块中), 357

abstract base class -- 抽象基类, 1759

AbstractBasicAuthHandler (*urllib.request* 中的类), 1108

abstractclassmethod() (在 *abc* 模块中), 1582

AbstractContextManager (*contextlib* 中的类), 1568

AbstractDigestAuthHandler (*urllib.request* 中的类), 1108

AbstractEventLoop (*asyncio* 中的类), 853

AbstractEventLoopPolicy (*asyncio* 中的类), 867

AbstractFormatter (*formatter* 中的类), 1685

abstractmethod() (在 *abc* 模块中), 1581

abstractproperty() (在 *abc* 模块中), 1583

AbstractSet (*typing* 中的类), 1359

abstractstaticmethod() (在 *abc* 模块中), 1582

AbstractWriter (*formatter* 中的类), 1686

accept() (*asyncore.dispatcher* 方法), 915

accept() (*multiprocessing.connection.Listener* 方法), 748

accept() (*socket.socket* 方法), 799

access() (在 *os* 模块中), 524

accumulate() (在 *itertools* 模块中), 314

acos() (在 *cmath* 模块中), 270

acos() (在 *math* 模块中), 267

acosh() (在 *cmath* 模块中), 271

acosh() (在 *math* 模块中), 267

acquire() (*_thread.lock* 方法), 787

acquire() (*asyncio.Condition* 方法), 903

acquire() (*asyncio.Lock* 方法), 902

acquire() (*asyncio.Semaphore* 方法), 904

acquire() (*logging.Handler* 方法), 610

acquire() (*multiprocessing.Lock* 方法), 734

acquire() (*multiprocessing.RLock* 方法), 735

acquire() (*threading.Condition* 方法), 715

acquire() (*threading.Lock* 方法), 713

acquire() (*threading.RLock* 方法), 713

acquire() (*threading.Semaphore* 方法), 716

acquire_lock() (在 *imp* 模块中), 1754

Action (*argparse* 中的类), 592

action (*optparse.Option* 属性), 1738

ACTIONS (*optparse.Option* 属性), 1750

active_children() (在 *multiprocessing* 模块中), 731

active_count() (在 *threading* 模块中), 709

add() (*decimal.Context* 方法), 286

add() (*frozenset* 方法), 71

add() (*mailbox.Mailbox* 方法), 997

add() (*mailbox.Maildir* 方法), 1000

add() (*msilib.RadioButtonGroup* 方法), 1694

add() (*pstats.Stats* 方法), 1504

add() (*tarfile.TarFile* 方法), 455

add() (*tkinter.ttk.Notebook* 方法), 1323

add() (在 *audioop* 模块中), 1227

add() (在 *operator* 模块中), 334

add_alias() (在 *email.charset* 模块中), 982

add_alternative() (*email.message.EmailMessage* 方法), 937

add_argument() (*argparse.ArgumentParser* 方法), 582

add_argument_group() (*argparse.ArgumentParser* 方法), 599

add_attachment() (*email.message.EmailMessage* 方法), 937

add_cgi_vars() (*wsgiref.handlers.BaseHandler* 方法), 1101

add_charset() (在 *email.charset* 模块中), 981

add_codec() (在 *email.charset* 模块中), 982

add_cookie_header() (*http.cookiejar.CookieJar* 方法), 1194

add_data() (在 *msilib* 模块中), 1690

add_done_callback() (*asyncio.Future* 方法), 872

add_done_callback() (*concurrent.futures.Future* 方法), 765

add_fallback() (*gettext.NullTranslations* 方法), 1250

add_file() (*msilib.Directory* 方法), 1693

add_flag() (*mailbox.MaildirMessage* 方法), 1006

add_flag() (*mailbox.mboxMessage* 方法), 1007

- add_flag() (*mailbox.MMDFMessage* 方法), 1011
 add_flow_data() (*formatter.formatter* 方法), 1684
 add_folder() (*mailbox.Maildir* 方法), 1000
 add_folder() (*mailbox.MH* 方法), 1002
 add_get_handler()
 (*email.contentmanager.ContentManager* 方法), 958
 add_handler() (*urllib.request.OpenerDirector* 方法), 1111
 add_header() (*email.message.EmailMessage* 方法), 933
 add_header() (*email.message.Message* 方法), 970
 add_header() (*urllib.request.Request* 方法), 1110
 add_header() (*wsgiref.headers.Headers* 方法), 1097
 add_history() (在 *readline* 模块中), 135
 add_hor_rule() (*formatter.formatter* 方法), 1684
 add_label() (*mailbox.BabylMessage* 方法), 1010
 add_label_data() (*formatter.formatter* 方法), 1684
 add_line_break() (*formatter.formatter* 方法), 1684
 add_literal_data() (*formatter.formatter* 方法), 1684
 add_mutually_exclusive_group() (*argparse.ArgumentParser* 方法), 600
 add_option() (*optparse.OptionParser* 方法), 1737
 add_parent() (*urllib.request.BaseHandler* 方法), 1111
 add_password() (*urllib.request.HTTPPasswordMgr* 方法), 1114
 add_password() (*urllib.request.HTTPPasswordMgrWithPriorAuth* 方法), 1114
 add_reader() (*asyncio.AbstractEventLoop* 方法), 858
 add_related() (*email.message.EmailMessage* 方法), 937
 add_section() (*configparser.ConfigParser* 方法), 481
 add_section() (*configparser.RawConfigParser* 方法), 484
 add_sequence() (*mailbox.MHMessage* 方法), 1009
 add_set_handler()
 (*email.contentmanager.ContentManager* 方法), 958
 add_signal_handler() (*asyncio.AbstractEventLoop* 方法), 860
 add_stream() (在 *msilib* 模块中), 1690
 add_subparsers() (*argparse.ArgumentParser* 方法), 596
 add_tables() (在 *msilib* 模块中), 1690
 add_type() (在 *mimetypes* 模块中), 1015
 add_unredirected_header() (*urllib.request.Request* 方法), 1110
 add_writer() (*asyncio.AbstractEventLoop* 方法), 858
 addch() (*curses.window* 方法), 651
 addCleanup() (*unittest.TestCase* 方法), 1408
 addcomponent() (*turtle.Shape* 方法), 1291
 addError() (*unittest.TestResult* 方法), 1413
 addExpectedFailure() (*unittest.TestResult* 方法), 1413
 addFailure() (*unittest.TestResult* 方法), 1413
 addfile() (*tarfile.TarFile* 方法), 455
 addFilter() (*logging.Handler* 方法), 610
 addFilter() (*logging.Logger* 方法), 608
 addHandler() (*logging.Logger* 方法), 609
 addLevelName() (在 *logging* 模块中), 618
 addnstr() (*curses.window* 方法), 651
 AddPackagePath() (在 *modulefinder* 模块中), 1624
 addr(*smtplib.SMTPChannel* 属性), 1169
 addr_spec(*email.headerregistry.Address* 属性), 957
 Address(*email.headerregistry* 中的类), 956
 address(*email.headerregistry.SingleAddressHeader* 属性), 955
 address(*multiprocessing.connection.Listener* 属性), 749
 address(*multiprocessing.managers.BaseManager* 属性), 740
 address_exclude() (*ipaddress.IPv4Network* 方法), 1220
 address_exclude() (*ipaddress.IPv6Network* 方法), 1222
 address_family(*socketserver.BaseServer* 属性), 1178
 address_string() (*http.server.BaseHTTPRequestHandler* 方法), 1187
 addresses(*email.headerregistry.AddressHeader* 属性), 954
 addresses(*email.headerregistry.Group* 属性), 957
 AddressHeader(*email.headerregistry* 中的类), 954
 addressof() (在 *ctypes* 模块中), 700
 AddressValueError, 1226
 addshape() (在 *turtle* 模块中), 1289
 addsitedir() (在 *site* 模块中), 1612
 addSkip() (*unittest.TestResult* 方法), 1413
 addstr() (*curses.window* 方法), 651
 addSubTest() (*unittest.TestResult* 方法), 1413
 addSuccess() (*unittest.TestResult* 方法), 1413
 addTest() (*unittest.TestSuite* 方法), 1409
 addTests() (*unittest.TestSuite* 方法), 1409
 addTypeEqualityFunc() (*unittest.TestCase* 方法), 1406
 addUnexpectedSuccess() (*unittest.TestResult* 方法), 1413
 adjusted() (*decimal.Decimal* 方法), 278
 Adler32() (在 *zlib* 模块中), 429
 ADPCM, Intel/DVI, 1227
 adpcm2lin() (在 *audioop* 模块中), 1227
 AF_ALG() (在 *socket* 模块中), 794
 AF_CAN() (在 *socket* 模块中), 793
 AF_INET() (在 *socket* 模块中), 792
 AF_INET6() (在 *socket* 模块中), 792
 AF_LINK() (在 *socket* 模块中), 794

- AF_PACKET() (在 *socket* 模块中), 794
 AF_RDS() (在 *socket* 模块中), 794
 AF_UNIX() (在 *socket* 模块中), 792
 aifc (模块), 1230
 aifc() (*aifc.aifc* 方法), 1232
 AIFF, 1230, 1238
 aiff() (*aifc.aifc* 方法), 1232
 AIFF-C, 1230, 1238
 alarm() (在 *signal* 模块中), 921
 A-LAW, 1232, 1240
 a-LAW, 1227
 alaw2lin() (在 *audioop* 模块中), 1227
 ALERT_DESCRIPTION_HANDSHAKE_FAILURE() (在 *ssl* 模块中), 821
 ALERT_DESCRIPTION_INTERNAL_ERROR() (在 *ssl* 模块中), 821
 AlertDescription(*ssl* 中的类), 821
 algorithms_available() (在 *hashlib* 模块中), 494
 algorithms_guaranteed() (在 *hashlib* 模块中), 494
 alias(*pdb* command), 1499
 alignment() (在 *ctypes* 模块中), 700
 alive(*weakref.finalize* 属性), 226
 all() (设置函数), 5
 all_errors() (在 *ftplib* 模块中), 1141
 all_features() (在 *xml.sax.handler* 模块中), 1066
 all_frames(*tracemalloc.Filter* 属性), 1521
 all_properties() (在 *xml.sax.handler* 模块中), 1067
 all_suffixes() (在 *importlib.machinery* 模块中), 1635
 all_tasks() (*asyncio.Task* 类方法), 874
 allocate_lock() (在 *_thread* 模块中), 786
 allow_reuse_address(*socketserver.BaseServer* 属性), 1178
 allowed_domains() (*http.cookiejar.DefaultCookiePolicy* 方法), 1198
 alt() (在 *curses.ascii* 模块中), 664
 ALT_DIGITS() (在 *locale* 模块中), 1258
 altsep() (在 *os* 模块中), 552
 altzone() (在 *time* 模块中), 573
 ALWAYS_TYPED_ACTIONS (*optparse.Option* 属性), 1750
 AMPER() (在 *token* 模块中), 1657
 AMPEREQUAL() (在 *token* 模块中), 1657
 and
 运算符, 27, 28
 and_() (在 *operator* 模块中), 334
 annotation -- 标注, 1759
 annotation(*inspect.Parameter* 属性), 1601
 answer_challenge() (在 *multiprocessing.connection* 模块中), 748
 anticipate_failure() (在 *test.support* 模块中), 1483
 any() (设置函数), 5
 Any() (在 *typing* 模块中), 1364
 ANY() (在 *unittest.mock* 模块中), 1447
 AnyStr() (在 *typing* 模块中), 1366
 api_version() (在 *sys* 模块中), 1558
 apop() (*poplib.POP3* 方法), 1146
 append() (*array.array* 方法), 221
 append() (*collections.deque* 方法), 200
 append() (*email.header.Header* 方法), 978
 append() (*imaplib.IMAP4* 方法), 1149
 append() (*msilib.CAB* 方法), 1692
 append() (*pipes.Template* 方法), 1718
 append() (*sequence method*), 37
 append() (*xml.etree.ElementTree.Element* 方法), 1041
 append_history_file() (在 *readline* 模块中), 134
 appendChild() (*xml.dom.Node* 方法), 1050
 appendleft() (*collections.deque* 方法), 200
 application_uri() (在 *wsgiref.util* 模块中), 1095
 apply (2to3 fixer), 1474
 apply() (*multiprocessing.pool.Pool* 方法), 746
 apply_async() (*multiprocessing.pool.Pool* 方法), 746
 apply_defaults() (*inspect.BoundArguments* 方法), 1603
 architecture() (在 *platform* 模块中), 666
 archive(*zipimport.zipimporter* 属性), 1620
 aRepr() (在 *reprlib* 模块中), 240
 argparse (模块), 573
 args (*BaseException* 属性), 80
 args (*functools.partial* 属性), 333
 args (*inspect.BoundArguments* 属性), 1603
 args (*pdb* command), 1498
 args (*subprocess.CompletedProcess* 属性), 768
 args (*subprocess.Popen* 属性), 774
 argtypes (*ctypes._FuncPtr* 属性), 697
 argument -- 参数, 1759
 ArgumentDefaultsHelpFormatter (*argparse* 中的类), 578
 ArgumentError, 697
 ArgumentParser (*argparse* 中的类), 575
 arguments (*inspect.BoundArguments* 属性), 1603
 argv() (在 *sys* 模块中), 1543
 arithmetic, 29
 ArithmeticError, 80
 array
 模块, 49
 array (*array* 中的类), 220
 Array (*ctypes* 中的类), 707
 array (模块), 220
 Array() (*multiprocessing.managers.SyncManager* 方法), 741
 Array() (在 *multiprocessing* 模块中), 736
 Array() (在 *multiprocessing.sharedctypes* 模块中), 737

- arrays, 220
- arraysize (*sqlite3.Cursor* 属性), 420
- article() (*nnplib.NNTP* 方法), 1158
- as_bytes() (*email.message.EmailMessage* 方法), 931
- as_bytes() (*email.message.Message* 方法), 968
- as_completed() (在 *asyncio* 模块中), 875
- as_completed() (在 *concurrent.futures* 模块中), 766
- as_integer_ratio() (*decimal.Decimal* 方法), 278
- as_integer_ratio() (*float* 方法), 32
- AS_IS() (在 *formatter* 模块中), 1683
- as_posix() (*pathlib.PurePath* 方法), 348
- as_string() (*email.message.EmailMessage* 方法), 931
- as_string() (*email.message.Message* 方法), 967
- as_tuple() (*decimal.Decimal* 方法), 278
- as_uri() (*pathlib.PurePath* 方法), 348
- ascii() (设置函数), 6
- ascii() (在 *curses.ascii* 模块中), 664
- ASCII() (在 *re* 模块中), 104
- ascii_letters() (在 *string* 模块中), 89
- ascii_lowercase() (在 *string* 模块中), 89
- ascii_uppercase() (在 *string* 模块中), 89
- asctime() (在 *time* 模块中), 567
- asin() (在 *cmath* 模块中), 270
- asin() (在 *math* 模块中), 267
- asinh() (在 *cmath* 模块中), 271
- asinh() (在 *math* 模块中), 267
- assert
 - 语句, 80
- assert_any_call() (*unittest.mock.Mock* 方法), 1422
- assert_called() (*unittest.mock.Mock* 方法), 1422
- assert_called_once() (*unittest.mock.Mock* 方法), 1422
- assert_called_once_with() (*unittest.mock.Mock* 方法), 1422
- assert_called_with() (*unittest.mock.Mock* 方法), 1422
- assert_has_calls() (*unittest.mock.Mock* 方法), 1423
- assert_line_data() (*formatter.formatter* 方法), 1685
- assert_not_called() (*unittest.mock.Mock* 方法), 1423
- assertAlmostEqual() (*unittest.TestCase* 方法), 1405
- assertCountEqual() (*unittest.TestCase* 方法), 1405
- assertDictEqual() (*unittest.TestCase* 方法), 1407
- assertEqual() (*unittest.TestCase* 方法), 1401
- assertFalse() (*unittest.TestCase* 方法), 1401
- assertGreater() (*unittest.TestCase* 方法), 1405
- assertGreaterEqual() (*unittest.TestCase* 方法), 1405
- assertIn() (*unittest.TestCase* 方法), 1402
- AssertionError, 80
- assertIs() (*unittest.TestCase* 方法), 1402
- assertIsInstance() (*unittest.TestCase* 方法), 1402
- assertIsNone() (*unittest.TestCase* 方法), 1402
- assertIsNot() (*unittest.TestCase* 方法), 1402
- assertIsNotNone() (*unittest.TestCase* 方法), 1402
- assertLess() (*unittest.TestCase* 方法), 1405
- assertLessEqual() (*unittest.TestCase* 方法), 1405
- assertListEqual() (*unittest.TestCase* 方法), 1406
- assertLogs() (*unittest.TestCase* 方法), 1404
- assertMultiLineEqual() (*unittest.TestCase* 方法), 1406
- assertNotAlmostEqual() (*unittest.TestCase* 方法), 1405
- assertNotEqual() (*unittest.TestCase* 方法), 1401
- assertNotIn() (*unittest.TestCase* 方法), 1402
- assertNotIsInstance() (*unittest.TestCase* 方法), 1402
- assertNotRegex() (*unittest.TestCase* 方法), 1405
- assertRaises() (*unittest.TestCase* 方法), 1402
- assertRaisesRegex() (*unittest.TestCase* 方法), 1403
- assertRegex() (*unittest.TestCase* 方法), 1405
- asserts (*2to3 fixer*), 1474
- assertSequenceEqual() (*unittest.TestCase* 方法), 1406
- assertSetEqual() (*unittest.TestCase* 方法), 1406
- assertTrue() (*unittest.TestCase* 方法), 1401
- assertTupleEqual() (*unittest.TestCase* 方法), 1406
- assertWarns() (*unittest.TestCase* 方法), 1403
- assertWarnsRegex() (*unittest.TestCase* 方法), 1404
- assignment
 - slice, 37
 - subscript, 37
- AST (*ast* 中的类), 1649
- ast (模块), 1649
- astimezone() (*datetime.datetime* 方法), 174
- async() (在 *asyncio* 模块中), 876
- ASYNC() (在 *token* 模块中), 1657
- async_chat (*asynchat* 中的类), 917
- async_chat.ac_in_buffer_size() (在 *asynchat* 模块中), 917
- async_chat.ac_out_buffer_size() (在 *asynchat* 模块中), 917
- AsyncContextManager (*typing* 中的类), 1361
- AsyncGenerator (*collections.abc* 中的类), 213
- AsyncGenerator (*typing* 中的类), 1361
- AsyncGeneratorType() (在 *types* 模块中), 231
- asynchat (模块), 917
- asynchronous context manager -- 异步上下文管理器, 1760
- asynchronous generator -- 异步生成器, 1760
- asynchronous generator iterator -- 异步生成器迭代器, 1760

- asynchronous iterable -- 异步可迭代对象, **1760**
- asynchronous iterator -- 异步迭代器, **1760**
- asyncio (模块), **852**
- asyncio.subprocess.DEVNULL() (在 *asyncio* 模块中), **897**
- asyncio.subprocess.PIPE() (在 *asyncio* 模块中), **897**
- asyncio.subprocess.Process (在 *asyncio* 中的类), **897**
- asyncio.subprocess.STDOUT() (在 *asyncio* 模块中), **897**
- AsyncIterable (*collections.abc* 中的类), **213**
- AsyncIterable (*typing* 中的类), **1360**
- AsyncIterator (*collections.abc* 中的类), **213**
- AsyncIterator (*typing* 中的类), **1360**
- asyncore (模块), **912**
- AsyncResult (*multiprocessing.pool* 中的类), **747**
- AT() (在 *token* 模块中), **1657**
- at_eof() (*asyncio.StreamReader* 方法), **890**
- atan() (在 *cmath* 模块中), **270**
- atan() (在 *math* 模块中), **267**
- atan2() (在 *math* 模块中), **267**
- atanh() (在 *cmath* 模块中), **271**
- atanh() (在 *math* 模块中), **267**
- ATEQUAL() (在 *token* 模块中), **1657**
- atexit (*weakref.finalize* 属性), **226**
- atexit (模块), **1584**
- atof() (在 *locale* 模块中), **1260**
- atoi() (在 *locale* 模块中), **1260**
- attach() (*email.message.Message* 方法), **968**
- attach_mock() (*unittest.mock.Mock* 方法), **1424**
- AttlistDeclHandler() (*xml.parsers.expat.xmlparser* 方法), **1078**
- attrgetter() (在 *operator* 模块中), **336**
- attrib (*xml.etree.ElementTree.Element* 属性), **1041**
- attribute -- 属性, **1760**
- AttributeError, **80**
- attributes (*xml.dom.Node* 属性), **1049**
- AttributesImpl (*xml.sax.xmlreader* 中的类), **1072**
- AttributesNSImpl (*xml.sax.xmlreader* 中的类), **1072**
- attroff() (*curses.window* 方法), **651**
- attron() (*curses.window* 方法), **651**
- attrset() (*curses.window* 方法), **651**
- Audio Interchange File Format, **1230, 1238**
- AUDIO_FILE_ENCODING_ADPCM_G721() (在 *sunau* 模块中), **1233**
- AUDIO_FILE_ENCODING_ADPCM_G722() (在 *sunau* 模块中), **1233**
- AUDIO_FILE_ENCODING_ADPCM_G723_3() (在 *sunau* 模块中), **1233**
- AUDIO_FILE_ENCODING_ADPCM_G723_5() (在 *sunau* 模块中), **1233**
- AUDIO_FILE_ENCODING_ALAW_8() (在 *sunau* 模块中), **1233**
- AUDIO_FILE_ENCODING_DOUBLE() (在 *sunau* 模块中), **1233**
- AUDIO_FILE_ENCODING_FLOAT() (在 *sunau* 模块中), **1233**
- AUDIO_FILE_ENCODING_LINEAR_8() (在 *sunau* 模块中), **1233**
- AUDIO_FILE_ENCODING_LINEAR_16() (在 *sunau* 模块中), **1233**
- AUDIO_FILE_ENCODING_LINEAR_24() (在 *sunau* 模块中), **1233**
- AUDIO_FILE_ENCODING_LINEAR_32() (在 *sunau* 模块中), **1233**
- AUDIO_FILE_ENCODING_MULAW_8() (在 *sunau* 模块中), **1233**
- AUDIO_FILE_MAGIC() (在 *sunau* 模块中), **1233**
- AUDIODEV, **1241**
- audiop (模块), **1227**
- auth() (*ftplib.FTP_TLS* 方法), **1144**
- auth() (*smtplib.SMTP* 方法), **1164**
- authenticate() (*imaplib.IMAP4* 方法), **1149**
- AuthenticationError, **728**
- authenticators() (*netrc.netrc* 方法), **485**
- authkey (*multiprocessing.Process* 属性), **727**
- auto (*enum* 中的类), **242**
- autorange() (*timeit.Timer* 方法), **1509**
- avg() (在 *audiop* 模块中), **1228**
- avgpp() (在 *audiop* 模块中), **1228**
- avoids_symlink_attacks (*shutil.rmtree* 属性), **380**
- AWAIT() (在 *token* 模块中), **1657**
- awaitable -- 可等待对象, **1760**
- Awaitable (*collections.abc* 中的类), **212**
- Awaitable (*typing* 中的类), **1360**
- ## B
- b
compileall 命令行选项, **1666**
- b, --buffer
unittest 命令行选项, **1393**
- b2a_base64() (在 *binascii* 模块中), **1020**
- b2a_hex() (在 *binascii* 模块中), **1021**
- b2a_hqx() (在 *binascii* 模块中), **1021**
- b2a_qp() (在 *binascii* 模块中), **1021**
- b2a_uu() (在 *binascii* 模块中), **1020**
- b16decode() (在 *base64* 模块中), **1018**
- b16encode() (在 *base64* 模块中), **1018**
- b32decode() (在 *base64* 模块中), **1018**
- b32encode() (在 *base64* 模块中), **1018**
- b64decode() (在 *base64* 模块中), **1017**
- b64encode() (在 *base64* 模块中), **1017**
- b85decode() (在 *base64* 模块中), **1018**
- b85encode() (在 *base64* 模块中), **1018**
- Babyl (*mailbox* 中的类), **1003**

- BabylMessage (*mailbox* 中的类), 1009
- back() (在 *turtle* 模块中), 1268
- backslashreplace_errors() (在 *codecs* 模块中), 149
- backward() (在 *turtle* 模块中), 1268
- BadStatusLine, 1135
- BadZipFile, 443
- BadZipfile, 443
- Balloon (*tkinter.tix* 中的类), 1336
- Barrier (*multiprocessing* 中的类), 734
- Barrier (*threading* 中的类), 718
- Barrier() (*multiprocessing.managers.SyncManager* 方法), 740
- base64
 - encoding, 1017
 - 模块, 1020
- base64 (模块), 1017
- base_exec_prefix() (在 *sys* 模块中), 1543
- base_prefix() (在 *sys* 模块中), 1543
- BaseCGIHandler (*wsgiref.handlers* 中的类), 1100
- BaseCookie (*http.cookies* 中的类), 1189
- BaseEventLoop (*asyncio* 中的类), 852
- BaseException, 80
- BaseHandler (*urllib.request* 中的类), 1107
- BaseHandler (*wsgiref.handlers* 中的类), 1101
- BaseHeader (*email.headerregistry* 中的类), 953
- BaseHTTPRequestHandler (*http.server* 中的类), 1184
- BaseManager (*multiprocessing.managers* 中的类), 739
- basename() (在 *os.path* 模块中), 357
- BaseProxy (*multiprocessing.managers* 中的类), 744
- BaseRequestHandler (*socketserver* 中的类), 1180
- BaseRotatingHandler (*logging.handlers* 中的类), 633
- BaseSelector (*selectors* 中的类), 849
- BaseServer (*socketserver* 中的类), 1178
- basestring (2to3 fixer), 1474
- BaseSubprocessTransport (*asyncio* 中的类), 881
- BaseTransport (*asyncio* 中的类), 879
- basicConfig() (在 *logging* 模块中), 619
- BasicContext (*decimal* 中的类), 283
- BasicInterpolation (*configparser* 中的类), 471
- baudrate() (在 *curses* 模块中), 645
- bbox() (*tkinter.ttk.Treeview* 方法), 1328
- BDADDR_ANY() (在 *socket* 模块中), 794
- BDADDR_LOCAL() (在 *socket* 模块中), 794
- bdb
 - 模块, 1494
- Bdb (*bdb* 中的类), 1488
- bdb (模块), 1487
- BdbQuit, 1487
- BDFL, 1760
- beep() (在 *curses* 模块中), 645
- Beep() (在 *winsound* 模块中), 1704
- BEFORE_ASYNC_WITH (*opcode*), 1674
- begin_fill() (在 *turtle* 模块中), 1277
- begin_poly() (在 *turtle* 模块中), 1282
- below() (*curses.panel.Panel* 方法), 665
- Benchmarking, 1507
- benchmarking, 567
- betavariate() (在 *random* 模块中), 303
- bgcolor() (在 *turtle* 模块中), 1284
- bgpic() (在 *turtle* 模块中), 1284
- bias() (在 *audioop* 模块中), 1228
- bidirectional() (在 *unicodedata* 模块中), 131
- BigEndianStructure (*ctypes* 中的类), 705
- bin() (置函数), 6
- binary
 - data, packing, 139
 - literals, 29
- Binary (*msilib* 中的类), 1690
- Binary (*xmlrpc.client* 中的类), 1204
- binary file -- 二进制文件, 1760
- binary mode, 17
- binary semaphores, 786
- BINARY_ADD (*opcode*), 1672
- BINARY_AND (*opcode*), 1673
- BINARY_FLOOR_DIVIDE (*opcode*), 1672
- BINARY_LSHIFT (*opcode*), 1672
- BINARY_MATRIX_MULTIPLY (*opcode*), 1672
- BINARY_MODULO (*opcode*), 1672
- BINARY_MULTIPLY (*opcode*), 1672
- BINARY_OR (*opcode*), 1673
- BINARY_POWER (*opcode*), 1672
- BINARY_RSHIFT (*opcode*), 1672
- BINARY_SUBSCR (*opcode*), 1672
- BINARY_SUBTRACT (*opcode*), 1672
- BINARY_TRUE_DIVIDE (*opcode*), 1672
- BINARY_XOR (*opcode*), 1673
- BinaryIO (*typing* 中的类), 1362
- binascii (模块), 1020
- bind (widgets), 1316
- bind() (*asyncore.dispatcher* 方法), 915
- bind() (*inspect.Signature* 方法), 1601
- bind() (*socket.socket* 方法), 800
- bind_partial() (*inspect.Signature* 方法), 1601
- bind_port() (在 *test.support* 模块中), 1483
- bind_textdomain_codeset() (在 *gettext* 模块中), 1248
- bindtextdomain() (在 *gettext* 模块中), 1247
- bindtextdomain() (在 *locale* 模块中), 1261
- binhex
 - 模块, 1020
- binhex (模块), 1019
- binhex() (在 *binhex* 模块中), 1020
- bisect (模块), 218
- bisect() (在 *bisect* 模块中), 218
- bisect_left() (在 *bisect* 模块中), 218

- `bisect_right()` (在 *bisect* 模块中), 218
- `bit_length()` (*int* 方法), 30
- `bitmap()` (*msilib.Dialog* 方法), 1694
- `bitwise`
 - `operations`, 30
- `bk()` (在 *turtle* 模块中), 1268
- `bkgd()` (*curses.window* 方法), 651
- `bkgdset()` (*curses.window* 方法), 651
- `blake2b()` (在 *hashlib* 模块中), 497
- `blake2b`, `blake2s`, 496
- `blake2b.MAX_DIGEST_SIZE()` (在 *hashlib* 模块中), 498
- `blake2b.MAX_KEY_SIZE()` (在 *hashlib* 模块中), 498
- `blake2b.PERSON_SIZE()` (在 *hashlib* 模块中), 498
- `blake2b.SALT_SIZE()` (在 *hashlib* 模块中), 497
- `blake2s()` (在 *hashlib* 模块中), 497
- `blake2s.MAX_DIGEST_SIZE()` (在 *hashlib* 模块中), 498
- `blake2s.MAX_KEY_SIZE()` (在 *hashlib* 模块中), 498
- `blake2s.PERSON_SIZE()` (在 *hashlib* 模块中), 498
- `blake2s.SALT_SIZE()` (在 *hashlib* 模块中), 497
- `block_size` (*hmac.HMAC* 属性), 504
- `blocked_domains()`
 - (*http.cookiejar.DefaultCookiePolicy* 方法), 1198
- `BlockingIOError`, 84, 555
- `body()` (*nnplib.NNTP* 方法), 1159
- `body_encode()` (*email.charset.Charset* 方法), 981
- `body_encoding` (*email.charset.Charset* 属性), 980
- `body_line_iterator()` (在 *email.iterators* 模块中), 985
- `BOM()` (在 *codecs* 模块中), 147
- `BOM_BE()` (在 *codecs* 模块中), 147
- `BOM_LE()` (在 *codecs* 模块中), 147
- `BOM_UTF8()` (在 *codecs* 模块中), 147
- `BOM_UTF16()` (在 *codecs* 模块中), 147
- `BOM_UTF16_BE()` (在 *codecs* 模块中), 147
- `BOM_UTF16_LE()` (在 *codecs* 模块中), 147
- `BOM_UTF32()` (在 *codecs* 模块中), 147
- `BOM_UTF32_BE()` (在 *codecs* 模块中), 147
- `BOM_UTF32_LE()` (在 *codecs* 模块中), 147
- `bool` (☐置类), 6
- `Boolean`
 - `operations`, 27, 28
 - `type`, 6
 - `values`, 77
 - 对象, 29
- `BOOLEAN_STATES` (*configparser.ConfigParser* 属性), 477
- `bootstrap()` (在 *ensurepip* 模块中), 1527
- `border()` (*curses.window* 方法), 651
- `bottom()` (*curses.panel.Panel* 方法), 665
- `bottom_panel()` (在 *curses.panel* 模块中), 665
- `BoundArguments` (*inspect* 中的类), 1603
- `BoundaryError`, 951
- `BoundedSemaphore` (*asyncio* 中的类), 904
- `BoundedSemaphore` (*multiprocessing* 中的类), 734
- `BoundedSemaphore` (*threading* 中的类), 716
- `BoundedSemaphore()`
 - (*multiprocessing.managers.SyncManager* 方法), 740
- `box()` (*curses.window* 方法), 652
- `bpformat()` (*bdb.Breakpoint* 方法), 1488
- `bpprint()` (*bdb.Breakpoint* 方法), 1488
- `break` (*pdb* command), 1496
- `break_anywhere()` (*bdb.Bdb* 方法), 1489
- `break_here()` (*bdb.Bdb* 方法), 1489
- `break_long_words` (*textwrap.TextWrapper* 属性), 130
- `BREAK_LOOP` (*opcode*), 1674
- `break_on_hyphens` (*textwrap.TextWrapper* 属性), 130
- `Breakpoint` (*bdb* 中的类), 1487
- `breakpoints`, 1343
- `broadcast_address` (*ipaddress.IPv4Network* 属性), 1219
- `broadcast_address` (*ipaddress.IPv6Network* 属性), 1222
- `broken` (*threading.Barrier* 属性), 719
- `BrokenBarrierError`, 719
- `BrokenPipeError`, 85
- `BrokenProcessPool`, 766
- `BROWSER`, 1085, 1086
- `BsdDbShelf` (*shelve* 中的类), 402
- `buffer` (*2to3* fixer), 1474
- `buffer` (*io.TextIOBase* 属性), 563
- `buffer` (*unittest.TestResult* 属性), 1412
- `buffer protocol`
 - binary sequence types, 49
 - str* (*built-in class*), 41
- `buffer size`, I/O, 17
- `buffer_info()` (*array.array* 方法), 221
- `buffer_size` (*xml.parsers.expat.xmlparser* 属性), 1077
- `buffer_text` (*xml.parsers.expat.xmlparser* 属性), 1077
- `buffer_used` (*xml.parsers.expat.xmlparser* 属性), 1077
- `BufferedIOBase` (*io* 中的类), 559
- `BufferedRandom` (*io* 中的类), 562
- `BufferedReader` (*io* 中的类), 561
- `BufferedRWPair` (*io* 中的类), 562
- `BufferedWriter` (*io* 中的类), 562
- `BufferError`, 80
- `BufferingHandler` (*logging.handlers* 中的类), 640
- `BufferTooShort`, 728
- `bufsize()` (*ossaudiodev.oss_audio_device* 方法), 1244
- `BUILD_CONST_KEY_MAP` (*opcode*), 1676
- `BUILD_LIST` (*opcode*), 1676
- `BUILD_LIST_UNPACK` (*opcode*), 1677

- BUILD_MAP (*opcode*), 1676
 - BUILD_MAP_UNPACK (*opcode*), 1677
 - BUILD_MAP_UNPACK_WITH_CALL (*opcode*), 1677
 - build_opener() (在 *urllib.request* 模块中), 1105
 - BUILD_SET (*opcode*), 1676
 - BUILD_SET_UNPACK (*opcode*), 1677
 - BUILD_SLICE (*opcode*), 1679
 - BUILD_STRING (*opcode*), 1676
 - BUILD_TUPLE (*opcode*), 1676
 - BUILD_TUPLE_UNPACK (*opcode*), 1676
 - BUILD_TUPLE_UNPACK_WITH_CALL (*opcode*), 1677
 - built-in
 - types, 27
 - builtin_module_names() (在 *sys* 模块中), 1544
 - BuiltinFunctionType() (在 *types* 模块中), 231
 - BuiltinImporter (*importlib.machinery* 中的类), 1635
 - BuiltinMethodType() (在 *types* 模块中), 231
 - builtins (模块), 1562
 - ButtonBox (*tkinter.tix* 中的类), 1336
 - bye() (在 *turtle* 模块中), 1289
 - byref() (在 *ctypes* 模块中), 700
 - bytearray
 - formatting, 60
 - interpolation, 60
 - methods, 51
 - 对象, 36, 49, 51
 - bytearray (☐置类), 51
 - byte-code
 - file, 1664, 1752
 - bytecode -- 字节码, 1760
 - Bytecode (*dis* 中的类), 1668
 - BYTECODE_SUFFIXES() (在 *importlib.machinery* 模块中), 1635
 - Bytecode.codeobj() (在 *dis* 模块中), 1669
 - Bytecode.first_line() (在 *dis* 模块中), 1669
 - byteorder() (在 *sys* 模块中), 1544
 - bytes
 - formatting, 60
 - interpolation, 60
 - methods, 51
 - str (*built-in class*), 41
 - 对象, 49
 - bytes (*uuid.UUID* 属性), 1173
 - bytes (☐置类), 49
 - bytes_le (*uuid.UUID* 属性), 1173
 - BytesFeedParser (*email.parser* 中的类), 939
 - BytesGenerator (*email.generator* 中的类), 942
 - BytesHeaderParser (*email.parser* 中的类), 940
 - BytesIO (*io* 中的类), 561
 - bytes-like object -- 字节类对象, 1760
 - BytesParser (*email.parser* 中的类), 940
 - ByteString (*collections.abc* 中的类), 212
 - ByteString (*typing* 中的类), 1359
 - byteswap() (*array.array* 方法), 221
 - byteswap() (在 *audioop* 模块中), 1228
 - BytesWarning, 86
 - bz2 (模块), 435
 - BZ2Compressor (*bz2* 中的类), 436
 - BZ2Decompressor (*bz2* 中的类), 436
 - BZ2File (*bz2* 中的类), 435
- ## C
- ### C
- language, 29
 - structures, 139
 - c <tarfile> <source1> ... <sourceN>
 - tarfile 命令行选项, 458
 - c <zipfile> <source1> ... <sourceN>
 - zipfile 命令行选项, 450
 - c, --catch
 - unittest 命令行选项, 1393
 - c, --clock
 - timeit 命令行选项, 1510
 - c, --count
 - trace 命令行选项, 1513
 - C, --coverdir=<dir>
 - trace 命令行选项, 1513
 - c_bool (*ctypes* 中的类), 705
 - C_BUILTIN() (在 *imp* 模块中), 1755
 - c_byte (*ctypes* 中的类), 703
 - c_char (*ctypes* 中的类), 703
 - c_char_p (*ctypes* 中的类), 703
 - c_contiguous (*memoryview* 属性), 69
 - c_double (*ctypes* 中的类), 704
 - C_EXTENSION() (在 *imp* 模块中), 1755
 - c_float (*ctypes* 中的类), 704
 - c_int (*ctypes* 中的类), 704
 - c_int8 (*ctypes* 中的类), 704
 - c_int16 (*ctypes* 中的类), 704
 - c_int32 (*ctypes* 中的类), 704
 - c_int64 (*ctypes* 中的类), 704
 - c_long (*ctypes* 中的类), 704
 - c_longdouble (*ctypes* 中的类), 704
 - c_longlong (*ctypes* 中的类), 704
 - c_short (*ctypes* 中的类), 704
 - c_size_t (*ctypes* 中的类), 704
 - c_ssize_t (*ctypes* 中的类), 704
 - c_ubyte (*ctypes* 中的类), 704
 - c_uint (*ctypes* 中的类), 704
 - c_uint8 (*ctypes* 中的类), 704
 - c_uint16 (*ctypes* 中的类), 704
 - c_uint32 (*ctypes* 中的类), 704
 - c_uint64 (*ctypes* 中的类), 704
 - c_ulong (*ctypes* 中的类), 705
 - c_ulonglong (*ctypes* 中的类), 705
 - c_ushort (*ctypes* 中的类), 705
 - c_void_p (*ctypes* 中的类), 705

- `c_wchar` (*ctypes* 中的类), 705
- `c_wchar_p` (*ctypes* 中的类), 705
- `CAB` (*msilib* 中的类), 1692
- `cache_from_source()` (在 *imp* 模块中), 1754
- `cache_from_source()` (在 *importlib.util* 模块中), 1639
- `cached` (*importlib.machinery.ModuleSpec* 属性), 1639
- `CacheFTPHandler` (*urllib.request* 中的类), 1109
- `calcszize()` (在 *struct* 模块中), 140
- `Calendar` (*calendar* 中的类), 192
- `calendar` (模块), 192
- `calendar()` (在 *calendar* 模块中), 194
- `call()` (在 *subprocess* 模块中), 776
- `call()` (在 *unittest.mock* 模块中), 1445
- `call_args` (*unittest.mock.Mock* 属性), 1426
- `call_args_list` (*unittest.mock.Mock* 属性), 1427
- `call_at()` (*asyncio.AbstractEventLoop* 方法), 854
- `call_count` (*unittest.mock.Mock* 属性), 1424
- `call_exception_handler()` (*asyncio.AbstractEventLoop* 方法), 861
- `CALL_FUNCTION` (*opcode*), 1679
- `CALL_FUNCTION_EX` (*opcode*), 1679
- `CALL_FUNCTION_KW` (*opcode*), 1679
- `call_later()` (*asyncio.AbstractEventLoop* 方法), 854
- `call_list()` (*unittest.mock.call* 方法), 1446
- `call_soon()` (*asyncio.AbstractEventLoop* 方法), 854
- `call_soon_threadsafe()` (*asyncio.AbstractEventLoop* 方法), 854
- `call_tracing()` (在 *sys* 模块中), 1544
- `Callable` (*collections.abc* 中的类), 211
- `callable()` (E 置函数), 7
- `Callable()` (在 *typing* 模块中), 1365
- `CallableProxyType()` (在 *weakref* 模块中), 226
- `callback` (*optparse.Option* 属性), 1738
- `callback()` (*contextlib.ExitStack* 方法), 1573
- `callback_args` (*optparse.Option* 属性), 1738
- `callback_kwargs` (*optparse.Option* 属性), 1738
- `callbacks()` (在 *gc* 模块中), 1594
- `called` (*unittest.mock.Mock* 属性), 1424
- `CalledProcessError`, 769
- `CAN_BCM()` (在 *socket* 模块中), 793
- `can_change_color()` (在 *curses* 模块中), 645
- `can_fetch()` (*urllib.robotparser.RobotFileParser* 方法), 1130
- `CAN_RAW_FD_FRAMES()` (在 *socket* 模块中), 793
- `can_symlink()` (在 *test.support* 模块中), 1483
- `can_write_eof()` (*asyncio.StreamWriter* 方法), 891
- `can_write_eof()` (*asyncio.WriteTransport* 方法), 880
- `cancel()` (*asyncio.Future* 方法), 872
- `cancel()` (*asyncio.Handle* 方法), 862
- `cancel()` (*asyncio.Task* 方法), 874
- `cancel()` (*concurrent.futures.Future* 方法), 765
- `cancel()` (*sched.scheduler* 方法), 783
- `cancel()` (*threading.Timer* 方法), 718
- `cancel_dump_traceback_later()` (在 *fault-handler* 模块中), 1493
- `cancel_join_thread()` (*multiprocessing.Queue* 方法), 730
- `cancelled()` (*asyncio.Future* 方法), 872
- `cancelled()` (*concurrent.futures.Future* 方法), 765
- `CancelledError`, 766
- `CannotSendHeader`, 1134
- `CannotSendRequest`, 1134
- `canonic()` (*bdb.Bdb* 方法), 1488
- `canonical()` (*decimal.Context* 方法), 286
- `canonical()` (*decimal.Decimal* 方法), 278
- `capa()` (*poplib.POP3* 方法), 1145
- `capitalize()` (*bytearray* 方法), 56
- `capitalize()` (*bytes* 方法), 56
- `capitalize()` (*str* 方法), 41
- `captured_stderr()` (在 *test.support* 模块中), 1482
- `captured_stdin()` (在 *test.support* 模块中), 1482
- `captured_stdout()` (在 *test.support* 模块中), 1482
- `captureWarnings()` (在 *logging* 模块中), 620
- `capwords()` (在 *string* 模块中), 98
- `casefold()` (*str* 方法), 41
- `cast()` (*memoryview* 方法), 66
- `cast()` (在 *ctypes* 模块中), 700
- `cast()` (在 *typing* 模块中), 1363
- `cat()` (在 *nis* 模块中), 1722
- `catch_warnings` (*warnings* 中的类), 1567
- `category()` (在 *unicodedata* 模块中), 131
- `cbreak()` (在 *curses* 模块中), 645
- `ccc()` (*ftplib.FTP_TLS* 方法), 1144
- `C-contiguous`, 1761
- `CDLL` (*ctypes* 中的类), 695
- `ceil()` (in module *math*), 29
- `ceil()` (在 *math* 模块中), 264
- `center()` (*bytearray* 方法), 54
- `center()` (*bytes* 方法), 54
- `center()` (*str* 方法), 41
- `CERT_NONE()` (在 *ssl* 模块中), 816
- `CERT_OPTIONAL()` (在 *ssl* 模块中), 817
- `CERT_REQUIRED()` (在 *ssl* 模块中), 817
- `cert_store_stats()` (*ssl.SSLContext* 方法), 826
- `cert_time_to_seconds()` (在 *ssl* 模块中), 815
- `CertificateError`, 812
- `certificates`, 832
- `CFUNCTYPE()` (在 *ctypes* 模块中), 698
- `CGI`
 - debugging, 1093
 - exceptions, 1094
 - protocol, 1087
 - security, 1092
 - tracebacks, 1094
- `cgi` (模块), 1087

- `cgi_directories` (`http.server.CGIHTTPRequestHandler` 属性), 1189
- `CGIHandler` (`wsgiref.handlers` 中的类), 1100
- `CGIHTTPRequestHandler` (`http.server` 中的类), 1188
- `cgitb` (模块), 1094
- `CGIXMLRPCRequestHandler` (`xmlrpc.server` 中的类), 1209
- `chain()` (在 `itertools` 模块中), 315
- `chaining`
 - `comparisons`, 28
- `ChainMap` (`collections` 中的类), 195
- `ChainMap` (`typing` 中的类), 1361
- `change_cwd()` (在 `test.support` 模块中), 1482
- `CHANNEL_BINDING_TYPES()` (在 `ssl` 模块中), 821
- `channel_class` (`smtpd.SMTPServer` 属性), 1168
- `channels()` (`ossaudiodev.oss_audio_device` 方法), 1243
- `CHAR_MAX()` (在 `locale` 模块中), 1260
- `character`, 130
- `CharacterDataHandler()`
 - (`xml.parsers.expat.xmlparser` 方法), 1079
- `characters()` (`xml.sax.handler.ContentHandler` 方法), 1068
- `characters_written` (`BlockingIOError` 属性), 84
- `Charset` (`email.charset` 中的类), 980
- `charset()` (`gettext.NullTranslations` 方法), 1250
- `chdir()` (在 `os` 模块中), 525
- `check` (`lzma.LZMADecompressor` 属性), 441
- `check()` (`imaplib.IMAP4` 方法), 1149
- `check()` (在 `tabnanny` 模块中), 1663
- `check__all__()` (在 `test.support` 模块中), 1484
- `check_call()` (在 `subprocess` 模块中), 776
- `check_hostname` (`ssl.SSLContext` 属性), 831
- `check_output()` (`doctest.OutputChecker` 方法), 1387
- `check_output()` (在 `subprocess` 模块中), 777
- `check_returncode()` (`subprocess.CompletedProcess` 方法), 768
- `check_unused_args()` (`string.Formatter` 方法), 91
- `check_warnings()` (在 `test.support` 模块中), 1481
- `checkbox()` (`msilib.Dialog` 方法), 1694
- `checkcache()` (在 `linecache` 模块中), 377
- `checkfuncname()` (在 `bdb` 模块中), 1491
- `CheckList` (`tkinter.tix` 中的类), 1337
- `checksum`
 - Cyclic Redundancy Check, 430
- `chflags()` (在 `os` 模块中), 525
- `chgat()` (`curses.window` 方法), 652
- `childNodes` (`xml.dom.Node` 属性), 1050
- `ChildProcessError`, 84
- `chmod()` (`pathlib.Path` 方法), 352
- `chmod()` (在 `os` 模块中), 525
- `choice()` (在 `random` 模块中), 302
- `choice()` (在 `secrets` 模块中), 505
- `choices` (`optparse.Option` 属性), 1738
- `choices()` (在 `random` 模块中), 302
- `chown()` (在 `os` 模块中), 526
- `chown()` (在 `shutil` 模块中), 381
- `chr()` (内置函数), 7
- `chroot()` (在 `os` 模块中), 526
- `Chunk` (`chunk` 中的类), 1238
- `chunk` (模块), 1238
- `cipher`
 - DES, 1710
- `cipher()` (`ssl.SSLSocket` 方法), 824
- `circle()` (在 `turtle` 模块中), 1270
- `CIRCUMFLEX()` (在 `token` 模块中), 1657
- `CIRCUMFLEXEQUAL()` (在 `token` 模块中), 1657
- `Clamped` (`decimal` 中的类), 290
- `class` -- 类, 1761
- `Class` (`symtable` 中的类), 1656
- `Class browser`, 1341
- `class variable` -- 类变量, 1761
- `classmethod()` (内置函数), 7
- `ClassVar()` (在 `typing` 模块中), 1366
- `CLD_CONTINUED()` (在 `os` 模块中), 548
- `CLD_DUMPED()` (在 `os` 模块中), 548
- `CLD_EXITED()` (在 `os` 模块中), 548
- `CLD_TRAPPED()` (在 `os` 模块中), 548
- `clean()` (`mailbox.Maildir` 方法), 1000
- `cleandoc()` (在 `inspect` 模块中), 1599
- `clear` (`pdb command`), 1497
- `Clear Breakpoint`, 1343
- `clear()` (`asyncio.Event` 方法), 902
- `clear()` (`collections.deque` 方法), 200
- `clear()` (`curses.window` 方法), 652
- `clear()` (`dict` 方法), 72
- `clear()` (`email.message.EmailMessage` 方法), 937
- `clear()` (`frozenset` 方法), 71
- `clear()` (`http.cookiejar.CookieJar` 方法), 1195
- `clear()` (`mailbox.Mailbox` 方法), 999
- `clear()` (`sequence method`), 37
- `clear()` (`threading.Event` 方法), 717
- `clear()` (在 `turtle` 模块中), 1278, 1284
- `clear()` (`xml.etree.ElementTree.Element` 方法), 1041
- `clear_all_breaks()` (`bdb.Bdb` 方法), 1490
- `clear_all_file_breaks()` (`bdb.Bdb` 方法), 1490
- `clear_bpbynumber()` (`bdb.Bdb` 方法), 1490
- `clear_break()` (`bdb.Bdb` 方法), 1490
- `clear_cache()` (在 `filecmp` 模块中), 369
- `clear_content()` (`email.message.EmailMessage` 方法), 937
- `clear_flags()` (`decimal.Context` 方法), 284
- `clear_frames()` (在 `traceback` 模块中), 1587
- `clear_history()` (在 `readline` 模块中), 135
- `clear_session_cookies()`
 - (`http.cookiejar.CookieJar` 方法), 1195
- `clear_traces()` (在 `tracemalloc` 模块中), 1519

- `clear_traps()` (*decimal.Context* 方法), 285
- `clearcache()` (在 *linecache* 模块中), 377
- `ClearData()` (*msilib.Record* 方法), 1692
- `clearok()` (*curses.window* 方法), 652
- `clearscreen()` (在 *turtle* 模块中), 1284
- `clearstamp()` (在 *turtle* 模块中), 1271
- `clearstamps()` (在 *turtle* 模块中), 1271
- `Client()` (在 *multiprocessing.connection* 模块中), 748
- `client_address` (*http.server.BaseHTTPRequestHandler* 属性), 1184
- `clock()` (在 *time* 模块中), 567
- `clock_getres()` (在 *time* 模块中), 567
- `clock_gettime()` (在 *time* 模块中), 567
- `CLOCK_HIGHRES()` (在 *time* 模块中), 572
- `CLOCK_MONOTONIC()` (在 *time* 模块中), 572
- `CLOCK_MONOTONIC_RAW()` (在 *time* 模块中), 572
- `CLOCK_PROCESS_CPUTIME_ID()` (在 *time* 模块中), 572
- `CLOCK_REALTIME()` (在 *time* 模块中), 572
- `clock_settime()` (在 *time* 模块中), 567
- `CLOCK_THREAD_CPUTIME_ID()` (在 *time* 模块中), 572
- `clone()` (*email.generator.BytesGenerator* 方法), 943
- `clone()` (*email.generator.Generator* 方法), 944
- `clone()` (*email.policy.Policy* 方法), 947
- `clone()` (*pipes.Template* 方法), 1717
- `clone()` (在 *turtle* 模块中), 1282
- `cloneNode()` (*xml.dom.Node* 方法), 1051
- `close()` (*aifc.aifc* 方法), 1231, 1232
- `close()` (*asyncio.AbstractEventLoop* 方法), 853
- `close()` (*asyncio.BaseSubprocessTransport* 方法), 882
- `close()` (*asyncio.BaseTransport* 方法), 879
- `close()` (*asyncio.Server* 方法), 862
- `close()` (*asyncio.StreamWriter* 方法), 891
- `close()` (*asyncore.dispatcher* 方法), 915
- `close()` (*chunk.Chunk* 方法), 1238
- `close()` (*contextlib.ExitStack* 方法), 1573
- `close()` (*dbm.dumb.dumbdbm* 方法), 408
- `close()` (*dbm.gnu.gdbm* 方法), 407
- `close()` (*dbm.ndbm.ndbm* 方法), 407
- `close()` (*email.parser.BytesFeedParser* 方法), 939
- `close()` (*ftplib.FTP* 方法), 1144
- `close()` (*html.parser.HTMLParser* 方法), 1027
- `close()` (*http.client.HTTPConnection* 方法), 1136
- `close()` (*imaplib.IMAP4* 方法), 1149
- `close()` (*io.IOBase* 方法), 557
- `close()` (*logging.FileHandler* 方法), 632
- `close()` (*logging.Handler* 方法), 610
- `close()` (*logging.handlers.MemoryHandler* 方法), 641
- `close()` (*logging.handlers.NTEventLogHandler* 方法), 639
- `close()` (*logging.handlers.SocketHandler* 方法), 636
- `close()` (*logging.handlers.SysLogHandler* 方法), 637
- `close()` (*mailbox.Mailbox* 方法), 999
- `close()` (*mailbox.Maildir* 方法), 1001
- `close()` (*mailbox.MH* 方法), 1003
- `close()` (*mmap.mmap* 方法), 926
- `Close()` (*msilib.View* 方法), 1691
- `close()` (*multiprocessing.connection.Connection* 方法), 732
- `close()` (*multiprocessing.connection.Listener* 方法), 749
- `close()` (*multiprocessing.pool.Pool* 方法), 747
- `close()` (*multiprocessing.Queue* 方法), 730
- `close()` (*ossaudiodev.oss_audio_device* 方法), 1242
- `close()` (*ossaudiodev.oss_mixer_device* 方法), 1244
- `close()` (*os.scandir* 方法), 531
- `close()` (*select.devpoll* 方法), 843
- `close()` (*select.epoll* 方法), 845
- `close()` (*select.kqueue* 方法), 846
- `close()` (*selectors.BaseSelector* 方法), 850
- `close()` (*shelve.Shelf* 方法), 401
- `close()` (*socket.socket* 方法), 800
- `close()` (*sqlite3.Connection* 方法), 412
- `close()` (*sqlite3.Cursor* 方法), 419
- `close()` (*sunau.AU_read* 方法), 1234
- `close()` (*sunau.AU_write* 方法), 1235
- `close()` (*tarfile.TarFile* 方法), 456
- `close()` (*telnetlib.Telnet* 方法), 1172
- `close()` (*urllib.request.BaseHandler* 方法), 1111
- `close()` (在 *fileinput* 模块中), 363
- `close()` (在 *os* 模块中), 516
- `close()` (*wave.Wave_read* 方法), 1236
- `close()` (*wave.Wave_write* 方法), 1237
- `Close()` (*winreg.PyHKEY* 方法), 1704
- `close()` (*xml.etree.ElementTree.TreeBuilder* 方法), 1044
- `close()` (*xml.etree.ElementTree.XMLParser* 方法), 1045
- `close()` (*xml.etree.ElementTree.XMLPullParser* 方法), 1046
- `close()` (*xml.sax.xmlreader.IncrementalParser* 方法), 1073
- `close()` (*zipfile.ZipFile* 方法), 445
- `close_connection` (*http.server.BaseHTTPRequestHandler* 属性), 1184
- `close_when_done()` (*asynchat.async_chat* 方法), 917
- `closed` (*http.client.HTTPResponse* 属性), 1138
- `closed` (*io.IOBase* 属性), 557
- `closed` (*mmap.mmap* 属性), 926
- `closed` (*ossaudiodev.oss_audio_device* 属性), 1244
- `closed` (*select.devpoll* 属性), 843
- `closed` (*select.epoll* 属性), 845
- `closed` (*select.kqueue* 属性), 846
- `CloseKey()` (在 *winreg* 模块中), 1696
- `closelog()` (在 *syslog* 模块中), 1723
- `closerange()` (在 *os* 模块中), 516
- `closing()` (在 *contextlib* 模块中), 1569
- `clrtoebot()` (*curses.window* 方法), 652

- `clrtoeol()` (*curses.window* 方法), 652
- `cmath` (模块), 269
- `cmd`
 - 模块, 1494
- `Cmd` (*cmd* 中的类), 1296
- `cmd` (*subprocess.CalledProcessError* 属性), 769
- `cmd` (*subprocess.TimeoutExpired* 属性), 768
- `cmd` (模块), 1296
- `cmdloop()` (*cmd.Cmd* 方法), 1296
- `cmdqueue` (*cmd.Cmd* 属性), 1297
- `cmp()` (在 *filecmp* 模块中), 369
- `cmp_op()` (在 *dis* 模块中), 1680
- `cmp_to_key()` (在 *functools* 模块中), 327
- `cmpfiles()` (在 *filecmp* 模块中), 369
- `CMSG_LEN()` (在 *socket* 模块中), 798
- `CMSG_SPACE()` (在 *socket* 模块中), 798
- `CO_ASYNC_GENERATOR()` (在 *inspect* 模块中), 1609
- `CO_COROUTINE()` (在 *inspect* 模块中), 1609
- `CO_GENERATOR()` (在 *inspect* 模块中), 1609
- `CO_ITERABLE_COROUTINE()` (在 *inspect* 模块中), 1609
- `CO_NESTED()` (在 *inspect* 模块中), 1609
- `CO_NEWLOCALS()` (在 *inspect* 模块中), 1609
- `CO_NOFREE()` (在 *inspect* 模块中), 1609
- `CO_OPTIMIZED()` (在 *inspect* 模块中), 1609
- `CO_VARARGS()` (在 *inspect* 模块中), 1609
- `CO_VARKEYWORDS()` (在 *inspect* 模块中), 1609
- `code` (*SystemExit* 属性), 83
- `code` (*urllib.error.HTTPError* 属性), 1130
- `code` (模块), 1615
- `code` (*xml.etree.ElementTree.ParseError* 属性), 1047
- `code` (*xml.parsers.expat.ExpatError* 属性), 1080
- `code object`, 76, 403
- `code_info()` (在 *dis* 模块中), 1669
- `CodecInfo` (*codecs* 中的类), 145
- `Codecs`, 145
 - `decode`, 145
 - `encode`, 145
- `codecs` (模块), 145
- `coded_value` (*http.cookies.Morsel* 属性), 1191
- `codeop` (模块), 1617
- `codepoint2name()` (在 *html.entities* 模块中), 1030
- `codes()` (在 *xml.parsers.expat.errors* 模块中), 1082
- `CODESET()` (在 *locale* 模块中), 1257
- `CodeType()` (在 *types* 模块中), 231
- `coercion` -- 强制类型转换, 1761
- `col_offset` (*ast.AST* 属性), 1649
- `collapse_addresses()` (在 *ipaddress* 模块中), 1225
- `collapse_rfc2231_value()` (在 *email.utils* 模块中), 985
- `collect()` (在 *gc* 模块中), 1593
- `collect_incoming_data()` (*asynchat.async_chat* 方法), 917
- `Collection` (*collections.abc* 中的类), 212
- `Collection` (*typing* 中的类), 1359
- `collections` (模块), 195
- `collections.abc` (模块), 210
- `colno` (*json.JSONDecodeError* 属性), 993
- `colno` (*re.error* 属性), 108
- `COLON()` (在 *token* 模块中), 1657
- `color()` (在 *turtle* 模块中), 1276
- `color_content()` (在 *curses* 模块中), 645
- `color_pair()` (在 *curses* 模块中), 645
- `colormode()` (在 *turtle* 模块中), 1288
- `colorsys` (模块), 1239
- `COLS`, 650
- `column()` (*tkinter.ttk.Treeview* 方法), 1328
- `COLUMNS`, 650
- `columns` (*os.terminal_size* 属性), 523
- `combinations()` (在 *itertools* 模块中), 316
- `combinations_with_replacement()` (在 *itertools* 模块中), 316
- `combine()` (*datetime.datetime* 类方法), 172
- `combining()` (在 *unicodedata* 模块中), 131
- `ComboBox` (*tkinter.tix* 中的类), 1336
- `Combobox` (*tkinter.ttk* 中的类), 1322
- `COMMA()` (在 *token* 模块中), 1657
- `command` (*http.server.BaseHTTPRequestHandler* 属性), 1185
- `CommandCompiler` (*codeop* 中的类), 1618
- `commands` (*pdb command*), 1497
- `comment` (*http.cookiejar.Cookie* 属性), 1200
- `comment` (*zipfile.ZipFile* 属性), 447
- `comment` (*zipfile.ZipInfo* 属性), 449
- `COMMENT()` (在 *tokenize* 模块中), 1659
- `Comment()` (在 *xml.etree.ElementTree* 模块中), 1039
- `comment_url` (*http.cookiejar.Cookie* 属性), 1200
- `commenters` (*shlex.shlex* 属性), 1303
- `CommentHandler()` (*xml.parsers.expat.xmlparser* 方法), 1079
- `commit()` (*msilib.CAB* 方法), 1692
- `Commit()` (*msilib.Database* 方法), 1690
- `commit()` (*sqlite3.Connection* 方法), 412
- `common` (*filecmp.dircmp* 属性), 370
- `Common Gateway Interface`, 1087
- `common_dirs` (*filecmp.dircmp* 属性), 370
- `common_files` (*filecmp.dircmp* 属性), 370
- `common_funny` (*filecmp.dircmp* 属性), 370
- `common_types()` (在 *mimetypes* 模块中), 1015
- `commonpath()` (在 *os.path* 模块中), 357
- `commonprefix()` (在 *os.path* 模块中), 357
- `communicate()` (*asyncio.asyncio.subprocess.Process* 方法), 897
- `communicate()` (*subprocess.Popen* 方法), 773
- `compare()` (*decimal.Context* 方法), 286
- `compare()` (*decimal.Decimal* 方法), 278
- `compare()` (*difflib.Differ* 方法), 124

- `compare_digest()` (在 *hmac* 模块中), 504
- `compare_digest()` (在 *secrets* 模块中), 506
- `compare_networks()` (*ipaddress.IPv4Network* 方法), 1221
- `compare_networks()` (*ipaddress.IPv6Network* 方法), 1222
- `COMPARE_OP` (*opcode*), 1677
- `compare_signal()` (*decimal.Context* 方法), 286
- `compare_signal()` (*decimal.Decimal* 方法), 278
- `compare_to()` (*tracemalloc.Snapshot* 方法), 1522
- `compare_total()` (*decimal.Context* 方法), 286
- `compare_total()` (*decimal.Decimal* 方法), 278
- `compare_total_mag()` (*decimal.Context* 方法), 286
- `compare_total_mag()` (*decimal.Decimal* 方法), 278
- `comparing`
 - objects, 28
- `comparison`
 - operator, 28
- `COMPARISON_FLAGS()` (在 *doctest* 模块中), 1376
- `comparisons`
 - chaining, 28
- `Compat32` (*email.policy* 中的类), 950
- `compat32()` (在 *email.policy* 模块中), 951
- `compile`
 - ☐置函数, 76, 231, 1647
- `Compile` (*codeop* 中的类), 1618
- `compile()` (*parser.ST* 方法), 1648
- `compile()` (☐置函数), 7
- `compile()` (在 *py_compile* 模块中), 1664
- `compile()` (在 *re* 模块中), 103
- `compile_command()` (在 *code* 模块中), 1616
- `compile_command()` (在 *codeop* 模块中), 1617
- `compile_dir()` (在 *compileall* 模块中), 1666
- `compile_file()` (在 *compileall* 模块中), 1667
- `compile_path()` (在 *compileall* 模块中), 1667
- `compileall` (模块), 1665
- `compileall` 命令行选项
 - b, 1666
 - d *destdir*, 1665
 - directory ..., 1665
 - f, 1665
 - file ..., 1665
 - i *list*, 1666
 - j *N*, 1666
 - l, 1665
 - q, 1665
 - r, 1666
 - x *regex*, 1666
- `compilest()` (在 *parser* 模块中), 1647
- `complete()` (*rlcompleter.Completer* 方法), 138
- `complete_statement()` (在 *sqlite3* 模块中), 411
- `completedefault()` (*cmd.Cmd* 方法), 1297
- `CompletedProcess` (*subprocess* 中的类), 768
- `complex`
 - ☐置函数, 29
- `Complex` (*numbers* 中的类), 261
- `complex` (☐置类), 8
- `complex number`
 - literals, 29
 - 对象, 29
- `complex number -- 复数`, 1761
- `compress()` (*bz2.BZ2Compressor* 方法), 436
- `compress()` (*lzma.LZMACompressor* 方法), 440
- `compress()` (在 *bz2* 模块中), 437
- `compress()` (在 *gzip* 模块中), 434
- `compress()` (在 *itertools* 模块中), 317
- `compress()` (在 *lzma* 模块中), 441
- `compress()` (在 *zlib* 模块中), 429
- `compress()` (*zlib.Compress* 方法), 431
- `compress_size` (*zipfile.ZipInfo* 属性), 450
- `compress_type` (*zipfile.ZipInfo* 属性), 449
- `compressed` (*ipaddress.IPv4Address* 属性), 1215
- `compressed` (*ipaddress.IPv4Network* 属性), 1220
- `compressed` (*ipaddress.IPv6Address* 属性), 1216
- `compressed` (*ipaddress.IPv6Network* 属性), 1222
- `compression()` (*ssl.SSLSocket* 方法), 824
- `CompressionError`, 452
- `compressobj()` (在 *zlib* 模块中), 430
- `COMSPEC`, 547, 771
- `concat()` (在 *operator* 模块中), 336
- `concatenation`
 - operation, 35
- `concurrent.futures` (模块), 761
- `Condition` (*asyncio* 中的类), 903
- `Condition` (*multiprocessing* 中的类), 734
- `condition` (*pdb command*), 1497
- `Condition` (*threading* 中的类), 715
- `condition()` (*msilib.Control* 方法), 1693
- `Condition()` (*multiprocessing.managers.SyncManager* 方法), 740
- `ConfigParser` (*configparser* 中的类), 480
- `configparser` (模块), 468
- `configuration`
 - file, 468
 - file, debugger, 1496
 - file, path, 1610
- `configuration information`, 1558
- `configure()` (*tkinter.ttk.Style* 方法), 1332
- `configure_mock()` (*unittest.mock.Mock* 方法), 1424
- `confstr()` (在 *os* 模块中), 551
- `confstr_names()` (在 *os* 模块中), 551
- `conjugate()` (*complex number method*), 29
- `conjugate()` (*decimal.Decimal* 方法), 279
- `conjugate()` (*numbers.Complex* 方法), 261
- `conn` (*smtpd.SMTPChannel* 属性), 1169
- `connect()` (*asyncore.dispatcher* 方法), 914
- `connect()` (*ftplib.FTP* 方法), 1141
- `connect()` (*http.client.HTTPConnection* 方法), 1136

- `connect()` (*multiprocessing.managers.BaseManager* 方法), 739
- `connect()` (*smtplib.SMTP* 方法), 1162
- `connect()` (*socket.socket* 方法), 800
- `connect()` (在 *sqlite3* 模块中), 410
- `connect_accepted_socket()` (*asyncio.BaseEventLoop* 方法), 858
- `connect_ex()` (*socket.socket* 方法), 800
- `connect_read_pipe()` (*asyncio.AbstractEventLoop* 方法), 860
- `connect_write_pipe()` (*asyncio.AbstractEventLoop* 方法), 860
- `Connection` (*multiprocessing.connection* 中的类), 732
- `Connection` (*sqlite3* 中的类), 412
- `connection` (*sqlite3.Cursor* 属性), 420
- `connection_lost()` (*asyncio.BaseProtocol* 方法), 883
- `connection_made()` (*asyncio.BaseProtocol* 方法), 882
- `ConnectionAbortedError`, 85
- `ConnectionError`, 84
- `ConnectionRefusedError`, 85
- `ConnectionResetError`, 85
- `ConnectRegistry()` (在 *winreg* 模块中), 1696
- `const` (*optparse.Option* 属性), 1738
- `constructor()` (在 *copyreg* 模块中), 400
- `consumed` (*asyncio.LimitOverrunError* 属性), 892
- `container`
iteration over, 34
- `Container` (*collections.abc* 中的类), 211
- `Container` (*typing* 中的类), 1359
- `contains()` (在 *operator* 模块中), 336
- `content type`
MIME, 1014
- `content_manager` (*email.policy.EmailPolicy* 属性), 949
- `content_type` (*email.headerregistry.ContentTypeHeader* 属性), 955
- `ContentDispositionHeader` (*email.headerregistry* 中的类), 955
- `ContentHandler` (*xml.sax.handler* 中的类), 1065
- `ContentManager` (*email.contentmanager* 中的类), 958
- `contents` (*ctypes._Pointer* 属性), 707
- `ContentTooShortError`, 1130
- `ContentTransferEncoding` (*email.headerregistry* 中的类), 955
- `ContentTypeHeader` (*email.headerregistry* 中的类), 955
- `Context` (*decimal* 中的类), 284
- `context` (*ssl.SSLSocket* 属性), 825
- `context management protocol`, 74
- `context manager`, 74
- `context manager` -- 上下文管理器, 1761
- `context_diff()` (在 *difflib* 模块中), 118
- `ContextDecorator` (*contextlib* 中的类), 1570
- `contextlib` (模块), 1568
- `ContextManager` (*typing* 中的类), 1360
- `contextmanager()` (在 *contextlib* 模块中), 1568
- `contiguous` -- 连续, 1761
- `contiguous` (*memoryview* 属性), 69
- `continue` (*pdb command*), 1498
- `CONTINUE_LOOP` (*opcode*), 1674
- `Control` (*msilib* 中的类), 1693
- `Control` (*tkinter.tix* 中的类), 1336
- `control()` (*msilib.Dialog* 方法), 1694
- `control()` (*select.kqueue* 方法), 846
- `controlnames()` (在 *curses.ascii* 模块中), 664
- `controls()` (*ossaudiodev.oss_mixer_device* 方法), 1244
- `ConversionError`, 488
- `conversions`
numeric, 29
- `convert_arg_line_to_args()` (*argparse.ArgumentParser* 方法), 602
- `convert_field()` (*string.Formatter* 方法), 91
- `Cookie` (*http.cookiejar* 中的类), 1194
- `CookieError`, 1189
- `CookieJar` (*http.cookiejar* 中的类), 1193
- `cookiejar` (*urllib.request.HTTPCookieProcessor* 属性), 1113
- `CookiePolicy` (*http.cookiejar* 中的类), 1193
- `Coordinated Universal Time`, 566
- `Copy`, 1343
- `copy`
protocol, 393
模块, 400
- `copy` (模块), 234
- `copy()` (*collections.deque* 方法), 200
- `copy()` (*decimal.Context* 方法), 285
- `copy()` (*dict* 方法), 72
- `copy()` (*frozenset* 方法), 70
- `copy()` (*hashlib.hash* 方法), 495
- `copy()` (*hmac.HMAC* 方法), 504
- `copy()` (*http.cookies.Morsel* 方法), 1191
- `copy()` (*imaplib.IMAP4* 方法), 1149
- `copy()` (*pipes.Template* 方法), 1718
- `copy()` (*sequence method*), 37
- `copy()` (*types.MappingProxyType* 方法), 232
- `copy()` (在 *copy* 模块中), 234
- `copy()` (在 *multiprocessing.sharedctypes* 模块中), 738
- `copy()` (在 *shutil* 模块中), 379
- `copy()` (*zlib.Compress* 方法), 431
- `copy()` (*zlib.Decompress* 方法), 432
- `copy2()` (在 *shutil* 模块中), 379
- `copy_abs()` (*decimal.Context* 方法), 286
- `copy_abs()` (*decimal.Decimal* 方法), 279
- `copy_decimal()` (*decimal.Context* 方法), 285
- `copy_location()` (在 *ast* 模块中), 1653

- `copy_negate()` (*decimal.Context* 方法), 286
- `copy_negate()` (*decimal.Decimal* 方法), 279
- `copy_sign()` (*decimal.Context* 方法), 286
- `copy_sign()` (*decimal.Decimal* 方法), 279
- `copyfile()` (在 *shutil* 模块中), 378
- `copyfileobj()` (在 *shutil* 模块中), 378
- copying files, 378
- `copymode()` (在 *shutil* 模块中), 378
- `copyreg` (模块), 400
- `copyright` (环境变量), 26
- `copyright()` (在 *sys* 模块中), 1544
- `copysign()` (在 *math* 模块中), 264
- `copystat()` (在 *shutil* 模块中), 379
- `copytree()` (在 *shutil* 模块中), 380
- `coroutine` -- 协程, 1761
- `Coroutine` (*collections.abc* 中的类), 213
- `Coroutine` (*typing* 中的类), 1360
- `coroutine function` -- 协程函数, 1761
- `coroutine()` (在 *asyncio* 模块中), 869
- `coroutine()` (在 *types* 模块中), 233
- `CoroutineType()` (在 *types* 模块中), 231
- `cos()` (在 *cmath* 模块中), 270
- `cos()` (在 *math* 模块中), 267
- `cosh()` (在 *cmath* 模块中), 271
- `cosh()` (在 *math* 模块中), 267
- `count` (*tracemalloc.Statistic* 属性), 1523
- `count` (*tracemalloc.StatisticDiff* 属性), 1523
- `count()` (*array.array* 方法), 221
- `count()` (*bytearray* 方法), 52
- `count()` (*bytes* 方法), 52
- `count()` (*collections.deque* 方法), 200
- `count()` (*sequence method*), 35
- `count()` (*str* 方法), 41
- `count()` (在 *itertools* 模块中), 317
- `count_diff` (*tracemalloc.StatisticDiff* 属性), 1523
- `Counter` (*collections* 中的类), 197
- `Counter` (*typing* 中的类), 1361
- `countOf()` (在 *operator* 模块中), 336
- `countTestCases()` (*unittest.TestCase* 方法), 1407
- `countTestCases()` (*unittest.TestSuite* 方法), 1409
- `CoverageResults` (*trace* 中的类), 1514
- `cProfile` (模块), 1502
- CPU time, 567
- `cpu_count()` (在 *multiprocessing* 模块中), 731
- `cpu_count()` (在 *os* 模块中), 551
- CPython, 1761
- `crawl_delay()` (*urllib.robotparser.RobotFileParser* 方法), 1130
- CRC (*zipfile.ZipInfo* 属性), 450
- `crc32()` (在 *binascii* 模块中), 1021
- `crc32()` (在 *zlib* 模块中), 430
- `crc_hqx()` (在 *binascii* 模块中), 1021
- create <tarfile> <source1> ...
 <sourceN>
- tarfile 命令行选项, 458
- `create()` (*imaplib.IMAP4* 方法), 1149
- `create()` (*venv.EnvBuilder* 方法), 1530
- `create()` (在 *venv* 模块中), 1531
- `create_aggregate()` (*sqlite3.Connection* 方法), 413
- `create_archive()` (在 *zipapp* 模块中), 1536
- `create_autospec()` (在 *unittest.mock* 模块中), 1447
- `create_collation()` (*sqlite3.Connection* 方法), 413
- `create_configuration()` (*venv.EnvBuilder* 方法), 1531
- `create_connection()` (*asyncio.AbstractEventLoop* 方法), 855
- `create_connection()` (在 *socket* 模块中), 795
- `create_datagram_endpoint()` (*asyncio.AbstractEventLoop* 方法), 856
- `create_decimal()` (*decimal.Context* 方法), 285
- `create_decimal_from_float()` (*decimal.Context* 方法), 285
- `create_default_context()` (在 *ssl* 模块中), 813
- `create_function()` (*sqlite3.Connection* 方法), 413
- `create_future()` (*asyncio.AbstractEventLoop* 方法), 855
- `create_module()` (*importlib.abc.Loader* 方法), 1631
- `create_module()` (*importlib.machinery.ExtensionFileLoader* 方法), 1638
- `CREATE_NEW_CONSOLE()` (在 *subprocess* 模块中), 776
- `CREATE_NEW_PROCESS_GROUP()` (在 *subprocess* 模块中), 776
- `create_server()` (*asyncio.AbstractEventLoop* 方法), 857
- `create_socket()` (*asyncore.dispatcher* 方法), 914
- `create_stats()` (*profile.Profile* 方法), 1503
- `create_string_buffer()` (在 *ctypes* 模块中), 700
- `create_subprocess_exec()` (在 *asyncio* 模块中), 895
- `create_subprocess_shell()` (在 *asyncio* 模块中), 895
- `create_system` (*zipfile.ZipInfo* 属性), 449
- `create_task()` (*asyncio.AbstractEventLoop* 方法), 855
- `create_unicode_buffer()` (在 *ctypes* 模块中), 700
- `create_unix_connection()` (*asyncio.AbstractEventLoop* 方法), 857
- `create_unix_server()` (*asyncio.AbstractEventLoop* 方法), 858
- `create_version` (*zipfile.ZipInfo* 属性), 449
- `createAttribute()` (*xml.dom.Document* 方法), 1052
- `createAttributeNS()` (*xml.dom.Document* 方法), 1052
- `createComment()` (*xml.dom.Document* 方法), 1052

- createDocument() (*xml.dom.DOMImplementation* 方法), 1049
- createDocumentType() (*xml.dom.DOMImplementation* 方法), 1049
- createElement() (*xml.dom.Document* 方法), 1052
- createElementNS() (*xml.dom.Document* 方法), 1052
- createfilehandler() (*tkinter.Widget.tk* 方法), 1317
- CreateKey() (在 *winreg* 模块中), 1696
- CreateKeyEx() (在 *winreg* 模块中), 1697
- createLock() (*logging.Handler* 方法), 610
- createLock() (*logging.NullHandler* 方法), 632
- createProcessingInstruction() (*xml.dom.Document* 方法), 1052
- CreateRecord() (在 *msilib* 模块中), 1690
- createSocket() (*logging.handlers.SocketHandler* 方法), 636
- createTextNode() (*xml.dom.Document* 方法), 1052
- credits (设置变量), 26
- critical() (*logging.Logger* 方法), 608
- critical() (在 *logging* 模块中), 618
- CRNCYSTR() (在 *locale* 模块中), 1258
- cross() (在 *audioop* 模块中), 1228
- crypt
模块, 1708
- crypt (模块), 1710
- crypt() (在 *crypt* 模块中), 1711
- crypt(3), 1710, 1711
- cryptography, 493
- csv, 461
- csv (模块), 461
- cte (*email.headerregistry.ContentTransferEncoding* 属性), 956
- cte_type (*email.policy.Policy* 属性), 946
- ctermid() (在 *os* 模块中), 510
- ctime() (*datetime.date* 方法), 169
- ctime() (*datetime.datetime* 方法), 177
- ctime() (在 *time* 模块中), 567
- ctrl() (在 *curses.ascii* 模块中), 664
- CTRL_BREAK_EVENT() (在 *signal* 模块中), 920
- CTRL_C_EVENT() (在 *signal* 模块中), 920
- ctypes (模块), 675
- curdir() (在 *os* 模块中), 552
- currency() (在 *locale* 模块中), 1259
- current() (*tkinter.ttk.Combobox* 方法), 1322
- current_process() (在 *multiprocessing* 模块中), 731
- current_task() (*asyncio.Task* 类方法), 874
- current_thread() (在 *threading* 模块中), 709
- CurrentByteIndex (*xml.parsers.expat.xmlparser* 属性), 1078
- CurrentColumnNumber (*xml.parsers.expat.xmlparser* 属性), 1078
- currentframe() (在 *inspect* 模块中), 1607
- CurrentLineNumber (*xml.parsers.expat.xmlparser* 属性), 1078
- curs_set() (在 *curses* 模块中), 645
- curses (模块), 644
- curses.ascii (模块), 662
- curses.panel (模块), 665
- curses.textpad (模块), 661
- Cursor (*sqlite3* 中的类), 417
- cursor() (*sqlite3.Connection* 方法), 412
- cursyncup() (*curses.window* 方法), 652
- Cut, 1343
- cwd() (*ftplib.FTP* 方法), 1143
- cwd() (*pathlib.Path* 类方法), 351
- cycle() (在 *itertools* 模块中), 318
- Cyclic Redundancy Check, 430
- ## D
- d destdir
compileall 命令行选项, 1665
- D_FMT() (在 *locale* 模块中), 1257
- D_T_FMT() (在 *locale* 模块中), 1257
- daemon (*multiprocessing.Process* 属性), 727
- daemon (*threading.Thread* 属性), 712
- data
packingbinary, 139
tabular, 461
- data (*collections.UserDict* 属性), 209
- data (*collections.UserList* 属性), 209
- Data (*plistlib* 中的类), 491
- data (*select.kevent* 属性), 848
- data (*selectors.SelectorKey* 属性), 849
- data (*urllib.request.Request* 属性), 1109
- data (*xml.dom.Comment* 属性), 1054
- data (*xml.dom.ProcessingInstruction* 属性), 1055
- data (*xml.dom.Text* 属性), 1054
- data (*xmlrpc.client.Binary* 属性), 1204
- data() (*xml.etree.ElementTree.TreeBuilder* 方法), 1045
- data_open() (*urllib.request.DataHandler* 方法), 1115
- data_received() (*asyncio.Protocol* 方法), 883
- database
Unicode, 130
- DatabaseError, 421
- databases, 408
- datagram_received() (*asyncio.DatagramProtocol* 方法), 884
- DatagramHandler (*logging.handlers* 中的类), 637
- DatagramProtocol (*asyncio* 中的类), 882
- DatagramRequestHandler (*socketserver* 中的类), 1180
- DataHandler (*urllib.request* 中的类), 1109
- date (*datetime* 中的类), 167
- date() (*datetime.datetime* 方法), 174
- date() (*nntplib.NNTP* 方法), 1159

- date_time (*zipfile.ZipInfo* 属性), 449
date_time_string()
 (*http.server.BaseHTTPRequestHandler* 方法), 1187
DateHeader (*email.headerregistry* 中的类), 954
datetime (*datetime* 中的类), 171
datetime (*email.headerregistry.DateHeader* 属性), 954
datetime (模块), 163
DateTime (*xmlrpc.client* 中的类), 1204
day (*datetime.date* 属性), 168
day (*datetime.datetime* 属性), 173
day_abbr() (在 *calendar* 模块中), 194
day_name() (在 *calendar* 模块中), 194
Daylight Saving Time, 566
daylight() (在 *time* 模块中), 573
DbfilenameShelf (*shelve* 中的类), 402
dbm (模块), 404
dbm.dumb (模块), 408
dbm.gnu
 模块, 402
dbm.gnu (模块), 406
dbm.ndbm
 模块, 402
dbm.ndbm (模块), 407
dcgettext() (在 *locale* 模块中), 1261
debug (*imaplib.IMAP4* 属性), 1153
debug (*shlex.shlex* 属性), 1303
debug (*zipfile.ZipFile* 属性), 447
debug() (*logging.Logger* 方法), 607
debug() (*pipes.Template* 方法), 1717
debug() (*unittest.TestCase* 方法), 1401
debug() (*unittest.TestSuite* 方法), 1409
debug() (在 *doctest* 模块中), 1389
debug() (在 *logging* 模块中), 617
DEBUG() (在 *re* 模块中), 104
DEBUG_BYTECODE_SUFFIXES() (在 *importlib.machinery* 模块中), 1635
DEBUG_COLLECTABLE() (在 *gc* 模块中), 1595
DEBUG_LEAK() (在 *gc* 模块中), 1595
DEBUG_SAVEALL() (在 *gc* 模块中), 1595
debug_src() (在 *doctest* 模块中), 1389
DEBUG_STATS() (在 *gc* 模块中), 1595
DEBUG_UNCOLLECTABLE() (在 *gc* 模块中), 1595
debugger, 1342, 1549, 1555
 configuration file, 1496
debugging, 1494
 CGI, 1093
DebuggingServer (*smtpd* 中的类), 1168
debuglevel (*http.client.HTTPResponse* 属性), 1138
DebugRunner (*doctest* 中的类), 1389
Decimal (*decimal* 中的类), 276
decimal (模块), 272
decimal() (在 *unicodedata* 模块中), 131
DecimalException (*decimal* 中的类), 290
decode
 Codecs, 145
decode (*codecs.CodecInfo* 属性), 145
decode() (*bytearray* 方法), 52
decode() (*bytes* 方法), 52
decode() (*codecs.Codec* 方法), 149
decode() (*codecs.IncrementalDecoder* 方法), 151
decode() (*json.JSONDecoder* 方法), 991
decode() (在 *base64* 模块中), 1019
decode() (在 *codecs* 模块中), 145
decode() (在 *quopri* 模块中), 1022
decode() (在 *uu* 模块中), 1023
decode() (*xmlrpc.client.Binary* 方法), 1205
decode() (*xmlrpc.client.DateTime* 方法), 1204
decode_header() (在 *email.header* 模块中), 979
decode_header() (在 *nntplib* 模块中), 1160
decode_params() (在 *email.utils* 模块中), 985
decode_rfc2231() (在 *email.utils* 模块中), 985
decode_source() (在 *importlib.util* 模块中), 1640
decodebytes() (在 *base64* 模块中), 1019
DecodedGenerator (*email.generator* 中的类), 944
decodestring() (在 *base64* 模块中), 1019
decodestring() (在 *quopri* 模块中), 1022
decomposition() (在 *unicodedata* 模块中), 131
decompress() (*bz2.BZ2Decompressor* 方法), 437
decompress() (*Izma.LZMADecompressor* 方法), 440
decompress() (在 *bz2* 模块中), 437
decompress() (在 *gzip* 模块中), 434
decompress() (在 *Izma* 模块中), 441
decompress() (在 *zlib* 模块中), 430
decompress() (*zlib.Decompress* 方法), 432
decompressobj() (在 *zlib* 模块中), 431
decorator -- 装饰器, 1761
dedent() (在 *textwrap* 模块中), 128
DEDENT() (在 *token* 模块中), 1657
deepcopy() (在 *copy* 模块中), 234
def_prog_mode() (在 *curses* 模块中), 645
def_shell_mode() (在 *curses* 模块中), 645
default (*inspect.Parameter* 属性), 1601
default (*optparse.Option* 属性), 1738
default() (*cmd.Cmd* 方法), 1297
default() (*json.JSONEncoder* 方法), 992
default() (在 *email.policy* 模块中), 950
DEFAULT() (在 *unittest.mock* 模块中), 1445
DEFAULT_BUFFER_SIZE() (在 *io* 模块中), 555
default_bufsize() (在 *xml.dom.pulldom* 模块中), 1063
default_exception_handler() (asyn-
 cio.AbstractEventLoop 方法), 861
default_factory (*collections.defaultdict* 属性), 203
DEFAULT_FORMAT() (在 *tarfile* 模块中), 453
DEFAULT_IGNORES() (在 *filecmp* 模块中), 370
default_open() (*urllib.request.BaseHandler* 方法), 1112

- DEFAULT_PROTOCOL() (在 *pickle* 模块中), 389
 default_timer() (在 *timeit* 模块中), 1508
 DefaultContext (*decimal* 中的类), 284
 DefaultCookiePolicy (*http.cookiejar* 中的类), 1193
 defaultdict (*collections* 中的类), 203
 DefaultDict (*typing* 中的类), 1361
 DefaultHandler() (*xml.parsers.expat.xmlparser* 方法), 1079
 DefaultHandlerExpand() (*xml.parsers.expat.xmlparser* 方法), 1079
 defaults() (*configparser.ConfigParser* 方法), 481
 DefaultSelector (*selectors* 中的类), 850
 defaultTestLoader() (在 *unittest* 模块中), 1414
 defaultTestResult() (*unittest.TestCase* 方法), 1407
 defects (*email.headerregistry.BaseHeader* 属性), 953
 defects (*email.message.EmailMessage* 属性), 938
 defects (*email.message.Message* 属性), 974
 defpath() (在 *os* 模块中), 552
 DefragResult (*urllib.parse* 中的类), 1127
 DefragResultBytes (*urllib.parse* 中的类), 1127
 degrees() (在 *math* 模块中), 267
 degrees() (在 *turtle* 模块中), 1273
 del
 语句, 37, 71
 del_param() (*email.message.EmailMessage* 方法), 934
 del_param() (*email.message.Message* 方法), 972
 delattr() (设置函数), 8
 delay() (在 *turtle* 模块中), 1285
 delay_output() (在 *curses* 模块中), 645
 delayload (*http.cookiejar.FileCookieJar* 属性), 1196
 delch() (*curses.window* 方法), 652
 dele() (*poplib.POP3* 方法), 1146
 delete() (*ftplib.FTP* 方法), 1143
 delete() (*imaplib.IMAP4* 方法), 1149
 delete() (*tkinter.ttk.Treeview* 方法), 1329
 DELETE_ATTR (*opcode*), 1676
 DELETE_DEREF (*opcode*), 1679
 DELETE_FAST (*opcode*), 1678
 DELETE_GLOBAL (*opcode*), 1676
 DELETE_NAME (*opcode*), 1676
 DELETE_SUBSCR (*opcode*), 1673
 deleteacl() (*imaplib.IMAP4* 方法), 1150
 deletefilehandler() (*tkinter.Widget.tk* 方法), 1318
 DeleteKey() (在 *winreg* 模块中), 1697
 DeleteKeyEx() (在 *winreg* 模块中), 1697
 deleteln() (*curses.window* 方法), 652
 deleteMe() (*bdb.Breakpoint* 方法), 1487
 DeleteValue() (在 *winreg* 模块中), 1698
 delimiter (*csv.Dialect* 属性), 465
 delitem() (在 *operator* 模块中), 336
 deliver_challenge() (在 *multiprocessing.connection* 模块中), 748
 delocalize() (在 *locale* 模块中), 1259
 demo_app() (在 *wsgiref.simple_server* 模块中), 1098
 denominator (*fractions.Fraction* 属性), 299
 denominator (*numbers.Rational* 属性), 262
 DeprecationWarning, 85
 deque (*collections* 中的类), 200
 Deque (*typing* 中的类), 1360
 dequeue() (*logging.handlers.QueueListener* 方法), 643
 DER_cert_to_PEM_cert() (在 *ssl* 模块中), 815
 derwin() (*curses.window* 方法), 652
 DES
 cipher, 1710
 description (*sqlite3.Cursor* 属性), 420
 description() (*nntplib.NNTP* 方法), 1157
 descriptions() (*nntplib.NNTP* 方法), 1157
 descriptor -- 描述器, 1761
 dest (*optparse.Option* 属性), 1738
 detach() (*io.BufferedIOBase* 方法), 559
 detach() (*io.TextIOBase* 方法), 563
 detach() (*socket.socket* 方法), 800
 detach() (*tkinter.ttk.Treeview* 方法), 1329
 detach() (*weakref.finalize* 方法), 226
 Detach() (*winreg.PyHKEY* 方法), 1704
 --details
 inspect 命令行选项, 1610
 detect_api_mismatch() (在 *test.support* 模块中), 1484
 detect_encoding() (在 *tokenize* 模块中), 1660
 deterministic profiling, 1500
 device_encoding() (在 *os* 模块中), 516
 devnull() (在 *os* 模块中), 552
 DEVNULL() (在 *subprocess* 模块中), 768
 devpoll() (在 *select* 模块中), 842
 DevpollSelector (*selectors* 中的类), 851
 dgettext() (在 *gettext* 模块中), 1248
 dgettext() (在 *locale* 模块中), 1261
 Dialect (*csv* 中的类), 464
 dialect (*csv.csvreader* 属性), 466
 dialect (*csv.csvwriter* 属性), 466
 Dialog (*msilib* 中的类), 1694
 dict (2to3 fixer), 1474
 Dict (*typing* 中的类), 1361
 dict (设置类), 71
 dict() (*multiprocessing.managers.SyncManager* 方法), 741
 dictConfig() (在 *logging.config* 模块中), 621
 dictionary
 type, operations on, 71
 对象, 71
 dictionary -- 字典, 1761
 dictionary view -- 字典视图, 1762
 DictReader (*csv* 中的类), 463

- DictWriter (csv 中的类), 463
- diff_bytes() (在 *difflib* 模块中), 120
- diff_files (filecmp.dircmp 属性), 370
- Differ (difflib 中的类), 117, 124
- difference() (frozenset 方法), 70
- difference_update() (frozenset 方法), 71
- difflib (模块), 117
- digest() (hashlib.hash 方法), 495
- digest() (hashlib.shake 方法), 495
- digest() (hmac.HMAC 方法), 504
- digest_size (hmac.HMAC 属性), 504
- digit() (在 unicodedata 模块中), 131
- digits() (在 string 模块中), 89
- dir() (ftplib.FTP 方法), 1143
- dir() (设置函数), 8
- dircmp (filecmp 中的类), 369
- directory
 - changing, 525
 - creating, 528
 - deleting, 380, 530
 - site-packages, 1610
 - traversal, 539, 540
 - walking, 539, 540
- directory ...
 - compileall 命令行选项, 1665
- Directory (msilib 中的类), 1692
- DirEntry (os 中的类), 532
- DirList (tkinter.tix 中的类), 1337
- dirname() (在 os.path 模块中), 358
- DirSelectBox (tkinter.tix 中的类), 1337
- DirSelectDialog (tkinter.tix 中的类), 1337
- DirTree (tkinter.tix 中的类), 1337
- dis (模块), 1668
- dis() (dis.Bytecode 方法), 1669
- dis() (在 dis 模块中), 1669
- dis() (在 pickletools 模块中), 1682
- disable (pdb command), 1497
- disable() (bdb.Breakpoint 方法), 1488
- disable() (profile.Profile 方法), 1503
- disable() (在 faulthandler 模块中), 1492
- disable() (在 gc 模块中), 1592
- disable() (在 logging 模块中), 618
- disable_interspersed_args() (opt-parse.OptionParser 方法), 1743
- DisableReflectionKey() (在 winreg 模块中), 1701
- disassemble() (在 dis 模块中), 1670
- discard (http.cookiejar.Cookie 属性), 1199
- discard() (frozenset 方法), 71
- discard() (mailbox.Mailbox 方法), 997
- discard() (mailbox.MH 方法), 1003
- discard_buffers() (asynchat.async_chat 方法), 917
- disco() (在 dis 模块中), 1670
- discover() (unittest.TestLoader 方法), 1411
- disk_usage() (在 shutil 模块中), 381
- dispatch_call() (bdb.Bdb 方法), 1489
- dispatch_exception() (bdb.Bdb 方法), 1489
- dispatch_line() (bdb.Bdb 方法), 1489
- dispatch_return() (bdb.Bdb 方法), 1489
- dispatch_table (pickle.Pickler 属性), 391
- dispatcher (asyncore 中的类), 913
- dispatcher_with_send (asyncore 中的类), 915
- display (pdb command), 1498
- display_name (email.headerregistry.Address 属性), 957
- display_name (email.headerregistry.Group 属性), 957
- displayhook() (在 sys 模块中), 1544
- dist() (在 platform 模块中), 669
- distance() (在 turtle 模块中), 1273
- distb() (在 dis 模块中), 1670
- distutils (模块), 1525
- divide() (decimal.Context 方法), 286
- divide_int() (decimal.Context 方法), 286
- DivisionByZero (decimal 中的类), 290
- divmod() (decimal.Context 方法), 286
- divmod() (设置函数), 9
- DllCanUnloadNow() (在 ctypes 模块中), 701
- DllGetClassObject() (在 ctypes 模块中), 701
- dllhandle() (在 sys 模块中), 1544
- dngettext() (在 gettext 模块中), 1248
- do_clear() (bdb.Bdb 方法), 1489
- do_command() (curses.textpad.Textbox 方法), 662
- do_GET() (http.server.SimpleHTTPRequestHandler 方法), 1188
- do_handshake() (ssl.SSLSocket 方法), 823
- do_HEAD() (http.server.SimpleHTTPRequestHandler 方法), 1187
- do_POST() (http.server.CGIHTTPRequestHandler 方法), 1189
- doc (json.JSONDecodeError 属性), 992
- doc_header (cmd.Cmd 属性), 1298
- DocCGIXMLRPCRequestHandler (xmlrpc.server 中的类), 1213
- DocFileSuite() (在 doctest 模块中), 1380
- doCleanups() (unittest.TestCase 方法), 1408
- docmd() (smtplib.SMTP 方法), 1162
- docstring -- 文档字符串, 1762
- docstring (doctest.DocTest 属性), 1383
- DocTest (doctest 中的类), 1383
- doctest (模块), 1368
- DocTestFailure, 1389
- DocTestFinder (doctest 中的类), 1384
- DocTestParser (doctest 中的类), 1385
- DocTestRunner (doctest 中的类), 1385
- DocTestSuite() (在 doctest 模块中), 1381
- doctype() (xml.etree.ElementTree.TreeBuilder 方法), 1045

- doctype() (*xml.etree.ElementTree.XMLParser* 方法), 1045
- documentation
 generation, 1366
 online, 1366
- documentElement (*xml.dom.Document* 属性), 1052
- DocXMLRPCRequestHandler (*xmlrpc.server* 中的类), 1213
- DocXMLRPCServer (*xmlrpc.server* 中的类), 1213
- domain (*email.headerregistry.Address* 属性), 957
- domain (*tracemalloc.DomainFilter* 属性), 1520
- domain (*tracemalloc.Filter* 属性), 1521
- domain_initial_dot (*http.cookiejar.Cookie* 属性), 1200
- domain_return_ok() (*http.cookiejar.CookiePolicy* 方法), 1196
- domain_specified (*http.cookiejar.Cookie* 属性), 1200
- DomainFilter (*tracemalloc* 中的类), 1520
- DomainLiberal (*http.cookiejar.DefaultCookiePolicy* 属性), 1199
- DomainRFC2965Match
 (*http.cookiejar.DefaultCookiePolicy* 属性), 1199
- DomainStrict (*http.cookiejar.DefaultCookiePolicy* 属性), 1199
- DomainStrictNoDots
 (*http.cookiejar.DefaultCookiePolicy* 属性), 1199
- DomainStrictNonDomain
 (*http.cookiejar.DefaultCookiePolicy* 属性), 1199
- DOMEventStream (*xml.dom.pulldom* 中的类), 1063
- DOMException, 1055
- DomstringSizeErr, 1055
- done() (*asyncio.Future* 方法), 872
- done() (*concurrent.futures.Future* 方法), 765
- done() (在 *turtle* 模块中), 1287
- done() (*xdrlib.Unpacker* 方法), 487
- DONT_ACCEPT_BLANKLINE() (在 *doctest* 模块中), 1375
- DONT_ACCEPT_TRUE_FOR_1() (在 *doctest* 模块中), 1375
- dont_write_bytecode() (在 *sys* 模块中), 1545
- doRollover() (*logging.handlers.RotatingFileHandler* 方法), 634
- doRollover() (*logging.handlers.TimedRotatingFileHandler* 方法), 635
- DOT() (在 *token* 模块中), 1657
- dot() (在 *turtle* 模块中), 1270
- DOTALL() (在 *re* 模块中), 104
- doublequote (*csv.Dialect* 属性), 465
- DOUBLESASH() (在 *token* 模块中), 1657
- DOUBLESASHEQUAL() (在 *token* 模块中), 1657
- DOUBLESTAR() (在 *token* 模块中), 1657
- DOUBLESTAREQUAL() (在 *token* 模块中), 1657
- doupdate() (在 *curses* 模块中), 645
- down (*pdb command*), 1496
- down() (在 *turtle* 模块中), 1274
- drain() (*asyncio.StreamWriter* 方法), 891
- drop_whitespace (*textwrap.TextWrapper* 属性), 129
- dropwhile() (在 *itertools* 模块中), 318
- dst() (*datetime.datetime* 方法), 175
- dst() (*datetime.time* 方法), 181
- dst() (*datetime.timezone* 方法), 189
- dst() (*datetime.tzinfo* 方法), 183
- DTDHandler (*xml.sax.handler* 中的类), 1065
- duck-typing -- 鸭子类型, 1762
- DumbWriter (*formatter* 中的类), 1686
- dummy_threading (模块), 785
- dump() (*pickle.Pickler* 方法), 390
- dump() (*tracemalloc.Snapshot* 方法), 1522
- dump() (在 *ast* 模块中), 1654
- dump() (在 *json* 模块中), 988
- dump() (在 *marshal* 模块中), 404
- dump() (在 *pickle* 模块中), 389
- dump() (在 *plistlib* 模块中), 490
- dump() (在 *xml.etree.ElementTree* 模块中), 1039
- dump_stats() (*profile.Profile* 方法), 1503
- dump_stats() (*pstats.Stats* 方法), 1504
- dump_traceback() (在 *faulthandler* 模块中), 1492
- dump_traceback_later() (在 *faulthandler* 模块中), 1492
- dumps() (在 *json* 模块中), 989
- dumps() (在 *marshal* 模块中), 404
- dumps() (在 *pickle* 模块中), 389
- dumps() (在 *plistlib* 模块中), 490
- dumps() (在 *xmlrpc.client* 模块中), 1207
- dup() (*socket.socket* 方法), 800
- dup() (在 *os* 模块中), 516
- dup2() (在 *os* 模块中), 516
- DUP_TOP (*opcode*), 1671
- DUP_TOP_TWO (*opcode*), 1671
- DuplicateOptionError, 484
- DuplicateSectionError, 484
- dwFlags (*subprocess.STARTUPINFO* 属性), 775
- DynamicClassAttribute() (在 *types* 模块中), 233
- ## E
- e <tarfile> [<output_dir>]
 tarfile 命令行选项, 458
- e <zipfile> <output_dir>
 zipfile 命令行选项, 450
- e() (在 *cmath* 模块中), 272
- e() (在 *math* 模块中), 268
- e, --exact
 tokenize 命令行选项, 1661
- E2BIG() (在 *errno* 模块中), 670

- EACCES() (在 *errno* 模块中), 670
 EADDRINUSE() (在 *errno* 模块中), 674
 EADDRNOTAVAIL() (在 *errno* 模块中), 674
 EADV() (在 *errno* 模块中), 672
 EAFNOSUPPORT() (在 *errno* 模块中), 674
 EAFP, 1762
 EAGAIN() (在 *errno* 模块中), 670
 EALREADY() (在 *errno* 模块中), 675
 east_asian_width() (在 *unicodedata* 模块中), 131
 EBADF() (在 *errno* 模块中), 672
 EBADF() (在 *errno* 模块中), 670
 EBADFD() (在 *errno* 模块中), 673
 EBADMSG() (在 *errno* 模块中), 673
 EBADR() (在 *errno* 模块中), 672
 EBADRQC() (在 *errno* 模块中), 672
 EBADSLT() (在 *errno* 模块中), 672
 EBFONT() (在 *errno* 模块中), 672
 EBUSY() (在 *errno* 模块中), 670
 ECHILD() (在 *errno* 模块中), 670
 echo() (在 *curses* 模块中), 646
 echochar() (*curses.window* 方法), 653
 ECHRNA() (在 *errno* 模块中), 671
 ECOMM() (在 *errno* 模块中), 673
 ECONNABORTED() (在 *errno* 模块中), 674
 ECONNREFUSED() (在 *errno* 模块中), 675
 ECONNRESET() (在 *errno* 模块中), 674
 EDEADLK() (在 *errno* 模块中), 671
 EDEADLOCK() (在 *errno* 模块中), 672
 EDESTADDRREQ() (在 *errno* 模块中), 673
 edit() (*curses.textpad.Textbox* 方法), 662
 EDOM() (在 *errno* 模块中), 671
 EDOTDOT() (在 *errno* 模块中), 673
 EDQUOT() (在 *errno* 模块中), 675
 EEXIST() (在 *errno* 模块中), 670
 EFAULT() (在 *errno* 模块中), 670
 EFBIG() (在 *errno* 模块中), 671
 effective() (在 *bdb* 模块中), 1491
 ehlo() (*smtpplib.SMTP* 方法), 1163
 ehlo_or_helo_if_needed() (*smtpplib.SMTP* 方法), 1163
 EHOSTDOWN() (在 *errno* 模块中), 675
 EHOSTUNREACH() (在 *errno* 模块中), 675
 EIDRM() (在 *errno* 模块中), 671
 EILSEQ() (在 *errno* 模块中), 673
 EINPROGRESS() (在 *errno* 模块中), 675
 EINTR() (在 *errno* 模块中), 669
 EINVAL() (在 *errno* 模块中), 670
 EIO() (在 *errno* 模块中), 670
 EISCONN() (在 *errno* 模块中), 674
 EISDIR() (在 *errno* 模块中), 670
 EISNAM() (在 *errno* 模块中), 675
 EL2HLT() (在 *errno* 模块中), 672
 EL2NSYNC() (在 *errno* 模块中), 671
 EL3HLT() (在 *errno* 模块中), 671
 EL3RST() (在 *errno* 模块中), 671
 Element (*xml.etree.ElementTree* 中的类), 1041
 element_create() (*tkinter.ttk.Style* 方法), 1333
 element_names() (*tkinter.ttk.Style* 方法), 1334
 element_options() (*tkinter.ttk.Style* 方法), 1334
 ElementDeclHandler()
 (*xml.parsers.expat.xmlparser* 方法), 1078
 elements() (*collections.Counter* 方法), 198
 ElementTree (*xml.etree.ElementTree* 中的类), 1043
 ELIBACC() (在 *errno* 模块中), 673
 ELIBBAD() (在 *errno* 模块中), 673
 ELIBEXEC() (在 *errno* 模块中), 673
 ELIBMAX() (在 *errno* 模块中), 673
 ELIBSCN() (在 *errno* 模块中), 673
 Ellinghouse, Lance, 1023
 Ellipsis (置变量), 25
 ELLIPSIS() (在 *doctest* 模块中), 1375
 ELLIPSIS() (在 *token* 模块中), 1657
 ELNRNG() (在 *errno* 模块中), 672
 ELOOP() (在 *errno* 模块中), 671
 email (模块), 929
 email.charset (模块), 980
 email.contentmanager (模块), 957
 email.encoders (模块), 982
 email.errors (模块), 951
 email.generator (模块), 942
 email.header (模块), 977
 email.headerregistry (模块), 952
 email.iterators (模块), 985
 EmailMessage (*email.message* 中的类), 931
 email.message (模块), 930
 email.mime (模块), 975
 email.parser (模块), 938
 EmailPolicy (*email.policy* 中的类), 948
 email.policy (模块), 944
 email.utils (模块), 983
 EMFILE() (在 *errno* 模块中), 670
 emit() (*logging.FileHandler* 方法), 632
 emit() (*logging.Handler* 方法), 611
 emit() (*logging.handlers.BufferingHandler* 方法), 640
 emit() (*logging.handlers.DatagramHandler* 方法), 637
 emit() (*logging.handlers.HTTPHandler* 方法), 641
 emit() (*logging.handlers.NTEventLogHandler* 方法), 639
 emit() (*logging.handlers.QueueHandler* 方法), 642
 emit() (*logging.handlers.RotatingFileHandler* 方法), 634
 emit() (*logging.handlers.SMTPHandler* 方法), 640
 emit() (*logging.handlers.SocketHandler* 方法), 636
 emit() (*logging.handlers.SysLogHandler* 方法), 637
 emit() (*logging.handlers.TimedRotatingFileHandler* 方法), 635
 emit() (*logging.handlers.WatchedFileHandler* 方法), 633

- `emit()` (`logging.NullHandler` 方法), 632
- `emit()` (`logging.StreamHandler` 方法), 631
- `EMLINK()` (在 `errno` 模块中), 671
- `Empty`, 784
- `empty` (`inspect.Parameter` 属性), 1601
- `empty` (`inspect.Signature` 属性), 1600
- `empty()` (`asyncio.Queue` 方法), 905
- `empty()` (`multiprocessing.Queue` 方法), 729
- `empty()` (`multiprocessing.SimpleQueue` 方法), 730
- `empty()` (`queue.Queue` 方法), 784
- `empty()` (`sched.scheduler` 方法), 783
- `EMPTY_NAMESPACE()` (在 `xml.dom` 模块中), 1048
- `emptyline()` (`cmd.Cmd` 方法), 1297
- `EMSGSIZE()` (在 `errno` 模块中), 674
- `EMULTIHOP()` (在 `errno` 模块中), 673
- `enable` (`pdb` command), 1497
- `enable()` (`bdb.Breakpoint` 方法), 1487
- `enable()` (`imaplib.IMAP4` 方法), 1150
- `enable()` (`profile.Profile` 方法), 1503
- `enable()` (在 `cgiib` 模块中), 1094
- `enable()` (在 `faulthandler` 模块中), 1492
- `enable()` (在 `gc` 模块中), 1592
- `enable_callback_tracebacks()` (在 `sqlite3` 模块中), 412
- `enable_interspersed_args()` (`opt-parse.OptionParser` 方法), 1743
- `enable_load_extension()` (`sqlite3.Connection` 方法), 415
- `enable_traversal()` (`tkinter.ttk.Notebook` 方法), 1324
- `ENABLE_USER_SITE()` (在 `site` 模块中), 1611
- `EnableReflectionKey()` (在 `winreg` 模块中), 1701
- `ENAMETOOLONG()` (在 `errno` 模块中), 671
- `ENAVAIL()` (在 `errno` 模块中), 675
- `enclose()` (`curses.window` 方法), 653
- `encode`
 - `Codecs`, 145
- `encode` (`codecs.CodecInfo` 属性), 145
- `encode()` (`codecs.Codec` 方法), 149
- `encode()` (`codecs.IncrementalEncoder` 方法), 150
- `encode()` (`email.header.Header` 方法), 979
- `encode()` (`json.JSONEncoder` 方法), 992
- `encode()` (`str` 方法), 41
- `encode()` (在 `base64` 模块中), 1019
- `encode()` (在 `codecs` 模块中), 145
- `encode()` (在 `quopri` 模块中), 1022
- `encode()` (在 `uu` 模块中), 1023
- `encode()` (`xmlrpc.client.Binary` 方法), 1205
- `encode()` (`xmlrpc.client.DateTime` 方法), 1204
- `encode_7or8bit()` (在 `email.encoders` 模块中), 982
- `encode_base64()` (在 `email.encoders` 模块中), 982
- `encode_noop()` (在 `email.encoders` 模块中), 982
- `encode_quopri()` (在 `email.encoders` 模块中), 982
- `encode_rfc2231()` (在 `email.utils` 模块中), 985
- `encodebytes()` (在 `base64` 模块中), 1019
- `EncodedFile()` (在 `codecs` 模块中), 147
- `encodePriority()` (`logging.handlers.SysLogHandler` 方法), 638
- `encodestring()` (在 `base64` 模块中), 1019
- `encodestring()` (在 `quopri` 模块中), 1022
- `encoding`
 - `base64`, 1017
 - `quoted-printable`, 1022
- `encoding` (`curses.window` 属性), 653
- `encoding` (`io.TextIOBase` 属性), 563
- `encoding` (`UnicodeError` 属性), 84
- `ENCODING()` (在 `tarfile` 模块中), 452
- `ENCODING()` (在 `tokenize` 模块中), 1659
- `encodings_map` (`mimetypes.MimeTypes` 属性), 1016
- `encodings_map()` (在 `mimetypes` 模块中), 1015
- `encodings.idna` (模块), 160
- `encodings.mbcx` (模块), 161
- `encodings.utf_8_sig` (模块), 161
- `end` (`UnicodeError` 属性), 84
- `end()` (`re.match` 方法), 111
- `end()` (`xml.etree.ElementTree.TreeBuilder` 方法), 1045
- `end_fill()` (在 `turtle` 模块中), 1277
- `END_FINALLY` (`opcode`), 1675
- `end_headers()` (`http.server.BaseHTTPRequestHandler` 方法), 1186
- `end_paragraph()` (`formatter.formatter` 方法), 1684
- `end_poly()` (在 `turtle` 模块中), 1282
- `EndCdataSectionHandler()`
 - (`xml.parsers.expat.xmlparser` 方法), 1079
- `EndDoctypeDeclHandler()`
 - (`xml.parsers.expat.xmlparser` 方法), 1078
- `endDocument()` (`xml.sax.handler.ContentHandler` 方法), 1067
- `endElement()` (`xml.sax.handler.ContentHandler` 方法), 1068
- `EndElementHandler()` (`xml.parsers.expat.xmlparser` 方法), 1079
- `endElementNS()` (`xml.sax.handler.ContentHandler` 方法), 1068
- `endheaders()` (`http.client.HTTPConnection` 方法), 1137
- `ENDMARKER()` (在 `token` 模块中), 1657
- `EndNamespaceDeclHandler()`
 - (`xml.parsers.expat.xmlparser` 方法), 1079
- `endpos` (`re.match` 属性), 111
- `endPrefixMapping()`
 - (`xml.sax.handler.ContentHandler` 方法), 1068
- `endswith()` (`bytearray` 方法), 52
- `endswith()` (`bytes` 方法), 52
- `endswith()` (`str` 方法), 41
- `endwin()` (在 `curses` 模块中), 646
- `ENETDOWN()` (在 `errno` 模块中), 674
- `ENETRESET()` (在 `errno` 模块中), 674

- ENETUNREACH() (在 *errno* 模块中), 674
- ENFILE() (在 *errno* 模块中), 670
- ENOANO() (在 *errno* 模块中), 672
- ENOBUFFS() (在 *errno* 模块中), 674
- ENOCSS() (在 *errno* 模块中), 672
- ENODATA() (在 *errno* 模块中), 672
- ENODEV() (在 *errno* 模块中), 670
- ENOENT() (在 *errno* 模块中), 669
- ENOEXEC() (在 *errno* 模块中), 670
- ENOLCK() (在 *errno* 模块中), 671
- ENOLINK() (在 *errno* 模块中), 672
- ENOMEM() (在 *errno* 模块中), 670
- ENOMSG() (在 *errno* 模块中), 671
- ENONET() (在 *errno* 模块中), 672
- ENOPKG() (在 *errno* 模块中), 672
- ENOPROTOOPT() (在 *errno* 模块中), 674
- ENOSPC() (在 *errno* 模块中), 671
- ENOSR() (在 *errno* 模块中), 672
- ENOSTR() (在 *errno* 模块中), 672
- ENOSYS() (在 *errno* 模块中), 671
- ENOTBLK() (在 *errno* 模块中), 670
- ENOTCONN() (在 *errno* 模块中), 674
- ENOTDIR() (在 *errno* 模块中), 670
- ENOTEMPTY() (在 *errno* 模块中), 671
- ENOTNAM() (在 *errno* 模块中), 675
- ENOTSOCK() (在 *errno* 模块中), 673
- ENOTTY() (在 *errno* 模块中), 670
- ENOTUNIQ() (在 *errno* 模块中), 673
- enqueue() (*logging.handlers.QueueHandler* 方法), 642
- enqueue_sentinel() (*logging.handlers.QueueListener* 方法), 643
- ensure_directories() (*venv.EnvBuilder* 方法), 1531
- ensure_future() (在 *asyncio* 模块中), 876
- ensurepip(模块), 1526
- enter() (*sched.scheduler* 方法), 782
- enter_context() (*contextlib.ExitStack* 方法), 1572
- enterabs() (*sched.scheduler* 方法), 782
- entities (*xml.dom.DocumentType* 属性), 1052
- EntityDeclHandler() (*xml.parsers.expat.xmlparser* 方法), 1079
- entitydefs() (在 *html.entities* 模块中), 1030
- EntityResolver (*xml.sax.handler* 中的类), 1065
- Enum (*enum* 中的类), 242
- enum(模块), 242
- enum_certificates() (在 *ssl* 模块中), 816
- enum_crls() (在 *ssl* 模块中), 816
- enumerate() (内置函数), 9
- enumerate() (在 *threading* 模块中), 709
- EnumKey() (在 *winreg* 模块中), 1698
- EnumValue() (在 *winreg* 模块中), 1698
- EnvBuilder (*venv* 中的类), 1530
- environ() (在 *os* 模块中), 510
- environ() (在 *posix* 模块中), 1708
- environb() (在 *os* 模块中), 511
- environment variables
 - deleting, 515
 - setting, 513
- EnvironmentError, 84
- Environments
 - virtual, 1527
- EnvironmentVarGuard (*test.support* 中的类), 1485
- ENXIO() (在 *errno* 模块中), 670
- eof (*bz2.BZ2Decompressor* 属性), 437
- eof (*lzma.LZMADecompressor* 属性), 441
- eof (*shlex.shlex* 属性), 1304
- eof (*ssl.MemoryBIO* 属性), 839
- eof (*zlib.Decompress* 属性), 431
- eof_received() (*asyncio.Protocol* 方法), 883
- EOFError, 80
- EOPNOTSUPP() (在 *errno* 模块中), 674
- EOVERFLOW() (在 *errno* 模块中), 673
- EPERM() (在 *errno* 模块中), 669
- EPFNOSUPPORT() (在 *errno* 模块中), 674
- epilogue (*email.message.EmailMessage* 属性), 938
- epilogue (*email.message.Message* 属性), 974
- EPIPE() (在 *errno* 模块中), 671
- epoch, 566
- epoll() (在 *select* 模块中), 842
- EpollSelector (*selectors* 中的类), 851
- EPROTO() (在 *errno* 模块中), 673
- EPROTONOSUPPORT() (在 *errno* 模块中), 674
- EPROTOTYPE() (在 *errno* 模块中), 674
- eq() (在 *operator* 模块中), 334
- EQUAL() (在 *token* 模块中), 1657
- EQUAL() (在 *token* 模块中), 1657
- ERA() (在 *locale* 模块中), 1258
- ERA_D_FMT() (在 *locale* 模块中), 1258
- ERA_D_T_FMT() (在 *locale* 模块中), 1258
- ERA_T_FMT() (在 *locale* 模块中), 1258
- ERANGE() (在 *errno* 模块中), 671
- erase() (*curses.window* 方法), 653
- erasechar() (在 *curses* 模块中), 646
- EREMCHG() (在 *errno* 模块中), 673
- EREMOTE() (在 *errno* 模块中), 672
- EREMOTEIO() (在 *errno* 模块中), 675
- ERESTART() (在 *errno* 模块中), 673
- erf() (在 *math* 模块中), 268
- erfc() (在 *math* 模块中), 268
- EROFS() (在 *errno* 模块中), 671
- ERR() (在 *curses* 模块中), 657
- errcheck (*ctypes.FuncPtr* 属性), 697
- errcode (*xmlrpc.client.ProtocolError* 属性), 1206
- errmsg (*xmlrpc.client.ProtocolError* 属性), 1206
- errno
 - 模块, 81
- errno (*OSError* 属性), 82
- errno(模块), 669

- Error, 381, 421, 465, 484, 488, 1012, 1020, 1022, 1023, 1085, 1233, 1236, 1255
- error, 107, 140, 234, 404, 406408, 429, 509, 604, 645, 786, 792, 842, 1075, 1227, 1718, 1722
- error() (*argparse.ArgumentParser* 方法), 602
- error() (*logging.Logger* 方法), 608
- error() (*urllib.request.OpenerDirector* 方法), 1111
- error() (在 *logging* 模块中), 618
- error() (*xml.sax.handler.ErrorHandler* 方法), 1070
- error_body (*wsgiref.handlers.BaseHandler* 属性), 1102
- error_content_type
(*http.server.BaseHTTPRequestHandler* 属性), 1185
- error_headers (*wsgiref.handlers.BaseHandler* 属性), 1102
- error_leader() (*shlex.shlex* 方法), 1302
- error_message_format
(*http.server.BaseHTTPRequestHandler* 属性), 1185
- error_output() (*wsgiref.handlers.BaseHandler* 方法), 1102
- error_perm, 1141
- error_proto, 1141, 1145
- error_received() (*asyncio.DatagramProtocol* 方法), 884
- error_reply, 1141
- error_status (*wsgiref.handlers.BaseHandler* 属性), 1102
- error_temp, 1141
- ErrorByteIndex (*xml.parsers.expat.xmlparser* 属性), 1077
- ErrorCode (*xml.parsers.expat.xmlparser* 属性), 1078
- errorcode() (在 *errno* 模块中), 669
- ErrorColumnNumber (*xml.parsers.expat.xmlparser* 属性), 1078
- ErrorHandler (*xml.sax.handler* 中的类), 1065
- ErrorLineNumber (*xml.parsers.expat.xmlparser* 属性), 1078
- Errors
logging, 605
- errors (*io.TextIOBase* 属性), 563
- errors (*unittest.TestLoader* 属性), 1410
- errors (*unittest.TestResult* 属性), 1412
- ErrorString() (在 *xml.parsers.expat* 模块中), 1075
- ERRORTOKEN() (在 *token* 模块中), 1657
- escape (*shlex.shlex* 属性), 1303
- escape() (在 *cgi* 模块中), 1091
- escape() (在 *glob* 模块中), 375
- escape() (在 *html* 模块中), 1025
- escape() (在 *re* 模块中), 107
- escape() (在 *xml.sax.saxutils* 模块中), 1070
- escapechar (*csv.Dialect* 属性), 465
- escapedquotes (*shlex.shlex* 属性), 1303
- ESHUTDOWN() (在 *errno* 模块中), 674
- ESOCKTNOSUPPORT() (在 *errno* 模块中), 674
- ESPIPE() (在 *errno* 模块中), 671
- ESRCH() (在 *errno* 模块中), 669
- ESRMNT() (在 *errno* 模块中), 673
- ESTALE() (在 *errno* 模块中), 675
- ESTRPIPE() (在 *errno* 模块中), 673
- ETIME() (在 *errno* 模块中), 672
- ETIMEDOUT() (在 *errno* 模块中), 674
- Etiny() (*decimal.Context* 方法), 285
- ETOOMANYREFS() (在 *errno* 模块中), 674
- Etop() (*decimal.Context* 方法), 285
- ETXTBSY() (在 *errno* 模块中), 670
- EUCLEAN() (在 *errno* 模块中), 675
- EUNATCH() (在 *errno* 模块中), 672
- EUSERS() (在 *errno* 模块中), 673
- eval
F置函数, 76, 236, 237, 1647
- eval() (F置函数), 9
- Event (*asyncio* 中的类), 902
- Event (*multiprocessing* 中的类), 734
- Event (*threading* 中的类), 717
- event scheduling, 782
- event() (*msilib.Control* 方法), 1693
- Event() (*multiprocessing.managers.SyncManager* 方法), 740
- events (*selectors.SelectorKey* 属性), 849
- events (widgets), 1316
- EWouldBlock() (在 *errno* 模块中), 671
- EX_CANTCREAT() (在 *os* 模块中), 544
- EX_CONFIG() (在 *os* 模块中), 544
- EX_DATAERR() (在 *os* 模块中), 543
- EX_IOERR() (在 *os* 模块中), 544
- EX_NOHOST() (在 *os* 模块中), 543
- EX_NOINPUT() (在 *os* 模块中), 543
- EX_NOPERM() (在 *os* 模块中), 544
- EX_NOTFOUND() (在 *os* 模块中), 544
- EX_NOUSER() (在 *os* 模块中), 543
- EX_OK() (在 *os* 模块中), 543
- EX_OSERR() (在 *os* 模块中), 544
- EX_OSFILE() (在 *os* 模块中), 544
- EX_PROTOCOL() (在 *os* 模块中), 544
- EX_SOFTWARE() (在 *os* 模块中), 543
- EX_TEMPFAIL() (在 *os* 模块中), 544
- EX_UNAVAILABLE() (在 *os* 模块中), 543
- EX_USAGE() (在 *os* 模块中), 543
- Example (*doctest* 中的类), 1384
- example (*doctest.DocTestFailure* 属性), 1390
- example (*doctest.UnexpectedException* 属性), 1390
- examples (*doctest.DocTest* 属性), 1383
- exc_info (*doctest.UnexpectedException* 属性), 1390
- exc_info() (在 *sys* 模块中), 1545
- exc_msg (*doctest.Example* 属性), 1384
- exc_type (*traceback.TracebackException* 属性), 1587

- excel (csv 中的类), 464
- excel_tab (csv 中的类), 464
- except
 - 语句, 79
- except (2to3 fixer), 1474
- excepthook () (in module sys), 1094
- excepthook () (在 sys 模块中), 1545
- Exception, 80
- exception () (asyncio.Future 方法), 872
- exception () (asyncio.StreamReader 方法), 889
- exception () (concurrent.futures.Future 方法), 765
- exception () (logging.Logger 方法), 608
- exception () (在 logging 模块中), 618
- EXCEPTION () (在 tkinter 模块中), 1318
- exceptions
 - in CGI scripts, 1094
- EXDEV () (在 errno 模块中), 670
- exec
 - ⓕ置函数, 10, 76, 1647
- exec (2to3 fixer), 1474
- exec () (ⓕ置函数), 10
- exec_module () (importlib.abc.InspectLoader 方法), 1633
- exec_module () (importlib.abc.Loader 方法), 1631
- exec_module () (importlib.abc.SourceLoader 方法), 1634
- exec_module () (importlib.machinery.ExtensionFileLoader 方法), 1638
- exec_prefix () (在 sys 模块中), 1545
- execfile (2to3 fixer), 1474
- execl () (在 os 模块中), 542
- execle () (在 os 模块中), 542
- execlp () (在 os 模块中), 542
- execlpe () (在 os 模块中), 542
- Executable Zip Files, 1535
- executable () (在 sys 模块中), 1546
- Execute () (msilib.View 方法), 1691
- execute () (sqlite3.Connection 方法), 412
- execute () (sqlite3.Cursor 方法), 417
- executemany () (sqlite3.Connection 方法), 412
- executemany () (sqlite3.Cursor 方法), 417
- executescript () (sqlite3.Connection 方法), 412
- executescript () (sqlite3.Cursor 方法), 418
- ExecutionLoader (importlib.abc 中的类), 1633
- Executor (concurrent.futures 中的类), 761
- execv () (在 os 模块中), 542
- execve () (在 os 模块中), 542
- execvp () (在 os 模块中), 542
- execvpe () (在 os 模块中), 542
- ExFileSelectBox (tkinter.tix 中的类), 1337
- EXFULL () (在 errno 模块中), 672
- exists () (pathlib.Path 方法), 352
- exists () (tkinter.ttk.Treeview 方法), 1329
- exists () (在 os.path 模块中), 358
- exit (ⓕ置变量), 26
- exit () (argparse.ArgumentParser 方法), 602
- exit () (在 _thread 模块中), 786
- exit () (在 sys 模块中), 1546
- exitcode (multiprocessing.Process 属性), 727
- exitfunc (2to3 fixer), 1475
- exitonclick () (在 turtle 模块中), 1289
- ExitStack (contextlib 中的类), 1572
- exp () (decimal.Context 方法), 286
- exp () (decimal.Decimal 方法), 279
- exp () (在 cmath 模块中), 270
- exp () (在 math 模块中), 266
- expand () (re.match 方法), 109
- expand_tabs (textwrap.TextWrapper 属性), 129
- ExpandEnvironmentStrings () (在 winreg 模块中), 1698
- expandNode () (xml.dom.pulldom.DOMEventStream 方法), 1063
- expandtabs () (bytearray 方法), 56
- expandtabs () (bytes 方法), 56
- expandtabs () (str 方法), 42
- expanduser () (pathlib.Path 方法), 352
- expanduser () (在 os.path 模块中), 358
- expandvars () (在 os.path 模块中), 358
- Expat, 1075
- ExpatError, 1075
- expect () (telnetlib.Telnet 方法), 1172
- expected (asyncio.IncompleteReadError 属性), 892
- expectedFailure () (在 unittest 模块中), 1398
- expectedFailures (unittest.TestResult 属性), 1412
- expires (http.cookiejar.Cookie 属性), 1199
- exploded (ipaddress.IPv4Address 属性), 1215
- exploded (ipaddress.IPv4Network 属性), 1220
- exploded (ipaddress.IPv6Address 属性), 1217
- exploded (ipaddress.IPv6Network 属性), 1222
- expm1 () (在 math 模块中), 266
- expovariate () (在 random 模块中), 303
- expr () (在 parser 模块中), 1646
- expression -- 表达式, 1762
- expunge () (imaplib.IMAP4 方法), 1150
- extend () (array.array 方法), 221
- extend () (collections.deque 方法), 200
- extend () (sequence method), 37
- extend () (xml.etree.ElementTree.Element 方法), 1042
- extend_path () (在 pkgutil 模块中), 1621
- EXTENDED_ARG (opcode), 1679
- ExtendedContext (decimal 中的类), 284
- ExtendedInterpolation (configparser 中的类), 472
- extendleft () (collections.deque 方法), 200
- extension module -- 扩展模块, 1762
- EXTENSION_SUFFIXES () (在 importlib.machinery 模块中), 1635

- ExtensionFileLoader (*importlib.machinery* 中的类), 1638
- extensions_map (*http.server.SimpleHTTPRequestHandler* 属性), 1187
- External Data Representation, 388, 486
- external_attr (*zipfile.ZipInfo* 属性), 450
- ExternalClashError, 1012
- ExternalEntityParserCreate() (*xml.parsers.expat.xmlparser* 方法), 1077
- ExternalEntityRefHandler() (*xml.parsers.expat.xmlparser* 方法), 1080
- extra (*zipfile.ZipInfo* 属性), 449
- extract <tarfile> [<output_dir>] tarfile 命令行选项, 458
- extract() (*tarfile.TarFile* 方法), 455
- extract() (*traceback.StackSummary* 类方法), 1588
- extract() (*zipfile.ZipFile* 方法), 446
- extract_cookies() (*http.cookiejar.CookieJar* 方法), 1194
- extract_stack() (在 *traceback* 模块中), 1586
- extract_tb() (在 *traceback* 模块中), 1586
- extract_version (*zipfile.ZipInfo* 属性), 449
- extractall() (*tarfile.TarFile* 方法), 454
- extractall() (*zipfile.ZipFile* 方法), 446
- ExtractError, 452
- extractfile() (*tarfile.TarFile* 方法), 455
- extsep() (在 *os* 模块中), 552
- ## F
- f compileall 命令行选项, 1665
- f, --failfast unittest 命令行选项, 1394
- f, --file=<file> trace 命令行选项, 1513
- f_contiguous (*memoryview* 属性), 69
- F_LOCK() (在 *os* 模块中), 518
- F_OK() (在 *os* 模块中), 524
- F_TEST() (在 *os* 模块中), 518
- F_TLOCK() (在 *os* 模块中), 518
- F_ULOCK() (在 *os* 模块中), 518
- fabs() (在 *math* 模块中), 264
- factorial() (在 *math* 模块中), 264
- factory() (*importlib.util.LazyLoader* 类方法), 1642
- fail() (*unittest.TestCase* 方法), 1407
- FAIL_FAST() (在 *doctest* 模块中), 1376
- failfast (*unittest.TestResult* 属性), 1412
- failureException (*unittest.TestCase* 属性), 1407
- failures (*unittest.TestResult* 属性), 1412
- FakePath (*test.support* 中的类), 1485
- False, 27, 77
- false, 27
- False (*Built-in object*), 27
- False (F 置变量), 25
- family (*socket.socket* 属性), 806
- FancyURLopener (*urllib.request* 中的类), 1120
- fast (*pickle.Pickler* 属性), 391
- fatalError() (*xml.sax.handler.ErrorHandler* 方法), 1070
- Fault (*xmlrpc.client* 中的类), 1205
- faultCode (*xmlrpc.client.Fault* 属性), 1205
- faulthandler (模块), 1491
- faultString (*xmlrpc.client.Fault* 属性), 1205
- fchdir() (在 *os* 模块中), 526
- fchmod() (在 *os* 模块中), 516
- fchown() (在 *os* 模块中), 517
- FCICreate() (在 *msilib* 模块中), 1689
- fcntl (模块), 1715
- fcntl() (在 *fcntl* 模块中), 1715
- fd (*selectors.SelectorKey* 属性), 849
- fd() (在 *turtle* 模块中), 1268
- fdatasync() (在 *os* 模块中), 517
- fdopen() (在 *os* 模块中), 516
- Feature (*msilib* 中的类), 1693
- feature_external_ges() (在 *xml.sax.handler* 模块中), 1066
- feature_external_pes() (在 *xml.sax.handler* 模块中), 1066
- feature_namespace_prefixes() (在 *xml.sax.handler* 模块中), 1066
- feature_namespaces() (在 *xml.sax.handler* 模块中), 1065
- feature_string_interning() (在 *xml.sax.handler* 模块中), 1066
- feature_validation() (在 *xml.sax.handler* 模块中), 1066
- feed() (*email.parser.BytesFeedParser* 方法), 939
- feed() (*html.parser.HTMLParser* 方法), 1027
- feed() (*xml.etree.ElementTree.XMLParser* 方法), 1045
- feed() (*xml.etree.ElementTree.XMLPullParser* 方法), 1046
- feed() (*xml.sax.xmlreader.IncrementalParser* 方法), 1073
- feed_data() (*asyncio.StreamReader* 方法), 890
- feed_eof() (*asyncio.StreamReader* 方法), 890
- FeedParser (*email.parser* 中的类), 939
- fetch() (*imaplib.IMAP4* 方法), 1150
- Fetch() (*msilib.View* 方法), 1691
- fetchall() (*sqlite3.Cursor* 方法), 419
- fetchmany() (*sqlite3.Cursor* 方法), 419
- fetchone() (*sqlite3.Cursor* 方法), 419
- fflags (*select.kevent* 属性), 847
- field_size_limit() (在 *csv* 模块中), 463
- fieldnames (*csv.csvreader* 属性), 466
- fields (*uuid.UUID* 属性), 1173
- file
- byte-code, 1664, 1752
 - configuration, 468

- copying, 378
- debugger configuration, 1496
- .ini, 468
- large files, 1707
- mime.types, 1015
- modes, 15
- path configuration, 1610
- .pdbrc, 1496
- plist, 489
- temporary, 371
- file ...
 - compileall 命令行选项, 1665
- file (*pyclbr.Class* 属性), 1664
- file (*pyclbr.Function* 属性), 1664
- file control
 - UNIX, 1715
- file name
 - temporary, 371
- file object
 - io module, 554
 - open() built-in function, 15
- file object -- 文件对象, 1762
- FILE_ATTRIBUTE_ARCHIVE() (在 *stat* 模块中), 368
- FILE_ATTRIBUTE_COMPRESSED() (在 *stat* 模块中), 368
- FILE_ATTRIBUTE_DEVICE() (在 *stat* 模块中), 368
- FILE_ATTRIBUTE_DIRECTORY() (在 *stat* 模块中), 368
- FILE_ATTRIBUTE_ENCRYPTED() (在 *stat* 模块中), 368
- FILE_ATTRIBUTE_HIDDEN() (在 *stat* 模块中), 368
- FILE_ATTRIBUTE_INTEGRITY_STREAM() (在 *stat* 模块中), 368
- FILE_ATTRIBUTE_NO_SCRUB_DATA() (在 *stat* 模块中), 368
- FILE_ATTRIBUTE_NORMAL() (在 *stat* 模块中), 368
- FILE_ATTRIBUTE_NOT_CONTENT_INDEXED() (在 *stat* 模块中), 368
- FILE_ATTRIBUTE_OFFLINE() (在 *stat* 模块中), 368
- FILE_ATTRIBUTE_READONLY() (在 *stat* 模块中), 368
- FILE_ATTRIBUTE_REPARSE_POINT() (在 *stat* 模块中), 368
- FILE_ATTRIBUTE_SPARSE_FILE() (在 *stat* 模块中), 368
- FILE_ATTRIBUTE_SYSTEM() (在 *stat* 模块中), 368
- FILE_ATTRIBUTE_TEMPORARY() (在 *stat* 模块中), 368
- FILE_ATTRIBUTE_VIRTUAL() (在 *stat* 模块中), 368
- file_dispatcher (*asyncore* 中的类), 915
- file_open() (*urllib.request.FileHandler* 方法), 1115
- file_size (*zipfile.ZipInfo* 属性), 450
- file_wrapper (*asyncore* 中的类), 915
- filecmp (模块), 369
- fileConfig() (在 *logging.config* 模块中), 622
- FileCookieJar (*http.cookiejar* 中的类), 1193
- FileEntry (*tkinter.tix* 中的类), 1337
- FileExistsError, 85
- FileFinder (*importlib.machinery* 中的类), 1636
- FileHandler (*logging* 中的类), 632
- FileHandler (*urllib.request* 中的类), 1109
- FileInput (*fileinput* 中的类), 363
- fileinput (模块), 361
- FileIO (*io* 中的类), 560
- file-like object -- 文件类对象, 1762
- filelineno() (在 *fileinput* 模块中), 362
- FileLoader (*importlib.abc* 中的类), 1633
- filemode() (在 *stat* 模块中), 365
- filename (*doctest.DocTest* 属性), 1383
- filename (*http.cookiejar.FileCookieJar* 属性), 1196
- filename (*OSError* 属性), 82
- filename (*traceback.TracebackException* 属性), 1587
- filename (*tracemalloc.Frame* 属性), 1521
- filename (*zipfile.ZipFile* 属性), 447
- filename (*zipfile.ZipInfo* 属性), 449
- filename() (在 *fileinput* 模块中), 362
- filename2 (*OSError* 属性), 82
- filename_only() (在 *tabnanny* 模块中), 1663
- filename_pattern (*tracemalloc.Filter* 属性), 1521
- filenames
 - pathname expansion, 375
 - wildcard expansion, 376
- fileno() (*http.client.HTTPResponse* 方法), 1137
- fileno() (*io.IOBase* 方法), 557
- fileno() (*multiprocessing.connection.Connection* 方法), 732
- fileno() (*ossaudiodev.oss_audio_device* 方法), 1242
- fileno() (*ossaudiodev.oss_mixer_device* 方法), 1244
- fileno() (*select.devpoll* 方法), 844
- fileno() (*select.epoll* 方法), 845
- fileno() (*select.kqueue* 方法), 846
- fileno() (*selectors.DevpollSelector* 方法), 851
- fileno() (*selectors.EpollSelector* 方法), 851
- fileno() (*selectors.KqueueSelector* 方法), 851
- fileno() (*socketserver.BaseServer* 方法), 1178
- fileno() (*socket.socket* 方法), 800
- fileno() (*telnetlib.Telnet* 方法), 1172
- fileno() (在 *fileinput* 模块中), 362
- FileNotFoundError, 85
- fileobj (*selectors.SelectorKey* 属性), 849
- FileSelectBox (*tkinter.tix* 中的类), 1337
- FileType (*argparse* 中的类), 599
- FileWrapper (*wsgiref.util* 中的类), 1096
- fill() (*textwrap.TextWrapper* 方法), 130
- fill() (在 *textwrap* 模块中), 127
- fillcolor() (在 *turtle* 模块中), 1276
- filling() (在 *turtle* 模块中), 1277
- filter (*2to3 fixer*), 1475

- Filter (*logging* 中的类), 612
- filter (*select.kevent* 属性), 847
- Filter (*tracemalloc* 中的类), 1521
- filter() (*logging.Filter* 方法), 612
- filter() (*logging.Handler* 方法), 610
- filter() (*logging.Logger* 方法), 609
- filter() (Ⓕ置函数), 10
- filter() (在 *curses* 模块中), 646
- filter() (在 *fnmatch* 模块中), 377
- FILTER_DIR() (在 *unittest.mock* 模块中), 1448
- filter_traces() (*tracemalloc.Snapshot* 方法), 1522
- filterfalse() (在 *itertools* 模块中), 318
- filterwarnings() (在 *warnings* 模块中), 1567
- finalize (*weakref* 中的类), 225
- find() (*bytearray* 方法), 52
- find() (*bytes* 方法), 52
- find() (*doctest.DocTestFinder* 方法), 1384
- find() (*mmap.mmap* 方法), 926
- find() (*str* 方法), 42
- find() (在 *gettext* 模块中), 1249
- find() (*xml.etree.ElementTree.Element* 方法), 1042
- find() (*xml.etree.ElementTree.ElementTree* 方法), 1043
- find_class() (*pickle.protocol*), 398
- find_class() (*pickle.Unpickler* 方法), 391
- find_library() (在 *ctypes.util* 模块中), 701
- find_loader() (*importlib.abc.PathEntryFinder* 方法), 1631
- find_loader() (*importlib.machinery.FileFinder* 方法), 1637
- find_loader() (在 *importlib* 模块中), 1628
- find_loader() (在 *pkgutil* 模块中), 1622
- find_longest_match() (*difflib.SequenceMatcher* 方法), 122
- find_module() (*imp.NullImporter* 方法), 1755
- find_module() (*importlib.abc.Finder* 方法), 1630
- find_module() (*importlib.abc.MetaPathFinder* 方法), 1630
- find_module() (*importlib.abc.PathEntryFinder* 方法), 1631
- find_module() (*importlib.machinery.PathFinder* 类方法), 1636
- find_module() (在 *imp* 模块中), 1752
- find_module() (*zipimport.zipimporter* 方法), 1620
- find_msvcr() (在 *ctypes.util* 模块中), 701
- find_spec() (*importlib.abc.MetaPathFinder* 方法), 1630
- find_spec() (*importlib.abc.PathEntryFinder* 方法), 1630
- find_spec() (*importlib.machinery.FileFinder* 方法), 1637
- find_spec() (*importlib.machinery.PathFinder* 类方法), 1636
- find_spec() (在 *importlib.util* 模块中), 1640
- find_unused_port() (在 *test.support* 模块中), 1484
- find_user_password() (*url-lib.request.HTTPPasswordMgr* 方法), 1114
- findall() (*re.regex* 方法), 109
- findall() (在 *re* 模块中), 106
- findall() (*xml.etree.ElementTree.Element* 方法), 1042
- findall() (*xml.etree.ElementTree.ElementTree* 方法), 1043
- findCaller() (*logging.Logger* 方法), 609
- finder -- 查找器, 1762
- Finder (*importlib.abc* 中的类), 1630
- findfactor() (在 *audioop* 模块中), 1228
- findfile() (在 *test.support* 模块中), 1480
- findfit() (在 *audioop* 模块中), 1228
- finditer() (*re.regex* 方法), 109
- finditer() (在 *re* 模块中), 106
- findlabels() (在 *dis* 模块中), 1670
- findlinestarts() (在 *dis* 模块中), 1670
- findmatch() (在 *mailcap* 模块中), 995
- findmax() (在 *audioop* 模块中), 1228
- findtext() (*xml.etree.ElementTree.Element* 方法), 1042
- findtext() (*xml.etree.ElementTree.ElementTree* 方法), 1043
- finish() (*socketserver.BaseRequestHandler* 方法), 1180
- finish_request() (*socketserver.BaseServer* 方法), 1179
- firstChild (*xml.dom.Node* 属性), 1050
- firstkey() (*dbm.gnu.gdbm* 方法), 406
- firstweekday() (在 *calendar* 模块中), 194
- fix_missing_locations() (在 *ast* 模块中), 1653
- fix_sentence_endings (*textwrap.TextWrapper* 属性), 130
- Flag (*enum* 中的类), 242
- flag_bits (*zipfile.ZipInfo* 属性), 449
- flags (*re.regex* 属性), 109
- flags (*select.kevent* 属性), 847
- flags() (在 *sys* 模块中), 1546
- flash() (在 *curses* 模块中), 646
- flatten() (*email.generator.BytesGenerator* 方法), 942
- flatten() (*email.generator.Generator* 方法), 943
- flattening
 - objects, 387
- float
 - Ⓕ置函数, 29
- float (Ⓕ置类), 10
- float_info() (在 *sys* 模块中), 1546
- float_repr_style() (在 *sys* 模块中), 1547
- floating point
 - literals, 29
 - 对象, 29
- FloatingPointError, 80, 1613
- FloatOperation (*decimal* 中的类), 291
- flock() (在 *fcntl* 模块中), 1716

- floor division -- 向下取整除法, 1762
- floor() (in module *math*), 29
- floor() (在 *math* 模块中), 264
- floordiv() (在 *operator* 模块中), 334
- flush() (*bz2.BZ2Compressor* 方法), 436
- flush() (*formatter.writer* 方法), 1685
- flush() (*io.BufferedWriter* 方法), 562
- flush() (*io.IOBase* 方法), 557
- flush() (*logging.Handler* 方法), 610
- flush() (*logging.handlers.BufferingHandler* 方法), 640
- flush() (*logging.handlers.MemoryHandler* 方法), 641
- flush() (*logging.StreamHandler* 方法), 631
- flush() (*lzma.LZMACompressor* 方法), 440
- flush() (*mailbox.Mailbox* 方法), 999
- flush() (*mailbox.Maildir* 方法), 1001
- flush() (*mailbox.MH* 方法), 1003
- flush() (*mmap.mmap* 方法), 926
- flush() (*zlib.Compress* 方法), 431
- flush() (*zlib.Decompress* 方法), 432
- flush_headers() (*http.server.BaseHTTPRequestHandler* 方法), 1187
- flush_softspace() (*formatter.formatter* 方法), 1684
- flushinp() (在 *curses* 模块中), 646
- FlushKey() (在 *winreg* 模块中), 1698
- fma() (*decimal.Context* 方法), 286
- fma() (*decimal.Decimal* 方法), 279
- fmod() (在 *math* 模块中), 265
- FMT_BINARY() (在 *plistlib* 模块中), 491
- FMT_XML() (在 *plistlib* 模块中), 491
- fnmatch (模块), 376
- fnmatch() (在 *fnmatch* 模块中), 376
- fnmatchcase() (在 *fnmatch* 模块中), 376
- focus() (*tkinter.ttk.Treeview* 方法), 1329
- fold(*datetime.datetime* 属性), 173
- fold(*datetime.time* 属性), 180
- fold() (*email.headerregistry.BaseHeader* 方法), 953
- fold() (*email.policy.Compat32* 方法), 951
- fold() (*email.policy.EmailPolicy* 方法), 949
- fold() (*email.policy.Policy* 方法), 948
- fold_binary() (*email.policy.Compat32* 方法), 951
- fold_binary() (*email.policy.EmailPolicy* 方法), 950
- fold_binary() (*email.policy.Policy* 方法), 948
- FOR_ITER (opcode), 1678
- forget() (*tkinter.ttk.Notebook* 方法), 1324
- forget() (在 *test.support* 模块中), 1480
- fork() (在 *os* 模块中), 544
- fork() (在 *pty* 模块中), 1714
- ForkingMixIn (*socketserver* 中的类), 1177
- ForkingTCPServer (*socketserver* 中的类), 1177
- ForkingUDPServer (*socketserver* 中的类), 1177
- forkpty() (在 *os* 模块中), 544
- Form (*tkinter.tix* 中的类), 1338
- format (*memoryview* 属性), 68
- format (*struct.Struct* 属性), 144
- format() (*logging.Formatter* 方法), 611
- format() (*logging.Handler* 方法), 611
- format() (*pprint.PrettyPrinter* 方法), 237
- format() (*str* 方法), 42
- format() (*string.Formatter* 方法), 90
- format() (*traceback.StackSummary* 方法), 1588
- format() (*traceback.TracebackException* 方法), 1588
- format() (*tracemalloc.Traceback* 方法), 1524
- format() (置函数), 11
- format() (在 *locale* 模块中), 1259
- format_datetime() (在 *email.utils* 模块中), 984
- format_exc() (在 *traceback* 模块中), 1586
- format_exception() (在 *traceback* 模块中), 1586
- format_exception_only() (*traceback.TracebackException* 方法), 1588
- format_exception_only() (在 *traceback* 模块中), 1586
- format_field() (*string.Formatter* 方法), 91
- format_help() (*argparse.ArgumentParser* 方法), 601
- format_list() (在 *traceback* 模块中), 1586
- format_map() (*str* 方法), 42
- format_stack() (在 *traceback* 模块中), 1587
- format_stack_entry() (*bdb.Bdb* 方法), 1491
- format_string() (在 *locale* 模块中), 1259
- format_tb() (在 *traceback* 模块中), 1586
- format_usage() (*argparse.ArgumentParser* 方法), 601
- FORMAT_VALUE (opcode), 1680
- formataddr() (在 *email.utils* 模块中), 983
- formatargspec() (在 *inspect* 模块中), 1604
- formatargvalues() (在 *inspect* 模块中), 1605
- formatdate() (在 *email.utils* 模块中), 984
- FormatError, 1013
- FormatError() (在 *ctypes* 模块中), 701
- FormatException() (*logging.Formatter* 方法), 612
- formatmonth() (*calendar.HTMLCalendar* 方法), 193
- formatmonth() (*calendar.TextCalendar* 方法), 193
- formatStack() (*logging.Formatter* 方法), 612
- Formatter (*logging* 中的类), 611
- Formatter (*string* 中的类), 90
- formatter (模块), 1683
- formatTime() (*logging.Formatter* 方法), 611
- formatting
 - bytearray (%), 60
 - bytes (%), 60
- formatting, string (%), 48
- formatwarning() (在 *warnings* 模块中), 1567
- formatyear() (*calendar.HTMLCalendar* 方法), 193
- formatyear() (*calendar.TextCalendar* 方法), 193
- formatyearpage() (*calendar.HTMLCalendar* 方法), 193
- Fortran contiguous, 1761
- forward() (在 *turtle* 模块中), 1268

- found_terminator() (*asynchat.async_chat* 方法), 917
 fpathconf() (在 *os* 模块中), 517
 fpectl (模块), 1613
 fqdn (*smtpd.SMTPChannel* 属性), 1169
 Fraction (*fractions* 中的类), 298
 fractions (模块), 298
 frame (*tkinter.scrolledtext.ScrolledText* 属性), 1340
 Frame (*tracemalloc* 中的类), 1521
 FrameSummary (*traceback* 中的类), 1589
 FrameType() (在 *types* 模块中), 232
 freeze_support() (在 *multiprocessing* 模块中), 731
 frexp() (在 *math* 模块中), 265
 from_address() (*ctypes._CData* 方法), 702
 from_buffer() (*ctypes._CData* 方法), 702
 from_buffer_copy() (*ctypes._CData* 方法), 702
 from_bytes() (*int* 类方法), 31
 from_callable() (*inspect.Signature* 类方法), 1601
 from_decimal() (*fractions.Fraction* 方法), 299
 from_exception() (*traceback.TracebackException* 类方法), 1588
 from_file() (*zipfile.ZipInfo* 类方法), 448
 from_float() (*decimal.Decimal* 方法), 279
 from_float() (*fractions.Fraction* 方法), 299
 from_iterable() (*itertools.chain* 类方法), 316
 from_list() (*traceback.StackSummary* 类方法), 1588
 from_param() (*ctypes._CData* 方法), 702
 from_traceback() (*dis.Bytecode* 类方法), 1669
 frombuf() (*tarfile.TarInfo* 类方法), 456
 frombytes() (*array.array* 方法), 221
 fromfd() (*select.epoll* 方法), 845
 fromfd() (*select.kqueue* 方法), 846
 fromfd() (在 *socket* 模块中), 795
 fromfile() (*array.array* 方法), 221
 fromhex() (*bytearray* 类方法), 51
 fromhex() (*bytes* 类方法), 50
 fromhex() (*float* 类方法), 32
 fromkeys() (*collections.Counter* 方法), 198
 fromkeys() (*dict* 类方法), 73
 fromlist() (*array.array* 方法), 222
 fromordinal() (*datetime.date* 类方法), 168
 fromordinal() (*datetime.datetime* 类方法), 172
 fromshare() (在 *socket* 模块中), 795
 fromstring() (*array.array* 方法), 222
 fromstring() (在 *xml.etree.ElementTree* 模块中), 1039
 fromstringlist() (在 *xml.etree.ElementTree* 模块中), 1039
 fromtarfile() (*tarfile.TarInfo* 类方法), 456
 fromtimestamp() (*datetime.date* 类方法), 167
 fromtimestamp() (*datetime.datetime* 类方法), 171
 fromunicode() (*array.array* 方法), 222
 fromutc() (*datetime.timezone* 方法), 189
 fromutc() (*datetime.tzinfo* 方法), 184
 FrozenImporter (*importlib.machinery* 中的类), 1636
 FrozenSet (*typing* 中的类), 1360
 frozenset (F 置类), 69
 fsdecode() (在 *os* 模块中), 511
 fsencode() (在 *os* 模块中), 511
 fspath() (在 *os* 模块中), 511
 fstat() (在 *os* 模块中), 517
 fstatvfs() (在 *os* 模块中), 517
 f-string -- f-字符串, 1762
 fsum() (在 *math* 模块中), 265
 fsync() (在 *os* 模块中), 517
 FTP, 1122
 ftplib (*standard module*), 1139
 protocol, 1121, 1139
 FTP (*ftplib* 中的类), 1140
 ftp_open() (*urllib.request.FTPHandler* 方法), 1116
 FTP-TLS (*ftplib* 中的类), 1140
 FTPHandler (*urllib.request* 中的类), 1109
 ftplib (模块), 1139
 ftruncate() (在 *os* 模块中), 517
 Full, 784
 full() (*asyncio.Queue* 方法), 905
 full() (*multiprocessing.Queue* 方法), 729
 full() (*queue.Queue* 方法), 784
 full_url (*urllib.request.Request* 属性), 1109
 fullmatch() (*re.regex* 方法), 108
 fullmatch() (在 *re* 模块中), 105
 func (*functools.partial* 属性), 333
 funcattrs (*2to3 fixer*), 1475
 function -- 函数, 1762
 Function (*symtable* 中的类), 1655
 function annotation -- 函数标注, 1762
 FunctionTestCase (*unittest* 中的类), 1408
 FunctionType() (在 *types* 模块中), 231
 functools (模块), 327
 funny_files (*filecmp.dircmp* 属性), 370
 future (*2to3 fixer*), 1475
 Future (*asyncio* 中的类), 871
 Future (*concurrent.futures* 中的类), 765
 FutureWarning, 86
 fwalk() (在 *os* 模块中), 540
- ## G
- g, --timing
 trace 命令行选项, 1513
 G.722, 1232
 gaierror, 792
 gamma() (在 *math* 模块中), 268
 gammavariate() (在 *random* 模块中), 303
 garbage collection -- 垃圾回收, 1763
 garbage() (在 *gc* 模块中), 1594
 gather() (*curses.textpad.Textbox* 方法), 662
 gather() (在 *asyncio* 模块中), 876
 gauss() (在 *random* 模块中), 303

- gc (模块), 1592
- gcd() (在 *fractions* 模块中), 300
- gcd() (在 *math* 模块中), 265
- ge() (在 *operator* 模块中), 334
- gen_uuid() (在 *msilib* 模块中), 1690
- generator, 1763
- generator -- 生成器, 1763
- Generator (*collections.abc* 中的类), 212
- Generator (*email.generator* 中的类), 943
- Generator (*typing* 中的类), 1361
- generator expression, 1763
- generator expression -- 生成器表达式, 1763
- generator iterator -- 生成器迭代器, 1763
- GeneratorExit, 80
- GeneratorType() (在 *types* 模块中), 231
- Generic (*typing* 中的类), 1358
- generic function -- 泛型函数, 1763
- generic_visit() (*ast.NodeVisitor* 方法), 1654
- genops() (在 *pickletools* 模块中), 1682
- get() (*asyncio.Queue* 方法), 905
- get() (*configparser.ConfigParser* 方法), 482
- get() (*dict* 方法), 73
- get() (*email.message.EmailMessage* 方法), 933
- get() (*email.message.Message* 方法), 970
- get() (*mailbox.Mailbox* 方法), 998
- get() (*multiprocessing.pool.AsyncResult* 方法), 747
- get() (*multiprocessing.Queue* 方法), 730
- get() (*multiprocessing.SimpleQueue* 方法), 730
- get() (*ossaudiodev.oss_mixer_device* 方法), 1245
- get() (*queue.Queue* 方法), 784
- get() (*tkinter.ttk.Combobox* 方法), 1322
- get() (*types.MappingProxyType* 方法), 232
- get() (在 *webbrowser* 模块中), 1086
- get() (*xml.etree.ElementTree.Element* 方法), 1041
- GET_AITER (*opcode*), 1674
- get_all() (*email.message.EmailMessage* 方法), 933
- get_all() (*email.message.Message* 方法), 970
- get_all() (*wsgiref.headers.Headers* 方法), 1097
- get_all_breaks() (*bdb.Bdb* 方法), 1490
- get_all_start_methods() (在 *multiprocessing* 模块中), 731
- GET_ANEXT (*opcode*), 1674
- get_app() (*wsgiref.simple_server.WSGIServer* 方法), 1098
- get_archive_formats() (在 *shutil* 模块中), 383
- get_asyncgen_hooks() (在 *sys* 模块中), 1550
- GET_AWAITABLE (*opcode*), 1674
- get_begidx() (在 *readline* 模块中), 136
- get_blocking() (在 *os* 模块中), 518
- get_body() (*email.message.EmailMessage* 方法), 936
- get_body_encoding() (*email.charset.Charset* 方法), 981
- get_boundary() (*email.message.EmailMessage* 方法), 934
- get_boundary() (*email.message.Message* 方法), 973
- get_bpbynumber() (*bdb.Bdb* 方法), 1490
- get_break() (*bdb.Bdb* 方法), 1490
- get_breaks() (*bdb.Bdb* 方法), 1490
- get_buffer() (*xdrlib.Packer* 方法), 486
- get_buffer() (*xdrlib.Unpacker* 方法), 487
- get_bytes() (*mailbox.Mailbox* 方法), 998
- get_ca_certs() (*ssl.SSLContext* 方法), 827
- get_cache_token() (在 *abc* 模块中), 1583
- get_channel_binding() (*ssl.SSLSocket* 方法), 824
- get_charset() (*email.message.Message* 方法), 969
- get_charsets() (*email.message.EmailMessage* 方法), 935
- get_charsets() (*email.message.Message* 方法), 973
- get_children() (*symtable.SymbolTable* 方法), 1655
- get_children() (*tkinter.ttk.Treeview* 方法), 1328
- get_ciphers() (*ssl.SSLContext* 方法), 827
- get_clock_info() (在 *time* 模块中), 567
- get_close_matches() (在 *difflib* 模块中), 119
- get_code() (*importlib.abc.InspectLoader* 方法), 1632
- get_code() (*importlib.abc.SourceLoader* 方法), 1634
- get_code() (*importlib.machinery.ExtensionFileLoader* 方法), 1638
- get_code() (*importlib.machinery.SourcelessFileLoader* 方法), 1638
- get_code() (*zipimport.zipimporter* 方法), 1620
- get_completer() (在 *readline* 模块中), 136
- get_completer_delims() (在 *readline* 模块中), 136
- get_completion_type() (在 *readline* 模块中), 136
- get_config_h_filename() (在 *sysconfig* 模块中), 1561
- get_config_var() (在 *sysconfig* 模块中), 1559
- get_config_vars() (在 *sysconfig* 模块中), 1559
- get_content() (*email.contentmanager.ContentManager* 方法), 958
- get_content() (*email.message.EmailMessage* 方法), 937
- get_content() (在 *email.contentmanager* 模块中), 959
- get_content_charset() (*email.message.EmailMessage* 方法), 935
- get_content_charset() (*email.message.Message* 方法), 973
- get_content_disposition() (*email.message.EmailMessage* 方法), 935
- get_content_disposition() (*email.message.Message* 方法), 973
- get_content_maintype() (*email.message.EmailMessage* 方法), 934
- get_content_maintype() (*email.message.Message* 方法), 971
- get_content_subtype() (*email.message.EmailMessage* 方法), 934

- `get_content_subtype()` (*email.message.Message* 方法), 971
`get_content_type()` (*email.message.EmailMessage* 方法), 933
`get_content_type()` (*email.message.Message* 方法), 971
`get_context()` (在 *multiprocessing* 模块中), 731
`get_coroutine_wrapper()` (在 *sys* 模块中), 1550
`get_count()` (在 *gc* 模块中), 1593
`get_current_history_length()` (在 *readline* 模块中), 135
`get_data()` (*importlib.abc.FileLoader* 方法), 1634
`get_data()` (*importlib.abc.ResourceLoader* 方法), 1632
`get_data()` (在 *pkgutil* 模块中), 1623
`get_data()` (*zipimport.zipimporter* 方法), 1620
`get_date()` (*mailbox.MaildirMessage* 方法), 1006
`get_debug()` (*asyncio.AbstractEventLoop* 方法), 862
`get_debug()` (在 *gc* 模块中), 1593
`get_default()` (*argparse.ArgumentParser* 方法), 601
`get_default_domain()` (在 *nis* 模块中), 1722
`get_default_type()` (*email.message.EmailMessage* 方法), 934
`get_default_type()` (*email.message.Message* 方法), 971
`get_default_verify_paths()` (在 *ssl* 模块中), 816
`get_dialect()` (在 *csv* 模块中), 462
`get_docstring()` (在 *ast* 模块中), 1653
`get_doctest()` (*doctest.DocTestParser* 方法), 1385
`get_endidx()` (在 *readline* 模块中), 136
`get_environ()` (*wsgiref.simple_server.WSGIRequestHandler* 方法), 1099
`get_errno()` (在 *ctypes* 模块中), 701
`get_event_loop()` (*asyncio.AbstractEventLoopPolicy* 方法), 867
`get_event_loop()` (在 *asyncio* 模块中), 865
`get_event_loop_policy()` (在 *asyncio* 模块中), 867
`get_examples()` (*doctest.DocTestParser* 方法), 1385
`get_exception_handler()` (*asyncio.AbstractEventLoop* 方法), 861
`get_exec_path()` (在 *os* 模块中), 512
`get_extra_info()` (*asyncio.BaseTransport* 方法), 879
`get_extra_info()` (*asyncio.StreamWriter* 方法), 891
`get_field()` (*string.Formatter* 方法), 90
`get_file()` (*mailbox.Babyl* 方法), 1004
`get_file()` (*mailbox.Mailbox* 方法), 998
`get_file()` (*mailbox.Maildir* 方法), 1001
`get_file()` (*mailbox.mbox* 方法), 1001
`get_file()` (*mailbox.MH* 方法), 1003
`get_file()` (*mailbox.MMDF* 方法), 1004
`get_file_breaks()` (*bdb.Bdb* 方法), 1490
`get_filename()` (*email.message.EmailMessage* 方法), 934
`get_filename()` (*email.message.Message* 方法), 973
`get_filename()` (*importlib.abc.ExecutionLoader* 方法), 1633
`get_filename()` (*importlib.abc.FileLoader* 方法), 1634
`get_filename()` (*importlib.machinery.ExtensionFileLoader* 方法), 1638
`get_filename()` (*zipimport.zipimporter* 方法), 1620
`get_flags()` (*mailbox.MaildirMessage* 方法), 1006
`get_flags()` (*mailbox.mboxMessage* 方法), 1007
`get_flags()` (*mailbox.MMDFMessage* 方法), 1011
`get_folder()` (*mailbox.Maildir* 方法), 1000
`get_folder()` (*mailbox.MH* 方法), 1002
`get_frees()` (*symtable.Function* 方法), 1655
`get_from()` (*mailbox.mboxMessage* 方法), 1007
`get_from()` (*mailbox.MMDFMessage* 方法), 1011
`get_full_url()` (*urllib.request.Request* 方法), 1110
`get_globals()` (*symtable.Function* 方法), 1655
`get_grouped_opcodes()` (*difflib.SequenceMatcher* 方法), 123
`get_handle_inheritable()` (在 *os* 模块中), 523
`get_header()` (*urllib.request.Request* 方法), 1110
`get_history_item()` (在 *readline* 模块中), 135
`get_history_length()` (在 *readline* 模块中), 135
`get_id()` (*symtable.SymbolTable* 方法), 1655
`get_ident()` (在 *_thread* 模块中), 786
`get_ident()` (在 *threading* 模块中), 709
`get_identifiers()` (*symtable.SymbolTable* 方法), 1655
`get_importer()` (在 *pkgutil* 模块中), 1622
`get_info()` (*mailbox.MaildirMessage* 方法), 1006
`get_inheritable()` (*socket.socket* 方法), 800
`get_inheritable()` (在 *os* 模块中), 523
`get_instructions()` (在 *dis* 模块中), 1670
`get_interpreter()` (在 *zipapp* 模块中), 1537
`GET_ITER(opcode)`, 1672
`get_key()` (*selectors.BaseSelector* 方法), 850
`get_labels()` (*mailbox.Babyl* 方法), 1003
`get_labels()` (*mailbox.BabylMessage* 方法), 1010
`get_last_error()` (在 *ctypes* 模块中), 701
`get_line_buffer()` (在 *readline* 模块中), 134
`get_lineno()` (*symtable.SymbolTable* 方法), 1655
`get_loader()` (在 *pkgutil* 模块中), 1622
`get_locals()` (*symtable.Function* 方法), 1655
`get_logger()` (在 *multiprocessing* 模块中), 751
`get_magic()` (在 *imp* 模块中), 1752
`get_makefile_filename()` (在 *sysconfig* 模块中), 1561
`get_map()` (*selectors.BaseSelector* 方法), 850

- get_matching_blocks() (*difflib.SequenceMatcher* 方法), 122
- get_message() (*mailbox.Mailbox* 方法), 998
- get_method() (*urllib.request.Request* 方法), 1110
- get_methods() (*symtable.Class* 方法), 1656
- get_mixed_type_key() (在 *ipaddress* 模块中), 1226
- get_name() (*symtable.Symbol* 方法), 1656
- get_name() (*symtable.SymbolTable* 方法), 1655
- get_namespace() (*symtable.Symbol* 方法), 1656
- get_namespaces() (*symtable.Symbol* 方法), 1656
- get_nonstandard_attr() (*http.cookiejar.Cookie* 方法), 1200
- get_nowait() (*asyncio.Queue* 方法), 905
- get_nowait() (*multiprocessing.Queue* 方法), 730
- get_nowait() (*queue.Queue* 方法), 784
- get_object_traceback() (在 *tracemalloc* 模块中), 1519
- get_objects() (在 *gc* 模块中), 1593
- get_opcodes() (*difflib.SequenceMatcher* 方法), 122
- get_option() (*optparse.OptionParser* 方法), 1743
- get_option_group() (*optparse.OptionParser* 方法), 1734
- get_osfhandle() (在 *msvcrt* 模块中), 1695
- get_output_charset() (*email.charset.Charset* 方法), 981
- get_param() (*email.message.Message* 方法), 972
- get_parameters() (*symtable.Function* 方法), 1655
- get_params() (*email.message.Message* 方法), 971
- get_path() (在 *sysconfig* 模块中), 1560
- get_path_names() (在 *sysconfig* 模块中), 1560
- get_paths() (在 *sysconfig* 模块中), 1560
- get_payload() (*email.message.Message* 方法), 968
- get_pid() (*asyncio.BaseSubprocessTransport* 方法), 881
- get_pipe_transport() (*asyncio.BaseSubprocessTransport* 方法), 881
- get_platform() (在 *sysconfig* 模块中), 1560
- get_poly() (在 *turtle* 模块中), 1282
- get_position() (*xdrlib.Unpacker* 方法), 487
- get_protocol() (*asyncio.BaseTransport* 方法), 879
- get_python_version() (在 *sysconfig* 模块中), 1560
- get_recsrc() (*ossaudiodev.oss_mixer_device* 方法), 1245
- get_referents() (在 *gc* 模块中), 1593
- get_referrers() (在 *gc* 模块中), 1593
- get_request() (*socketserver.BaseServer* 方法), 1179
- get_returncode() (*asyncio.BaseSubprocessTransport* 方法), 881
- get_scheme() (*wsgiref.handlers.BaseHandler* 方法), 1102
- get_scheme_names() (在 *sysconfig* 模块中), 1560
- get_sequences() (*mailbox.MH* 方法), 1002
- get_sequences() (*mailbox.MHMessage* 方法), 1009
- get_server() (*multiprocessing.managers.BaseManager* 方法), 739
- get_server_certificate() (在 *ssl* 模块中), 815
- get_shapepoly() (在 *turtle* 模块中), 1281
- get_socket() (*telnetlib.Telnet* 方法), 1172
- get_source() (*importlib.abc.InspectLoader* 方法), 1632
- get_source() (*importlib.abc.SourceLoader* 方法), 1635
- get_source() (*importlib.machinery.ExtensionFileLoader* 方法), 1638
- get_source() (*importlib.machinery.SourcelessFileLoader* 方法), 1638
- get_source() (*zipimport.zipimporter* 方法), 1620
- get_stack() (*asyncio.Task* 方法), 874
- get_stack() (*bdb.Bdb* 方法), 1491
- get_start_method() (在 *multiprocessing* 模块中), 732
- get_starttag_text() (*html.parser.HTMLParser* 方法), 1027
- get_stats() (在 *gc* 模块中), 1593
- get_stderr() (*wsgiref.handlers.BaseHandler* 方法), 1101
- get_stderr() (*wsgiref.simple_server.WSGIRequestHandler* 方法), 1099
- get_stdin() (*wsgiref.handlers.BaseHandler* 方法), 1101
- get_string() (*mailbox.Mailbox* 方法), 998
- get_subdir() (*mailbox.MaildirMessage* 方法), 1005
- get_suffixes() (在 *imp* 模块中), 1752
- get_symbols() (*symtable.SymbolTable* 方法), 1655
- get_tag() (在 *imp* 模块中), 1754
- get_task_factory() (*asyncio.AbstractEventLoop* 方法), 855
- get_terminal_size() (在 *os* 模块中), 522
- get_terminal_size() (在 *shutil* 模块中), 385
- get_terminator() (*asynchat.async_chat* 方法), 918
- get_threshold() (在 *gc* 模块中), 1593
- get_token() (*shlex.shlex* 方法), 1302
- get_traceback_limit() (在 *tracemalloc* 模块中), 1519
- get_traced_memory() (在 *tracemalloc* 模块中), 1519
- get_tracemalloc_memory() (在 *tracemalloc* 模块中), 1519
- get_type() (*symtable.SymbolTable* 方法), 1655
- get_type_hints() (在 *typing* 模块中), 1363
- get_unixfrom() (*email.message.EmailMessage* 方法), 932
- get_unixfrom() (*email.message.Message* 方法), 968
- get_unpack_formats() (在 *shutil* 模块中), 384
- get_usage() (*optparse.OptionParser* 方法), 1744
- get_value() (*string.Formatter* 方法), 90

- `get_version()` (*optparse.OptionParser* 方法), 1734
`get_visible()` (*mailbox.BabylMessage* 方法), 1010
`get_wch()` (*curses.window* 方法), 653
`get_write_buffer_limits()` (*asyncio.WriteTransport* 方法), 880
`get_write_buffer_size()` (*asyncio.WriteTransport* 方法), 880
`GET_YIELD_FROM_ITER` (*opcode*), 1672
`getacl()` (*imaplib.IMAP4* 方法), 1150
`getaddresses()` (在 *email.utils* 模块中), 983
`getaddrinfo()` (*asyncio.AbstractEventLoop* 方法), 859
`getaddrinfo()` (在 *socket* 模块中), 796
`getallocatedblocks()` (在 *sys* 模块中), 1547
`getannotation()` (*imaplib.IMAP4* 方法), 1150
`getargspec()` (在 *inspect* 模块中), 1604
`getargvalues()` (在 *inspect* 模块中), 1604
`getatime()` (在 *os.path* 模块中), 358
`getattr()` (设置函数), 11
`getattr_static()` (在 *inspect* 模块中), 1607
`getAttribute()` (*xml.dom.Element* 方法), 1053
`getAttributeNode()` (*xml.dom.Element* 方法), 1053
`getAttributeNodeNS()` (*xml.dom.Element* 方法), 1053
`getAttributeNS()` (*xml.dom.Element* 方法), 1053
`GetBase()` (*xml.parsers.expat.xmlparser* 方法), 1076
`getbegyx()` (*curses.window* 方法), 653
`getbkgd()` (*curses.window* 方法), 653
`getboolean()` (*configparser.ConfigParser* 方法), 482
`getbuffer()` (*io.BytesIO* 方法), 561
`getByteStream()` (*xml.sax.xmlreader.InputSource* 方法), 1074
`getcallargs()` (在 *inspect* 模块中), 1605
`getcanvas()` (在 *turtle* 模块中), 1288
`getcapabilities()` (*nntplib.NNTP* 方法), 1156
`getcaps()` (在 *mailcap* 模块中), 996
`getch()` (*curses.window* 方法), 653
`getch()` (在 *msvcrt* 模块中), 1695
`getCharacterStream()` (*xml.sax.xmlreader.InputSource* 方法), 1074
`getche()` (在 *msvcrt* 模块中), 1695
`getcheckinterval()` (在 *sys* 模块中), 1548
`getChild()` (*logging.Logger* 方法), 607
`getchildren()` (*xml.etree.ElementTree.Element* 方法), 1042
`getclasstree()` (在 *inspect* 模块中), 1603
`getclosurevars()` (在 *inspect* 模块中), 1605
`GetColumnInfo()` (*msilib.View* 方法), 1691
`getColumnNumber()` (*xml.sax.xmlreader.Locator* 方法), 1073
`getcomments()` (在 *inspect* 模块中), 1599
`getcompname()` (*aifc.aifc* 方法), 1231
`getcompname()` (*sunau.AU_read* 方法), 1234
`getcompname()` (*wave.Wave_read* 方法), 1236
`getcomptype()` (*aifc.aifc* 方法), 1231
`getcomptype()` (*sunau.AU_read* 方法), 1234
`getcomptype()` (*wave.Wave_read* 方法), 1236
`getContentHandler()` (*xml.sax.xmlreader.XMLReader* 方法), 1072
`getcontext()` (在 *decimal* 模块中), 283
`getcoroutinelocals()` (在 *inspect* 模块中), 1609
`getcoroutinestate()` (在 *inspect* 模块中), 1608
`getctime()` (在 *os.path* 模块中), 359
`getcwd()` (在 *os* 模块中), 526
`getcwdb()` (在 *os* 模块中), 526
`getcwdu (2to3 fixer)`, 1475
`getdecoder()` (在 *codecs* 模块中), 146
`getdefaultencoding()` (在 *sys* 模块中), 1548
`getdefaultlocale()` (在 *locale* 模块中), 1258
`getdefaulttimeout()` (在 *socket* 模块中), 799
`getdlopenflags()` (在 *sys* 模块中), 1548
`getdoc()` (在 *inspect* 模块中), 1599
`getDOMImplementation()` (在 *xml.dom* 模块中), 1048
`getDTDHandler()` (*xml.sax.xmlreader.XMLReader* 方法), 1072
`getEffectiveLevel()` (*logging.Logger* 方法), 607
`getegid()` (在 *os* 模块中), 512
`getElementsByTagName()` (*xml.dom.Document* 方法), 1052
`getElementsByTagName()` (*xml.dom.Element* 方法), 1053
`getElementsByTagNameNS()` (*xml.dom.Document* 方法), 1052
`getElementsByTagNameNS()` (*xml.dom.Element* 方法), 1053
`getencoder()` (在 *codecs* 模块中), 146
`getEncoding()` (*xml.sax.xmlreader.InputSource* 方法), 1074
`getEntityResolver()` (*xml.sax.xmlreader.XMLReader* 方法), 1072
`getenv()` (在 *os* 模块中), 511
`getenvb()` (在 *os* 模块中), 511
`getErrorHandler()` (*xml.sax.xmlreader.XMLReader* 方法), 1072
`geteuid()` (在 *os* 模块中), 512
`getEvent()` (*xml.dom.pulldom.DOMEEventStream* 方法), 1063
`getEventCategory()` (*logging.handlers.NTEventLogHandler* 方法), 639
`getEventType()` (*logging.handlers.NTEventLogHandler* 方法), 640
`getException()` (*xml.sax.SAXException* 方法), 1065
`getFeature()` (*xml.sax.xmlreader.XMLReader* 方法), 1073
`GetFieldCount()` (*msilib.Record* 方法), 1692

- getfile() (在 *inspect* 模块中), 1599
 getfilesystemcodeerrors() (在 *sys* 模块中), 1548
 getfilesystemencoding() (在 *sys* 模块中), 1548
 getfirst() (*cgi.FieldStorage* 方法), 1090
 getfloat() (*configparser.ConfigParser* 方法), 482
 getfmts() (*ossaudiodev.oss_audio_device* 方法), 1242
 getfqdn() (在 *socket* 模块中), 796
 getframeinfo() (在 *inspect* 模块中), 1606
 getframerate() (*aifc.aifc* 方法), 1231
 getframerate() (*sunau.AU_read* 方法), 1234
 getframerate() (*wave.Wave_read* 方法), 1236
 getfullargspec() (在 *inspect* 模块中), 1604
 getgeneratorlocals() (在 *inspect* 模块中), 1608
 getgeneratorstate() (在 *inspect* 模块中), 1608
 getgid() (在 *os* 模块中), 512
 getgrall() (在 *grp* 模块中), 1710
 getgrgid() (在 *grp* 模块中), 1710
 getgrnam() (在 *grp* 模块中), 1710
 getgrouplist() (在 *os* 模块中), 512
 getgroups() (在 *os* 模块中), 512
 getheader() (*http.client.HTTPResponse* 方法), 1137
 getheaders() (*http.client.HTTPResponse* 方法), 1137
 gethostbyaddr() (in module *socket*), 515
 gethostbyaddr() (在 *socket* 模块中), 797
 gethostbyname() (在 *socket* 模块中), 796
 gethostbyname_ex() (在 *socket* 模块中), 796
 gethostname() (in module *socket*), 515
 gethostname() (在 *socket* 模块中), 797
 getincrementaldecoder() (在 *codecs* 模块中), 146
 getincrementalencoder() (在 *codecs* 模块中), 146
 getinfo() (*zipfile.ZipFile* 方法), 445
 getinnerframes() (在 *inspect* 模块中), 1606
 GetInputContext() (*xml.parsers.expat.xmlparser* 方法), 1076
 getint() (*configparser.ConfigParser* 方法), 482
 GetInteger() (*msilib.Record* 方法), 1692
 getitem() (在 *operator* 模块中), 336
 getiterator() (*xml.etree.ElementTree.Element* 方法), 1042
 getiterator() (*xml.etree.ElementTree.ElementTree* 方法), 1043
 getitimer() (在 *signal* 模块中), 922
 getkey() (*curses.window* 方法), 653
 GetLastError() (在 *ctypes* 模块中), 701
 getLength() (*xml.sax.xmlreader.Attributes* 方法), 1074
 getLevelName() (在 *logging* 模块中), 618
 getline() (在 *linecache* 模块中), 377
 getLineNumber() (*xml.sax.xmlreader.Locator* 方法), 1073
 getlist() (*cgi.FieldStorage* 方法), 1090
 getloadavg() (在 *os* 模块中), 552
 getlocale() (在 *locale* 模块中), 1258
 getLogger() (在 *logging* 模块中), 616
 getLoggerClass() (在 *logging* 模块中), 616
 getlogin() (在 *os* 模块中), 512
 getLogRecordFactory() (在 *logging* 模块中), 616
 getmark() (*aifc.aifc* 方法), 1231
 getmark() (*sunau.AU_read* 方法), 1234
 getmark() (*wave.Wave_read* 方法), 1236
 getmarkers() (*aifc.aifc* 方法), 1231
 getmarkers() (*sunau.AU_read* 方法), 1234
 getmarkers() (*wave.Wave_read* 方法), 1236
 getmaxyx() (*curses.window* 方法), 653
 getmember() (*tarfile.TarFile* 方法), 454
 getmembers() (*tarfile.TarFile* 方法), 454
 getmembers() (在 *inspect* 模块中), 1597
 getMessage() (*logging.LogRecord* 方法), 613
 getMessage() (*xml.sax.SAXException* 方法), 1065
 getMessageID() (*logging.handlers.NTEventLogHandler* 方法), 640
 getmodule() (在 *inspect* 模块中), 1599
 getmodulename() (在 *inspect* 模块中), 1597
 getmouse() (在 *curses* 模块中), 646
 getmro() (在 *inspect* 模块中), 1605
 getmtime() (在 *os.path* 模块中), 359
 getname() (*chunk.Chunk* 方法), 1238
 getName() (*threading.Thread* 方法), 712
 getNameByQName() (*xml.sax.xmlreader.AttributesNS* 方法), 1075
 getnameinfo() (*asyncio.AbstractEventLoop* 方法), 859
 getnameinfo() (在 *socket* 模块中), 797
 getnames() (*tarfile.TarFile* 方法), 454
 getNames() (*xml.sax.xmlreader.Attributes* 方法), 1074
 getnchannels() (*aifc.aifc* 方法), 1231
 getnchannels() (*sunau.AU_read* 方法), 1234
 getnchannels() (*wave.Wave_read* 方法), 1236
 getnframes() (*aifc.aifc* 方法), 1231
 getnframes() (*sunau.AU_read* 方法), 1234
 getnframes() (*wave.Wave_read* 方法), 1236
 getnode, 1174
 getnode() (在 *uuid* 模块中), 1174
 getopt (模块), 603
 getopt() (在 *getopt* 模块中), 603
 GetoptError, 604
 getouterframes() (在 *inspect* 模块中), 1606
 getoutput() (在 *subprocess* 模块中), 781
 getpagesize() (在 *resource* 模块中), 1721
 getparams() (*aifc.aifc* 方法), 1231
 getparams() (*sunau.AU_read* 方法), 1234
 getparams() (*wave.Wave_read* 方法), 1236
 getparyx() (*curses.window* 方法), 653
 getpass (模块), 643

- getpass() (在 *getpass* 模块中), 643
 GetPassWarning, 644
 getpeercert() (*ssl.SSLSocket* 方法), 823
 getpeername() (*socket.socket* 方法), 801
 getpen() (在 *turtle* 模块中), 1282
 getpgid() (在 *os* 模块中), 512
 getpgrp() (在 *os* 模块中), 512
 getpid() (在 *os* 模块中), 512
 getpos() (*html.parser.HTMLParser* 方法), 1027
 getppid() (在 *os* 模块中), 513
 getpreferredencoding() (在 *locale* 模块中), 1259
 getpriority() (在 *os* 模块中), 513
 getprofile() (在 *sys* 模块中), 1549
 GetProperty() (*msilib.SummaryInformation* 方法), 1691
 getProperty() (*xml.sax.xmlreader.XMLReader* 方法), 1073
 GetPropertyCount() (*msilib.SummaryInformation* 方法), 1691
 getprotobyname() (在 *socket* 模块中), 797
 getproxies() (在 *urllib.request* 模块中), 1106
 getPublicId() (*xml.sax.xmlreader.InputSource* 方法), 1074
 getPublicId() (*xml.sax.xmlreader.Locator* 方法), 1073
 getpwall() (在 *pwd* 模块中), 1709
 getpwnam() (在 *pwd* 模块中), 1709
 getpwuid() (在 *pwd* 模块中), 1709
 getQNameByName() (*xml.sax.xmlreader.AttributesNS* 方法), 1075
 getQNames() (*xml.sax.xmlreader.AttributesNS* 方法), 1075
 getquota() (*imaplib.IMAP4* 方法), 1150
 getquotaroot() (*imaplib.IMAP4* 方法), 1150
 getrandbits() (在 *random* 模块中), 301
 getRandom() (在 *os* 模块中), 553
 getreader() (在 *codecs* 模块中), 146
 getrecursionlimit() (在 *sys* 模块中), 1548
 getrefcount() (在 *sys* 模块中), 1548
 getresgid() (在 *os* 模块中), 513
 getresponse() (*http.client.HTTPConnection* 方法), 1136
 getresuid() (在 *os* 模块中), 513
 getrlimit() (在 *resource* 模块中), 1718
 getroot() (*xml.etree.ElementTree.ElementTree* 方法), 1043
 getrusage() (在 *resource* 模块中), 1721
 getsample() (在 *audioop* 模块中), 1228
 getsampwidth() (*aifc.aifc* 方法), 1231
 getsampwidth() (*sunau.AU_read* 方法), 1234
 getsampwidth() (*wave.Wave_read* 方法), 1236
 getscreen() (在 *turtle* 模块中), 1283
 getservbyname() (在 *socket* 模块中), 797
 getservbyport() (在 *socket* 模块中), 797
 GetSetDescriptorType() (在 *types* 模块中), 232
 getshapes() (在 *turtle* 模块中), 1289
 getsid() (在 *os* 模块中), 515
 getsignal() (在 *signal* 模块中), 921
 getsitepackages() (在 *site* 模块中), 1612
 getsize() (*chunk.Chunk* 方法), 1238
 getsize() (在 *os.path* 模块中), 359
 getsizeof() (在 *sys* 模块中), 1548
 getsockname() (*socket.socket* 方法), 801
 getsockopt() (*socket.socket* 方法), 801
 getsource() (在 *inspect* 模块中), 1599
 getsourcefile() (在 *inspect* 模块中), 1599
 getsourcelines() (在 *inspect* 模块中), 1599
 getspall() (在 *spwd* 模块中), 1709
 getspnam() (在 *spwd* 模块中), 1709
 getstate() (*codecs.IncrementalDecoder* 方法), 151
 getstate() (*codecs.IncrementalEncoder* 方法), 150
 getstate() (在 *random* 模块中), 301
 getstatusoutput() (在 *subprocess* 模块中), 780
 getstr() (*curses.window* 方法), 653
 GetString() (*msilib.Record* 方法), 1692
 getSubject() (*logging.handlers.SMTPHandler* 方法), 640
 GetSummaryInformation() (*msilib.Database* 方法), 1690
 getswitchinterval() (在 *sys* 模块中), 1549
 getSystemId() (*xml.sax.xmlreader.InputSource* 方法), 1074
 getSystemId() (*xml.sax.xmlreader.Locator* 方法), 1073
 getsyx() (在 *curses* 模块中), 646
 gettarinfo() (*tarfile.TarFile* 方法), 455
 gettempdir() (在 *tempfile* 模块中), 373
 gettempdirb() (在 *tempfile* 模块中), 373
 gettempprefix() (在 *tempfile* 模块中), 373
 gettempprefixb() (在 *tempfile* 模块中), 373
 getTestCaseNames() (*unittest.TestLoader* 方法), 1411
 gettext (模块), 1247
 gettext() (*gettext.GNUTranslations* 方法), 1251
 gettext() (*gettext.NullTranslations* 方法), 1250
 gettext() (在 *gettext* 模块中), 1248
 gettext() (在 *locale* 模块中), 1261
 gettimeout() (*socket.socket* 方法), 801
 gettrace() (在 *sys* 模块中), 1549
 getturtle() (在 *turtle* 模块中), 1282
 getType() (*xml.sax.xmlreader.Attributes* 方法), 1074
 getuid() (在 *os* 模块中), 513
 geturl() (*urllib.parse.urllib.parse.SplitResult* 方法), 1126
 getuser() (在 *getpass* 模块中), 644
 getuserbase() (在 *site* 模块中), 1612
 getusersitepackages() (在 *site* 模块中), 1612

getvalue() (*io.BytesIO* 方法), 561
 getvalue() (*io.StringIO* 方法), 564
 getValue() (*xml.sax.xmlreader.Attributes* 方法), 1074
 getValueByQName() (*xml.sax.xmlreader.AttributesNS* 方法), 1075
 getwch() (在 *msvcrt* 模块中), 1695
 getwche() (在 *msvcrt* 模块中), 1696
 getweakrefcount() (在 *weakref* 模块中), 224
 getweakrefs() (在 *weakref* 模块中), 224
 getwelcome() (*ftplib.FTP* 方法), 1142
 getwelcome() (*nnplib.NNTP* 方法), 1156
 getwelcome() (*poplib.POP3* 方法), 1145
 getwin() (在 *curses* 模块中), 646
 getwindowsversion() (在 *sys* 模块中), 1549
 getwriter() (在 *codecs* 模块中), 146
 getxattr() (在 *os* 模块中), 541
 getyx() (*curses.window* 方法), 653
 gid (*tarfile.TarInfo* 属性), 457
 GIL, 1763
 glob
 模块, 376
 glob (模块), 375
 glob() (*msilib.Directory* 方法), 1693
 glob() (*pathlib.Path* 方法), 352
 glob() (在 *glob* 模块中), 375
 global interpreter lock -- 全局解释器锁, 1763
 globals() (☐置函数), 11
 globs (*doctest.DocTest* 属性), 1383
 gmtime() (在 *time* 模块中), 568
 gname (*tarfile.TarInfo* 属性), 457
 GNOME, 1252
 GNU_FORMAT() (在 *tarfile* 模块中), 453
 gnu_getopt() (在 *getopt* 模块中), 604
 GNUTranslations (*gettext* 中的类), 1251
 got (*doctest.DocTestFailure* 属性), 1390
 goto() (在 *turtle* 模块中), 1268
 Graphical User Interface, 1307
 GREATER() (在 *token* 模块中), 1657
 GREATEREQUAL() (在 *token* 模块中), 1657
 Greenwich Mean Time, 566
 GRND_NONBLOCK() (在 *os* 模块中), 554
 GRND_RANDOM() (在 *os* 模块中), 554
 Group (*email.headerregistry* 中的类), 957
 group() (*nnplib.NNTP* 方法), 1157
 group() (*pathlib.Path* 方法), 353
 group() (*re.match* 方法), 109
 groupby() (在 *itertools* 模块中), 318
 groupdict() (*re.match* 方法), 111
 groupindex (*re.regex* 属性), 109
 groups (*email.headerregistry.AddressHeader* 属性), 954
 groups (*re.regex* 属性), 109
 groups() (*re.match* 方法), 110
 grp (模块), 1710

gt() (在 *operator* 模块中), 334
 guess_all_extensions() (*mimetypes.MimeTypes* 方法), 1016
 guess_all_extensions() (在 *mimetypes* 模块中), 1014
 guess_extension() (*mimetypes.MimeTypes* 方法), 1016
 guess_extension() (在 *mimetypes* 模块中), 1014
 guess_scheme() (在 *wsgiref.util* 模块中), 1095
 guess_type() (*mimetypes.MimeTypes* 方法), 1016
 guess_type() (在 *mimetypes* 模块中), 1014
 GUI, 1307
 gzip (模块), 432
 GzipFile (*gzip* 中的类), 433

H

-h, --help
 json.tool 命令行选项, 995
 timeit 命令行选项, 1510
 tokenize 命令行选项, 1660
 zipapp 命令行选项, 1536
 halfdelay() (在 *curses* 模块中), 647
 Handle (*asyncio* 中的类), 862
 handle() (*http.server.BaseHTTPRequestHandler* 方法), 1186
 handle() (*logging.Handler* 方法), 610
 handle() (*logging.handlers.QueueListener* 方法), 643
 handle() (*logging.Logger* 方法), 609
 handle() (*logging.NullHandler* 方法), 632
 handle() (*socketserver.BaseRequestHandler* 方法), 1180
 handle() (*wsgiref.simple_server.WSGIRequestHandler* 方法), 1099
 handle_accept() (*asyncore.dispatcher* 方法), 914
 handle_accepted() (*asyncore.dispatcher* 方法), 914
 handle_charref() (*html.parser.HTMLParser* 方法), 1027
 handle_close() (*asyncore.dispatcher* 方法), 914
 handle_comment() (*html.parser.HTMLParser* 方法), 1027
 handle_connect() (*asyncore.dispatcher* 方法), 914
 handle_data() (*html.parser.HTMLParser* 方法), 1027
 handle_decl() (*html.parser.HTMLParser* 方法), 1028
 handle_defect() (*email.policy.Policy* 方法), 947
 handle_endtag() (*html.parser.HTMLParser* 方法), 1027
 handle_entityref() (*html.parser.HTMLParser* 方法), 1027
 handle_error() (*asyncore.dispatcher* 方法), 914
 handle_error() (*socketserver.BaseServer* 方法), 1179
 handle_expect_100()
 (*http.server.BaseHTTPRequestHandler* 方法), 1186

- [handle_expt\(\)](#) (*asyncore.dispatcher* 方法), 914
[handle_one_request\(\)](#) (*http.server.BaseHTTPRequestHandler* 方法), 1186
[handle_pi\(\)](#) (*html.parser.HTMLParser* 方法), 1028
[handle_read\(\)](#) (*asyncore.dispatcher* 方法), 913
[handle_request\(\)](#) (*socketserver.BaseServer* 方法), 1178
[handle_request\(\)](#) (*xmlrpc.server.CGIXMLRPCRequestHandler* 方法), 1212
[handle_startendtag\(\)](#) (*html.parser.HTMLParser* 方法), 1027
[handle_starttag\(\)](#) (*html.parser.HTMLParser* 方法), 1027
[handle_timeout\(\)](#) (*socketserver.BaseServer* 方法), 1179
[handle_write\(\)](#) (*asyncore.dispatcher* 方法), 913
[handleError\(\)](#) (*logging.Handler* 方法), 610
[handleError\(\)](#) (*logging.handlers.SocketHandler* 方法), 636
[Handler](#) (*logging* 中的类), 610
[handler\(\)](#) (在 *cgiib* 模块中), 1094
[harmonic_mean\(\)](#) (在 *statistics* 模块中), 308
[HAS_ALPN\(\)](#) (在 *ssl* 模块中), 820
[has_children\(\)](#) (*syntable.SymbolTable* 方法), 1655
[has_colors\(\)](#) (在 *curses* 模块中), 646
[HAS_ECDH\(\)](#) (在 *ssl* 模块中), 820
[has_exec\(\)](#) (*syntable.SymbolTable* 方法), 1655
[has_extn\(\)](#) (*smtpplib.SMTP* 方法), 1163
[has_header\(\)](#) (*csv.Sniffer* 方法), 464
[has_header\(\)](#) (*urllib.request.Request* 方法), 1110
[has_ic\(\)](#) (在 *curses* 模块中), 646
[has_il\(\)](#) (在 *curses* 模块中), 646
[has_ipv6\(\)](#) (在 *socket* 模块中), 794
[has_key\(2to3 fixer\)](#), 1475
[has_key\(\)](#) (在 *curses* 模块中), 646
[has_location](#) (*importlib.machinery.ModuleSpec* 属性), 1639
[has_nonstandard_attr\(\)](#) (*http.cookiejar.Cookie* 方法), 1200
[HAS_NPN\(\)](#) (在 *ssl* 模块中), 820
[has_option\(\)](#) (*configparser.ConfigParser* 方法), 481
[has_option\(\)](#) (*optparse.OptionParser* 方法), 1743
[has_section\(\)](#) (*configparser.ConfigParser* 方法), 481
[HAS_SNI\(\)](#) (在 *ssl* 模块中), 820
[has_ticket](#) (*ssl.SSLSession* 属性), 839
[HAS_TLSv1_3\(\)](#) (在 *ssl* 模块中), 821
[hasattr\(\)](#) (内置函数), 11
[hasAttribute\(\)](#) (*xml.dom.Element* 方法), 1053
[hasAttributeNS\(\)](#) (*xml.dom.Element* 方法), 1053
[hasAttributes\(\)](#) (*xml.dom.Node* 方法), 1050
[hasChildNodes\(\)](#) (*xml.dom.Node* 方法), 1050
[hascompare\(\)](#) (在 *dis* 模块中), 1680
[hasconst\(\)](#) (在 *dis* 模块中), 1680
[hasFeature\(\)](#) (*xml.dom.DOMImplementation* 方法), 1049
[hasfree\(\)](#) (在 *dis* 模块中), 1680
[hash](#) 内置函数, 36
[hash\(\)](#) (内置函数), 12
[hash_info\(\)](#) (在 *sys* 模块中), 1550
[hashable -- 可哈希](#), 1763
[Hashable](#) (*collections.abc* 中的类), 211
[Hashable](#) (*typing* 中的类), 1359
[hasHandlers\(\)](#) (*logging.Logger* 方法), 609
[hash.block_size\(\)](#) (在 *hashlib* 模块中), 495
[hash.digest_size\(\)](#) (在 *hashlib* 模块中), 495
[hashlib](#) (模块), 493
[hasjabs\(\)](#) (在 *dis* 模块中), 1680
[hasjrel\(\)](#) (在 *dis* 模块中), 1680
[haslocal\(\)](#) (在 *dis* 模块中), 1680
[hasname\(\)](#) (在 *dis* 模块中), 1680
[HAVE_ARGUMENT](#) (*opcode*), 1680
[HAVE_THREADS\(\)](#) (在 *decimal* 模块中), 289
[HCI_DATA_DIR\(\)](#) (在 *socket* 模块中), 794
[HCI_FILTER\(\)](#) (在 *socket* 模块中), 794
[HCI_TIME_STAMP\(\)](#) (在 *socket* 模块中), 794
[head\(\)](#) (*nntplib.NNTP* 方法), 1159
[Header](#) (*email.header* 中的类), 978
[header_encode\(\)](#) (*email.charset.Charset* 方法), 981
[header_encode_lines\(\)](#) (*email.charset.Charset* 方法), 981
[header_encoding](#) (*email.charset.Charset* 属性), 980
[header_factory](#) (*email.policy.EmailPolicy* 属性), 949
[header_fetch_parse\(\)](#) (*email.policy.Compat32* 方法), 951
[header_fetch_parse\(\)](#) (*email.policy.EmailPolicy* 方法), 949
[header_fetch_parse\(\)](#) (*email.policy.Policy* 方法), 948
[header_items\(\)](#) (*urllib.request.Request* 方法), 1110
[header_max_count\(\)](#) (*email.policy.EmailPolicy* 方法), 949
[header_max_count\(\)](#) (*email.policy.Policy* 方法), 947
[header_offset](#) (*zipfile.ZipInfo* 属性), 450
[header_source_parse\(\)](#) (*email.policy.Compat32* 方法), 951
[header_source_parse\(\)](#) (*email.policy.EmailPolicy* 方法), 949
[header_source_parse\(\)](#) (*email.policy.Policy* 方法), 947
[header_store_parse\(\)](#) (*email.policy.Compat32* 方法), 951
[header_store_parse\(\)](#) (*email.policy.EmailPolicy* 方法), 949

- header_store_parse() (*email.policy.Policy* 方法), 948
- HeaderError, 452
- HeaderParseError, 951
- HeaderParser (*email.parser* 中的类), 940
- HeaderRegistry (*email.headerregistry* 中的类), 956
- headers
- MIME, 1014, 1087
- headers (*http.server.BaseHTTPRequestHandler* 属性), 1185
- headers (*urllib.error.HTTPError* 属性), 1130
- Headers (*wsgiref.headers* 中的类), 1097
- headers (*xmlrpc.client.ProtocolError* 属性), 1206
- heading() (*tkinter.ttk.Treeview* 方法), 1329
- heading() (在 *turtle* 模块中), 1273
- heapify() (在 *heapq* 模块中), 215
- heapmin() (在 *msvcrt* 模块中), 1696
- heappop() (在 *heapq* 模块中), 214
- heappush() (在 *heapq* 模块中), 214
- heappushpop() (在 *heapq* 模块中), 214
- heapq (模块), 214
- heapreplace() (在 *heapq* 模块中), 215
- helo() (*smtpplib.SMTP* 方法), 1163
- help
- online, 1366
- help
- trace 命令行选项, 1512
- help (*optparse.Option* 属性), 1739
- help (*pdb command*), 1496
- help() (*nnplib.NNTP* 方法), 1158
- help() (设置函数), 12
- herror, 792
- hex (*uuid.UUID* 属性), 1174
- hex() (*bytearray* 方法), 51
- hex() (*bytes* 方法), 50
- hex() (*float* 方法), 32
- hex() (*memoryview* 方法), 65
- hex() (设置函数), 12
- hexadecimal
- literals, 29
- hexbin() (在 *binhex* 模块中), 1020
- hexdigest() (*hashlib.hash* 方法), 495
- hexdigest() (*hashlib.shake* 方法), 495
- hexdigest() (*hmac.HMAC* 方法), 504
- hexdigits() (在 *string* 模块中), 89
- hexlify() (在 *binascii* 模块中), 1021
- hexversion() (在 *sys* 模块中), 1550
- hidden() (*curses.panel.Panel* 方法), 665
- hide() (*curses.panel.Panel* 方法), 665
- hide() (*tkinter.ttk.Notebook* 方法), 1324
- hide_cookie2 (*http.cookiejar.CookiePolicy* 属性), 1197
- hideturtle() (在 *turtle* 模块中), 1278
- HierarchyRequestErr, 1055
- HIGHEST_PROTOCOL() (在 *pickle* 模块中), 389
- HKEY_CLASSES_ROOT() (在 *winreg* 模块中), 1701
- HKEY_CURRENT_CONFIG() (在 *winreg* 模块中), 1701
- HKEY_CURRENT_USER() (在 *winreg* 模块中), 1701
- HKEY_DYN_DATA() (在 *winreg* 模块中), 1702
- HKEY_LOCAL_MACHINE() (在 *winreg* 模块中), 1701
- HKEY_PERFORMANCE_DATA() (在 *winreg* 模块中), 1701
- HKEY_USERS() (在 *winreg* 模块中), 1701
- hline() (*curses.window* 方法), 653
- HList (*tkinter.tix* 中的类), 1337
- hls_to_rgb() (在 *coloursys* 模块中), 1239
- hmac (模块), 503
- HOME, 358
- home() (*pathlib.Path* 类方法), 351
- home() (在 *turtle* 模块中), 1270
- HOMEDRIVE, 358
- HOMEPATH, 358
- hook_compressed() (在 *fileinput* 模块中), 363
- hook_encoded() (在 *fileinput* 模块中), 363
- host (*urllib.request.Request* 属性), 1109
- hostmask (*ipaddress.IPv4Network* 属性), 1220
- hostmask (*ipaddress.IPv6Network* 属性), 1222
- hosts (*netrc.netrc* 属性), 485
- hosts() (*ipaddress.IPv4Network* 方法), 1220
- hosts() (*ipaddress.IPv6Network* 方法), 1222
- hour (*datetime.datetime* 属性), 173
- hour (*datetime.time* 属性), 180
- HRESULT (*ctypes* 中的类), 705
- hStdError (*subprocess.STARTUPINFO* 属性), 775
- hStdInput (*subprocess.STARTUPINFO* 属性), 775
- hStdOutput (*subprocess.STARTUPINFO* 属性), 775
- hsv_to_rgb() (在 *coloursys* 模块中), 1239
- ht() (在 *turtle* 模块中), 1278
- HTML, 1026, 1121
- html (模块), 1025
- html5() (在 *html.entities* 模块中), 1030
- HTMLCalendar (*calendar* 中的类), 193
- HtmlDiff (*difflib* 中的类), 117
- html.entities (模块), 1030
- HTMLParser (*html.parser* 中的类), 1026
- html.parser (模块), 1026
- htonl() (在 *socket* 模块中), 797
- htons() (在 *socket* 模块中), 797
- HTTP
- http (*standard module*), 1131
 - http.client (*standard module*), 1133
 - protocol, 1087, 1121, 1131, 1133, 1184
- http (模块), 1131
- HTTP() (在 *email.policy* 模块中), 950
- http_error_301() (*urllib.request.HTTPRedirectHandler* 方法), 1113
- http_error_302() (*urllib.request.HTTPRedirectHandler* 方法), 1113

- [http_error_303\(\)](#) ([url-lib.request.HTTPRedirectHandler](#) 方法), 1113
[http_error_307\(\)](#) ([url-lib.request.HTTPRedirectHandler](#) 方法), 1113
[http_error_401\(\)](#) ([url-lib.request.HTTPBasicAuthHandler](#) 方法), 1114
[http_error_401\(\)](#) ([url-lib.request.HTTPDigestAuthHandler](#) 方法), 1115
[http_error_407\(\)](#) ([url-lib.request.ProxyBasicAuthHandler](#) 方法), 1115
[http_error_407\(\)](#) ([url-lib.request.ProxyDigestAuthHandler](#) 方法), 1115
[http_error_auth_reged\(\)](#) ([url-lib.request.AbstractBasicAuthHandler](#) 方法), 1114
[http_error_auth_reged\(\)](#) ([url-lib.request.AbstractDigestAuthHandler](#) 方法), 1115
[http_error_default\(\)](#) ([urllib.request.BaseHandler](#) 方法), 1112
[http_error_nnn\(\)](#) ([urllib.request.BaseHandler](#) 方法), 1112
[http_open\(\)](#) ([urllib.request.HTTPHandler](#) 方法), 1115
[HTTP_PORT\(\)](#) (在 [http.client](#) 模块中), 1135
[http_proxy](#), 1105, 1118
[http_response\(\)](#) ([urllib.request.HTTPErrorProcessor](#) 方法), 1116
[http_version](#) ([wsgiref.handlers.BaseHandler](#) 属性), 1103
[HTTPBasicAuthHandler](#) ([urllib.request](#) 中的类), 1108
[http.client](#) (模块), 1133
[HTTPConnection](#) ([http.client](#) 中的类), 1133
[http.cookiejar](#) (模块), 1192
[HTTPCookieProcessor](#) ([urllib.request](#) 中的类), 1107
[http.cookies](#) (模块), 1189
[httpd](#), 1184
[HTTPDefaultErrorHandler](#) ([urllib.request](#) 中的类), 1107
[HTTPDigestAuthHandler](#) ([urllib.request](#) 中的类), 1108
[HTTPError](#), 1129
[HTTPErrorProcessor](#) ([urllib.request](#) 中的类), 1109
[HTTPException](#), 1134
[HTTPHandler](#) ([logging.handlers](#) 中的类), 641
[HTTPHandler](#) ([urllib.request](#) 中的类), 1109
[HTTPPasswordMgr](#) ([urllib.request](#) 中的类), 1108
[HTTPPasswordMgrWithDefaultRealm](#) ([url-lib.request](#) 中的类), 1108
[HTTPPasswordMgrWithPriorAuth](#) ([urllib.request](#) 中的类), 1108
[HTTPRedirectHandler](#) ([urllib.request](#) 中的类), 1107
[HTTPResponse](#) ([http.client](#) 中的类), 1134
[https_open\(\)](#) ([urllib.request.HTTPSHandler](#) 方法), 1115
[HTTPS_PORT\(\)](#) (在 [http.client](#) 模块中), 1135
[https_response\(\)](#) ([url-lib.request.HTTPErrorProcessor](#) 方法), 1116
[HTTPSConnection](#) ([http.client](#) 中的类), 1134
[HTTPServer](#) ([http.server](#) 中的类), 1184
[http.server](#) (模块), 1184
[HTTPSHandler](#) ([urllib.request](#) 中的类), 1109
[HTTPStatus](#) ([http](#) 中的类), 1131
[hypot\(\)](#) (在 [math](#) 模块中), 267

|
-i list
 [compileall](#) 命令行选项, 1666
[I\(\)](#) (在 [re](#) 模块中), 104
I/O control
 [buffering](#), 17, 801
 [POSIX](#), 1712
 [tty](#), 1712
 [UNIX](#), 1715
[iadd\(\)](#) (在 [operator](#) 模块中), 339
[iand\(\)](#) (在 [operator](#) 模块中), 339
[iconcat\(\)](#) (在 [operator](#) 模块中), 339
[id](#) ([ssl.SSLSession](#) 属性), 839
[id\(\)](#) ([unittest.TestCase](#) 方法), 1407
[id\(\)](#) (设置函数), 12
[idcok\(\)](#) ([curses.window](#) 方法), 653
[ident](#) ([select.kevent](#) 属性), 847
[ident](#) ([threading.Thread](#) 属性), 712
[identchars](#) ([cmd.Cmd](#) 属性), 1297
[identify\(\)](#) ([tkinter.ttk.Notebook](#) 方法), 1324
[identify\(\)](#) ([tkinter.ttk.Treeview](#) 方法), 1329
[identify\(\)](#) ([tkinter.ttk.Widget](#) 方法), 1321
[identify_column\(\)](#) ([tkinter.ttk.Treeview](#) 方法), 1329
[identify_element\(\)](#) ([tkinter.ttk.Treeview](#) 方法), 1330
[identify_region\(\)](#) ([tkinter.ttk.Treeview](#) 方法), 1330
[identify_row\(\)](#) ([tkinter.ttk.Treeview](#) 方法), 1329
[idioms](#) ([2to3 fixer](#)), 1475
[IDLE](#), 1340, 1763
[IDLESTARTUP](#), 1346
[idlok\(\)](#) ([curses.window](#) 方法), 654
[IEEE-754](#), 1613
if
 语句, 27
[if_indextoname\(\)](#) (在 [socket](#) 模块中), 799
[if_nameindex\(\)](#) (在 [socket](#) 模块中), 799
[if_nametoindex\(\)](#) (在 [socket](#) 模块中), 799

- `ifloordiv()` (在 *operator* 模块中), 339
- `iglob()` (在 *glob* 模块中), 375
- `ignorableWhitespace()`
(*xml.sax.handler.ContentHandler* 方法), 1069
- `ignore` (*pdb* command), 1497
- `ignore_errors()` (在 *codecs* 模块中), 149
- `IGNORE_EXCEPTION_DETAIL()` (在 *doctest* 模块中), 1375
- `ignore_patterns()` (在 *shutil* 模块中), 379
- `IGNORECASE()` (在 *re* 模块中), 104
- `--ignore-dir=<dir>`
 trace 命令行选项, 1513
- `--ignore-module=<mod>`
 trace 命令行选项, 1513
- `ihave()` (*nntplib.NNTP* 方法), 1159
- IISCGIHandler* (*wsgiref.handlers* 中的类), 1100
- `ilshift()` (在 *operator* 模块中), 339
- `imag` (*numbers.Complex* 属性), 261
- `imap()` (*multiprocessing.pool.Pool* 方法), 746
- IMAP4*
 protocol, 1147
- IMAP4* (*imaplib* 中的类), 1147
- IMAP4_SSL*
 protocol, 1147
- IMAP4_SSL* (*imaplib* 中的类), 1148
- IMAP4_stream*
 protocol, 1147
- IMAP4_stream* (*imaplib* 中的类), 1148
- IMAP4.abort*, 1148
- IMAP4.error*, 1148
- IMAP4.readonly*, 1148
- `imap_unordered()` (*multiprocessing.pool.Pool* 方法), 746
- imaplib* (模块), 1147
- `imatmul()` (在 *operator* 模块中), 339
- imgchr* (模块), 1240
- `immedok()` (*curses.window* 方法), 654
- immutable*
 sequence types, 36
- immutable* -- 不可变, 1764
- `imod()` (在 *operator* 模块中), 339
- imp*
 模块, 22
- imp* (模块), 1751
- ImpImporter* (*pkgutil* 中的类), 1621
- `implementation()` (在 *sys* 模块中), 1551
- ImpLoader* (*pkgutil* 中的类), 1622
- import*
 语句, 22, 1610, 1751
- `import` (*2to3 fixer*), 1475
- `import path` -- 导入路径, 1764
- `import_fresh_module()` (在 *test.support* 模块中), 1483
- `IMPORT_FROM` (*opcode*), 1677
- `import_module()` (在 *importlib* 模块中), 1628
- `import_module()` (在 *test.support* 模块中), 1483
- `IMPORT_NAME` (*opcode*), 1677
- `IMPORT_STAR` (*opcode*), 1675
- importer* -- 导入器, 1764
- ImportError*, 81
- `importing` -- 导入, 1764
- importlib* (模块), 1627
- importlib.abc* (模块), 1629
- importlib.machinery* (模块), 1635
- importlib.util* (模块), 1639
- `imports` (*2to3 fixer*), 1475
- `imports2` (*2to3 fixer*), 1475
- ImportWarning*, 86
- ImproperConnectionState*, 1134
- `imul()` (在 *operator* 模块中), 339
- in*
 运算符, 28, 35
- `in_dll()` (*ctypes._CData* 方法), 703
- `in_table_a1()` (在 *stringprep* 模块中), 132
- `in_table_b1()` (在 *stringprep* 模块中), 133
- `in_table_c3()` (在 *stringprep* 模块中), 133
- `in_table_c4()` (在 *stringprep* 模块中), 133
- `in_table_c5()` (在 *stringprep* 模块中), 133
- `in_table_c6()` (在 *stringprep* 模块中), 133
- `in_table_c7()` (在 *stringprep* 模块中), 133
- `in_table_c8()` (在 *stringprep* 模块中), 133
- `in_table_c9()` (在 *stringprep* 模块中), 133
- `in_table_c11()` (在 *stringprep* 模块中), 133
- `in_table_c11_c12()` (在 *stringprep* 模块中), 133
- `in_table_c12()` (在 *stringprep* 模块中), 133
- `in_table_c21()` (在 *stringprep* 模块中), 133
- `in_table_c21_c22()` (在 *stringprep* 模块中), 133
- `in_table_c22()` (在 *stringprep* 模块中), 133
- `in_table_d1()` (在 *stringprep* 模块中), 133
- `in_table_d2()` (在 *stringprep* 模块中), 133
- `in_transaction` (*sqlite3.Connection* 属性), 412
- `inch()` (*curses.window* 方法), 654
- inclusive* (*tracemalloc.DomainFilter* 属性), 1520
- inclusive* (*tracemalloc.Filter* 属性), 1521
- Incomplete*, 1022
- IncompleteRead*, 1134
- IncompleteReadError*, 892
- `increment_lineno()` (在 *ast* 模块中), 1653
- IncrementalDecoder* (*codecs* 中的类), 150
- incrementaldecoder* (*codecs.CodecInfo* 属性), 145
- IncrementalEncoder* (*codecs* 中的类), 150
- incrementalencoder* (*codecs.CodecInfo* 属性), 145
- IncrementalNewlineDecoder* (*io* 中的类), 565
- IncrementalParser* (*xml.sax.xmlreader* 中的类), 1071
- indent* (*doctest.Example* 属性), 1384
- `indent()` (在 *textwrap* 模块中), 128
- `INDENT()` (在 *token* 模块中), 1657

- IndentationError, 83
- index() (array.array 方法), 222
- index() (bytearray 方法), 53
- index() (bytes 方法), 53
- index() (collections.deque 方法), 200
- index() (sequence method), 35
- index() (str 方法), 43
- index() (tkinter.ttk.Notebook 方法), 1324
- index() (tkinter.ttk.Treeview 方法), 1330
- index() (在 operator 模块中), 335
- IndexError, 81
- indexOf() (在 operator 模块中), 336
- IndexSizeErr, 1055
- inet_aton() (在 socket 模块中), 797
- inet_ntoa() (在 socket 模块中), 798
- inet_ntop() (在 socket 模块中), 798
- inet_pton() (在 socket 模块中), 798
- Inexact (decimal 中的类), 290
- inf() (在 cmath 模块中), 272
- inf() (在 math 模块中), 268
- infile
 - json.tool 命令行选项, 995
- infile (shlex.shlex 属性), 1303
- Infinity, 10
- infj() (在 cmath 模块中), 272
- info
 - zipapp 命令行选项, 1536
- info() (dis.Bytecode 方法), 1669
- info() (gettext.NullTranslations 方法), 1250
- info() (logging.Logger 方法), 608
- info() (在 logging 模块中), 617
- infolist() (zipfile.ZipFile 方法), 445
- .ini
 - file, 468
- ini file, 468
- init() (在 mimetypes 模块中), 1015
- init_color() (在 curses 模块中), 647
- init_database() (在 msilib 模块中), 1690
- init_pair() (在 curses 模块中), 647
- inited() (在 mimetypes 模块中), 1015
- initgroups() (在 os 模块中), 513
- initial_indent (textwrap.TextWrapper 属性), 129
- initscr() (在 curses 模块中), 647
- inode() (os.DirEntry 方法), 532
- INPLACE_ADD (opcode), 1673
- INPLACE_AND (opcode), 1673
- INPLACE_FLOOR_DIVIDE (opcode), 1673
- INPLACE_LSHIFT (opcode), 1673
- INPLACE_MATRIX_MULTIPLY (opcode), 1673
- INPLACE_MODULO (opcode), 1673
- INPLACE_MULTIPLY (opcode), 1673
- INPLACE_OR (opcode), 1673
- INPLACE_POWER (opcode), 1673
- INPLACE_RSHIFT (opcode), 1673
- INPLACE_SUBTRACT (opcode), 1673
- INPLACE_TRUE_DIVIDE (opcode), 1673
- INPLACE_XOR (opcode), 1673
- input (2to3 fixer), 1475
- input() (设置函数), 12
- input() (在 fileinput 模块中), 362
- input_charset (email.charset.Charset 属性), 980
- input_codec (email.charset.Charset 属性), 980
- InputOnly (tkinter.tix 中的类), 1338
- InputSource (xml.sax.xmlreader 中的类), 1071
- insch() (curses.window 方法), 654
- insdelln() (curses.window 方法), 654
- insert() (array.array 方法), 222
- insert() (collections.deque 方法), 200
- insert() (sequence method), 37
- insert() (tkinter.ttk.Notebook 方法), 1324
- insert() (tkinter.ttk.Treeview 方法), 1330
- insert() (xml.etree.ElementTree.Element 方法), 1042
- insert_text() (在 readline 模块中), 134
- insertBefore() (xml.dom.Node 方法), 1050
- insertln() (curses.window 方法), 654
- insnstr() (curses.window 方法), 654
- insort() (在 bisect 模块中), 218
- insort_left() (在 bisect 模块中), 218
- insort_right() (在 bisect 模块中), 218
- inspect (模块), 1595
- inspect 命令行选项
 - details, 1610
- InspectLoader (importlib.abc 中的类), 1632
- insstr() (curses.window 方法), 654
- install() (gettext.NullTranslations 方法), 1250
- install() (在 gettext 模块中), 1249
- install_opener() (在 urllib.request 模块中), 1105
- install_scripts() (venv.EnvBuilder 方法), 1531
- installHandler() (在 unittest 模块中), 1418
- instate() (tkinter.ttk.Widget 方法), 1321
- instr() (curses.window 方法), 654
- instream (shlex.shlex 属性), 1303
- Instruction (dis 中的类), 1671
- Instruction.arg() (在 dis 模块中), 1671
- Instruction.argrepr() (在 dis 模块中), 1671
- Instruction.argval() (在 dis 模块中), 1671
- Instruction.is_jump_target() (在 dis 模块中), 1671
- Instruction.offset() (在 dis 模块中), 1671
- Instruction.opcode() (在 dis 模块中), 1671
- Instruction.opname() (在 dis 模块中), 1671
- Instruction.starts_line() (在 dis 模块中), 1671
- int
 - 设置函数, 29
- int (uuid.UUID 属性), 1174
- int (设置类), 13
- Int2AP() (在 imaplib 模块中), 1148

- `int_info()` (在 `sys` 模块中), 1551
- `integer`
 - `literals`, 29
 - `types, operations on`, 30
 - 对象, 29
- `Integral` (`numbers` 中的类), 262
- `Integrated Development Environment`, 1340
- `IntegrityError`, 421
- `Intel/DVI ADPCM`, 1227
- `IntEnum` (`enum` 中的类), 242
- `interact` (`pdb` command), 1499
- `interact()` (`code.InteractiveConsole` 方法), 1617
- `interact()` (`telnetlib.Telnet` 方法), 1172
- `interact()` (在 `code` 模块中), 1615
- `interactive -- 交互`, 1764
- `InteractiveConsole` (`code` 中的类), 1615
- `InteractiveInterpreter` (`code` 中的类), 1615
- `intern` (`2to3` fixer), 1475
- `intern()` (在 `sys` 模块中), 1551
- `internal_attr` (`zipfile.ZipInfo` 属性), 450
- `Internaldate2tuple()` (在 `imaplib` 模块中), 1148
- `internalSubset` (`xml.dom.DocumentType` 属性), 1051
- `Internet`, 1085
- `interpolation`
 - `bytearray (%)`, 60
 - `bytes (%)`, 60
- `interpolation, string (%)`, 48
- `InterpolationDepthError`, 484
- `InterpolationError`, 484
- `InterpolationMissingOptionError`, 485
- `InterpolationSyntaxError`, 485
- `interpreted -- 解释型`, 1764
- `interpreter prompts`, 1554
- `interpreter shutdown -- 解释器关闭`, 1764
- `interrupt()` (`sqlite3.Connection` 方法), 414
- `interrupt_main()` (在 `_thread` 模块中), 786
- `InterruptedError`, 85
- `intersection()` (`frozenset` 方法), 70
- `intersection_update()` (`frozenset` 方法), 71
- `IntFlag` (`enum` 中的类), 242
- `intro` (`cmd.Cmd` 属性), 1297
- `InuseAttributeErr`, 1055
- `inv()` (在 `operator` 模块中), 335
- `InvalidAccessErr`, 1055
- `invalidate_caches()`
 - (`importlib.abc.MetaPathFinder` 方法), 1630
- `invalidate_caches()`
 - (`importlib.abc.PathEntryFinder` 方法), 1631
- `invalidate_caches()`
 - (`importlib.machinery.FileFinder` 方法), 1637
- `invalidate_caches()`
 - (`importlib.machinery.PathFinder` 类方法), 1636
- `invalidate_caches()` (在 `importlib` 模块中), 1628
- `InvalidCharacterErr`, 1055
- `InvalidModificationErr`, 1055
- `InvalidOperation` (`decimal` 中的类), 290
- `InvalidStateErr`, 1055
- `InvalidStateError`, 871
- `InvalidURL`, 1134
- `invert()` (在 `operator` 模块中), 335
- `IO` (`typing` 中的类), 1362
- `io` (模块), 554
- `IOBase` (`io` 中的类), 556
- `ioctl()` (`socket.socket` 方法), 801
- `ioctl()` (在 `fcntl` 模块中), 1715
- `IOError`, 84
- `ior()` (在 `operator` 模块中), 340
- `io.StringIO`
 - 对象, 40
- `ip` (`ipaddress.IPv4Interface` 属性), 1224
- `ip` (`ipaddress.IPv6Interface` 属性), 1224
- `ip_address()` (在 `ipaddress` 模块中), 1214
- `ip_interface()` (在 `ipaddress` 模块中), 1215
- `ip_network()` (在 `ipaddress` 模块中), 1214
- `ipaddress` (模块), 1214
- `ipow()` (在 `operator` 模块中), 340
- `ipv4_mapped` (`ipaddress.IPv6Address` 属性), 1217
- `IPv4Address` (`ipaddress` 中的类), 1215
- `IPv4Interface` (`ipaddress` 中的类), 1224
- `IPv4Network` (`ipaddress` 中的类), 1219
- `IPv6Address` (`ipaddress` 中的类), 1216
- `IPv6Interface` (`ipaddress` 中的类), 1224
- `IPv6Network` (`ipaddress` 中的类), 1221
- `irshift()` (在 `operator` 模块中), 340
- `is`
 - 运算符, 28
- `is not`
 - 运算符, 28
- `is_()` (在 `operator` 模块中), 334
- `is_absolute()` (`pathlib.PurePath` 方法), 348
- `is_alive()` (`multiprocessing.Process` 方法), 727
- `is_alive()` (`threading.Thread` 方法), 712
- `is_assigned()` (`symtable.Symbol` 方法), 1656
- `is_attachment()` (`email.message.EmailMessage` 方法), 935
- `is_authenticated()`
 - (`url-lib.request.HTTPPasswordMgrWithPriorAuth` 方法), 1114
- (`im-`) `is_block_device()` (`pathlib.Path` 方法), 353
- (`im-`) `is_blocked()` (`http.cookiejar.DefaultCookiePolicy` 方法), 1198
- (`im-`) `is_canonical()` (`decimal.Context` 方法), 286
- (`im-`) `is_canonical()` (`decimal.Decimal` 方法), 279
- (`im-`) `is_char_device()` (`pathlib.Path` 方法), 353
- (`im-`) `IS_CHARACTER_JUNK()` (在 `difflib` 模块中), 121
- (`im-`) `is_check_supported()` (在 `lzma` 模块中), 441
- (`im-`) `is_closed()` (`asyncio.AbstractEventLoop` 方法), 853

- `is_closing()` (*asyncio.BaseTransport* 方法), 879
- `is_declared_global()` (*symtable.Symbol* 方法), 1656
- `is_dir()` (*os.DirEntry* 方法), 532
- `is_dir()` (*pathlib.Path* 方法), 353
- `is_dir()` (*zipfile.ZipInfo* 方法), 449
- `is_enabled()` (在 *faulthandler* 模块中), 1492
- `is_expired()` (*http.cookiejar.Cookie* 方法), 1200
- `is_fifo()` (*pathlib.Path* 方法), 353
- `is_file()` (*os.DirEntry* 方法), 532
- `is_file()` (*pathlib.Path* 方法), 353
- `is_finalizing()` (在 *sys* 模块中), 1551
- `is_finite()` (*decimal.Context* 方法), 286
- `is_finite()` (*decimal.Decimal* 方法), 280
- `is_free()` (*symtable.Symbol* 方法), 1656
- `is_global` (*ipaddress.IPv4Address* 属性), 1216
- `is_global` (*ipaddress.IPv6Address* 属性), 1217
- `is_global()` (*symtable.Symbol* 方法), 1656
- `is_hop_by_hop()` (在 *wsgiref.util* 模块中), 1096
- `is_imported()` (*symtable.Symbol* 方法), 1656
- `is_infinite()` (*decimal.Context* 方法), 286
- `is_infinite()` (*decimal.Decimal* 方法), 280
- `is_integer()` (*float* 方法), 32
- `is_jython()` (在 *test.support* 模块中), 1480
- `IS_LINE_JUNK()` (在 *difflib* 模块中), 121
- `is_linetouched()` (*curses.window* 方法), 654
- `is_link_local` (*ipaddress.IPv4Address* 属性), 1216
- `is_link_local` (*ipaddress.IPv4Network* 属性), 1219
- `is_link_local` (*ipaddress.IPv6Address* 属性), 1217
- `is_link_local` (*ipaddress.IPv6Network* 属性), 1222
- `is_local()` (*symtable.Symbol* 方法), 1656
- `is_loopback` (*ipaddress.IPv4Address* 属性), 1216
- `is_loopback` (*ipaddress.IPv4Network* 属性), 1219
- `is_loopback` (*ipaddress.IPv6Address* 属性), 1217
- `is_loopback` (*ipaddress.IPv6Network* 属性), 1222
- `is_multicast` (*ipaddress.IPv4Address* 属性), 1216
- `is_multicast` (*ipaddress.IPv4Network* 属性), 1219
- `is_multicast` (*ipaddress.IPv6Address* 属性), 1217
- `is_multicast` (*ipaddress.IPv6Network* 属性), 1222
- `is_multipart()` (*email.message.EmailMessage* 方法), 931
- `is_multipart()` (*email.message.Message* 方法), 968
- `is_namespace()` (*symtable.Symbol* 方法), 1656
- `is_nan()` (*decimal.Context* 方法), 286
- `is_nan()` (*decimal.Decimal* 方法), 280
- `is_nested()` (*symtable.SymbolTable* 方法), 1655
- `is_normal()` (*decimal.Context* 方法), 286
- `is_normal()` (*decimal.Decimal* 方法), 280
- `is_not()` (在 *operator* 模块中), 334
- `is_not_allowed()` (*http.cookiejar.DefaultCookiePolicy* 方法), 1198
- `is_optimized()` (*symtable.SymbolTable* 方法), 1655
- `is_package()` (*importlib.abc.InspectLoader* 方法), 1633
- `is_package()` (*importlib.abc.SourceLoader* 方法), 1635
- `is_package()` (*importlib.machinery.ExtensionFileLoader* 方法), 1638
- `is_package()` (*importlib.machinery.SourceFileLoader* 方法), 1637
- `is_package()` (*importlib.machinery.SourcelessFileLoader* 方法), 1638
- `is_package()` (*zipimport.zipimporter* 方法), 1620
- `is_parameter()` (*symtable.Symbol* 方法), 1656
- `is_private` (*ipaddress.IPv4Address* 属性), 1216
- `is_private` (*ipaddress.IPv4Network* 属性), 1219
- `is_private` (*ipaddress.IPv6Address* 属性), 1217
- `is_private` (*ipaddress.IPv6Network* 属性), 1222
- `is_python_build()` (在 *sysconfig* 模块中), 1561
- `is_qnan()` (*decimal.Context* 方法), 286
- `is_qnan()` (*decimal.Decimal* 方法), 280
- `is_referenced()` (*symtable.Symbol* 方法), 1656
- `is_reserved` (*ipaddress.IPv4Address* 属性), 1216
- `is_reserved` (*ipaddress.IPv4Network* 属性), 1219
- `is_reserved` (*ipaddress.IPv6Address* 属性), 1217
- `is_reserved` (*ipaddress.IPv6Network* 属性), 1222
- `is_reserved()` (*pathlib.PurePath* 方法), 349
- `is_resource_enabled()` (在 *test.support* 模块中), 1480
- `is_running()` (*asyncio.AbstractEventLoop* 方法), 853
- `is_set()` (*asyncio.Event* 方法), 902
- `is_set()` (*threading.Event* 方法), 717
- `is_signed()` (*decimal.Context* 方法), 286
- `is_signed()` (*decimal.Decimal* 方法), 280
- `is_site_local` (*ipaddress.IPv6Address* 属性), 1217
- `is_site_local` (*ipaddress.IPv6Network* 属性), 1222
- `is_snan()` (*decimal.Context* 方法), 287
- `is_snan()` (*decimal.Decimal* 方法), 280
- `is_socket()` (*pathlib.Path* 方法), 353
- `is_subnormal()` (*decimal.Context* 方法), 287
- `is_subnormal()` (*decimal.Decimal* 方法), 280
- `is_symlink()` (*os.DirEntry* 方法), 533
- `is_symlink()` (*pathlib.Path* 方法), 353
- `is_tarfile()` (在 *tarfile* 模块中), 452
- `is_term_resized()` (在 *curses* 模块中), 647
- `is_tracing()` (在 *tracemalloc* 模块中), 1519
- `is_tracked()` (在 *gc* 模块中), 1594
- `is_unspecified` (*ipaddress.IPv4Address* 属性), 1216
- `is_unspecified` (*ipaddress.IPv4Network* 属性), 1219
- `is_unspecified` (*ipaddress.IPv6Address* 属性), 1217
- `is_unspecified` (*ipaddress.IPv6Network* 属性), 1222
- `is_wintouched()` (*curses.window* 方法), 654
- `is_zero()` (*decimal.Context* 方法), 287
- `is_zero()` (*decimal.Decimal* 方法), 280
- `is_zipfile()` (在 *zipfile* 模块中), 444
- `isabs()` (在 *os.path* 模块中), 359
- `isabstract()` (在 *inspect* 模块中), 1598
- `IsADirectoryError`, 85

- isalnum() (bytearray 方法), 57
- isalnum() (bytes 方法), 57
- isalnum() (str 方法), 43
- isalnum() (在 *curses.ascii* 模块中), 663
- isalpha() (bytearray 方法), 57
- isalpha() (bytes 方法), 57
- isalpha() (str 方法), 43
- isalpha() (在 *curses.ascii* 模块中), 663
- isascii() (在 *curses.ascii* 模块中), 663
- isasynccgen() (在 *inspect* 模块中), 1598
- isasynccgenfunction() (在 *inspect* 模块中), 1598
- isatty() (*chunk.Chunk* 方法), 1238
- isatty() (*io.IOBBase* 方法), 557
- isatty() (在 *os* 模块中), 518
- isawaitable() (在 *inspect* 模块中), 1597
- isblank() (在 *curses.ascii* 模块中), 663
- isblk() (*tarfile.TarInfo* 方法), 457
- isbuiltin() (在 *inspect* 模块中), 1598
- ischr() (*tarfile.TarInfo* 方法), 457
- isclass() (在 *inspect* 模块中), 1597
- isclose() (在 *cmath* 模块中), 271
- isclose() (在 *math* 模块中), 265
- iscntrl() (在 *curses.ascii* 模块中), 664
- iscode() (在 *inspect* 模块中), 1598
- iscoroutine() (在 *asyncio* 模块中), 876
- iscoroutine() (在 *inspect* 模块中), 1597
- iscoroutinefunction() (在 *asyncio* 模块中), 876
- iscoroutinefunction() (在 *inspect* 模块中), 1597
- isctrl() (在 *curses.ascii* 模块中), 664
- isDaemon() (*threading.Thread* 方法), 712
- isdatadescriptor() (在 *inspect* 模块中), 1598
- isdecimal() (str 方法), 43
- isdev() (*tarfile.TarInfo* 方法), 457
- isdigit() (bytearray 方法), 57
- isdigit() (bytes 方法), 57
- isdigit() (str 方法), 43
- isdigit() (在 *curses.ascii* 模块中), 664
- isdir() (*tarfile.TarInfo* 方法), 457
- isdir() (在 *os.path* 模块中), 359
- isdisjoint() (*frozenset* 方法), 69
- isdown() (在 *turtle* 模块中), 1275
- iselement() (在 *xml.etree.ElementTree* 模块中), 1039
- isenabled() (在 *gc* 模块中), 1593
- isEnabledFor() (*logging.Logger* 方法), 607
- isendwin() (在 *curses* 模块中), 647
- ISEOF() (在 *token* 模块中), 1657
- isexpr() (*parser.ST* 方法), 1648
- isexpr() (在 *parser* 模块中), 1647
- isfifo() (*tarfile.TarInfo* 方法), 457
- isfile() (*tarfile.TarInfo* 方法), 457
- isfile() (在 *os.path* 模块中), 359
- isfinite() (在 *cmath* 模块中), 271
- isfinite() (在 *math* 模块中), 265
- isfirstline() (在 *fileinput* 模块中), 362
- isframe() (在 *inspect* 模块中), 1598
- isfunction() (在 *inspect* 模块中), 1597
- isgenerator() (在 *inspect* 模块中), 1597
- isgeneratorfunction() (在 *inspect* 模块中), 1597
- isgetsetdescriptor() (在 *inspect* 模块中), 1599
- isgraph() (在 *curses.ascii* 模块中), 664
- isidentifier() (str 方法), 43
- isinf() (在 *cmath* 模块中), 271
- isinf() (在 *math* 模块中), 265
- isinstance(2to3 fixer), 1475
- isinstance() (设置函数), 13
- iskeyword() (在 *keyword* 模块中), 1658
- isleap() (在 *calendar* 模块中), 194
- islice() (在 *itertools* 模块中), 319
- islink() (在 *os.path* 模块中), 359
- islnk() (*tarfile.TarInfo* 方法), 457
- islower() (bytearray 方法), 57
- islower() (bytes 方法), 57
- islower() (str 方法), 43
- islower() (在 *curses.ascii* 模块中), 664
- ismemberdescriptor() (在 *inspect* 模块中), 1599
- ismeta() (在 *curses.ascii* 模块中), 664
- ismethod() (在 *inspect* 模块中), 1597
- ismethoddescriptor() (在 *inspect* 模块中), 1598
- ismodule() (在 *inspect* 模块中), 1597
- ismount() (在 *os.path* 模块中), 359
- isnan() (在 *cmath* 模块中), 271
- isnan() (在 *math* 模块中), 266
- ISNONTERMINAL() (在 *token* 模块中), 1657
- isnumeric() (str 方法), 43
- isocalendar() (*datetime.date* 方法), 169
- isocalendar() (*datetime.datetime* 方法), 176
- isoformat() (*datetime.date* 方法), 169
- isoformat() (*datetime.datetime* 方法), 176
- isoformat() (*datetime.time* 方法), 180
- isolation_level (*sqlite3.Connection* 属性), 412
- isowekday() (*datetime.date* 方法), 169
- isowekday() (*datetime.datetime* 方法), 176
- isprint() (在 *curses.ascii* 模块中), 664
- isprintable() (str 方法), 43
- ispunct() (在 *curses.ascii* 模块中), 664
- isreadable() (*pprint.PrettyPrinter* 方法), 236
- isreadable() (在 *pprint* 模块中), 236
- isrecursive() (*pprint.PrettyPrinter* 方法), 237
- isrecursive() (在 *pprint* 模块中), 236
- isreg() (*tarfile.TarInfo* 方法), 457
- isReservedKey() (*http.cookies.Morsel* 方法), 1191
- isroutine() (在 *inspect* 模块中), 1598
- isSameNode() (*xml.dom.Node* 方法), 1050
- isspace() (bytearray 方法), 58
- isspace() (bytes 方法), 58
- isspace() (str 方法), 43
- isspace() (在 *curses.ascii* 模块中), 664
- isstdin() (在 *fileinput* 模块中), 363

issubclass() (设置函数), 13
 issubset() (frozenset 方法), 70
 issuite() (parser.ST 方法), 1648
 issuite() (在 parser 模块中), 1647
 issuperset() (frozenset 方法), 70
 issym() (tarfile.TarInfo 方法), 457
 ISTERMINAL() (在 token 模块中), 1657
 istitle() (bytearray 方法), 58
 istitle() (bytes 方法), 58
 istitle() (str 方法), 43
 istraceback() (在 inspect 模块中), 1598
 isub() (在 operator 模块中), 340
 isupper() (bytearray 方法), 58
 isupper() (bytes 方法), 58
 isupper() (str 方法), 44
 isupper() (在 curses.ascii 模块中), 664
 isvisible() (在 turtle 模块中), 1278
 isxdigit() (在 curses.ascii 模块中), 664
 item() (tkinter.ttk.Treeview 方法), 1330
 item() (xml.dom.NamedNodeMap 方法), 1054
 item() (xml.dom.NodeList 方法), 1051
 itemgetter() (在 operator 模块中), 337
 items() (configparser.ConfigParser 方法), 482
 items() (dict 方法), 73
 items() (email.message.EmailMessage 方法), 933
 items() (email.message.Message 方法), 970
 items() (mailbox.Mailbox 方法), 998
 items() (types.MappingProxyType 方法), 233
 items() (xml.etree.ElementTree.Element 方法), 1041
 itemsize (array.array 属性), 221
 itemsize (memoryview 属性), 68
 ItemsView (collections.abc 中的类), 212
 ItemsView (typing 中的类), 1360
 iter() (设置函数), 13
 iter() (xml.etree.ElementTree.Element 方法), 1042
 iter() (xml.etree.ElementTree.ElementTree 方法), 1043
 iter_attachments() (email.message.EmailMessage 方法), 936
 iter_child_nodes() (在 ast 模块中), 1653
 iter_fields() (在 ast 模块中), 1653
 iter_importers() (在 pkgutil 模块中), 1622
 iter_modules() (在 pkgutil 模块中), 1622
 iter_parts() (email.message.EmailMessage 方法), 936
 iter_unpack() (struct.Struct 方法), 144
 iter_unpack() (在 struct 模块中), 140
 iterable -- 可迭代对象, 1764
 Iterable (collections.abc 中的类), 211
 Iterable (typing 中的类), 1359
 iterator -- 迭代器, 1764
 Iterator (collections.abc 中的类), 212
 Iterator (typing 中的类), 1359
 iterator protocol, 34
 iterdecode() (在 codecs 模块中), 147

iterdir() (pathlib.Path 方法), 353
 iterdump() (sqlite3.Connection 方法), 416
 iterencode() (json.JSONEncoder 方法), 992
 iterencode() (在 codecs 模块中), 147
 iterfind() (xml.etree.ElementTree.Element 方法), 1042
 iterfind() (xml.etree.ElementTree.ElementTree 方法), 1043
 iteritems() (mailbox.Mailbox 方法), 998
 iterkeys() (mailbox.Mailbox 方法), 998
 itermonthdates() (calendar.Calendar 方法), 192
 itermonthdays() (calendar.Calendar 方法), 192
 itermonthdays2() (calendar.Calendar 方法), 192
 iterparse() (在 xml.etree.ElementTree 模块中), 1039
 itertext() (xml.etree.ElementTree.Element 方法), 1042
 itertools (2to3 fixer), 1475
 itertools (模块), 313
 itertools_imports (2to3 fixer), 1475
 itervalues() (mailbox.Mailbox 方法), 998
 iterweekdays() (calendar.Calendar 方法), 192
 ITIMER_PROF() (在 signal 模块中), 920
 ITIMER_REAL() (在 signal 模块中), 920
 ITIMER_VIRTUAL() (在 signal 模块中), 920
 ItimerError, 921
 itruediv() (在 operator 模块中), 340
 ixor() (在 operator 模块中), 340

J

-j N
 compileall 命令行选项, 1666
 Jansen, Jack, 1023
 java_ver() (在 platform 模块中), 668
 join() (asyncio.Queue 方法), 905
 join() (bytearray 方法), 53
 join() (bytes 方法), 53
 join() (multiprocessing.JoinableQueue 方法), 730
 join() (multiprocessing.pool.Pool 方法), 747
 join() (multiprocessing.Process 方法), 726
 join() (queue.Queue 方法), 784
 join() (str 方法), 44
 join() (threading.Thread 方法), 712
 join() (在 os.path 模块中), 359
 join_thread() (multiprocessing.Queue 方法), 730
 JoinableQueue (multiprocessing 中的类), 730
 joinpath() (pathlib.PurePath 方法), 349
 js_output() (http.cookies.BaseCookie 方法), 1190
 js_output() (http.cookies.Morsel 方法), 1191
 json (模块), 986
 JSONDecodeError, 992
 JSONDecoder (json 中的类), 990
 JSONEncoder (json 中的类), 991
 json.tool (模块), 994
 json.tool 命令行选项

-h, --help, 995
 infile, 995
 outfile, 995
 --sort-keys, 995
 jump (*pdb* command), 1498
 JUMP_ABSOLUTE (*opcode*), 1678
 JUMP_FORWARD (*opcode*), 1677
 JUMP_IF_FALSE_OR_POP (*opcode*), 1678
 JUMP_IF_TRUE_OR_POP (*opcode*), 1678
K
 kbhit() (在 *msvcrt* 模块中), 1695
 KDEDIR, 1087
 kevent() (在 *select* 模块中), 843
 key (*http.cookies.Morsel* 属性), 1191
 key function -- 键函数, 1764
 KEY_ALL_ACCESS() (在 *winreg* 模块中), 1702
 KEY_CREATE_LINK() (在 *winreg* 模块中), 1702
 KEY_CREATE_SUB_KEY() (在 *winreg* 模块中), 1702
 KEY_ENUMERATE_SUB_KEYS() (在 *winreg* 模块中), 1702
 KEY_EXECUTE() (在 *winreg* 模块中), 1702
 KEY_NOTIFY() (在 *winreg* 模块中), 1702
 KEY_QUERY_VALUE() (在 *winreg* 模块中), 1702
 KEY_READ() (在 *winreg* 模块中), 1702
 KEY_SET_VALUE() (在 *winreg* 模块中), 1702
 KEY_WOW64_32KEY() (在 *winreg* 模块中), 1702
 KEY_WOW64_64KEY() (在 *winreg* 模块中), 1702
 KEY_WRITE() (在 *winreg* 模块中), 1702
 KeyboardInterrupt, 81
 KeyError, 81
 keyname() (在 *curses* 模块中), 647
 keypad() (*curses.window* 方法), 654
 keyrefs() (*weakref.WeakKeyDictionary* 方法), 224
 keys() (*dict* 方法), 73
 keys() (*email.message.EmailMessage* 方法), 933
 keys() (*email.message.Message* 方法), 970
 keys() (*mailbox.Mailbox* 方法), 998
 keys() (*sqlite3.Row* 方法), 420
 keys() (*types.MappingProxyType* 方法), 233
 keys() (*xml.etree.ElementTree.Element* 方法), 1041
 KeysView (*collections.abc* 中的类), 212
 KeysView (*typing* 中的类), 1360
 keyword (模块), 1658
 keyword argument -- 关键字参数, 1764
 keywords (*functools.partial* 属性), 333
 kill() (*asyncio.asyncio.subprocess.Process* 方法), 898
 kill() (*asyncio.BaseSubprocessTransport* 方法), 882
 kill() (*subprocess.Popen* 方法), 774
 kill() (在 *os* 模块中), 545
 killchar() (在 *curses* 模块中), 647
 killpg() (在 *os* 模块中), 545
 kind (*inspect.Parameter* 属性), 1602
 knownfiles() (在 *mimetypes* 模块中), 1015

kqueue() (在 *select* 模块中), 842
 KqueueSelector (*selectors* 中的类), 851
 kwargs (*inspect.BoundArguments* 属性), 1603
 kwlist() (在 *keyword* 模块中), 1659

L

-l
 compileall 命令行选项, 1665
 -l <tarfile>
 tarfile 命令行选项, 458
 -l <zipfile>
 zipfile 命令行选项, 450
 L() (在 *re* 模块中), 104
 -l, --indentlevel=<num>
 pickletools 命令行选项, 1681
 -l, --listfuncs
 trace 命令行选项, 1513
 LabelEntry (*tkinter.tix* 中的类), 1336
 LabelFrame (*tkinter.tix* 中的类), 1336
 lambda, 1765
 LambdaType() (在 *types* 模块中), 231
 LANG, 1247, 1249, 1256, 1258
 LANGUAGE, 1247, 1249
 language
 C, 29
 large files, 1707
 LargeZipFile, 444
 last() (*nntplib.NNTP* 方法), 1158
 last_accepted (*multiprocessing.connection.Listener* 属性), 749
 last_traceback() (在 *sys* 模块中), 1552
 last_type() (在 *sys* 模块中), 1552
 last_value() (在 *sys* 模块中), 1552
 lastChild (*xml.dom.Node* 属性), 1050
 lastcmd (*cmd.Cmd* 属性), 1297
 lastgroup (*re.match* 属性), 111
 lastindex (*re.match* 属性), 111
 lastResort() (在 *logging* 模块中), 620
 lastrowid (*sqlite3.Cursor* 属性), 419
 layout() (*tkinter.ttk.Style* 方法), 1333
 lazycache() (在 *linecache* 模块中), 377
 LazyLoader (*importlib.util* 中的类), 1641
 LBRACE() (在 *token* 模块中), 1657
 LBYL, 1765
 LC_ALL, 1247, 1249
 LC_ALL() (在 *locale* 模块中), 1260
 LC_COLLATE() (在 *locale* 模块中), 1260
 LC_CTYPE() (在 *locale* 模块中), 1260
 LC_MESSAGES, 1247, 1249
 LC_MESSAGES() (在 *locale* 模块中), 1260
 LC_MONETARY() (在 *locale* 模块中), 1260
 LC_NUMERIC() (在 *locale* 模块中), 1260
 LC_TIME() (在 *locale* 模块中), 1260
 lchflags() (在 *os* 模块中), 527

- lchmod() (*pathlib.Path* 方法), 354
- lchmod() (在 *os* 模块中), 527
- lchown() (在 *os* 模块中), 527
- ldexp() (在 *math* 模块中), 266
- ldgettext() (在 *gettext* 模块中), 1248
- ldngettext() (在 *gettext* 模块中), 1248
- le() (在 *operator* 模块中), 334
- leapdays() (在 *calendar* 模块中), 194
- leaveok() (*curses.window* 方法), 654
- left (*filecmp.dircmp* 属性), 370
- left() (在 *turtle* 模块中), 1268
- left_list (*filecmp.dircmp* 属性), 370
- left_only (*filecmp.dircmp* 属性), 370
- LEFTSHIFT() (在 *token* 模块中), 1657
- LEFTSHIFTEQUAL() (在 *token* 模块中), 1657
- len
 - ☐置函数, 35, 71
- len() (☐置函数), 13
- length (*xml.dom.NamedNodeMap* 属性), 1054
- length (*xml.dom.NodeList* 属性), 1051
- length_hint() (在 *operator* 模块中), 336
- LESS() (在 *token* 模块中), 1657
- LESSEQUAL() (在 *token* 模块中), 1657
- lexists() (在 *os.path* 模块中), 358
- lgamma() (在 *math* 模块中), 268
- lgettext() (*gettext.GNUTranslations* 方法), 1251
- lgettext() (*gettext.NullTranslations* 方法), 1250
- lgettext() (在 *gettext* 模块中), 1248
- lib2to3 (模块), 1477
- libc_ver() (在 *platform* 模块中), 669
- library (*ssl.SSLError* 属性), 811
- library() (在 *dbm.ndbm* 模块中), 407
- LibraryLoader (*ctypes* 中的类), 696
- license (☐置变量), 26
- LifoQueue (*asyncio* 中的类), 906
- LifoQueue (*queue* 中的类), 783
- light-weight processes, 786
- limit_denominator() (*fractions.Fraction* 方法), 300
- LimitOverrunError, 892
- lin2adpcm() (在 *audioop* 模块中), 1228
- lin2alaw() (在 *audioop* 模块中), 1228
- lin2lin() (在 *audioop* 模块中), 1228
- lin2ulaw() (在 *audioop* 模块中), 1229
- line() (*msilib.Dialog* 方法), 1694
- line_buffering (*io.TextIOWrapper* 属性), 564
- line_num (*csv.csvreader* 属性), 466
- line-buffered I/O, 17
- linecache (模块), 377
- lineno (*ast.AST* 属性), 1649
- lineno (*doctest.DocTest* 属性), 1383
- lineno (*doctest.Example* 属性), 1384
- lineno (*json.JSONDecodeError* 属性), 992
- lineno (*pyclbr.Class* 属性), 1664
- lineno (*pyclbr.Function* 属性), 1664
- lineno (*re.error* 属性), 108
- lineno (*shlex.shlex* 属性), 1304
- lineno (*traceback.TracebackException* 属性), 1587
- lineno (*tracemalloc.Filter* 属性), 1521
- lineno (*tracemalloc.Frame* 属性), 1521
- lineno (*xml.parsers.expat.ExpatError* 属性), 1080
- lineno() (在 *fileinput* 模块中), 362
- LINES, 646, 650
- lines (*os.terminal_size* 属性), 523
- linesep (*email.policy.Policy* 属性), 946
- linesep() (在 *os* 模块中), 552
- lineterminator (*csv.Dialect* 属性), 465
- LineTooLong, 1135
- link() (在 *os* 模块中), 527
- linkname (*tarfile.TarInfo* 属性), 456
- linux_distribution() (在 *platform* 模块中), 669
- list
 - type, operations on, 37
 - 对象, 36, 37
- list -- 列表, 1765
- list (*pdb command*), 1498
- List (*typing* 中的类), 1360
- list (☐置类), 37
- list <tarfile>
 - tarfile 命令行选项, 458
- list comprehension -- 列表推导式, 1765
- list() (*imaplib.IMAP4* 方法), 1150
- list() (*multiprocessing.managers.SyncManager* 方法), 741
- list() (*nntplib.NNTP* 方法), 1157
- list() (*poplib.POP3* 方法), 1146
- list() (*tarfile.TarFile* 方法), 454
- LIST_APPEND (*opcode*), 1674
- list_dialects() (在 *csv* 模块中), 463
- list_folders() (*mailbox.Maildir* 方法), 1000
- list_folders() (*mailbox.MH* 方法), 1002
- listdir() (在 *os* 模块中), 527
- listen() (*asyncore.dispatcher* 方法), 915
- listen() (*socket.socket* 方法), 801
- listen() (在 *logging.config* 模块中), 622
- listen() (在 *turtle* 模块中), 1286
- Listener (*multiprocessing.connection* 中的类), 748
- listMethods() (*xmlrpc.client.ServerProxy.system* 方法), 1203
- ListNoteBook (*tkinter.tix* 中的类), 1338
- listxattr() (在 *os* 模块中), 541
- literal_eval() (在 *ast* 模块中), 1653
- literals
 - binary, 29
 - complex number, 29
 - floating point, 29
 - hexadecimal, 29
 - integer, 29

- numeric, 29
- octal, 29
- LittleEndianStructure (ctypes 中的类), 705
- ljust() (bytearray 方法), 54
- ljust() (bytes 方法), 54
- ljust() (str 方法), 44
- LK_LOCK() (在 msvcrt 模块中), 1695
- LK_NBLCK() (在 msvcrt 模块中), 1695
- LK_NBLCK() (在 msvcrt 模块中), 1695
- LK_RLCK() (在 msvcrt 模块中), 1695
- LK_UNLCK() (在 msvcrt 模块中), 1695
- ll (pdb command), 1498
- LMTP (smtplib 中的类), 1161
- ln() (decimal.Context 方法), 287
- ln() (decimal.Decimal 方法), 280
- LNAME, 644
- lngettext() (gettext.GNUTranslations 方法), 1252
- lngettext() (gettext.NullTranslations 方法), 1250
- lngettext() (在 gettext 模块中), 1248
- load() (http.cookiejar.FileCookieJar 方法), 1195
- load() (http.cookies.BaseCookie 方法), 1190
- load() (pickle.Unpickler 方法), 391
- load() (tracemalloc.Snapshot 类方法), 1522
- load() (在 json 模块中), 989
- load() (在 marshal 模块中), 404
- load() (在 pickle 模块中), 389
- load() (在 plistlib 模块中), 489
- LOAD_ATTR (opcode), 1677
- LOAD_BUILD_CLASS (opcode), 1675
- load_cert_chain() (ssl.SSLContext 方法), 826
- LOAD_CLASSDEREF (opcode), 1678
- LOAD_CLOSURE (opcode), 1678
- LOAD_CONST (opcode), 1676
- load_default_certs() (ssl.SSLContext 方法), 827
- LOAD_DEREF (opcode), 1678
- load_dh_params() (ssl.SSLContext 方法), 830
- load_extension() (sqlite3.Connection 方法), 415
- LOAD_FAST (opcode), 1678
- LOAD_GLOBAL (opcode), 1678
- load_module() (importlib.abc.FileLoader 方法), 1633
- load_module() (importlib.abc.InspectLoader 方法), 1633
- load_module() (importlib.abc.Loader 方法), 1631
- load_module() (importlib.abc.SourceLoader 方法), 1634
- load_module() (importlib.machinery.SourceFileLoader 方法), 1637
- load_module() (importlib.machinery.SourcelessFileLoader 方法), 1638
- load_module() (在 imp 模块中), 1752
- load_module() (zipimport.zipimporter 方法), 1620
- LOAD_NAME (opcode), 1676
- load_package_tests() (在 test.support 模块中), 1484
- load_verify_locations() (ssl.SSLContext 方法), 827
- loader -- 加载器, 1765
- Loader (importlib.abc 中的类), 1631
- loader (importlib.machinery.ModuleSpec 属性), 1639
- loader_state (importlib.machinery.ModuleSpec 属性), 1639
- LoadError, 1193
- LoadKey() (在 winreg 模块中), 1699
- LoadLibrary() (ctypes.LibraryLoader 方法), 696
- loads() (在 json 模块中), 990
- loads() (在 marshal 模块中), 404
- loads() (在 pickle 模块中), 390
- loads() (在 plistlib 模块中), 489
- loads() (在 xmlrpc.client 模块中), 1208
- loadTestsFromModule() (unittest.TestLoader 方法), 1410
- loadTestsFromName() (unittest.TestLoader 方法), 1410
- loadTestsFromNames() (unittest.TestLoader 方法), 1411
- loadTestsFromTestCase() (unittest.TestLoader 方法), 1410
- local (threading 中的类), 710
- localcontext() (在 decimal 模块中), 283
- locale (模块), 1255
- LOCALE() (在 re 模块中), 104
- localeconv() (在 locale 模块中), 1256
- LocaleHTMLCalendar (calendar 中的类), 193
- LocaleTextCalendar (calendar 中的类), 193
- localName (xml.dom.Attr 属性), 1054
- localName (xml.dom.Node 属性), 1050
- locals
 - unittest 命令行选项, 1394
- locals() (设置函数), 14
- localtime() (在 email.utils 模块中), 983
- localtime() (在 time 模块中), 568
- Locator (xml.sax.xmlreader 中的类), 1071
- Lock (asyncio 中的类), 901
- Lock (multiprocessing 中的类), 734
- Lock (threading 中的类), 713
- lock() (mailbox.Babyl 方法), 1004
- lock() (mailbox.Mailbox 方法), 999
- lock() (mailbox.Maildir 方法), 1001
- lock() (mailbox.mbox 方法), 1001
- lock() (mailbox.MH 方法), 1003
- lock() (mailbox.MMDf 方法), 1004
- Lock() (multiprocessing.managers.SyncManager 方法), 740
- lock_held() (在 imp 模块中), 1754
- locked() (_thread.lock 方法), 787
- locked() (asyncio.Condition 方法), 903

- `locked()` (`asyncio.Lock` 方法), 902
- `locked()` (`asyncio.Semaphore` 方法), 904
- `lockf()` (在 `fcntl` 模块中), 1716
- `lockf()` (在 `os` 模块中), 518
- `locking()` (在 `msvcrt` 模块中), 1695
- `LockType()` (在 `_thread` 模块中), 786
- `log()` (`logging.Logger` 方法), 608
- `log()` (在 `cmath` 模块中), 270
- `log()` (在 `logging` 模块中), 618
- `log()` (在 `math` 模块中), 266
- `log1p()` (在 `math` 模块中), 266
- `log2()` (在 `math` 模块中), 266
- `log10()` (`decimal.Context` 方法), 287
- `log10()` (`decimal.Decimal` 方法), 280
- `log10()` (在 `cmath` 模块中), 270
- `log10()` (在 `math` 模块中), 266
- `log_date_time_string()`
(`http.server.BaseHTTPRequestHandler` 方法), 1187
- `log_error()` (`http.server.BaseHTTPRequestHandler` 方法), 1187
- `log_exception()` (`wsgiref.handlers.BaseHandler` 方法), 1102
- `log_message()` (`http.server.BaseHTTPRequestHandler` 方法), 1187
- `log_request()` (`http.server.BaseHTTPRequestHandler` 方法), 1187
- `log_to_stderr()` (在 `multiprocessing` 模块中), 751
- `logb()` (`decimal.Context` 方法), 287
- `logb()` (`decimal.Decimal` 方法), 280
- `Logger` (`logging` 中的类), 606
- `LoggerAdapter` (`logging` 中的类), 616
- `logging`
Errors, 605
- `logging` (模块), 605
- `logging.config` (模块), 621
- `logging.handlers` (模块), 631
- `logical_and()` (`decimal.Context` 方法), 287
- `logical_and()` (`decimal.Decimal` 方法), 280
- `logical_invert()` (`decimal.Context` 方法), 287
- `logical_invert()` (`decimal.Decimal` 方法), 280
- `logical_or()` (`decimal.Context` 方法), 287
- `logical_or()` (`decimal.Decimal` 方法), 280
- `logical_xor()` (`decimal.Context` 方法), 287
- `logical_xor()` (`decimal.Decimal` 方法), 280
- `login()` (`ftplib.FTP` 方法), 1142
- `login()` (`imaplib.IMAP4` 方法), 1150
- `login()` (`nntplib.NNTP` 方法), 1156
- `login()` (`smtplib.SMTP` 方法), 1163
- `login_cram_md5()` (`imaplib.IMAP4` 方法), 1150
- `LOGNAME`, 512, 644
- `lognormvariate()` (在 `random` 模块中), 303
- `logout()` (`imaplib.IMAP4` 方法), 1150
- `LogRecord` (`logging` 中的类), 613
- `long` (2to3 fixer), 1476
- `longMessage` (`unittest.TestCase` 属性), 1407
- `longname()` (在 `curses` 模块中), 647
- `lookup()` (`symtable.SymbolTable` 方法), 1655
- `lookup()` (`tkinter.ttk.Style` 方法), 1332
- `lookup()` (在 `codecs` 模块中), 145
- `lookup()` (在 `unicodedata` 模块中), 131
- `lookup_error()` (在 `codecs` 模块中), 149
- `LookupError`, 80
- `loop()` (在 `asyncore` 模块中), 913
- `lower()` (`bytearray` 方法), 58
- `lower()` (`bytes` 方法), 58
- `lower()` (`str` 方法), 44
- `LPAR()` (在 `token` 模块中), 1657
- `lru_cache()` (在 `functools` 模块中), 327
- `lseek()` (在 `os` 模块中), 518
- `lshift()` (在 `operator` 模块中), 335
- `LSQB()` (在 `token` 模块中), 1657
- `lstat()` (`pathlib.Path` 方法), 354
- `lstat()` (在 `os` 模块中), 527
- `lstrip()` (`bytearray` 方法), 54
- `lstrip()` (`bytes` 方法), 54
- `lstrip()` (`str` 方法), 44
- `lsub()` (`imaplib.IMAP4` 方法), 1150
- `lt()` (在 `operator` 模块中), 334
- `lt()` (在 `turtle` 模块中), 1268
- `LWPCookieJar` (`http.cookiejar` 中的类), 1196
- `lzma` (模块), 438
- `LZMACompressor` (`lzma` 中的类), 439
- `LZMADecompressor` (`lzma` 中的类), 440
- `LZMAError`, 438
- `LZMAFile` (`lzma` 中的类), 438

M

- `-m` `<mainfn>`, `--main=<mainfn>`
zipapp 命令行选项, 1536
- `M()` (在 `re` 模块中), 104
- `-m`, `--memo`
pickletools 命令行选项, 1681
- `-m`, `--missing`
trace 命令行选项, 1513
- `mac_ver()` (在 `platform` 模块中), 668
- `machine()` (在 `platform` 模块中), 666
- `macpath` (模块), 385
- `macros` (`netrc.netrc` 属性), 486
- `MAGIC_NUMBER()` (在 `importlib.util` 模块中), 1639
- `MagicMock` (`unittest.mock` 中的类), 1443
- `Mailbox` (`mailbox` 中的类), 997
- `mailbox` (模块), 996
- `mailcap` (模块), 995
- `Maildir` (`mailbox` 中的类), 1000
- `MaildirMessage` (`mailbox` 中的类), 1005
- `mailfrom` (`smtpd.SMTPChannel` 属性), 1169
- `MailmanProxy` (`smtpd` 中的类), 1168

- `main()` (在 *py_compile* 模块中), 1665
- `main()` (在 *site* 模块中), 1612
- `main()` (在 *unittest* 模块中), 1414
- `main_thread()` (在 *threading* 模块中), 710
- `mainloop()` (在 *turtle* 模块中), 1287
- `maintype` (*email.headerregistry.ContentTypeHeader* 属性), 955
- `major` (*email.headerregistry.MIMEVersionHeader* 属性), 955
- `major()` (在 *os* 模块中), 529
- `make_alternative()` (*email.message.EmailMessage* 方法), 937
- `make_archive()` (在 *shutil* 模块中), 383
- `make_bad_fd()` (在 *test.support* 模块中), 1483
- `make_cookies()` (*http.cookiejar.CookieJar* 方法), 1195
- `make_file()` (*difflib.HtmlDiff* 方法), 118
- `MAKE_FUNCTION` (*opcode*), 1679
- `make_header()` (在 *email.header* 模块中), 979
- `make_mixed()` (*email.message.EmailMessage* 方法), 937
- `make_msgid()` (在 *email.utils* 模块中), 983
- `make_parser()` (在 *xml.sax* 模块中), 1064
- `make_related()` (*email.message.EmailMessage* 方法), 937
- `make_server()` (在 *wsgiref.simple_server* 模块中), 1098
- `make_table()` (*difflib.HtmlDiff* 方法), 118
- `makedev()` (在 *os* 模块中), 529
- `makedirs()` (在 *os* 模块中), 528
- `makeelement()` (*xml.etree.ElementTree.Element* 方法), 1042
- `makefile()` (*socket.socket* 方法), 801
- `makeLogRecord()` (在 *logging* 模块中), 619
- `makePickle()` (*logging.handlers.SocketHandler* 方法), 636
- `makeRecord()` (*logging.Logger* 方法), 609
- `makeSocket()` (*logging.handlers.DatagramHandler* 方法), 637
- `makeSocket()` (*logging.handlers.SocketHandler* 方法), 636
- `maketrans()` (*bytearray* 静态方法), 53
- `maketrans()` (*bytes* 静态方法), 53
- `maketrans()` (*str* 静态方法), 44
- `mangle_from_` (*email.policy.Compat32* 属性), 951
- `mangle_from_` (*email.policy.Policy* 属性), 947
- `map` (*2to3 fixer*), 1476
- `map()` (*concurrent.futures.Executor* 方法), 762
- `map()` (*multiprocessing.pool.Pool* 方法), 746
- `map()` (*tkinter.ttk.Style* 方法), 1332
- `map()` (⌘置函数), 14
- `MAP_ADD` (*opcode*), 1674
- `map_async()` (*multiprocessing.pool.Pool* 方法), 746
- `map_table_b2()` (在 *stringprep* 模块中), 133
- `map_table_b3()` (在 *stringprep* 模块中), 133
- `map_to_type()` (*email.headerregistry.HeaderRegistry* 方法), 956
- `mapLogRecord()` (*logging.handlers.HTTPHandler* 方法), 641
- `mapping`
 - types, operations on, 71
 - 对象, 71
- `mapping -- 映射`, 1765
- `Mapping` (*collections.abc* 中的类), 212
- `Mapping` (*typing* 中的类), 1359
- `mapping()` (*msilib.Control* 方法), 1693
- `MappingProxyType` (*types* 中的类), 232
- `MapView` (*collections.abc* 中的类), 212
- `MapView` (*typing* 中的类), 1360
- `mapPriority()` (*logging.handlers.SysLogHandler* 方法), 639
- `maps` (*collections.ChainMap* 属性), 195
- `maps()` (在 *nis* 模块中), 1722
- `marshal` (模块), 403
- `marshalling`
 - objects, 387
- `masking`
 - operations, 30
- `Match` (*typing* 中的类), 1362
- `match()` (*pathlib.PurePath* 方法), 349
- `match()` (*re.regex* 方法), 108
- `match()` (在 *nis* 模块中), 1722
- `match()` (在 *re* 模块中), 105
- `match_hostname()` (在 *ssl* 模块中), 815
- `math`
 - 模块, 29, 272
- `math` (模块), 264
- `matmul()` (在 *operator* 模块中), 335
- `max`
 - ⌘置函数, 35
- `max` (*datetime.date* 属性), 168
- `max` (*datetime.datetime* 属性), 172
- `max` (*datetime.time* 属性), 179
- `max` (*datetime.timedelta* 属性), 165
- `max()` (*decimal.Context* 方法), 287
- `max()` (*decimal.Decimal* 方法), 280
- `max()` (⌘置函数), 14
- `max()` (在 *audioop* 模块中), 1229
- `max_count` (*email.headerregistry.BaseHeader* 属性), 953
- `MAX_EMAX()` (在 *decimal* 模块中), 289
- `MAX_INTERPOLATION_DEPTH()` (在 *configparser* 模块中), 483
- `max_line_length` (*email.policy.Policy* 属性), 946
- `max_lines` (*textwrap.TextWrapper* 属性), 130
- `max_mag()` (*decimal.Context* 方法), 287
- `max_mag()` (*decimal.Decimal* 方法), 280
- `MAX_PREC()` (在 *decimal* 模块中), 289

- `max_prefixlen` (`ipaddress.IPv4Address` 属性), 1215
- `max_prefixlen` (`ipaddress.IPv4Network` 属性), 1219
- `max_prefixlen` (`ipaddress.IPv6Address` 属性), 1217
- `max_prefixlen` (`ipaddress.IPv6Network` 属性), 1222
- `maxarray` (`reprlib.Repr` 属性), 241
- `maxdeque` (`reprlib.Repr` 属性), 241
- `maxdict` (`reprlib.Repr` 属性), 241
- `maxDiff` (`unittest.TestCase` 属性), 1407
- `maxfrozenset` (`reprlib.Repr` 属性), 241
- `maxlen` (`collections.deque` 属性), 201
- `maxlevel` (`reprlib.Repr` 属性), 241
- `maxlist` (`reprlib.Repr` 属性), 241
- `maxlong` (`reprlib.Repr` 属性), 241
- `maxother` (`reprlib.Repr` 属性), 241
- `maxpp` () (在 `audioop` 模块中), 1229
- `maxset` (`reprlib.Repr` 属性), 241
- `maxsize` (`asyncio.Queue` 属性), 906
- `maxsize` () (在 `sys` 模块中), 1552
- `maxstring` (`reprlib.Repr` 属性), 241
- `maxtuple` (`reprlib.Repr` 属性), 241
- `maxunicode` () (在 `sys` 模块中), 1552
- `MAXYEAR` () (在 `datetime` 模块中), 164
- `MB_ICONASTERISK` () (在 `winsound` 模块中), 1706
- `MB_ICONEXCLAMATION` () (在 `winsound` 模块中), 1706
- `MB_ICONHAND` () (在 `winsound` 模块中), 1706
- `MB_ICONQUESTION` () (在 `winsound` 模块中), 1706
- `MB_OK` () (在 `winsound` 模块中), 1706
- `mbox` (`mailbox` 中的类), 1001
- `mboxMessage` (`mailbox` 中的类), 1007
- `mean` () (在 `statistics` 模块中), 307
- `median` () (在 `statistics` 模块中), 308
- `median_grouped` () (在 `statistics` 模块中), 309
- `median_high` () (在 `statistics` 模块中), 309
- `median_low` () (在 `statistics` 模块中), 308
- `MemberDescriptorType` () (在 `types` 模块中), 232
- `memmove` () (在 `ctypes` 模块中), 701
- `MemoryBIO` (`ssl` 中的类), 839
- `MemoryError`, 81
- `MemoryHandler` (`logging.handlers` 中的类), 641
- `memoryview`
 - 对象, 49
- `memoryview` (☐置类), 62
- `memset` () (在 `ctypes` 模块中), 701
- `merge` () (在 `heapq` 模块中), 215
- `Message` (`email.message` 中的类), 967
- `Message` (`mailbox` 中的类), 1005
- `message digest`, MD5, 493
- `message_factory` (`email.policy.Policy` 属性), 947
- `message_from_binary_file` () (在 `email` 模块中), 941
- `message_from_bytes` () (在 `email` 模块中), 941
- `message_from_file` () (在 `email` 模块中), 941
- `message_from_string` () (在 `email` 模块中), 941
- `MessageBeep` () (在 `winsound` 模块中), 1704
- `MessageClass` (`http.server.BaseHTTPRequestHandler` 属性), 1185
- `MessageError`, 951
- `MessageParseError`, 951
- `messages` () (在 `xml.parsers.expat.errors` 模块中), 1082
- `meta path finder` -- 元路径查找器, 1765
- `meta` () (在 `curses` 模块中), 647
- `meta_path` () (在 `sys` 模块中), 1552
- `metaclass` -- 元类, 1765
- `metaclass` (`2to3 fixer`), 1476
- `MetaPathFinder` (`importlib.abc` 中的类), 1630
- `metavar` (`optparse.Option` 属性), 1739
- `MetavarTypeHelpFormatter` (`argparse` 中的类), 578
- `Meter` (`tkinter.tix` 中的类), 1336
- `method`
 - 对象, 76
- `method` -- 方法, 1765
- `method` (`urllib.request.Request` 属性), 1110
- `method resolution order` -- 方法解析顺序, 1765
- `method_calls` (`unittest.mock.Mock` 属性), 1427
- `METHOD_CRYPT` () (在 `crypt` 模块中), 1711
- `METHOD_MD5` () (在 `crypt` 模块中), 1711
- `METHOD_SHA256` () (在 `crypt` 模块中), 1711
- `METHOD_SHA512` () (在 `crypt` 模块中), 1711
- `methodattrs` (`2to3 fixer`), 1476
- `methodcaller` () (在 `operator` 模块中), 337
- `methodHelp` () (`xmlrpc.client.ServerProxy.system` 方法), 1203
- `methods`
 - `bytearray`, 51
 - `bytes`, 51
 - `string`, 41
- `methods` (`pyclbr.Class` 属性), 1664
- `methods` () (在 `crypt` 模块中), 1711
- `methodSignature` () (`xmlrpc.client.ServerProxy.system` 方法), 1203
- `MethodType` () (在 `types` 模块中), 231
- `MH` (`mailbox` 中的类), 1002
- `MHMessage` (`mailbox` 中的类), 1008
- `microsecond` (`datetime.datetime` 属性), 173
- `microsecond` (`datetime.time` 属性), 180
- `MIME`
 - `base64 encoding`, 1017
 - `content type`, 1014
 - `headers`, 1014, 1087
 - `quoted-printable encoding`, 1022
- `MIMEApplication` (`email.mime.application` 中的类), 976
- `MIMEAudio` (`email.mime.audio` 中的类), 976
- `MIMEBase` (`email.mime.base` 中的类), 975
- `MIMEImage` (`email.mime.image` 中的类), 976

- MIMEMessage (*email.mime.message* 中的类), 977
- MIMEMultipart (*email.mime.multipart* 中的类), 975
- MIMENonMultipart (*email.mime.nonmultipart* 中的类), 975
- MIMEPart (*email.message* 中的类), 938
- MIMEText (*email.mime.text* 中的类), 977
- MimeTypes (*mimetypes* 中的类), 1016
- mimetypes (模块), 1014
- MIMEVersionHeader (*email.headerregistry* 中的类), 955
- min
- Ⓡ置函数, 35
- min (*datetime.date* 属性), 168
- min (*datetime.datetime* 属性), 172
- min (*datetime.time* 属性), 179
- min (*datetime.timedelta* 属性), 165
- min() (*decimal.Context* 方法), 287
- min() (*decimal.Decimal* 方法), 281
- min() (Ⓡ置函数), 14
- MIN_EMIN() (在 *decimal* 模块中), 289
- MIN_ETINY() (在 *decimal* 模块中), 289
- min_mag() (*decimal.Context* 方法), 287
- min_mag() (*decimal.Decimal* 方法), 281
- MINEQUAL() (在 *token* 模块中), 1657
- minmax() (在 *audioop* 模块中), 1229
- minor (*email.headerregistry.MIMEVersionHeader* 属性), 955
- minor() (在 *os* 模块中), 529
- minus() (*decimal.Context* 方法), 287
- MINUS() (在 *token* 模块中), 1657
- minute (*datetime.datetime* 属性), 173
- minute (*datetime.time* 属性), 180
- MINYEAR() (在 *datetime* 模块中), 164
- mirrored() (在 *unicodedata* 模块中), 131
- misc_header (*cmd.Cmd* 属性), 1298
- MissingSectionHeaderError, 485
- MIXERDEV, 1242
- mkd() (*ftplib.FTP* 方法), 1143
- mkdir() (*pathlib.Path* 方法), 354
- mkdir() (在 *os* 模块中), 528
- mkdtemp() (在 *tempfile* 模块中), 373
- mkfifo() (在 *os* 模块中), 528
- mknod() (在 *os* 模块中), 529
- mksalt() (在 *crypt* 模块中), 1711
- mkstemp() (在 *tempfile* 模块中), 372
- mktemp() (在 *tempfile* 模块中), 374
- mktime() (在 *time* 模块中), 568
- mktime_tz() (在 *email.utils* 模块中), 984
- mlsd() (*ftplib.FTP* 方法), 1143
- mmap (*mmap* 中的类), 925
- mmap (模块), 924
- MMDF (*mailbox* 中的类), 1004
- MMDFMessage (*mailbox* 中的类), 1011
- Mock (*unittest.mock* 中的类), 1421
- mock_add_spec() (*unittest.mock.Mock* 方法), 1423
- mock_calls (*unittest.mock.Mock* 属性), 1427
- mock_open() (在 *unittest.mock* 模块中), 1448
- mod() (在 *operator* 模块中), 335
- mode (*io.FileIO* 属性), 560
- mode (*ossaudiodev.oss_audio_device* 属性), 1244
- mode (*tarfile.TarInfo* 属性), 456
- mode() (在 *statistics* 模块中), 309
- mode() (在 *turtle* 模块中), 1288
- modes
- file, 15
- modf() (在 *math* 模块中), 266
- modified() (*urllib.robotparser.RobotFileParser* 方法), 1130
- Modify() (*msilib.View* 方法), 1691
- modify() (*select.devpoll* 方法), 844
- modify() (*select.epoll* 方法), 845
- modify() (*selectors.BaseSelector* 方法), 850
- modify() (*select.poll* 方法), 846
- module
- search path, 377, 1552, 1610
- module -- 模块, 1765
- module (*pyclbr.Class* 属性), 1663
- module (*pyclbr.Function* 属性), 1664
- module spec -- 模块规格, 1765
- module_for_loader() (在 *importlib.util* 模块中), 1641
- module_from_spec() (在 *importlib.util* 模块中), 1640
- module_repr() (*importlib.abc.Loader* 方法), 1632
- ModuleFinder (*modulefinder* 中的类), 1624
- modulefinder (模块), 1624
- ModuleInfo (*pkgutil* 中的类), 1621
- ModuleNotFoundError, 81
- modules (*modulefinder.ModuleFinder* 属性), 1624
- modules() (在 *sys* 模块中), 1552
- ModuleSpec (*importlib.machinery* 中的类), 1638
- ModuleType (*types* 中的类), 231
- monotonic() (在 *time* 模块中), 568
- month (*datetime.date* 属性), 168
- month (*datetime.datetime* 属性), 173
- month() (在 *calendar* 模块中), 194
- month_abbr() (在 *calendar* 模块中), 194
- month_name() (在 *calendar* 模块中), 194
- monthcalendar() (在 *calendar* 模块中), 194
- monthdatescalendar() (*calendar.Calendar* 方法), 192
- monthdays2calendar() (*calendar.Calendar* 方法), 192
- monthdayscalendar() (*calendar.Calendar* 方法), 192
- monthrange() (在 *calendar* 模块中), 194
- Morsel (*http.cookies* 中的类), 1190
- most_common() (*collections.Counter* 方法), 198

`mouseinterval()` (在 *curses* 模块中), 647
`mousemask()` (在 *curses* 模块中), 647
`move()` (*curses.panel.Panel* 方法), 665
`move()` (*curses.window* 方法), 655
`move()` (*mmap.mmap* 方法), 926
`move()` (*tkinter.ttk.Treeview* 方法), 1330
`move()` (在 *shutil* 模块中), 381
`move_to_end()` (*collections.OrderedDict* 方法), 207
`MozillaCookieJar` (*http.cookiejar* 中的类), 1196
`MRO`, 1765
`mro()` (*class* 方法), 78
`msg` (*http.client.HTTPResponse* 属性), 1137
`msg` (*json.JSONDecodeError* 属性), 992
`msg` (*re.error* 属性), 107
`msg` (*traceback.TracebackException* 属性), 1588
`msg()` (*telnetlib.Telnet* 方法), 1171
`msi`, 1689
`msilib` (模块), 1689
`msvcrt` (模块), 1694
`mt_interact()` (*telnetlib.Telnet* 方法), 1172
`mtime` (*gzip.GzipFile* 属性), 434
`mtime` (*tarfile.TarInfo* 属性), 456
`mtime()` (*urllib.robotparser.RobotFileParser* 方法), 1130
`mul()` (在 *audioop* 模块中), 1229
`mul()` (在 *operator* 模块中), 335
`MultiCall` (*xmlrpc.client* 中的类), 1207
`MULTILINE()` (在 *re* 模块中), 104
`MultipartConversionError`, 952
`multiply()` (*decimal.Context* 方法), 287
`multiprocessing` (模块), 720
`multiprocessing.connection` (模块), 748
`multiprocessing.dummy` (模块), 752
`multiprocessing.Manager()` (在 *multiprocessing.sharedctypes* 模块中), 739
`multiprocessing.managers` (模块), 739
`multiprocessing.pool` (模块), 745
`multiprocessing.sharedctypes` (模块), 737
`mutable`
 sequence types, 36
`mutable` -- 可变, 1765
`MutableMapping` (*collections.abc* 中的类), 212
`MutableMapping` (*typing* 中的类), 1359
`MutableSequence` (*collections.abc* 中的类), 212
`MutableSequence` (*typing* 中的类), 1359
`MutableSet` (*collections.abc* 中的类), 212
`MutableSet` (*typing* 中的类), 1359
`mvderwin()` (*curses.window* 方法), 655
`mvwin()` (*curses.window* 方法), 655
`myrights()` (*imaplib.IMAP4* 方法), 1150

N

`-n N`, `--number=N`
 `timeit` 命令行选项, 1510
`N_TOKENS()` (在 *token* 模块中), 1657

`n_waiting` (*threading.Barrier* 属性), 719
`name` (*codecs.CodecInfo* 属性), 145
`name` (*doctest.DocTest* 属性), 1383
`name` (*email.headerregistry.BaseHeader* 属性), 953
`name` (*hashlib.hash* 属性), 495
`name` (*hmac.HMAC* 属性), 504
`name` (*http.cookiejar.Cookie* 属性), 1199
`name` (*importlib.abc.FileLoader* 属性), 1633
`name` (*importlib.machinery.ExtensionFileLoader* 属性), 1638
`name` (*importlib.machinery.ModuleSpec* 属性), 1639
`name` (*importlib.machinery.SourceFileLoader* 属性), 1637
`name` (*importlib.machinery.SourcelessFileLoader* 属性), 1637
`name` (*inspect.Parameter* 属性), 1601
`name` (*io.FileIO* 属性), 560
`name` (*multiprocessing.Process* 属性), 726
`name` (*os.DirEntry* 属性), 532
`name` (*ossaudiodev.oss_audio_device* 属性), 1244
`name` (*pyclbr.Class* 属性), 1663
`name` (*pyclbr.Function* 属性), 1664
`name` (*tarfile.TarInfo* 属性), 456
`name` (*threading.Thread* 属性), 712
`name` (*xml.dom.Attr* 属性), 1054
`name` (*xml.dom.DocumentType* 属性), 1052
`name()` (在 *os* 模块中), 510
`NAME()` (在 *token* 模块中), 1657
`name()` (在 *unicodedata* 模块中), 131
`name2codepoint()` (在 *html.entities* 模块中), 1030
`named tuple` -- 具名元组, 1765
`NamedTemporaryFile()` (在 *tempfile* 模块中), 371
`NamedTuple` (*typing* 中的类), 1362
`namedtuple()` (在 *collections* 模块中), 204
`NameError`, 81
`namelist()` (*zipfile.ZipFile* 方法), 445
`nameprep()` (在 *encodings.idna* 模块中), 161
`namer` (*logging.handlers.BaseRotatingHandler* 属性), 633
`namereplace_errors()` (在 *codecs* 模块中), 149
`namespace` -- 命名空间, 1766
`Namespace` (*argparse* 中的类), 595
`Namespace` (*multiprocessing.managers* 中的类), 741
`namespace package` -- 命名空间包, 1766
`namespace()` (*imaplib.IMAP4* 方法), 1150
`Namespace()` (*multiprocessing.managers.SyncManager* 方法), 741
`NAMESPACE_DNS()` (在 *uuid* 模块中), 1174
`NAMESPACE_OID()` (在 *uuid* 模块中), 1174
`NAMESPACE_URL()` (在 *uuid* 模块中), 1174
`NAMESPACE_X500()` (在 *uuid* 模块中), 1175
`NamespaceErr`, 1055
`namespaceURI` (*xml.dom.Node* 属性), 1050
`NaN`, 10
`nan()` (在 *cmath* 模块中), 272
`nan()` (在 *math* 模块中), 268

- `nanj()` (在 *cmath* 模块中), 272
- `NannyNag`, 1663
- `napms()` (在 *curses* 模块中), 648
- `nargs` (*optparse.Option* 属性), 1738
- `nbytes` (*memoryview* 属性), 67
- `ndiff()` (在 *difflib* 模块中), 119
- `ndim` (*memoryview* 属性), 68
- `ne` (2to3 fixer), 1476
- `ne()` (在 *operator* 模块中), 334
- `needs_input` (*bz2.BZ2Decompressor* 属性), 437
- `needs_input` (*lzma.LZMADecompressor* 属性), 441
- `neg()` (在 *operator* 模块中), 335
- nested scope -- 嵌套作用域, 1766
- `netmask` (*ipaddress.IPv4Network* 属性), 1220
- `netmask` (*ipaddress.IPv6Network* 属性), 1222
- `NetmaskValueError`, 1226
- `netrc` (*netrc* 中的类), 485
- `netrc` (模块), 485
- `NetrcParseError`, 485
- `netscape` (*http.cookiejar.CookiePolicy* 属性), 1197
- `network` (*ipaddress.IPv4Interface* 属性), 1224
- `network` (*ipaddress.IPv6Interface* 属性), 1224
- Network News Transfer Protocol, 1153
- `network_address` (*ipaddress.IPv4Network* 属性), 1219
- `network_address` (*ipaddress.IPv6Network* 属性), 1222
- `new()` (在 *hashlib* 模块中), 494
- `new()` (在 *hmac* 模块中), 503
- `new_alignment()` (*formatter.writer* 方法), 1685
- `new_child()` (*collections.ChainMap* 方法), 195
- `new_class()` (在 *types* 模块中), 230
- `new_event_loop()` (*asyncio.AbstractEventLoopPolicy* 方法), 867
- `new_event_loop()` (在 *asyncio* 模块中), 865
- `new_font()` (*formatter.writer* 方法), 1685
- `new_margin()` (*formatter.writer* 方法), 1685
- `new_module()` (在 *imp* 模块中), 1753
- `new_panel()` (在 *curses.panel* 模块中), 665
- `new_spacing()` (*formatter.writer* 方法), 1686
- `new_styles()` (*formatter.writer* 方法), 1686
- `newgroups()` (*nntplib.NNTP* 方法), 1156
- `NEWLINE()` (在 *token* 模块中), 1657
- `newlines` (*io.TextIOBase* 属性), 563
- `newnews()` (*nntplib.NNTP* 方法), 1157
- `newpad()` (在 *curses* 模块中), 648
- `new-style class` -- 新式类, 1766
- `NewType()` (在 *typing* 模块中), 1363
- `newwin()` (在 *curses* 模块中), 648
- `next` (2to3 fixer), 1476
- `next` (*pdb* command), 1497
- `next()` (*nntplib.NNTP* 方法), 1158
- `next()` (*tarfile.TarFile* 方法), 454
- `next()` (*tkinter.ttk.Treeview* 方法), 1330
- `next()` (赋值函数), 14
- `next_minus()` (*decimal.Context* 方法), 287
- `next_minus()` (*decimal.Decimal* 方法), 281
- `next_plus()` (*decimal.Context* 方法), 287
- `next_plus()` (*decimal.Decimal* 方法), 281
- `next_toward()` (*decimal.Context* 方法), 287
- `next_toward()` (*decimal.Decimal* 方法), 281
- `nextfile()` (在 *fileinput* 模块中), 363
- `nextkey()` (*dbm.gnu.gdbm* 方法), 406
- `nextSibling` (*xml.dom.Node* 属性), 1050
- `ngettext()` (*gettext.GNUTranslations* 方法), 1251
- `ngettext()` (*gettext.NullTranslations* 方法), 1250
- `ngettext()` (在 *gettext* 模块中), 1248
- `nice()` (在 *os* 模块中), 545
- `nis` (模块), 1722
- `nl()` (在 *curses* 模块中), 648
- `NL()` (在 *tokenize* 模块中), 1659
- `nl_langinfo()` (在 *locale* 模块中), 1257
- `nlargest()` (在 *heapq* 模块中), 215
- `nlst()` (*ftplib.FTP* 方法), 1143
- NNTP
 - protocol, 1153
- NNTP (*nntplib* 中的类), 1154
- `nntp_implementation` (*nntplib.NNTP* 属性), 1155
- NNTP_SSL (*nntplib* 中的类), 1154
- `nntp_version` (*nntplib.NNTP* 属性), 1155
- NNTPDataError, 1155
- NNTPError, 1155
- `nntplib` (模块), 1153
- NNTPPermanentError, 1155
- NNTPProtocolError, 1155
- NNTPReplyError, 1155
- NNTPTemporaryError, 1155
- `no_proxy`, 1107
- `no_type_check()` (在 *typing* 模块中), 1364
- `no_type_check_decorator()` (在 *typing* 模块中), 1364
- `nocbreak()` (在 *curses* 模块中), 648
- NoDataAllowedErr, 1056
- `node()` (在 *platform* 模块中), 666
- `nodelay()` (*curses.window* 方法), 655
- `nodeName` (*xml.dom.Node* 属性), 1050
- `NodeTransformer` (*ast* 中的类), 1654
- `nodeType` (*xml.dom.Node* 属性), 1049
- `nodeValue` (*xml.dom.Node* 属性), 1050
- `NodeVisitor` (*ast* 中的类), 1653
- `noecho()` (在 *curses* 模块中), 648
- NOEXPR() (在 *locale* 模块中), 1258
- NoModificationAllowedErr, 1056
- `nonblock()` (*ossaudiodev.oss_audio_device* 方法), 1242
- NonCallableMagicMock (*unittest.mock* 中的类), 1443
- NonCallableMock (*unittest.mock* 中的类), 1428

- None (*Built-in object*), 27
 - None (☐置变量), 25
 - nonl() (在 *curses* 模块中), 648
 - nonzero (*2to3 fixer*), 1476
 - noop() (*imaplib.IMAP4* 方法), 1150
 - noop() (*poplib.POP3* 方法), 1146
 - NoOptionError, 484
 - NOP (*opcode*), 1671
 - noqiflush() (在 *curses* 模块中), 648
 - noraw() (在 *curses* 模块中), 648
 - NoReturn() (在 *typing* 模块中), 1364
 - normalize() (*decimal.Context* 方法), 287
 - normalize() (*decimal.Decimal* 方法), 281
 - normalize() (在 *locale* 模块中), 1259
 - normalize() (在 *unicodedata* 模块中), 131
 - normalize() (*xml.dom.Node* 方法), 1051
 - NORMALIZE_WHITESPACE() (在 *doctest* 模块中), 1375
 - normalvariate() (在 *random* 模块中), 303
 - normcase() (在 *os.path* 模块中), 360
 - normpath() (在 *os.path* 模块中), 360
 - NoSectionError, 484
 - NoSuchMailboxError, 1012
 - not
 - 运算符, 28
 - not in
 - 运算符, 28, 35
 - not_() (在 *operator* 模块中), 334
 - NotADirectoryError, 85
 - notationDecl() (*xml.sax.handler.DTDHandler* 方法), 1069
 - NotationDeclHandler()
 - (*xml.parsers.expat.xmlparser* 方法), 1079
 - notations (*xml.dom.DocumentType* 属性), 1052
 - NotConnected, 1134
 - NoteBook (*tkinter.tix* 中的类), 1338
 - Notebook (*tkinter.ttk* 中的类), 1323
 - NotEmptyError, 1012
 - NOTEQUAL() (在 *token* 模块中), 1657
 - NotFoundErr, 1055
 - notify() (*asyncio.Condition* 方法), 903
 - notify() (*threading.Condition* 方法), 715
 - notify_all() (*asyncio.Condition* 方法), 903
 - notify_all() (*threading.Condition* 方法), 716
 - notimeout() (*curses.window* 方法), 655
 - NotImplemented (☐置变量), 25
 - NotImplementedError, 81
 - NotStandaloneHandler()
 - (*xml.parsers.expat.xmlparser* 方法), 1079
 - NotSupportedError, 1056
 - NotSupportedError, 421
 - noutrefresh() (*curses.window* 方法), 655
 - now() (*datetime.datetime* 类方法), 171
 - NSIG() (在 *signal* 模块中), 920
 - nsmallest() (在 *heapq* 模块中), 215
 - NT_OFFSET() (在 *token* 模块中), 1657
 - NTEventLogHandler (*logging.handlers* 中的类), 639
 - ntohl() (在 *socket* 模块中), 797
 - ntohs() (在 *socket* 模块中), 797
 - ntransfercmd() (*ftplib.FTP* 方法), 1143
 - NullFormatter (*formatter* 中的类), 1685
 - NullHandler (*logging* 中的类), 632
 - NullImporter (*imp* 中的类), 1755
 - NullTranslations (*gettext* 中的类), 1250
 - NullWriter (*formatter* 中的类), 1686
 - num_addresses (*ipaddress.IPv4Network* 属性), 1220
 - num_addresses (*ipaddress.IPv6Network* 属性), 1222
 - Number (*numbers* 中的类), 261
 - NUMBER() (在 *token* 模块中), 1657
 - number_class() (*decimal.Context* 方法), 287
 - number_class() (*decimal.Decimal* 方法), 281
 - numbers (模块), 261
 - numerator (*fractions.Fraction* 属性), 299
 - numerator (*numbers.Rational* 属性), 262
 - numeric
 - conversions, 29
 - literals, 29
 - object, 28
 - types, operations on, 29
 - 对象, 29
 - numeric() (在 *unicodedata* 模块中), 131
 - Numerical Python, 20
 - numinput() (在 *turtle* 模块中), 1288
 - numliterals (*2to3 fixer*), 1476
- ## O
- o <output>, --output=<output>
 - zipapp 命令行选项, 1536
 - o, --output=<file>
 - pickletools 命令行选项, 1681
 - O_APPEND() (在 *os* 模块中), 519
 - O_ASYNC() (在 *os* 模块中), 519
 - O_BINARY() (在 *os* 模块中), 519
 - O_CLOEXEC() (在 *os* 模块中), 519
 - O_CREAT() (在 *os* 模块中), 519
 - O_DIRECT() (在 *os* 模块中), 519
 - O_DIRECTORY() (在 *os* 模块中), 519
 - O_DSYNC() (在 *os* 模块中), 519
 - O_EXCL() (在 *os* 模块中), 519
 - O_EXLOCK() (在 *os* 模块中), 519
 - O_NDELAY() (在 *os* 模块中), 519
 - O_NOATIME() (在 *os* 模块中), 519
 - O_NOCTTY() (在 *os* 模块中), 519
 - O_NOFOLLOW() (在 *os* 模块中), 519
 - O_NOINHERIT() (在 *os* 模块中), 519
 - O_NONBLOCK() (在 *os* 模块中), 519
 - O_PATH() (在 *os* 模块中), 519
 - O_RANDOM() (在 *os* 模块中), 519

- `O_RDONLY()` (在 `os` 模块中), 519
- `O_RDWR()` (在 `os` 模块中), 519
- `O_RSYNC()` (在 `os` 模块中), 519
- `O_SEQUENTIAL()` (在 `os` 模块中), 519
- `O_SHLOCK()` (在 `os` 模块中), 519
- `O_SHORT_LIVED()` (在 `os` 模块中), 519
- `O_SYNC()` (在 `os` 模块中), 519
- `O_TEMPORARY()` (在 `os` 模块中), 519
- `O_TEXT()` (在 `os` 模块中), 519
- `O_TMPFILE()` (在 `os` 模块中), 519
- `O_TRUNC()` (在 `os` 模块中), 519
- `O_WRONLY()` (在 `os` 模块中), 519
- `obj` (`memoryview` 属性), 67
- `object`
 - `code`, 76, 403
 - `numeric`, 28
- `object` -- 对象, 1766
- `object` (`UnicodeError` 属性), 84
- `object` (☐置类), 14
- `objects`
 - `comparing`, 28
 - `flattening`, 387
 - `marshalling`, 387
 - `persistent`, 387
 - `pickling`, 387
 - `serializing`, 387
- `obufcount()` (`ossaudiodev.oss_audio_device` 方法), 1244
- `obuffree()` (`ossaudiodev.oss_audio_device` 方法), 1244
- `oct()` (☐置函数), 14
- `octal`
 - `literals`, 29
- `octdigits()` (在 `string` 模块中), 90
- `offset` (`traceback.TracebackException` 属性), 1588
- `offset` (`xml.parsers.expat.ExpatError` 属性), 1080
- `OK()` (在 `curses` 模块中), 657
- `OleDLL` (`ctypes` 中的类), 695
- `onclick()` (在 `turtle` 模块中), 1281, 1286
- `ondrag()` (在 `turtle` 模块中), 1282
- `onecmd()` (`cmd.Cmd` 方法), 1297
- `onkey()` (在 `turtle` 模块中), 1286
- `onkeypress()` (在 `turtle` 模块中), 1286
- `onkeyrelease()` (在 `turtle` 模块中), 1286
- `onrelease()` (在 `turtle` 模块中), 1281
- `onscreenclick()` (在 `turtle` 模块中), 1286
- `ontimer()` (在 `turtle` 模块中), 1287
- `OP()` (在 `token` 模块中), 1657
- `OP_ALL()` (在 `ssl` 模块中), 819
- `OP_CIPHER_SERVER_PREFERENCE()` (在 `ssl` 模块中), 819
- `OP_ENABLE_MIDDLEBOX_COMPAT()` (在 `ssl` 模块中), 820
- `OP_NO_COMPRESSION()` (在 `ssl` 模块中), 820
- `OP_NO_SSLv2()` (在 `ssl` 模块中), 819
- `OP_NO_SSLv3()` (在 `ssl` 模块中), 819
- `OP_NO_TICKET()` (在 `ssl` 模块中), 820
- `OP_NO_TLSv1()` (在 `ssl` 模块中), 819
- `OP_NO_TLSv1_1()` (在 `ssl` 模块中), 819
- `OP_NO_TLSv1_2()` (在 `ssl` 模块中), 819
- `OP_NO_TLSv1_3()` (在 `ssl` 模块中), 819
- `OP_SINGLE_DH_USE()` (在 `ssl` 模块中), 820
- `OP_SINGLE_ECDH_USE()` (在 `ssl` 模块中), 820
- `open()` (`imaplib.IMAP4` 方法), 1150
- `open()` (`pathlib.Path` 方法), 354
- `open()` (`pipes.Template` 方法), 1718
- `open()` (`tarfile.TarFile` 类方法), 454
- `open()` (`telnetlib.Telnet` 方法), 1171
- `open()` (`urllib.request.OpenerDirector` 方法), 1111
- `open()` (`urllib.request.URLOpener` 方法), 1120
- `open()` (☐置函数), 15
- `open()` (在 `aifc` 模块中), 1230
- `open()` (在 `bz2` 模块中), 435
- `open()` (在 `codecs` 模块中), 146
- `open()` (在 `dbm` 模块中), 405
- `open()` (在 `dbm.dumb` 模块中), 408
- `open()` (在 `dbm.gnu` 模块中), 406
- `open()` (在 `dbm.ndbm` 模块中), 407
- `open()` (在 `gzip` 模块中), 433
- `open()` (在 `io` 模块中), 555
- `open()` (在 `lzma` 模块中), 438
- `open()` (在 `os` 模块中), 518
- `open()` (在 `ossaudiodev` 模块中), 1241
- `open()` (在 `shelve` 模块中), 401
- `open()` (在 `sunau` 模块中), 1233
- `open()` (在 `tarfile` 模块中), 451
- `open()` (在 `tokenize` 模块中), 1660
- `open()` (在 `wave` 模块中), 1235
- `open()` (在 `webbrowser` 模块中), 1086
- `open()` (`webbrowser.controller` 方法), 1087
- `open()` (`zipfile.ZipFile` 方法), 445
- `open_connection()` (在 `asyncio` 模块中), 889
- `open_new()` (在 `webbrowser` 模块中), 1086
- `open_new()` (`webbrowser.controller` 方法), 1087
- `open_new_tab()` (在 `webbrowser` 模块中), 1086
- `open_new_tab()` (`webbrowser.controller` 方法), 1087
- `open_osfhandle()` (在 `msvcrt` 模块中), 1695
- `open_unix_connection()` (在 `asyncio` 模块中), 889
- `open_unknown()` (`urllib.request.URLOpener` 方法), 1120
- `OpenDatabase()` (在 `msilib` 模块中), 1689
- `OpenerDirector` (`urllib.request` 中的类), 1107
- `openfp()` (在 `sunau` 模块中), 1233
- `openfp()` (在 `wave` 模块中), 1236
- `OpenKey()` (在 `winreg` 模块中), 1699
- `OpenKeyEx()` (在 `winreg` 模块中), 1699
- `openlog()` (在 `syslog` 模块中), 1723

- `openmixer()` (在 *ossaudiodev* 模块中), 1242
 - `openpty()` (在 *os* 模块中), 519
 - `openpty()` (在 *pty* 模块中), 1714
 - `OpenSSL`
 - (use in module *hashlib*), 494
 - (use in module *ssl*), 810
 - `OPENSSL_VERSION()` (在 *ssl* 模块中), 821
 - `OPENSSL_VERSION_INFO()` (在 *ssl* 模块中), 821
 - `OPENSSL_VERSION_NUMBER()` (在 *ssl* 模块中), 821
 - `OpenView()` (*msilib.Database* 方法), 1690
 - `operation`
 - concatenation, 35
 - repetition, 35
 - slice, 35
 - subscript, 35
 - `OperationalError`, 421
 - `operations`
 - bitwise, 30
 - Boolean, 27, 28
 - masking, 30
 - shifting, 30
 - `operations on`
 - dictionary type, 71
 - integer types, 30
 - list type, 37
 - mapping types, 71
 - numeric types, 29
 - sequence types, 35, 37
 - `operator`
 - (*minus*), 29
 - + (*plus*), 29
 - comparison, 28
 - `operator (2to3 fixer)`, 1476
 - `operator` (模块), 333
 - `opmap()` (在 *dis* 模块中), 1680
 - `opname()` (在 *dis* 模块中), 1680
 - `optimize()` (在 *pickletools* 模块中), 1682
 - `OPTIMIZED_BYTECODE_SUFFIXES()` (在 *importlib.machinery* 模块中), 1635
 - `Optional()` (在 *typing* 模块中), 1365
 - `OptionGroup` (*optparse* 中的类), 1732
 - `OptionMenu` (*tkinter.tix* 中的类), 1336
 - `OptionParser` (*optparse* 中的类), 1735
 - `options` (*doctest.Example* 属性), 1384
 - `Options` (*ssl* 中的类), 820
 - `options` (*ssl.SSLContext* 属性), 831
 - `options()` (*configparser.ConfigParser* 方法), 481
 - `optionxform()` (*configparser.ConfigParser* 方法), 477, 483
 - `optparse` (模块), 1725
 - `or`
 - 运算符, 27, 28
 - `or_()` (在 *operator* 模块中), 335
 - `ord()` (内置函数), 17
 - `ordered_attributes` (*xml.parsers.expat.xmlparser* 属性), 1077
 - `OrderedDict` (*collections* 中的类), 207
 - `origin` (*importlib.machinery.ModuleSpec* 属性), 1639
 - `origin_req_host` (*urllib.request.Request* 属性), 1109
 - `origin_server` (*wsgiref.handlers.BaseHandler* 属性), 1103
 - `os`
 - 模块, 1707
 - `os` (模块), 509
 - `os_environ` (*wsgiref.handlers.BaseHandler* 属性), 1101
 - `OSError`, 81
 - `os.path` (模块), 357
 - `ossaudiodev` (模块), 1241
 - `OSSAudioError`, 1241
 - `outfile`
 - json.tool* 命令行选项, 995
 - `output` (*subprocess.CalledProcessError* 属性), 769
 - `output` (*subprocess.TimeoutExpired* 属性), 769
 - `output` (*unittest.TestCase* 属性), 1404
 - `output()` (*http.cookies.BaseCookie* 方法), 1190
 - `output()` (*http.cookies.Morsel* 方法), 1191
 - `output_charset` (*email.charset.Charset* 属性), 980
 - `output_charset()` (*gettext.NullTranslations* 方法), 1250
 - `output_codec` (*email.charset.Charset* 属性), 980
 - `output_difference()` (*doctest.OutputChecker* 方法), 1387
 - `OutputChecker` (*doctest* 中的类), 1387
 - `OutputString()` (*http.cookies.Morsel* 方法), 1191
 - `over()` (*nntplib.NNTP* 方法), 1157
 - `Overflow` (*decimal* 中的类), 290
 - `OverflowError`, 82
 - `overlaps()` (*ipaddress.IPv4Network* 方法), 1220
 - `overlaps()` (*ipaddress.IPv6Network* 方法), 1222
 - `overlay()` (*curses.window* 方法), 655
 - `overload()` (在 *typing* 模块中), 1363
 - `overwrite()` (*curses.window* 方法), 655
 - `owner()` (*pathlib.Path* 方法), 354
- ## P
- `p` (*pdb command*), 1498
 - `-p <interpreter>`,
 - python=<interpreter>
 - `zipapp` 命令行选项, 1536
 - `-p, --pattern pattern`
 - unittest-discover* 命令行选项, 1394
 - `-p, --preamble=<preamble>`
 - pickletools* 命令行选项, 1681
 - `-p, --process`
 - timeit* 命令行选项, 1510
 - `P_ALL()` (在 *os* 模块中), 548
 - `P_DETACH()` (在 *os* 模块中), 546

- P_NOWAIT() (在 *os* 模块中), 546
- P_NOWAITO() (在 *os* 模块中), 546
- P_OVERLAY() (在 *os* 模块中), 546
- P_PGID() (在 *os* 模块中), 548
- P_PID() (在 *os* 模块中), 548
- P_WAIT() (在 *os* 模块中), 546
- pack() (*mailbox.MH* 方法), 1002
- pack() (*struct.Struct* 方法), 144
- pack() (在 *struct* 模块中), 140
- pack_array() (*xdrlib.Packer* 方法), 487
- pack_bytes() (*xdrlib.Packer* 方法), 487
- pack_double() (*xdrlib.Packer* 方法), 486
- pack_farray() (*xdrlib.Packer* 方法), 487
- pack_float() (*xdrlib.Packer* 方法), 486
- pack_fopaque() (*xdrlib.Packer* 方法), 487
- pack_fstring() (*xdrlib.Packer* 方法), 487
- pack_into() (*struct.Struct* 方法), 144
- pack_into() (在 *struct* 模块中), 140
- pack_list() (*xdrlib.Packer* 方法), 487
- pack_opaque() (*xdrlib.Packer* 方法), 487
- pack_string() (*xdrlib.Packer* 方法), 487
- package, 1610
- package -- 包, 1766
- packed (*ipaddress.IPv4Address* 属性), 1216
- packed (*ipaddress.IPv6Address* 属性), 1217
- Packer (*xdrlib* 中的类), 486
- packing
 - binary data, 139
- packing (*widgets*), 1313
- PAGER, 1367
- pair_content() (在 *curses* 模块中), 648
- pair_number() (在 *curses* 模块中), 648
- PanedWindow (*tkinter.tix* 中的类), 1338
- parameter -- 形参, 1766
- Parameter (*inspect* 中的类), 1601
- ParameterizedMIMEHeader (*email.headerregistry* 中的类), 955
- parameters (*inspect.Signature* 属性), 1600
- params (*email.headerregistry.ParameterizedMIMEHeader* 属性), 955
- pardir() (在 *os* 模块中), 552
- paren (2to3 fixer), 1476
- parent (*importlib.machinery.ModuleSpec* 属性), 1639
- parent (*urllib.request.BaseHandler* 属性), 1112
- parent() (*tkinter.ttk.Treeview* 方法), 1330
- parentNode (*xml.dom.Node* 属性), 1049
- parents (*collections.ChainMap* 属性), 196
- paretovariate() (在 *random* 模块中), 303
- parse() (*doctest.DocTestParser* 方法), 1385
- parse() (*email.parser.BytesParser* 方法), 940
- parse() (*email.parser.Parser* 方法), 940
- parse() (*string.Formatter* 方法), 90
- parse() (*urllib.robotparser.RobotFileParser* 方法), 1130
- parse() (在 *ast* 模块中), 1653
- parse() (在 *cgi* 模块中), 1091
- parse() (在 *xml.dom.minidom* 模块中), 1058
- parse() (在 *xml.dom.pulldom* 模块中), 1063
- parse() (在 *xml.etree.ElementTree* 模块中), 1039
- parse() (在 *xml.sax* 模块中), 1064
- parse() (*xml.etree.ElementTree.ElementTree* 方法), 1043
- Parse() (*xml.parsers.expat.xmlparser* 方法), 1076
- parse() (*xml.sax.xmlreader.XMLReader* 方法), 1072
- parse_and_bind() (在 *readline* 模块中), 134
- parse_args() (*argparse.ArgumentParser* 方法), 592
- PARSE_COLNAMES() (在 *sqlite3* 模块中), 410
- parse_config_h() (在 *sysconfig* 模块中), 1561
- PARSE_DECLTYPES() (在 *sqlite3* 模块中), 410
- parse_header() (在 *cgi* 模块中), 1091
- parse_known_args() (*argparse.ArgumentParser* 方法), 601
- parse_multipart() (在 *cgi* 模块中), 1091
- parse_qs() (在 *cgi* 模块中), 1091
- parse_qs() (在 *urllib.parse* 模块中), 1124
- parse_qsl() (在 *cgi* 模块中), 1091
- parse_qsl() (在 *urllib.parse* 模块中), 1124
- parseaddr() (在 *email.utils* 模块中), 983
- parsebytes() (*email.parser.BytesParser* 方法), 940
- parsedate() (在 *email.utils* 模块中), 984
- parsedate_to_datetime() (在 *email.utils* 模块中), 984
- parsedate_tz() (在 *email.utils* 模块中), 984
- ParseError (*xml.etree.ElementTree* 中的类), 1047
- ParseFile() (*xml.parsers.expat.xmlparser* 方法), 1076
- ParseFlags() (在 *imaplib* 模块中), 1148
- Parser (*email.parser* 中的类), 940
- parser (模块), 1645
- ParserCreate() (在 *xml.parsers.expat* 模块中), 1075
- ParserError, 1648
- ParseResult (*urllib.parse* 中的类), 1127
- ParseResultBytes (*urllib.parse* 中的类), 1127
- parsestr() (*email.parser.Parser* 方法), 940
- parseString() (在 *xml.dom.minidom* 模块中), 1058
- parseString() (在 *xml.dom.pulldom* 模块中), 1063
- parseString() (在 *xml.sax* 模块中), 1064
- parsing
 - Python source code, 1645
 - URL, 1122
- ParsingError, 485
- partial (*asyncio.IncompleteReadError* 属性), 892
- partial() (*imaplib.IMAP4* 方法), 1151
- partial() (在 *functools* 模块中), 329
- partialmethod (*functools* 中的类), 329
- parties (*threading.Barrier* 属性), 719
- partition() (*bytearray* 方法), 53
- partition() (*bytes* 方法), 53
- partition() (*str* 方法), 44
- pass_() (*poplib.POP3* 方法), 1146

- Paste, 1343
- patch() (在 *unittest.mock* 模块中), 1433
- patch.dict() (在 *unittest.mock* 模块中), 1436
- patch.multiple() (在 *unittest.mock* 模块中), 1437
- patch.object() (在 *unittest.mock* 模块中), 1436
- patch.stopall() (在 *unittest.mock* 模块中), 1439
- PATH, 542, 546, 552, 1085, 1092, 1094
- path
 - configuration file, 1610
 - module search, 377, 1552, 1610
 - operations, 341, 357
- path (*http.cookiejar.Cookie* 属性), 1199
- path (*http.server.BaseHTTPRequestHandler* 属性), 1185
- path (*importlib.abc.FileLoader* 属性), 1633
- path (*importlib.machinery.ExtensionFileLoader* 属性), 1638
- path (*importlib.machinery.FileFinder* 属性), 1637
- path (*importlib.machinery.SourceFileLoader* 属性), 1637
- path (*importlib.machinery.SourcelessFileLoader* 属性), 1638
- path (*os.DirEntry* 属性), 532
- Path (*pathlib* 中的类), 350
- path based finder -- 基于路径的查找器, 1767
- Path browser, 1341
- path entry -- 路径入口, 1766
- path entry finder -- 路径入口查找器, 1766
- path entry hook -- 路径入口钩子, 1767
- path() (在 *sys* 模块中), 1552
- path_hook() (*importlib.machinery.FileFinder* 类方法), 1637
- path_hooks() (在 *sys* 模块中), 1553
- path_importer_cache() (在 *sys* 模块中), 1553
- path_mtime() (*importlib.abc.SourceLoader* 方法), 1634
- path_return_ok() (*http.cookiejar.CookiePolicy* 方法), 1197
- path_stats() (*importlib.abc.SourceLoader* 方法), 1634
- path_stats() (*importlib.machinery.SourceFileLoader* 方法), 1637
- pathconf() (在 *os* 模块中), 529
- pathconf_names() (在 *os* 模块中), 529
- PathEntryFinder (*importlib.abc* 中的类), 1630
- PathFinder (*importlib.machinery* 中的类), 1636
- pathlib (模块), 341
- PathLike (*os* 中的类), 511
- path-like object -- 路径类对象, 1767
- pathname2url() (在 *urllib.request* 模块中), 1106
- pathsep() (在 *os* 模块中), 552
- pattern (*re.error* 属性), 107
- pattern (*re.regex* 属性), 109
- Pattern (*typing* 中的类), 1362
- pause() (在 *signal* 模块中), 921
- pause_reading() (*asyncio.ReadTransport* 方法), 880
- pause_writing() (*asyncio.BaseProtocol* 方法), 884
- PAX_FORMAT() (在 *tarfile* 模块中), 453
- pax_headers (*tarfile.TarFile* 属性), 456
- pax_headers (*tarfile.TarInfo* 属性), 457
- pbkdf2_hmac() (在 *hashlib* 模块中), 496
- pd() (在 *turtle* 模块中), 1274
- Pdb (class in *pdb*), 1494
- Pdb (*pdb* 中的类), 1495
- pdb (模块), 1494
- .pdbrc
 - file, 1496
- peek() (*bz2.BZ2File* 方法), 436
- peek() (*gzip.GzipFile* 方法), 433
- peek() (*io.BufferedReader* 方法), 561
- peek() (*lzma.LZMAFile* 方法), 439
- peek() (*weakref.finalize* 方法), 226
- peer (*smtpd.SMTPChannel* 属性), 1169
- PEM_cert_to_DER_cert() (在 *ssl* 模块中), 816
- pen() (在 *turtle* 模块中), 1274
- pencolor() (在 *turtle* 模块中), 1275
- pending (*ssl.MemoryBIO* 属性), 839
- pending() (*ssl.SSLSocket* 方法), 825
- PendingDeprecationWarning, 86
- pendown() (在 *turtle* 模块中), 1274
- pensize() (在 *turtle* 模块中), 1274
- penup() (在 *turtle* 模块中), 1274
- PEP, 1767
- PERCENT() (在 *token* 模块中), 1657
- PERCENTEQUAL() (在 *token* 模块中), 1657
- perf_counter() (在 *time* 模块中), 568
- Performance, 1507
- PermissionError, 85
- permutations() (在 *itertools* 模块中), 320
- Persist() (*msilib.SummaryInformation* 方法), 1691
- persistence, 387
- persistent
 - objects, 387
- persistent_id (*pickle protocol*), 394
- persistent_id() (*pickle.Pickler* 方法), 390
- persistent_load (*pickle protocol*), 394
- persistent_load() (*pickle.Unpickler* 方法), 391
- PF_CAN() (在 *socket* 模块中), 793
- PF_PACKET() (在 *socket* 模块中), 794
- PF_RDS() (在 *socket* 模块中), 794
- pformat() (*pprint.PrettyPrinter* 方法), 236
- pformat() (在 *pprint* 模块中), 235
- phase() (在 *cmath* 模块中), 269
- pi() (在 *cmath* 模块中), 272
- pi() (在 *math* 模块中), 268
- pickle
 - 模块, 234, 400, 401, 403
- pickle (模块), 387
- pickle() (在 *copyreg* 模块中), 400

- PickleError, 390
- Pickler (*pickle* 中的类), 390
- pickletools (模块), 1681
- pickletools 命令行选项
 - a, --annotate, 1681
 - l, --indentlevel=<num>, 1681
 - m, --memo, 1681
 - o, --output=<file>, 1681
 - p, --preamble=<preamble>, 1681
- pickling
 - objects, 387
- PicklingError, 390
- pid (*asyncio.asyncio.subprocess.Process* 属性), 898
- pid (*multiprocessing.Process* 属性), 727
- pid (*subprocess.Popen* 属性), 775
- Pipe() (在 *multiprocessing* 模块中), 729
- pipe() (在 *os* 模块中), 520
- PIPE() (在 *subprocess* 模块中), 768
- pipe2() (在 *os* 模块中), 520
- PIPE_BUF() (在 *select* 模块中), 843
- pipe_connection_lost()
 - (*asyncio.SubprocessProtocol* 方法), 883
- pipe_data_received()
 - (*asyncio.SubprocessProtocol* 方法), 883
- pipes (模块), 1717
- PKG_DIRECTORY() (在 *imp* 模块中), 1755
- pkgutil (模块), 1621
- placeholder (*textwrap.TextWrapper* 属性), 130
- platform (模块), 666
- platform() (在 *platform* 模块中), 666
- platform() (在 *sys* 模块中), 1553
- PlaySound() (在 *winsound* 模块中), 1704
- plist
 - file, 489
- plistlib (模块), 489
- plock() (在 *os* 模块中), 545
- plus() (*decimal.Context* 方法), 288
- PLUS() (在 *token* 模块中), 1657
- PLUSEQUAL() (在 *token* 模块中), 1657
- pm() (在 *pdb* 模块中), 1495
- POINTER() (在 *ctypes* 模块中), 701
- pointer() (在 *ctypes* 模块中), 701
- polar() (在 *cmath* 模块中), 270
- Policy (*email.policy* 中的类), 946
- poll() (*multiprocessing.connection.Connection* 方法), 733
- poll() (*select.devpoll* 方法), 844
- poll() (*select.epoll* 方法), 845
- poll() (*select.poll* 方法), 846
- poll() (*subprocess.Popen* 方法), 773
- poll() (在 *select* 模块中), 842
- PollSelector (*selectors* 中的类), 851
- Pool (*multiprocessing.pool* 中的类), 745
- pop() (*array.array* 方法), 222
- pop() (*collections.deque* 方法), 201
- pop() (*dict* 方法), 73
- pop() (*frozenset* 方法), 71
- pop() (*mailbox.Mailbox* 方法), 999
- pop() (*sequence method*), 37
- POP3
 - protocol, 1144
- POP3 (*poplib* 中的类), 1145
- POP3_SSL (*poplib* 中的类), 1145
- pop_alignment() (*formatter.formatter* 方法), 1684
- pop_all() (*contextlib.ExitStack* 方法), 1573
- POP_BLOCK (*opcode*), 1675
- POP_EXCEPT (*opcode*), 1675
- pop_font() (*formatter.formatter* 方法), 1684
- POP_JUMP_IF_FALSE (*opcode*), 1677
- POP_JUMP_IF_TRUE (*opcode*), 1677
- pop_margin() (*formatter.formatter* 方法), 1685
- pop_source() (*shlex.shlex* 方法), 1302
- pop_style() (*formatter.formatter* 方法), 1685
- POP_TOP (*opcode*), 1671
- (*asyncio*) Popen (*subprocess* 中的类), 770
- popen() (in module *os*), 843
- (*asyncio*) popen() (在 *os* 模块中), 545
- popen() (在 *platform* 模块中), 668
- popitem() (*collections.OrderedDict* 方法), 207
- popitem() (*dict* 方法), 73
- popitem() (*mailbox.Mailbox* 方法), 999
- popleft() (*collections.deque* 方法), 201
- poplib (模块), 1144
- PopupMenu (*tkinter.tix* 中的类), 1336
- port (*http.cookiejar.Cookie* 属性), 1199
- port_specified (*http.cookiejar.Cookie* 属性), 1200
- portion -- 部分, 1767
- pos (*json.JSONDecodeError* 属性), 992
- pos (*re.error* 属性), 107
- pos (*re.match* 属性), 111
- pos() (在 *operator* 模块中), 335
- pos() (在 *turtle* 模块中), 1272
- position (*xml.etree.ElementTree.ParseError* 属性), 1047
- position() (在 *turtle* 模块中), 1272
- positional argument -- 位置参数, 1767
- POSIX
 - I/O control, 1712
 - threads, 786
- posix (模块), 1707
- POSIX_FADV_DONTNEED() (在 *os* 模块中), 520
- POSIX_FADV_NOREUSE() (在 *os* 模块中), 520
- POSIX_FADV_NORMAL() (在 *os* 模块中), 520
- POSIX_FADV_RANDOM() (在 *os* 模块中), 520
- POSIX_FADV_SEQUENTIAL() (在 *os* 模块中), 520
- POSIX_FADV_WILLNEED() (在 *os* 模块中), 520
- posix_fadvise() (在 *os* 模块中), 520
- posix_fallocate() (在 *os* 模块中), 520

- POSIXLY_CORRECT, 604
 PosixPath (*pathlib* 中的类), 351
 post() (*nntplib.NNTP* 方法), 1159
 post() (*ossaudiodev.oss_audio_device* 方法), 1243
 post_handshake_auth(*ssl.SSLContext* 属性), 831
 post_mortem() (在 *pdb* 模块中), 1495
 post_setup() (*venv.EnvBuilder* 方法), 1531
 postcmd() (*cmd.Cmd* 方法), 1297
 postloop() (*cmd.Cmd* 方法), 1297
 pow() (⌈置函数), 17
 pow() (在 *math* 模块中), 266
 pow() (在 *operator* 模块中), 335
 power() (*decimal.Context* 方法), 288
 pp (*pdb command*), 1498
 pprint (模块), 235
 pprint() (*pprint.PrettyPrinter* 方法), 236
 pprint() (在 *pprint* 模块中), 236
 prcal() (在 *calendar* 模块中), 194
 pread() (在 *os* 模块中), 520
 preamble(*email.message.EmailMessage* 属性), 938
 preamble(*email.message.Message* 属性), 974
 precmd() (*cmd.Cmd* 方法), 1297
 prefix(*xml.dom.Attr* 属性), 1054
 prefix(*xml.dom.Node* 属性), 1050
 prefix(*zipimport.zipimporter* 属性), 1620
 prefix() (在 *sys* 模块中), 1553
 PREFIXES() (在 *site* 模块中), 1611
 prefixlen(*ipaddress.IPv4Network* 属性), 1220
 prefixlen(*ipaddress.IPv6Network* 属性), 1222
 preloop() (*cmd.Cmd* 方法), 1297
 prepare() (*logging.handlers.QueueHandler* 方法), 642
 prepare() (*logging.handlers.QueueListener* 方法), 643
 prepare_class() (在 *types* 模块中), 230
 prepare_input_source() (在 *xml.sax.saxutils* 模块中), 1071
 prepend() (*pipes.Template* 方法), 1718
 PrettyPrinter (*pprint* 中的类), 235
 prev() (*tkinter.ttk.Treeview* 方法), 1330
 previousSibling(*xml.dom.Node* 属性), 1050
 print (2to3 fixer), 1476
 print() (⌈置函数), 17
 print_callees() (*pstats.Stats* 方法), 1505
 print_callers() (*pstats.Stats* 方法), 1505
 print_directory() (在 *cgi* 模块中), 1091
 print_envron() (在 *cgi* 模块中), 1091
 print_envron_usage() (在 *cgi* 模块中), 1091
 print_exc() (*timeit.Timer* 方法), 1509
 print_exc() (在 *traceback* 模块中), 1586
 print_exception() (在 *traceback* 模块中), 1585
 PRINT_EXPR(*opcode*), 1674
 print_form() (在 *cgi* 模块中), 1091
 print_help() (*argparse.ArgumentParser* 方法), 601
 print_last() (在 *traceback* 模块中), 1586
 print_stack() (*asyncio.Task* 方法), 874
 print_stack() (在 *traceback* 模块中), 1586
 print_stats() (*profile.Profile* 方法), 1503
 print_stats() (*pstats.Stats* 方法), 1505
 print_tb() (在 *traceback* 模块中), 1585
 print_usage() (*argparse.ArgumentParser* 方法), 601
 print_usage() (*optparse.OptionParser* 方法), 1744
 print_version() (*optparse.OptionParser* 方法), 1734
 printable() (在 *string* 模块中), 90
 printdir() (*zipfile.ZipFile* 方法), 446
 printf-style formatting, 48, 60
 PRIO_GRP() (在 *os* 模块中), 513
 PRIO_PROCESS() (在 *os* 模块中), 513
 PRIO_USER() (在 *os* 模块中), 513
 PriorityQueue (*asyncio* 中的类), 906
 PriorityQueue (*queue* 中的类), 783
 prlimit() (在 *resource* 模块中), 1719
 prmonth() (*calendar.TextCalendar* 方法), 193
 prmonth() (在 *calendar* 模块中), 194
 ProactorEventLoop (*asyncio* 中的类), 865
 process
 group, 512
 id, 513
 id of parent, 513
 killing, 545
 scheduling priority, 513, 514
 signalling, 545
 Process (*multiprocessing* 中的类), 726
 process() (*logging.LoggerAdapter* 方法), 616
 process_exited() (*asyncio.SubprocessProtocol* 方法), 883
 process_message() (*smtpd.SMTPServer* 方法), 1167
 process_request() (*socketserver.BaseServer* 方法), 1179
 process_time() (在 *time* 模块中), 568
 process_tokens() (在 *tabnanny* 模块中), 1663
 ProcessError, 728
 processes, light-weight, 786
 ProcessingInstruction() (在 *xml.etree.ElementTree* 模块中), 1040
 processingInstruction() (*xml.sax.handler.ContentHandler* 方法), 1069
 ProcessingInstructionHandler() (*xml.parsers.expat.xmlparser* 方法), 1079
 ProcessLookupError, 85
 processor time, 567
 processor() (在 *platform* 模块中), 667
 ProcessPoolExecutor (*concurrent.futures* 中的类), 764
 product() (在 *itertools* 模块中), 321
 Profile (*profile* 中的类), 1502
 profile (模块), 1502
 profile function, 710, 1549, 1554
 profiler, 1549, 1554

- profiling, deterministic, 1500
- ProgrammingError, 421
- Progressbar (*tkinter.ttk* 中的类), 1325
- prompt (*cmd.Cmd* 属性), 1297
- prompt_user_passwd() (*url-lib.request.FancyURLopener* 方法), 1121
- prompts, interpreter, 1554
- propagate (*logging.Logger* 属性), 606
- property (设置类), 18
- property list, 489
- property_declaration_handler() (在 *xml.sax.handler* 模块中), 1067
- property_dom_node() (在 *xml.sax.handler* 模块中), 1067
- property_lexical_handler() (在 *xml.sax.handler* 模块中), 1066
- property_xml_string() (在 *xml.sax.handler* 模块中), 1067
- PropertyMock (*unittest.mock* 中的类), 1429
- prot_c() (*ftplib.FTP_TLS* 方法), 1144
- prot_p() (*ftplib.FTP_TLS* 方法), 1144
- proto (*socket.socket* 属性), 806
- protocol
 - CGI, 1087
 - context management, 74
 - copy, 393
 - FTP, 1121, 1139
 - HTTP, 1087, 1121, 1131, 1133, 1184
 - IMAP4, 1147
 - IMAP4_SSL, 1147
 - IMAP4_stream, 1147
 - iterator, 34
 - NNTP, 1153
 - POP3, 1144
 - SMTP, 1160
 - Telnet, 1170
- Protocol (*asyncio* 中的类), 882
- protocol (*ssl.SSLContext* 属性), 831
- PROTOCOL_SSLv2() (在 *ssl* 模块中), 818
- PROTOCOL_SSLv3() (在 *ssl* 模块中), 818
- PROTOCOL_SSLv23() (在 *ssl* 模块中), 818
- PROTOCOL_TLS() (在 *ssl* 模块中), 818
- PROTOCOL_TLS_CLIENT() (在 *ssl* 模块中), 818
- PROTOCOL_TLS_SERVER() (在 *ssl* 模块中), 818
- PROTOCOL_TLSv1() (在 *ssl* 模块中), 818
- PROTOCOL_TLSv1_1() (在 *ssl* 模块中), 818
- PROTOCOL_TLSv1_2() (在 *ssl* 模块中), 819
- protocol_version (*http.server.BaseHTTPRequestHandler* 属性), 1185
- PROTOCOL_VERSION (*imaplib.IMAP4* 属性), 1153
- ProtocolError (*xmlrpc.client* 中的类), 1206
- provisional API -- 暂定 API, 1767
- provisional package -- 暂定包, 1767
- proxy() (在 *weakref* 模块中), 224
- proxyauth() (*imaplib.IMAP4* 方法), 1151
- ProxyBasicAuthHandler (*urllib.request* 中的类), 1108
- ProxyDigestAuthHandler (*urllib.request* 中的类), 1108
- ProxyHandler (*urllib.request* 中的类), 1107
- ProxyType() (在 *weakref* 模块中), 226
- ProxyTypes() (在 *weakref* 模块中), 226
- pryear() (*calendar.TextCalendar* 方法), 193
- ps1() (在 *sys* 模块中), 1554
- ps2() (在 *sys* 模块中), 1554
- pstats (模块), 1503
- pstdev() (在 *statistics* 模块中), 310
- pthread_kill() (在 *signal* 模块中), 921
- pthread_sigmask() (在 *signal* 模块中), 922
- pthreads, 786
- pty
 - 模块, 519
- pty (模块), 1714
- pu() (在 *turtle* 模块中), 1274
- publicId (*xml.dom.DocumentType* 属性), 1051
- PullDom (*xml.dom.pulldom* 中的类), 1062
- punctuation() (在 *string* 模块中), 90
- punctuation_chars (*shlex.shlex* 属性), 1304
- PurePath (*pathlib* 中的类), 343
- PurePath.anchor() (在 *pathlib* 模块中), 346
- PurePath.drive() (在 *pathlib* 模块中), 346
- PurePath.name() (在 *pathlib* 模块中), 347
- PurePath.parent() (在 *pathlib* 模块中), 347
- PurePath.parents() (在 *pathlib* 模块中), 347
- PurePath.parts() (在 *pathlib* 模块中), 346
- PurePath.root() (在 *pathlib* 模块中), 346
- PurePath.stem() (在 *pathlib* 模块中), 348
- PurePath.suffix() (在 *pathlib* 模块中), 348
- PurePath.suffixes() (在 *pathlib* 模块中), 348
- PurePosixPath (*pathlib* 中的类), 344
- PureProxy (*smtpd* 中的类), 1168
- PureWindowsPath (*pathlib* 中的类), 344
- purge() (在 *re* 模块中), 107
- Purpose.CLIENT_AUTH() (在 *ssl* 模块中), 821
- Purpose.SERVER_AUTH() (在 *ssl* 模块中), 821
- push() (*asynchat.async_chat* 方法), 918
- push() (*code.InteractiveConsole* 方法), 1617
- push() (*contextlib.ExitStack* 方法), 1572
- push_alignment() (*formatter.formatter* 方法), 1684
- push_font() (*formatter.formatter* 方法), 1684
- push_margin() (*formatter.formatter* 方法), 1684
- push_source() (*shlex.shlex* 方法), 1302
- push_style() (*formatter.formatter* 方法), 1685
- push_token() (*shlex.shlex* 方法), 1302
- push_with_producer() (*asynchat.async_chat* 方法), 918
- pushbutton() (*msilib.Dialog* 方法), 1694
- put() (*asyncio.Queue* 方法), 905

- put() (*multiprocessing.Queue* 方法), 729
 put() (*multiprocessing.SimpleQueue* 方法), 730
 put() (*queue.Queue* 方法), 784
 put_nowait() (*asyncio.Queue* 方法), 905
 put_nowait() (*multiprocessing.Queue* 方法), 729
 put_nowait() (*queue.Queue* 方法), 784
 putch() (在 *msvcrt* 模块中), 1696
 putenv() (在 *os* 模块中), 513
 putheader() (*http.client.HTTPConnection* 方法), 1137
 putp() (在 *curses* 模块中), 648
 putrequest() (*http.client.HTTPConnection* 方法), 1136
 putwch() (在 *msvcrt* 模块中), 1696
 putwin() (*curses.window* 方法), 655
 pvariance() (在 *statistics* 模块中), 310
 pwd
 模块, 358
 pwd(模块), 1708
 pwd() (*ftplib.FTP* 方法), 1143
 pwrite() (在 *os* 模块中), 520
 py_compile(模块), 1664
 PY_COMPILED() (在 *imp* 模块中), 1755
 PY_FROZEN() (在 *imp* 模块中), 1755
 py_object(*ctypes* 中的类), 705
 PY_SOURCE() (在 *imp* 模块中), 1755
 pyclbr(模块), 1663
 PyCompileError, 1664
 PyDLL(*ctypes* 中的类), 695
 pydoc(模块), 1366
 pyexpat
 模块, 1075
 PYFUNCTYPE() (在 *ctypes* 模块中), 698
 Python 3000, 1767
 Python Editor, 1340
 Python 提高建议
 PEP 1, 1767
 PEP 205, 226
 PEP 227, 1592
 PEP 235, 1627
 PEP 237, 49, 62
 PEP 238, 1592, 1762
 PEP 249, 408, 410
 PEP 255, 1592
 PEP 263, 1627, 1659, 1660
 PEP 273, 1619
 PEP 278, 1769
 PEP 282, 383, 621
 PEP 292, 97
 PEP 302, 23, 377, 1553, 1619, 1621, 1622, 1625, 1628, 1630, 1633, 1755, 1762, 1765
 PEP 307, 388
 PEP 324, 767
 PEP 328, 23, 1592, 1628
 PEP 338, 1627
 PEP 342, 212
 PEP 343, 1577, 1592, 1761
 PEP 362, 1603, 1760, 1766
 PEP 366, 1627, 1628
 PEP 370, 1612
 PEP 378, 93
 PEP 380, 852, 868
 PEP 383, 148, 790
 PEP 393, 153, 159, 1552
 PEP 397, 1530
 PEP 405, 1527
 PEP 411, 1550, 1556, 1767
 PEP 420, 1628, 1762, 1766, 1767
 PEP 421, 1551
 PEP 428, 342
 PEP 442, 1594
 PEP 443, 1763
 PEP 451, 1552, 1622, 1626, 1628, 1762
 PEP 453, 1526
 PEP 461, 62
 PEP 468, 208
 PEP 475, 17, 85, 519, 521, 522, 549, 569, 800, 802, 804, 843, 847, 850, 923, 924
 PEP 479, 1592
 PEP 483, 1351
 PEP 484, 1351, 1353, 1358, 1364, 1759, 1762, 1769
 PEP 485, 265, 271
 PEP 488, 1628, 1640, 1664
 PEP 489, 1628, 1636, 1638
 PEP 492, 213, 1550, 1556, 1609, 1760, 1761
 PEP 498, 1762
 PEP 506, 505
 PEP 515, 93
 PEP 519, 1767
 PEP 524, 553
 PEP 525, 213, 1550, 1555, 1609, 1760
 PEP 526, 1351, 1363, 1366, 1759, 1769
 PEP 529, 1548, 1556
 PEP 3101, 90
 PEP 3105, 1592
 PEP 3112, 1592
 PEP 3115, 231
 PEP 3116, 1769
 PEP 3118, 64
 PEP 3119, 214, 1579
 PEP 3120, 1628
 PEP 3141, 261, 1579
 PEP 3147, 1626, 1628, 1640, 1664, 1667, 1754
 PEP 3148, 766
 PEP 3149, 1543
 PEP 3151, 85, 792, 842, 1718
 PEP 3153, 912
 PEP 3154, 389

PEP 3155, 1767
 PEP 3156, 912
 PEP 3333, 10951099, 1102, 1103
 python_branch() (在 *platform* 模块中), 667
 python_build() (在 *platform* 模块中), 667
 python_compiler() (在 *platform* 模块中), 667
 PYTHON_DOM, 1048
 python_implementation() (在 *platform* 模块中), 667
 python_revision() (在 *platform* 模块中), 667
 python_version() (在 *platform* 模块中), 667
 python_version_tuple() (在 *platform* 模块中), 667
 PYTHONASYNCIODEBUG, 862, 907
 PYTHONDOCS, 1367
 PYTHONDONTWRITEBYTECODE, 1545
 PYTHONFAULTHANDLER, 1492
 Pythonic, 1767
 PYTHONIOENCODING, 1557
 PYTHONLEGACYWINDOWSFSENCODING, 1556
 PYTHONNOUSERSITE, 1611, 1612
 PYTHONPATH, 1092, 1552
 PYTHONSTARTUP, 136, 1346, 1551, 1611
 PYTHONTRACEMALLOC, 1515, 1520
 PYTHONUSERBASE, 1612
 PyZipFile (*zipfile* 中的类), 448

Q

-q
 compileall 命令行选项, 1665
 qiflush() (在 *curses* 模块中), 648
 QName (*xml.etree.ElementTree* 中的类), 1044
 qsize() (*asyncio.Queue* 方法), 906
 qsize() (*multiprocessing.Queue* 方法), 729
 qsize() (*queue.Queue* 方法), 784
 qualified name -- 限定名称, 1767
 quantize() (*decimal.Context* 方法), 288
 quantize() (*decimal.Decimal* 方法), 281
 QueryInfoKey() (在 *winreg* 模块中), 1699
 QueryReflectionKey() (在 *winreg* 模块中), 1701
 QueryValue() (在 *winreg* 模块中), 1699
 QueryValueEx() (在 *winreg* 模块中), 1700
 Queue (*asyncio* 中的类), 905
 Queue (*multiprocessing* 中的类), 729
 Queue (*queue* 中的类), 783
 queue (*sched.scheduler* 属性), 783
 queue (模块), 783
 Queue() (*multiprocessing.managers.SyncManager* 方法), 741
 QueueEmpty, 906
 QueueFull, 906
 QueueHandler (*logging.handlers* 中的类), 642
 QueueListener (*logging.handlers* 中的类), 642
 quick_ratio() (*difflib.SequenceMatcher* 方法), 123

quit (*pdb* command), 1499
 quit (设置变量), 26
 quit() (*ftplib.FTP* 方法), 1144
 quit() (*nntplib.NNTP* 方法), 1156
 quit() (*poplib.POP3* 方法), 1146
 quit() (*smtplib.SMTP* 方法), 1166
 quopri (模块), 1022
 quote() (在 *email.utils* 模块中), 983
 quote() (在 *shlex* 模块中), 1301
 quote() (在 *urllib.parse* 模块中), 1127
 QUOTE_ALL() (在 *csv* 模块中), 464
 quote_from_bytes() (在 *urllib.parse* 模块中), 1128
 QUOTE_MINIMAL() (在 *csv* 模块中), 464
 QUOTE_NONE() (在 *csv* 模块中), 464
 QUOTE_NONNUMERIC() (在 *csv* 模块中), 464
 quote_plus() (在 *urllib.parse* 模块中), 1128
 quoteattr() (在 *xml.sax.saxutils* 模块中), 1070
 quotechar (*csv.Dialect* 属性), 465
 quoted-printable
 encoding, 1022
 quotes (*shlex.shlex* 属性), 1303
 quoting (*csv.Dialect* 属性), 465

R

-r
 compileall 命令行选项, 1666
 -r N, --repeat=N
 timeit 命令行选项, 1510
 -R, --no-report
 trace 命令行选项, 1513
 -r, --report
 trace 命令行选项, 1513
 R_OK() (在 *os* 模块中), 524
 radians() (在 *math* 模块中), 267
 radians() (在 *turtle* 模块中), 1274
 RadioButtonGroup (*msilib* 中的类), 1693
 radiogroup() (*msilib.Dialog* 方法), 1694
 radix() (*decimal.Context* 方法), 288
 radix() (*decimal.Decimal* 方法), 282
 RADIXCHAR() (在 *locale* 模块中), 1257
 raise
 语句, 79
 raise (2to3 fixer), 1476
 raise_on_defect (*email.policy.Policy* 属性), 947
 RAISE_VARARGS (*opcode*), 1679
 RAND_add() (在 *ssl* 模块中), 814
 RAND_bytes() (在 *ssl* 模块中), 814
 RAND_egd() (在 *ssl* 模块中), 814
 RAND_pseudo_bytes() (在 *ssl* 模块中), 814
 RAND_status() (在 *ssl* 模块中), 814
 randbelow() (在 *secrets* 模块中), 505
 randbits() (在 *secrets* 模块中), 505
 randint() (在 *random* 模块中), 302
 random (模块), 300

- `random()` (在 `random` 模块中), 303
- `randrange()` (在 `random` 模块中), 302
- `range`
 - 对象, 38
- `range` (范围类), 38
- `RARROW()` (在 `token` 模块中), 1657
- `ratecv()` (在 `audioop` 模块中), 1229
- `ratio()` (`difflib.SequenceMatcher` 方法), 123
- `Rational` (`numbers` 中的类), 262
- `raw` (`io.BufferedIOBase` 属性), 559
- `raw()` (在 `curses` 模块中), 649
- `raw_data_manager()` (在 `email.contentmanager` 模块中), 958
- `raw_decode()` (`json.JSONDecoder` 方法), 991
- `raw_input` (2to3 fixer), 1476
- `raw_input()` (`code.InteractiveConsole` 方法), 1617
- `RawArray()` (在 `multiprocessing.sharedctypes` 模块中), 737
- `RawConfigParser` (`configparser` 中的类), 484
- `RawDescriptionHelpFormatter` (`argparse` 中的类), 578
- `RawIOBase` (`io` 中的类), 558
- `RawPen` (`turtle` 中的类), 1290
- `RawTextHelpFormatter` (`argparse` 中的类), 578
- `RawTurtle` (`turtle` 中的类), 1290
- `RawValue()` (在 `multiprocessing.sharedctypes` 模块中), 737
- `RBRACE()` (在 `token` 模块中), 1657
- `rcpttos` (`smtpd.SMTPChannel` 属性), 1169
- `re`
 - 模块, 41, 376
- `re` (`re.match` 属性), 111
- `re` (模块), 99
- `read()` (`asyncio.StreamReader` 方法), 890
- `read()` (`chunk.Chunk` 方法), 1239
- `read()` (`codecs.StreamReader` 方法), 152
- `read()` (`configparser.ConfigParser` 方法), 481
- `read()` (`http.client.HTTPResponse` 方法), 1137
- `read()` (`imaplib.IMAP4` 方法), 1151
- `read()` (`io.BufferedIOBase` 方法), 559
- `read()` (`io.BufferedReader` 方法), 561
- `read()` (`io.RawIOBase` 方法), 558
- `read()` (`io.TextIOBase` 方法), 563
- `read()` (`mimetypes.MimeTypes` 方法), 1016
- `read()` (`mmap.mmap` 方法), 927
- `read()` (`ossaudiodev.oss_audio_device` 方法), 1242
- `read()` (`ssl.MemoryBIO` 方法), 839
- `read()` (`ssl.SSLSocket` 方法), 822
- `read()` (`urllib.robotparser.RobotFileParser` 方法), 1130
- `read()` (在 `os` 模块中), 521
- `read()` (`zipfile.ZipFile` 方法), 446
- `read1()` (`io.BufferedIOBase` 方法), 559
- `read1()` (`io.BufferedReader` 方法), 561
- `read1()` (`io.BytesIO` 方法), 561
- `read_all()` (`telnetlib.Telnet` 方法), 1171
- `read_byte()` (`mmap.mmap` 方法), 927
- `read_bytes()` (`pathlib.Path` 方法), 354
- `read_dict()` (`configparser.ConfigParser` 方法), 482
- `read_eager()` (`telnetlib.Telnet` 方法), 1171
- `read_environ()` (在 `wsgiref.handlers` 模块中), 1103
- `read_events()` (`xml.etree.ElementTree.XMLPullParser` 方法), 1046
- `read_file()` (`configparser.ConfigParser` 方法), 482
- `read_history_file()` (在 `readline` 模块中), 134
- `read_init_file()` (在 `readline` 模块中), 134
- `read_lazy()` (`telnetlib.Telnet` 方法), 1171
- `read_mime_types()` (在 `mimetypes` 模块中), 1015
- `read_sb_data()` (`telnetlib.Telnet` 方法), 1171
- `read_some()` (`telnetlib.Telnet` 方法), 1171
- `read_string()` (`configparser.ConfigParser` 方法), 482
- `read_text()` (`pathlib.Path` 方法), 354
- `read_token()` (`shlex.shlex` 方法), 1302
- `read_until()` (`telnetlib.Telnet` 方法), 1171
- `read_very_eager()` (`telnetlib.Telnet` 方法), 1171
- `read_very_lazy()` (`telnetlib.Telnet` 方法), 1171
- `read_windows_registry()` (`mimetypes.MimeTypes` 方法), 1016
- `readable()` (`asyncore.dispatcher` 方法), 914
- `readable()` (`io.IOBase` 方法), 557
- `READABLE()` (在 `tkinter` 模块中), 1318
- `readall()` (`io.RawIOBase` 方法), 558
- `reader()` (在 `csv` 模块中), 461
- `ReadError`, 452
- `readexactly()` (`asyncio.StreamReader` 方法), 890
- `readfp()` (`configparser.ConfigParser` 方法), 483
- `readfp()` (`mimetypes.MimeTypes` 方法), 1016
- `readframes()` (`aifc.aifc` 方法), 1231
- `readframes()` (`sunau.AU_read` 方法), 1234
- `readframes()` (`wave.Wave_read` 方法), 1236
- `readinto()` (`http.client.HTTPResponse` 方法), 1137
- `readinto()` (`io.BufferedIOBase` 方法), 559
- `readinto()` (`io.RawIOBase` 方法), 558
- `readinto1()` (`io.BufferedIOBase` 方法), 559
- `readinto1()` (`io.BytesIO` 方法), 561
- `readline` (模块), 133
- `readline()` (`asyncio.StreamReader` 方法), 890
- `readline()` (`codecs.StreamReader` 方法), 152
- `readline()` (`imaplib.IMAP4` 方法), 1151
- `readline()` (`io.IOBase` 方法), 557
- `readline()` (`io.TextIOBase` 方法), 563
- `readline()` (`mmap.mmap` 方法), 927
- `readlines()` (`codecs.StreamReader` 方法), 152
- `readlines()` (`io.IOBase` 方法), 557
- `readlink()` (在 `os` 模块中), 529
- `readmodule()` (在 `pyclbr` 模块中), 1663
- `readmodule_ex()` (在 `pyclbr` 模块中), 1663
- `readonly` (`memoryview` 属性), 68
- `readPlist()` (在 `plistlib` 模块中), 490

- readPlistFromBytes() (在 *plistlib* 模块中), 490
- ReadTransport (*asyncio* 中的类), 880
- readuntil() (*asyncio.StreamReader* 方法), 890
- readv() (在 *os* 模块中), 521
- ready() (*multiprocessing.pool.AsyncResult* 方法), 747
- Real (*numbers* 中的类), 262
- real (*numbers.Complex* 属性), 261
- Real Media File Format, 1238
- real_quick_ratio() (*difflib.SequenceMatcher* 方法), 123
- realpath() (在 *os.path* 模块中), 360
- reason (*http.client.HTTPResponse* 属性), 1138
- reason (*ssl.SSLError* 属性), 811
- reason (*UnicodeError* 属性), 84
- reason (*urllib.error.HTTPError* 属性), 1130
- reason (*urllib.error.URLError* 属性), 1129
- reattach() (*tkinter.ttk.Treeview* 方法), 1330
- recontrols() (*ossaudiodev.oss_mixer_device* 方法), 1245
- received_data (*smtpd.SMTPChannel* 属性), 1169
- received_lines (*smtpd.SMTPChannel* 属性), 1169
- recent() (*imaplib.IMAP4* 方法), 1151
- records (*unittest.TestCase* 属性), 1404
- rect() (在 *cmath* 模块中), 270
- rectangle() (在 *curses.textpad* 模块中), 661
- RecursionError, 82
- recursive_repr() (在 *reprlib* 模块中), 240
- recv() (*asyncore.dispatcher* 方法), 915
- recv() (*multiprocessing.connection.Connection* 方法), 732
- recv() (*socket.socket* 方法), 801
- recv_bytes() (*multiprocessing.connection.Connection* 方法), 733
- recv_bytes_into() (*multiprocessing.connection.Connection* 方法), 733
- recv_into() (*socket.socket* 方法), 803
- recvfrom() (*socket.socket* 方法), 802
- recvfrom_into() (*socket.socket* 方法), 803
- recvmsg() (*socket.socket* 方法), 802
- recvmsg_into() (*socket.socket* 方法), 803
- redirect_request() (*url-lib.request.HTTPRedirectHandler* 方法), 1113
- redirect_stderr() (在 *contextlib* 模块中), 1570
- redirect_stdout() (在 *contextlib* 模块中), 1570
- redisplay() (在 *readline* 模块中), 134
- redrawln() (*curses.window* 方法), 655
- redrawwin() (*curses.window* 方法), 655
- reduce (2to3 fixer), 1476
- reduce() (在 *functools* 模块中), 330
- ref (*weakref* 中的类), 223
- reference count -- 引用计数, 1768
- ReferenceError, 82, 226
- ReferenceType() (在 *weakref* 模块中), 226
- refold_source (*email.policy.EmailPolicy* 属性), 949
- refresh() (*curses.window* 方法), 655
- REG_BINARY() (在 *winreg* 模块中), 1702
- REG_DWORD() (在 *winreg* 模块中), 1703
- REG_DWORD_BIG_ENDIAN() (在 *winreg* 模块中), 1703
- REG_DWORD_LITTLE_ENDIAN() (在 *winreg* 模块中), 1703
- REG_EXPAND_SZ() (在 *winreg* 模块中), 1703
- REG_FULL_RESOURCE_DESCRIPTOR() (在 *winreg* 模块中), 1703
- REG_LINK() (在 *winreg* 模块中), 1703
- REG_MULTI_SZ() (在 *winreg* 模块中), 1703
- REG_NONE() (在 *winreg* 模块中), 1703
- REG_QWORD() (在 *winreg* 模块中), 1703
- REG_QWORD_LITTLE_ENDIAN() (在 *winreg* 模块中), 1703
- REG_RESOURCE_LIST() (在 *winreg* 模块中), 1703
- REG_RESOURCE_REQUIREMENTS_LIST() (在 *winreg* 模块中), 1703
- REG_SZ() (在 *winreg* 模块中), 1703
- register() (*abc.ABCMeta* 方法), 1580
- register() (*multiprocessing.managers.BaseManager* 方法), 740
- register() (*select.devpoll* 方法), 844
- register() (*select.epoll* 方法), 845
- register() (*selectors.BaseSelector* 方法), 849
- register() (*select.poll* 方法), 845
- register() (在 *atexit* 模块中), 1584
- register() (在 *codecs* 模块中), 146
- register() (在 *faulthandler* 模块中), 1493
- register() (在 *webbrowser* 模块中), 1086
- register_adapter() (在 *sqlite3* 模块中), 411
- register_archive_format() (在 *shutil* 模块中), 384
- register_converter() (在 *sqlite3* 模块中), 411
- register_defect() (*email.policy.Policy* 方法), 947
- register_dialect() (在 *csv* 模块中), 462
- register_error() (在 *codecs* 模块中), 148
- register_function() (*xml-rpc.server.CGIXMLRPCRequestHandler* 方法), 1212
- register_function() (*xml-rpc.server.SimpleXMLRPCServer* 方法), 1209
- register_instance() (*xml-rpc.server.CGIXMLRPCRequestHandler* 方法), 1212
- register_instance() (*xml-rpc.server.SimpleXMLRPCServer* 方法), 1210
- register_introspection_functions() (*xml-rpc.server.CGIXMLRPCRequestHandler* 方法), 1212
- register_introspection_functions() (*xml-rpc.server.SimpleXMLRPCServer* 方法), 1210
- register_multicall_functions() (*xml-*

- `rpc.server.CGIXMLRPCRequestHandler` 方法), 1212
- `register_multicall_functions()` (`xmlrpc.server.SimpleXMLRPCServer` 方法), 1210
- `register_namespace()` (在 `xml.etree.ElementTree` 模块中), 1040
- `register_optionflag()` (在 `doctest` 模块中), 1377
- `register_shape()` (在 `turtle` 模块中), 1289
- `register_unpack_format()` (在 `shutil` 模块中), 384
- `registerDOMImplementation()` (在 `xml.dom` 模块中), 1048
- `registerResult()` (在 `unittest` 模块中), 1418
- `regular package` -- 正规包, 1768
- `relative`
URL, 1122
- `relative_to()` (`pathlib.PurePath` 方法), 350
- `release()` (`_thread.lock` 方法), 787
- `release()` (`asyncio.Condition` 方法), 903
- `release()` (`asyncio.Lock` 方法), 902
- `release()` (`asyncio.Semaphore` 方法), 904
- `release()` (`logging.Handler` 方法), 610
- `release()` (`memoryview` 方法), 65
- `release()` (`multiprocessing.Lock` 方法), 735
- `release()` (`multiprocessing.RLock` 方法), 735
- `release()` (`threading.Condition` 方法), 715
- `release()` (`threading.Lock` 方法), 713
- `release()` (`threading.RLock` 方法), 714
- `release()` (`threading.Semaphore` 方法), 716
- `release()` (在 `platform` 模块中), 667
- `release_lock()` (在 `imp` 模块中), 1755
- `reload(2to3 fixer)`, 1476
- `reload()` (在 `imp` 模块中), 1753
- `reload()` (在 `importlib` 模块中), 1629
- `relpath()` (在 `os.path` 模块中), 360
- `remainder()` (`decimal.Context` 方法), 288
- `remainder_near()` (`decimal.Context` 方法), 288
- `remainder_near()` (`decimal.Decimal` 方法), 282
- `RemoteDisconnected`, 1135
- `remove()` (`array.array` 方法), 222
- `remove()` (`collections.deque` 方法), 201
- `remove()` (`frozenset` 方法), 71
- `remove()` (`mailbox.Mailbox` 方法), 997
- `remove()` (`mailbox.MH` 方法), 1003
- `remove()` (`sequence method`), 37
- `remove()` (在 `os` 模块中), 530
- `remove()` (`xml.etree.ElementTree.Element` 方法), 1042
- `remove_done_callback()` (`asyncio.Future` 方法), 872
- `remove_flag()` (`mailbox.MaildirMessage` 方法), 1006
- `remove_flag()` (`mailbox.mboxMessage` 方法), 1007
- `remove_flag()` (`mailbox.MMDFMessage` 方法), 1011
- `remove_folder()` (`mailbox.Maildir` 方法), 1000
- `remove_folder()` (`mailbox.MH` 方法), 1002
- `remove_header()` (`urllib.request.Request` 方法), 1110
- `remove_history_item()` (在 `readline` 模块中), 135
- `remove_label()` (`mailbox.BabylMessage` 方法), 1010
- `remove_option()` (`configparser.ConfigParser` 方法), 483
- `remove_option()` (`optparse.OptionParser` 方法), 1743
- `remove_pyc()` (`msilib.Directory` 方法), 1693
- `remove_reader()` (`asyncio.AbstractEventLoop` 方法), 858
- `remove_section()` (`configparser.ConfigParser` 方法), 483
- `remove_sequence()` (`mailbox.MHMessage` 方法), 1009
- `remove_signal_handler()` (`asyncio.AbstractEventLoop` 方法), 860
- `remove_writer()` (`asyncio.AbstractEventLoop` 方法), 858
- `removeAttribute()` (`xml.dom.Element` 方法), 1053
- `removeAttributeNode()` (`xml.dom.Element` 方法), 1053
- `removeAttributeNS()` (`xml.dom.Element` 方法), 1053
- `removeChild()` (`xml.dom.Node` 方法), 1051
- `removedirs()` (在 `os` 模块中), 530
- `removeFilter()` (`logging.Handler` 方法), 610
- `removeFilter()` (`logging.Logger` 方法), 609
- `removeHandler()` (`logging.Logger` 方法), 609
- `removeHandler()` (在 `unittest` 模块中), 1418
- `removeResult()` (在 `unittest` 模块中), 1418
- `removexattr()` (在 `os` 模块中), 541
- `rename()` (`ftplib.FTP` 方法), 1143
- `rename()` (`imaplib.IMAP4` 方法), 1151
- `rename()` (`pathlib.Path` 方法), 355
- `rename()` (在 `os` 模块中), 530
- `renames(2to3 fixer)`, 1476
- `renames()` (在 `os` 模块中), 530
- `reopenIfNeeded()` (`logging.handlers.WatchedFileHandler` 方法), 633
- `reorganize()` (`dbm.gnu.gdbm` 方法), 407
- `repeat()` (`timeit.Timer` 方法), 1509
- `repeat()` (在 `itertools` 模块中), 321
- `repeat()` (在 `timeit` 模块中), 1508
- `repetition`
operation, 35
- `replace()` (`bytearray` 方法), 53
- `replace()` (`bytes` 方法), 53
- `replace()` (`curses.panel.Panel` 方法), 665
- `replace()` (`datetime.date` 方法), 169
- `replace()` (`datetime.datetime` 方法), 174
- `replace()` (`datetime.time` 方法), 180
- `replace()` (`inspect.Parameter` 方法), 1602
- `replace()` (`inspect.Signature` 方法), 1601

- `replace()` (*pathlib.Path* 方法), 355
- `replace()` (*str* 方法), 44
- `replace()` (在 *os* 模块中), 530
- `replace_errors()` (在 *codecs* 模块中), 149
- `replace_header()` (*email.message.EmailMessage* 方法), 933
- `replace_header()` (*email.message.Message* 方法), 971
- `replace_history_item()` (在 *readline* 模块中), 135
- `replace_whitespace` (*textwrap.TextWrapper* 属性), 129
- `replaceChild()` (*xml.dom.Node* 方法), 1051
- `ReplacePackage()` (在 *modulefinder* 模块中), 1624
- `report()` (*filecmp.dircmp* 方法), 369
- `report()` (*modulefinder.ModuleFinder* 方法), 1624
- `REPORT_CDIF` (在 *doctest* 模块中), 1376
- `report_failure()` (*doctest.DocTestRunner* 方法), 1386
- `report_full_closure()` (*filecmp.dircmp* 方法), 370
- `REPORT_NDIFF` (在 *doctest* 模块中), 1376
- `REPORT_ONLY_FIRST_FAILURE` (在 *doctest* 模块中), 1376
- `report_partial_closure()` (*filecmp.dircmp* 方法), 369
- `report_start()` (*doctest.DocTestRunner* 方法), 1386
- `report_success()` (*doctest.DocTestRunner* 方法), 1386
- `REPORT_UDIFF` (在 *doctest* 模块中), 1376
- `report_unexpected_exception()` (*doctest.DocTestRunner* 方法), 1386
- `REPORTING_FLAGS` (在 *doctest* 模块中), 1376
- `repr (2to3 fixer)`, 1476
- `Repr` (*reprlib* 中的类), 240
- `repr()` (*reprlib.Repr* 方法), 241
- `repr()` (设置函数), 19
- `repr()` (在 *reprlib* 模块中), 240
- `repr1()` (*reprlib.Repr* 方法), 241
- `reprlib` (模块), 240
- `Request` (*urllib.request* 中的类), 1106
- `request()` (*http.client.HTTPConnection* 方法), 1135
- `request_queue_size` (*socketserver.BaseServer* 属性), 1179
- `request_rate()` (*urllib.robotparser.RobotFileParser* 方法), 1131
- `request_uri()` (在 *wsgiref.util* 模块中), 1095
- `request_version` (*http.server.BaseHTTPRequestHandler* 属性), 1185
- `RequestHandlerClass` (*socketserver.BaseServer* 属性), 1178
- `requestline` (*http.server.BaseHTTPRequestHandler* 属性), 1185
- `requires()` (在 *test.support* 模块中), 1480
- `reserved` (*zipfile.ZipInfo* 属性), 449
- `RESERVED_FUTURE` (在 *uuid* 模块中), 1175
- `RESERVED_MICROSOFT` (在 *uuid* 模块中), 1175
- `RESERVED_NCS` (在 *uuid* 模块中), 1175
- `reset()` (*bdb.Bdb* 方法), 1488
- `reset()` (*codecs.IncrementalDecoder* 方法), 151
- `reset()` (*codecs.IncrementalEncoder* 方法), 150
- `reset()` (*codecs.StreamReader* 方法), 153
- `reset()` (*codecs.StreamWriter* 方法), 152
- `reset()` (*html.parser.HTMLParser* 方法), 1027
- `reset()` (*ossaudiodev.oss_audio_device* 方法), 1243
- `reset()` (*pipes.Template* 方法), 1717
- `reset()` (*threading.Barrier* 方法), 719
- `reset()` (在 *turtle* 模块中), 1277, 1284
- `reset()` (*xdrllib.Packer* 方法), 486
- `reset()` (*xdrllib.Unpacker* 方法), 487
- `reset()` (*xml.dom.pulldom.DOMEEventStream* 方法), 1063
- `reset()` (*xml.sax.xmlreader.IncrementalParser* 方法), 1073
- `reset_mock()` (*unittest.mock.Mock* 方法), 1423
- `reset_prog_mode()` (在 *curses* 模块中), 649
- `reset_shell_mode()` (在 *curses* 模块中), 649
- `resetbuffer()` (*code.InteractiveConsole* 方法), 1617
- `resetlocale()` (在 *locale* 模块中), 1259
- `resetscreen()` (在 *turtle* 模块中), 1284
- `resetty()` (在 *curses* 模块中), 649
- `resetwarnings()` (在 *warnings* 模块中), 1567
- `resize()` (*curses.window* 方法), 656
- `resize()` (*mmap.mmap* 方法), 927
- `resize()` (在 *ctypes* 模块中), 701
- `resize_term()` (在 *curses* 模块中), 649
- `resizemode()` (在 *turtle* 模块中), 1279
- `resizeterm()` (在 *curses* 模块中), 649
- `resolution` (*datetime.date* 属性), 168
- `resolution` (*datetime.datetime* 属性), 172
- `resolution` (*datetime.time* 属性), 179
- `resolution` (*datetime.timedelta* 属性), 165
- `resolve()` (*pathlib.Path* 方法), 355
- `resolve_name()` (在 *importlib.util* 模块中), 1640
- `resolveEntity()` (*xml.sax.handler.EntityResolver* 方法), 1069
- `resource` (模块), 1718
- `ResourceDenied`, 1480
- `ResourceLoader` (*importlib.abc* 中的类), 1632
- `ResourceWarning`, 86
- `response` (*nntplib.NNTPError* 属性), 1155
- `response()` (*imaplib.IMAP4* 方法), 1151
- `ResponseNotReady`, 1135
- `responses` (*http.server.BaseHTTPRequestHandler* 属性), 1185
- `responses()` (在 *http.client* 模块中), 1135
- `restart` (*pdb* command), 1499
- `restore()` (在 *difflib* 模块中), 120

- `restype` (`ctypes.FuncPtr` 属性), 697
- `result()` (`asyncio.Future` 方法), 872
- `result()` (`concurrent.futures.Future` 方法), 765
- `results()` (`trace.Trace` 方法), 1514
- `resume_reading()` (`asyncio.ReadTransport` 方法), 880
- `resume_writing()` (`asyncio.BaseProtocol` 方法), 884
- `retr()` (`poplib.POP3` 方法), 1146
- `retrbinary()` (`ftplib.FTP` 方法), 1142
- `retrieve()` (`urllib.request.URLopener` 方法), 1120
- `retrlines()` (`ftplib.FTP` 方法), 1142
- `return` (`pdb` command), 1498
- `return_annotation` (`inspect.Signature` 属性), 1601
- `return_ok()` (`http.cookiejar.CookiePolicy` 方法), 1196
- `RETURN_VALUE` (`opcode`), 1674
- `return_value` (`unittest.mock.Mock` 属性), 1425
- `returncode` (`asyncio.asyncio.subprocess.Process` 属性), 898
- `returncode` (`subprocess.CalledProcessError` 属性), 769
- `returncode` (`subprocess.CompletedProcess` 属性), 768
- `returncode` (`subprocess.Popen` 属性), 775
- `reverse()` (`array.array` 方法), 222
- `reverse()` (`collections.deque` 方法), 201
- `reverse()` (`sequence method`), 37
- `reverse()` (在 `audioop` 模块中), 1229
- `reverse_order()` (`pstats.Stats` 方法), 1504
- `reverse_pointer` (`ipaddress.IPv4Address` 属性), 1216
- `reverse_pointer` (`ipaddress.IPv6Address` 属性), 1217
- `reversed()` (`⌈置函数`), 19
- `Reversible` (`collections.abc` 中的类), 212
- `Reversible` (`typing` 中的类), 1359
- `revert()` (`http.cookiejar.FileCookieJar` 方法), 1195
- `rewind()` (`aifc.aifc` 方法), 1231
- `rewind()` (`sunau.AU_read` 方法), 1234
- `rewind()` (`wave.Wave_read` 方法), 1236
- RFC
 - RFC 821, 1160, 1162
 - RFC 822, 569, 960, 977, 1137, 1163, 1165, 1166, 1251
 - RFC 854, 1170, 1171
 - RFC 959, 1139, 1143
 - RFC 977, 1153
 - RFC 1014, 486
 - RFC 1123, 569
 - RFC 1321, 493
 - RFC 1422, 832, 841
 - RFC 1521, 1019, 1022
 - RFC 1522, 1021, 1022
 - RFC 1524, 995, 996
 - RFC 1730, 1147
 - RFC 1738, 1129
 - RFC 1750, 814
 - RFC 1766, 1258, 1259
 - RFC 1808, 1123, 1129
 - RFC 1832, 486
 - RFC 1869, 1160, 1162
 - RFC 1870, 1167, 1170
 - RFC 1939, 1144
 - RFC 2045, 929, 933, 934, 955, 956, 971, 972, 978, 1017, 1019
 - RFC 2045#section-6.8, 1205
 - RFC 2046, 929, 959, 978
 - RFC 2047, 929, 948, 953, 954, 978, 979, 983
 - RFC 2060, 1147, 1152
 - RFC 2068, 1189
 - RFC 2104, 503
 - RFC 2109, 11891191, 1193, 1194, 11981200
 - RFC 2183, 929, 935, 973
 - RFC 2231, 929, 933, 934, 971, 972, 978, 985
 - RFC 2295, 1133
 - RFC 2342, 1150
 - RFC 2368, 1129
 - RFC 2373, 1216
 - RFC 2396, 1125, 1129
 - RFC 2397, 1115
 - RFC 2449, 1146
 - RFC 2595, 1144, 1147
 - RFC 2616, 1096, 1099, 1113, 1121, 1130
 - RFC 2732, 1129
 - RFC 2818, 815
 - RFC 2821, 929
 - RFC 2822, 569, 969, 977979, 983, 984, 1005, 1185
 - RFC 2964, 1194
 - RFC 2965, 1107, 1110, 1193, 1194, 11961199, 1201
 - RFC 2980, 1153, 1159
 - RFC 3056, 1217
 - RFC 3171, 1216
 - RFC 3280, 823
 - RFC 3330, 1216
 - RFC 3454, 132
 - RFC 3490, 159161
 - RFC 3490#section-3.1, 161
 - RFC 3492, 159, 160
 - RFC 3493, 810
 - RFC 3501, 1152
 - RFC 3542, 798
 - RFC 3548, 1017, 1018
 - RFC 3659, 1143
 - RFC 3879, 1217
 - RFC 3927, 1216
 - RFC 3977, 1153, 11551157, 1159
 - RFC 3986, 1124, 1126, 1129
 - RFC 4086, 841
 - RFC 4122, 11731175
 - RFC 4180, 461

- RFC 4193, 1217
- RFC 4217, 1140
- RFC 4291, 1216
- RFC 4380, 1217
- RFC 4627, 986, 994
- RFC 4642, 1155
- RFC 4954, 1164
- RFC 5161, 1150
- RFC 5233, 929, 967
- RFC 5246, 821, 841
- RFC 5280, 815, 841
- RFC 5321, 957, 1167, 1168
- RFC 5322, 930, 940, 943, 944, 946, 948, 949, 951, 953, 954, 957, 1165
- RFC 5424, 638
- RFC 5735, 1216
- RFC 5929, 824
- RFC 6066, 820, 829, 841
- RFC 6125, 815
- RFC 6152, 1167
- RFC 6531, 931, 949, 1161, 1167, 1169
- RFC 6532, 929, 930, 940, 949
- RFC 6585, 1133
- RFC 6855, 1150
- RFC 6856, 1146
- RFC 7159, 986, 993, 994
- RFC 7230, 1106, 1137
- RFC 7231, 1132
- RFC 7238, 1132
- RFC 7301, 820, 829
- RFC 7525, 841
- RFC 7693, 496
- RFC 7914, 496
- rfc2109 (*http.cookiejar.Cookie* 属性), 1200
- rfc2109_as_netscape (*http.cookiejar.DefaultCookiePolicy* 属性), 1198
- rfc2965 (*http.cookiejar.CookiePolicy* 属性), 1197
- RFC_4122 () (在 *uuid* 模块中), 1175
- rfile (*http.server.BaseHTTPRequestHandler* 属性), 1185
- rfind () (*bytearray* 方法), 53
- rfind () (*bytes* 方法), 53
- rfind () (*mmap.mmap* 方法), 927
- rfind () (*str* 方法), 44
- rgb_to_hls () (在 *colorsys* 模块中), 1239
- rgb_to_hsv () (在 *colorsys* 模块中), 1239
- rgb_to_yiq () (在 *colorsys* 模块中), 1239
- rglob () (*pathlib.Path* 方法), 355
- right (*filecmp.dircmp* 属性), 370
- right () (在 *turtle* 模块中), 1268
- right_list (*filecmp.dircmp* 属性), 370
- right_only (*filecmp.dircmp* 属性), 370
- RIGHTSHIFT () (在 *token* 模块中), 1657
- RIGHTSHIFTEQUAL () (在 *token* 模块中), 1657
- rindex () (*bytearray* 方法), 53
- rindex () (*bytes* 方法), 53
- rindex () (*str* 方法), 44
- rjust () (*bytearray* 方法), 55
- rjust () (*bytes* 方法), 55
- rjust () (*str* 方法), 44
- rlcompleter (模块), 138
- rlencode_hqx () (在 *binascii* 模块中), 1021
- rldecode_hqx () (在 *binascii* 模块中), 1021
- RLIM_INFINITY () (在 *resource* 模块中), 1718
- RLIMIT_AS () (在 *resource* 模块中), 1720
- RLIMIT_CORE () (在 *resource* 模块中), 1719
- RLIMIT_CPU () (在 *resource* 模块中), 1719
- RLIMIT_DATA () (在 *resource* 模块中), 1719
- RLIMIT_FSIZE () (在 *resource* 模块中), 1719
- RLIMIT_MEMLOCK () (在 *resource* 模块中), 1720
- RLIMIT_MSGQUEUE () (在 *resource* 模块中), 1720
- RLIMIT_NICE () (在 *resource* 模块中), 1720
- RLIMIT_NOFILE () (在 *resource* 模块中), 1719
- RLIMIT_NPROC () (在 *resource* 模块中), 1719
- RLIMIT_NPTS () (在 *resource* 模块中), 1721
- RLIMIT_OFILE () (在 *resource* 模块中), 1720
- RLIMIT_RSS () (在 *resource* 模块中), 1719
- RLIMIT_RTPRIO () (在 *resource* 模块中), 1720
- RLIMIT_RTTIME () (在 *resource* 模块中), 1720
- RLIMIT_SBSIZE () (在 *resource* 模块中), 1720
- RLIMIT_SIGPENDING () (在 *resource* 模块中), 1720
- RLIMIT_STACK () (在 *resource* 模块中), 1719
- RLIMIT_SWAP () (在 *resource* 模块中), 1720
- RLIMIT_VMEM () (在 *resource* 模块中), 1720
- RLock (*multiprocessing* 中的类), 735
- RLock (*threading* 中的类), 713
- RLock () (*multiprocessing.managers.SyncManager* 方法), 741
- rmd () (*ftplib.FTP* 方法), 1143
- rmdir () (*pathlib.Path* 方法), 355
- rmdir () (在 *os* 模块中), 531
- RMFF, 1238
- rms () (在 *audioop* 模块中), 1229
- rmtree () (在 *shutil* 模块中), 380
- RobotFileParser (*urllib.robotparser* 中的类), 1130
- robots.txt, 1130
- rollback () (*sqlite3.Connection* 方法), 412
- ROT_THREE (*opcode*), 1671
- ROT_TWO (*opcode*), 1671
- rotate () (*collections.deque* 方法), 201
- rotate () (*decimal.Context* 方法), 288
- rotate () (*decimal.Decimal* 方法), 282
- rotate () (*logging.handlers.BaseRotatingHandler* 方法), 633
- RotatingFileHandler (*logging.handlers* 中的类), 634

rotation_filename() (logging.handlers.BaseRotatingHandler 方法), 633
 rotator (logging.handlers.BaseRotatingHandler 属性), 633
 round() (内置函数), 19
 ROUND_UP() (在 decimal 模块中), 290
 ROUND_CEILING() (在 decimal 模块中), 289
 ROUND_DOWN() (在 decimal 模块中), 289
 ROUND_FLOOR() (在 decimal 模块中), 289
 ROUND_HALF_DOWN() (在 decimal 模块中), 289
 ROUND_HALF_EVEN() (在 decimal 模块中), 289
 ROUND_HALF_UP() (在 decimal 模块中), 289
 ROUND_UP() (在 decimal 模块中), 289
 Rounded (decimal 中的类), 291
 Row (sqlite3 中的类), 420
 row_factory (sqlite3.Connection 属性), 415
 rowcount (sqlite3.Cursor 属性), 419
 RPAR() (在 token 模块中), 1657
 rpartition() (bytearray 方法), 54
 rpartition() (bytes 方法), 54
 rpartition() (str 方法), 44
 rpc_paths (xmlrpc.server.SimpleXMLRPCRequestHandler 属性), 1210
 rpop() (poplib.POP3 方法), 1146
 rset() (poplib.POP3 方法), 1146
 rshift() (在 operator 模块中), 335
 rsplit() (bytearray 方法), 55
 rsplit() (bytes 方法), 55
 rsplit() (str 方法), 44
 RSQB() (在 token 模块中), 1657
 rstrip() (bytearray 方法), 55
 rstrip() (bytes 方法), 55
 rstrip() (str 方法), 45
 rt() (在 turtle 模块中), 1268
 RTLD_DEEPBIND() (在 os 模块中), 553
 RTLD_GLOBAL() (在 os 模块中), 553
 RTLD_LAZY() (在 os 模块中), 553
 RTLD_LOCAL() (在 os 模块中), 553
 RTLD_NODELETE() (在 os 模块中), 553
 RTLD_NOLOAD() (在 os 模块中), 553
 RTLD_NOW() (在 os 模块中), 553
 ruler (cmd.Cmd 属性), 1298
 run (pdb command), 1499
 Run script, 1342
 run() (bdb.Bdb 方法), 1491
 run() (doctest.DocTestRunner 方法), 1386
 run() (multiprocessing.Process 方法), 726
 run() (pdb.Pdb 方法), 1496
 run() (profile.Profile 方法), 1503
 run() (sched.scheduler 方法), 783
 run() (threading.Thread 方法), 711
 run() (trace.Trace 方法), 1514
 run() (unittest.TestCase 方法), 1400
 run() (unittest.TestSuite 方法), 1409
 run() (unittest.TextTestRunner 方法), 1414
 run() (在 pdb 模块中), 1495
 run() (在 profile 模块中), 1502
 run() (在 subprocess 模块中), 767
 run() (wsgiref.handlers.BaseHandler 方法), 1101
 run_coroutine_threadsafe() (在 asyncio 模块中), 876
 run_docstring_examples() (在 doctest 模块中), 1380
 run_doctest() (在 test.support 模块中), 1481
 run_forever() (asyncio.AbstractEventLoop 方法), 853
 run_in_executor() (asyncio.AbstractEventLoop 方法), 860
 run_module() (在 runpy 模块中), 1625
 run_path() (在 runpy 模块中), 1626
 run_script() (modulefinder.ModuleFinder 方法), 1624
 run_unittest() (在 test.support 模块中), 1481
 run_until_complete() (asyncio.AbstractEventLoop 方法), 853
 run_with_locale() (在 test.support 模块中), 1483
 runcall() (bdb.Bdb 方法), 1491
 runcall() (pdb.Pdb 方法), 1496
 runcall() (profile.Profile 方法), 1503
 runcall() (在 pdb 模块中), 1495
 runcode() (code.InteractiveInterpreter 方法), 1616
 runctx() (bdb.Bdb 方法), 1491
 runctx() (profile.Profile 方法), 1503
 runctx() (trace.Trace 方法), 1514
 runctx() (在 profile 模块中), 1502
 runeval() (bdb.Bdb 方法), 1491
 runeval() (pdb.Pdb 方法), 1496
 runeval() (在 pdb 模块中), 1495
 runfunc() (trace.Trace 方法), 1514
 running() (concurrent.futures.Future 方法), 765
 runpy (模块), 1625
 runsource() (code.InteractiveInterpreter 方法), 1616
 RuntimeError, 82
 RuntimeWarning, 86
 RUSAGE_BOTH() (在 resource 模块中), 1722
 RUSAGE_CHILDREN() (在 resource 模块中), 1722
 RUSAGE_SELF() (在 resource 模块中), 1721
 RUSAGE_THREAD() (在 resource 模块中), 1722

S

-s S, --setup=S
 timeit 命令行选项, 1510
 S() (在 re 模块中), 104
 -s, --start-directory directory
 unittest-discover 命令行选项, 1394
 -s, --summary
 trace 命令行选项, 1513

- S_ENFMT() (在 *stat* 模块中), 367
- S_IEXEC() (在 *stat* 模块中), 367
- S_IFBLK() (在 *stat* 模块中), 366
- S_IFCHR() (在 *stat* 模块中), 366
- S_IFDIR() (在 *stat* 模块中), 366
- S_IFDOOR() (在 *stat* 模块中), 366
- S_IFIFO() (在 *stat* 模块中), 366
- S_IFLNK() (在 *stat* 模块中), 366
- S_IFMT() (在 *stat* 模块中), 364
- S_IFPORT() (在 *stat* 模块中), 366
- S_IFREG() (在 *stat* 模块中), 366
- S_IFSOCK() (在 *stat* 模块中), 366
- S_IFWHT() (在 *stat* 模块中), 366
- S_IMODE() (在 *stat* 模块中), 364
- S_IREAD() (在 *stat* 模块中), 367
- S_IRGRP() (在 *stat* 模块中), 367
- S_IROTH() (在 *stat* 模块中), 367
- S_IRUSR() (在 *stat* 模块中), 367
- S_IRWXG() (在 *stat* 模块中), 367
- S_IRWXO() (在 *stat* 模块中), 367
- S_IRWXU() (在 *stat* 模块中), 367
- S_ISBLK() (在 *stat* 模块中), 364
- S_ISCHR() (在 *stat* 模块中), 364
- S_ISDIR() (在 *stat* 模块中), 364
- S_ISDOOR() (在 *stat* 模块中), 364
- S_ISFIFO() (在 *stat* 模块中), 364
- S_ISGID() (在 *stat* 模块中), 367
- S_ISLNK() (在 *stat* 模块中), 364
- S_ISPORT() (在 *stat* 模块中), 364
- S_ISREG() (在 *stat* 模块中), 364
- S_ISSOCK() (在 *stat* 模块中), 364
- S_ISUID() (在 *stat* 模块中), 367
- S_ISVTX() (在 *stat* 模块中), 367
- S_ISWHT() (在 *stat* 模块中), 364
- S_IWGRP() (在 *stat* 模块中), 367
- S_IWOTH() (在 *stat* 模块中), 367
- S_IWRITE() (在 *stat* 模块中), 367
- S_IWUSR() (在 *stat* 模块中), 367
- S_IXGRP() (在 *stat* 模块中), 367
- S_IXOTH() (在 *stat* 模块中), 367
- S_IXUSR() (在 *stat* 模块中), 367
- safe_substitute() (*string.Template* 方法), 97
- saferepr() (在 *pprint* 模块中), 236
- same_files(*filecmp.dircmp* 属性), 370
- same_quantum() (*decimal.Context* 方法), 288
- same_quantum() (*decimal.Decimal* 方法), 282
- samefile() (*pathlib.Path* 方法), 355
- samefile() (在 *os.path* 模块中), 360
- SameFileError, 378
- sameopenfile() (在 *os.path* 模块中), 360
- samestat() (在 *os.path* 模块中), 360
- sample() (在 *random* 模块中), 302
- save() (*http.cookiejar.FileCookieJar* 方法), 1195
- SaveKey() (在 *winreg* 模块中), 1700
- savetty() (在 *curses* 模块中), 649
- SAX2DOM(*xml.dom.pulldom* 中的类), 1062
- SAXException, 1064
- SAXNotRecognizedException, 1064
- SAXNotSupportedException, 1065
- SAXParseException, 1064
- scaleb() (*decimal.Context* 方法), 288
- scaleb() (*decimal.Decimal* 方法), 282
- scandir() (在 *os* 模块中), 531
- scanf(), 112
- sched(模块), 782
- SCHED_BATCH() (在 *os* 模块中), 550
- SCHED_FIFO() (在 *os* 模块中), 550
- sched_get_priority_max() (在 *os* 模块中), 550
- sched_get_priority_min() (在 *os* 模块中), 550
- sched_getaffinity() (在 *os* 模块中), 551
- sched_getparam() (在 *os* 模块中), 551
- sched_getscheduler() (在 *os* 模块中), 551
- SCHED_IDLE() (在 *os* 模块中), 550
- SCHED_OTHER() (在 *os* 模块中), 550
- sched_param(*os* 中的类), 550
- sched_priority(*os.sched_param* 属性), 550
- SCHED_RESET_ON_FORK() (在 *os* 模块中), 550
- SCHED_RR() (在 *os* 模块中), 550
- sched_rr_get_interval() (在 *os* 模块中), 551
- sched_setaffinity() (在 *os* 模块中), 551
- sched_setparam() (在 *os* 模块中), 551
- sched_setscheduler() (在 *os* 模块中), 551
- SCHED_SPORADIC() (在 *os* 模块中), 550
- sched_yield() (在 *os* 模块中), 551
- scheduler(*sched* 中的类), 782
- schema() (在 *msilib* 模块中), 1694
- Screen(*turtle* 中的类), 1290
- screenize() (在 *turtle* 模块中), 1284
- script_from_examples() (在 *doctest* 模块中), 1388
- scroll() (*curses.window* 方法), 656
- ScrolledCanvas(*turtle* 中的类), 1290
- scrollok() (*curses.window* 方法), 656
- script() (在 *hashlib* 模块中), 496
- search
 - path, module, 377, 1552, 1610
- search() (*imaplib.IMAP4* 方法), 1151
- search() (*re.regex* 方法), 108
- search() (在 *re* 模块中), 105
- second(*datetime.datetime* 属性), 173
- second(*datetime.time* 属性), 180
- seconds since the epoch, 566
- secrets(模块), 505
- SECTCRE(*configparser.ConfigParser* 属性), 478
- sections() (*configparser.ConfigParser* 方法), 481
- secure(*http.cookiejar.CookieJar* 属性), 1199
- secure hash algorithm, SHA1, SHA224, SHA256, SHA384, SHA512, 493

- Secure Sockets Layer, 810
- security
CGI, 1092
- see() (*tkinter.ttk.Treeview* 方法), 1330
- seed() (在 *random* 模块中), 301
- seek() (*chunk.Chunk* 方法), 1238
- seek() (*io.IOBase* 方法), 557
- seek() (*io.TextIOBase* 方法), 563
- seek() (*mmap.mmap* 方法), 927
- SEEK_CUR() (在 *os* 模块中), 518
- SEEK_END() (在 *os* 模块中), 518
- SEEK_SET() (在 *os* 模块中), 518
- seekable() (*io.IOBase* 方法), 558
- seen_greeting (*smtpd.SMTPChannel* 属性), 1169
- Select (*tkinter.tix* 中的类), 1337
- select (模块), 841
- select() (*imaplib.IMAP4* 方法), 1151
- select() (*selectors.BaseSelector* 方法), 850
- select() (*tkinter.ttk.Notebook* 方法), 1324
- select() (在 *select* 模块中), 843
- selected_alpn_protocol() (*ssl.SSLSocket* 方法), 825
- selected_npn_protocol() (*ssl.SSLSocket* 方法), 825
- selection() (*tkinter.ttk.Treeview* 方法), 1331
- selection_add() (*tkinter.ttk.Treeview* 方法), 1331
- selection_remove() (*tkinter.ttk.Treeview* 方法), 1331
- selection_set() (*tkinter.ttk.Treeview* 方法), 1331
- selection_toggle() (*tkinter.ttk.Treeview* 方法), 1331
- selector (*urllib.request.Request* 属性), 1109
- SelectorEventLoop (*asyncio* 中的类), 865
- SelectorKey (*selectors* 中的类), 849
- selectors (模块), 848
- SelectSelector (*selectors* 中的类), 851
- Semaphore (*asyncio* 中的类), 904
- Semaphore (*multiprocessing* 中的类), 735
- Semaphore (*threading* 中的类), 716
- Semaphore() (*multiprocessing.managers.SyncManager* 方法), 741
- semaphores, binary, 786
- SEMI() (在 *token* 模块中), 1657
- send() (*asyncore.dispatcher* 方法), 914
- send() (*http.client.HTTPConnection* 方法), 1137
- send() (*imaplib.IMAP4* 方法), 1151
- send() (*logging.handlers.DatagramHandler* 方法), 637
- send() (*logging.handlers.SocketHandler* 方法), 636
- send() (*multiprocessing.connection.Connection* 方法), 732
- send() (*socket.socket* 方法), 803
- send_bytes() (*multiprocessing.connection.Connection* 方法), 733
- send_error() (*http.server.BaseHTTPRequestHandler* 方法), 1186
- send_flow_data() (*formatter.writer* 方法), 1686
- send_header() (*http.server.BaseHTTPRequestHandler* 方法), 1186
- send_hor_rule() (*formatter.writer* 方法), 1686
- send_label_data() (*formatter.writer* 方法), 1686
- send_line_break() (*formatter.writer* 方法), 1686
- send_literal_data() (*formatter.writer* 方法), 1686
- send_message() (*smtpplib.SMTP* 方法), 1165
- send_paragraph() (*formatter.writer* 方法), 1686
- send_response() (*http.server.BaseHTTPRequestHandler* 方法), 1186
- send_response_only() (*http.server.BaseHTTPRequestHandler* 方法), 1186
- send_signal() (*asyncio.asyncio.subprocess.Process* 方法), 898
- send_signal() (*asyncio.BaseSubprocessTransport* 方法), 882
- send_signal() (*subprocess.Popen* 方法), 774
- sendall() (*socket.socket* 方法), 804
- sendcmd() (*ftplib.FTP* 方法), 1142
- sendfile() (*socket.socket* 方法), 804
- sendfile() (在 *os* 模块中), 521
- sendfile() (*wsgiref.handlers.BaseHandler* 方法), 1103
- sendmail() (*smtpplib.SMTP* 方法), 1165
- sendmsg() (*socket.socket* 方法), 804
- sendmsg_afalg() (*socket.socket* 方法), 804
- sendto() (*asyncio.DatagramTransport* 方法), 881
- sendto() (*socket.socket* 方法), 804
- sentinel (*multiprocessing.Process* 属性), 727
- sentinel() (在 *unittest.mock* 模块中), 1445
- sep() (在 *os* 模块中), 552
- sequence
iteration, 34
types, immutable, 36
types, mutable, 36
types, operations on, 35, 37
对象, 35
- sequence -- 序列, 1768
- Sequence (*collections.abc* 中的类), 212
- Sequence (*typing* 中的类), 1359
- sequence() (在 *msilib* 模块中), 1694
- sequence2st() (在 *parser* 模块中), 1646
- SequenceMatcher (*difflib* 中的类), 117, 121
- serializing
objects, 387
- serve_forever() (*socketserver.BaseServer* 方法), 1178
- server

- WWW, 1087, 1184
- Server (*asyncio* 中的类), 862
- server (*http.server.BaseHTTPRequestHandler* 属性), 1184
- server_activate() (*socketserver.BaseServer* 方法), 1179
- server_address (*socketserver.BaseServer* 属性), 1178
- server_bind() (*socketserver.BaseServer* 方法), 1179
- server_close() (*socketserver.BaseServer* 方法), 1178
- server_hostname (*ssl.SSLSocket* 属性), 825
- server_side (*ssl.SSLSocket* 属性), 825
- server_software (*wsgiref.handlers.BaseHandler* 属性), 1102
- server_version (*http.server.BaseHTTPRequestHandler* 属性), 1185
- server_version (*http.server.SimpleHTTPRequestHandler* 属性), 1187
- ServerProxy (*xmlrpc.client* 中的类), 1201
- service_actions() (*socketserver.BaseServer* 方法), 1178
- session (*ssl.SSLSocket* 属性), 826
- session_reused (*ssl.SSLSocket* 属性), 826
- session_stats() (*ssl.SSLContext* 方法), 830
- set
对象, 69
- Set (*collections.abc* 中的类), 212
- Set (*typing* 中的类), 1360
- set (设置类), 69
- Set Breakpoint, 1343
- set() (*asyncio.Event* 方法), 902
- set() (*configparser.ConfigParser* 方法), 482
- set() (*configparser.RawConfigParser* 方法), 484
- set() (*http.cookies.Morsel* 方法), 1191
- set() (*ossaudiodev.oss_mixer_device* 方法), 1245
- set() (*test.support.EnvironmentVarGuard* 方法), 1485
- set() (*threading.Event* 方法), 717
- set() (*tkinter.ttk.Combobox* 方法), 1322
- set() (*tkinter.ttk.Treeview* 方法), 1331
- set() (*xml.etree.ElementTree.Element* 方法), 1041
- SET_ADD (*opcode*), 1674
- set_allowed_domains()
(*http.cookiejar.DefaultCookiePolicy* 方法), 1198
- set_alpn_protocols() (*ssl.SSLContext* 方法), 829
- set_app() (*wsgiref.simple_server.WSGIServer* 方法), 1098
- set_asyncgen_hooks() (在 *sys* 模块中), 1555
- set_authorizer() (*sqlite3.Connection* 方法), 414
- set_auto_history() (在 *readline* 模块中), 135
- set_blocked_domains()
(*http.cookiejar.DefaultCookiePolicy* 方法), 1198
- set_blocking() (在 *os* 模块中), 521
- set_boundary() (*email.message.EmailMessage* 方法), 934
- set_boundary() (*email.message.Message* 方法), 973
- set_break() (*bdb.Bdb* 方法), 1490
- set_charset() (*email.message.Message* 方法), 969
- set_children() (*tkinter.ttk.Treeview* 方法), 1328
- set_ciphers() (*ssl.SSLContext* 方法), 828
- set_completer() (在 *readline* 模块中), 136
- set_completer_delims() (在 *readline* 模块中), 136
- set_completion_display_matches_hook()
(在 *readline* 模块中), 136
- set_content() (*email.contentmanager.ContentManager* 方法), 958
- set_content() (*email.message.EmailMessage* 方法), 937
- set_content() (在 *email.contentmanager* 模块中), 959
- set_continue() (*bdb.Bdb* 方法), 1490
- set_cookie() (*http.cookiejar.CookieJar* 方法), 1195
- set_cookie_if_ok() (*http.cookiejar.CookieJar* 方法), 1195
- set_coroutine_wrapper() (在 *sys* 模块中), 1556
- set_current() (*msilib.Feature* 方法), 1693
- set_data() (*importlib.abc.SourceLoader* 方法), 1634
- set_data() (*importlib.machinery.SourceFileLoader* 方法), 1637
- set_date() (*mailbox.MaildirMessage* 方法), 1006
- set_debug() (*asyncio.AbstractEventLoop* 方法), 862
- set_debug() (在 *gc* 模块中), 1593
- set_debuglevel() (*ftplib.FTP* 方法), 1141
- set_debuglevel() (*http.client.HTTPConnection* 方法), 1136
- set_debuglevel() (*nntplib.NNTP* 方法), 1159
- set_debuglevel() (*poplib.POP3* 方法), 1145
- set_debuglevel() (*smtp.SMTP* 方法), 1162
- set_debuglevel() (*telnetlib.Telnet* 方法), 1171
- set_default_executor() (*asyncio.AbstractEventLoop* 方法), 861
- set_default_type() (*email.message.EmailMessage* 方法), 934
- set_default_type() (*email.message.Message* 方法), 971
- set_default_verify_paths() (*ssl.SSLContext* 方法), 828
- set_defaults() (*argparse.ArgumentParser* 方法), 600
- set_defaults() (*optparse.OptionParser* 方法), 1744
- set_ecdh_curve() (*ssl.SSLContext* 方法), 830
- set_errno() (在 *ctypes* 模块中), 702
- set_event_loop() (*asyncio.AbstractEventLoopPolicy* 方法), 867
- set_event_loop() (在 *asyncio* 模块中), 865

- `set_event_loop_policy()` (在 `asyncio` 模块中), 867
`set_exception()` (`asyncio.Future` 方法), 872
`set_exception()` (`asyncio.StreamReader` 方法), 890
`set_exception()` (`concurrent.futures.Future` 方法), 766
`set_exception_handler()` (`asyncio.AbstractEventLoop` 方法), 861
`set_executable()` (在 `multiprocessing` 模块中), 732
`set_flags()` (`mailbox.MaildirMessage` 方法), 1006
`set_flags()` (`mailbox.mboxMessage` 方法), 1007
`set_flags()` (`mailbox.MMDfMessage` 方法), 1011
`set_from()` (`mailbox.mboxMessage` 方法), 1007
`set_from()` (`mailbox.MMDfMessage` 方法), 1011
`set_handle_inheritable()` (在 `os` 模块中), 523
`set_history_length()` (在 `readline` 模块中), 135
`set_info()` (`mailbox.MaildirMessage` 方法), 1006
`set_inheritable()` (`socket.socket` 方法), 805
`set_inheritable()` (在 `os` 模块中), 523
`set_labels()` (`mailbox.BabylMessage` 方法), 1010
`set_last_error()` (在 `ctypes` 模块中), 702
`set_literal(2to3 fixer)`, 1477
`set_loader()` (在 `importlib.util` 模块中), 1641
`set_next()` (`bdb.Bdb` 方法), 1490
`set_nonstandard_attr()` (`http.cookiejar.Cookie` 方法), 1200
`set_npn_protocols()` (`ssl.SSLContext` 方法), 829
`set_ok()` (`http.cookiejar.CookiePolicy` 方法), 1196
`set_option_negotiation_callback()` (`telnetlib.Telnet` 方法), 1172
`set_output_charset()` (`gettext.NullTranslations` 方法), 1250
`set_package()` (在 `importlib.util` 模块中), 1641
`set_param()` (`email.message.EmailMessage` 方法), 934
`set_param()` (`email.message.Message` 方法), 972
`set_pasv()` (`ftplib.FTP` 方法), 1142
`set_payload()` (`email.message.Message` 方法), 969
`set_policy()` (`http.cookiejar.CookieJar` 方法), 1195
`set_position()` (`xdrlib.Unpacker` 方法), 487
`set_pre_input_hook()` (在 `readline` 模块中), 135
`set_progress_handler()` (`sqlite3.Connection` 方法), 414
`set_protocol()` (`asyncio.BaseTransport` 方法), 879
`set_proxy()` (`urllib.request.Request` 方法), 1110
`set_quit()` (`bdb.Bdb` 方法), 1490
`set_recsrc()` (`ossaudiodev.oss_mixer_device` 方法), 1245
`set_result()` (`asyncio.Future` 方法), 872
`set_result()` (`concurrent.futures.Future` 方法), 765
`set_return()` (`bdb.Bdb` 方法), 1490
`set_running_or_notify_cancel()` (`concurrent.futures.Future` 方法), 765
`set_seq1()` (`difflib.SequenceMatcher` 方法), 121
`set_seq2()` (`difflib.SequenceMatcher` 方法), 122
`set_seqs()` (`difflib.SequenceMatcher` 方法), 121
`set_sequences()` (`mailbox.MH` 方法), 1002
`set_sequences()` (`mailbox.MHMessage` 方法), 1009
`set_server_documentation()` (`xmlrpc.server.DocCGIXMLRPCRequestHandler` 方法), 1214
`set_server_documentation()` (`xmlrpc.server.DocXMLRPCServer` 方法), 1213
`set_server_name()` (`xmlrpc.server.DocCGIXMLRPCRequestHandler` 方法), 1214
`set_server_name()` (`xmlrpc.server.DocXMLRPCServer` 方法), 1213
`set_server_title()` (`xmlrpc.server.DocCGIXMLRPCRequestHandler` 方法), 1214
`set_server_title()` (`xmlrpc.server.DocXMLRPCServer` 方法), 1213
`set_servername_callback()` (`ssl.SSLContext` 方法), 829
`set_spacing()` (`formatter.formatter` 方法), 1685
`set_start_method()` (在 `multiprocessing` 模块中), 732
`set_startup_hook()` (在 `readline` 模块中), 135
`set_step()` (`bdb.Bdb` 方法), 1490
`set_subdir()` (`mailbox.MaildirMessage` 方法), 1005
`set_task_factory()` (`asyncio.AbstractEventLoop` 方法), 855
`set_terminator()` (`asynchat.async_chat` 方法), 918
`set_threshold()` (在 `gc` 模块中), 1593
`set_trace()` (`bdb.Bdb` 方法), 1490
`set_trace()` (`pdb.Pdb` 方法), 1496
`set_trace()` (在 `bdb` 模块中), 1491
`set_trace()` (在 `pdb` 模块中), 1495
`set_trace_callback()` (`sqlite3.Connection` 方法), 414
`set_transport()` (`asyncio.StreamReader` 方法), 890
`set_tunnel()` (`http.client.HTTPConnection` 方法), 1136
`set_type()` (`email.message.Message` 方法), 972
`set_unittest_reportflags()` (在 `doctest` 模块中), 1382
`set_unixfrom()` (`email.message.EmailMessage` 方法), 932
`set_unixfrom()` (`email.message.Message` 方法), 968
`set_until()` (`bdb.Bdb` 方法), 1490
`set_url()` (`urllib.robotparser.RobotFileParser` 方法), 1130
`set_usage()` (`optparse.OptionParser` 方法), 1744
`set_userptr()` (`curses.panel.Panel` 方法), 666
`set_visible()` (`mailbox.BabylMessage` 方法), 1010
`set_wakeup_fd()` (在 `signal` 模块中), 922
`set_write_buffer_limits()` (`asyncio.WriteTransport` 方法), 880

- `setacl()` (*imaplib.IMAP4* 方法), 1151
- `setannotation()` (*imaplib.IMAP4* 方法), 1151
- `setattr()` (`setattr` 函数), 19
- `setAttribute()` (*xml.dom.Element* 方法), 1053
- `setAttributeNode()` (*xml.dom.Element* 方法), 1053
- `setAttributeNodeNS()` (*xml.dom.Element* 方法), 1053
- `setAttributeNS()` (*xml.dom.Element* 方法), 1053
- `SetBase()` (*xml.parsers.expat.xmlparser* 方法), 1076
- `setblocking()` (*socket.socket* 方法), 805
- `setByteStream()` (*xml.sax.xmlreader.InputSource* 方法), 1074
- `setcbreak()` (在 *tty* 模块中), 1713
- `setCharacterStream()`
(*xml.sax.xmlreader.InputSource* 方法), 1074
- `setcheckinterval()` (在 *sys* 模块中), 1554
- `setcomptype()` (*aifc.aifc* 方法), 1232
- `setcomptype()` (*sunau.AU_write* 方法), 1235
- `setcomptype()` (*wave.Wave_write* 方法), 1237
- `setContentHandler()`
(*xml.sax.xmlreader.XMLReader* 方法), 1072
- `setcontext()` (在 *decimal* 模块中), 283
- `setDaemon()` (*threading.Thread* 方法), 712
- `setdefault()` (*dict* 方法), 73
- `setdefault()` (*http.cookies.Morsel* 方法), 1191
- `setdefaulttimeout()` (在 *socket* 模块中), 799
- `setdlopenflags()` (在 *sys* 模块中), 1554
- `setDocumentLocator()`
(*xml.sax.handler.ContentHandler* 方法), 1067
- `setDTDHandler()` (*xml.sax.xmlreader.XMLReader* 方法), 1072
- `setegid()` (在 *os* 模块中), 514
- `setEncoding()` (*xml.sax.xmlreader.InputSource* 方法), 1074
- `setEntityResolver()`
(*xml.sax.xmlreader.XMLReader* 方法), 1072
- `setErrorHandler()` (*xml.sax.xmlreader.XMLReader* 方法), 1072
- `seteuid()` (在 *os* 模块中), 514
- `setFeature()` (*xml.sax.xmlreader.XMLReader* 方法), 1073
- `setfirstweekday()` (在 *calendar* 模块中), 193
- `setfmt()` (*ossaudiodev.oss_audio_device* 方法), 1243
- `setFormatter()` (*logging.Handler* 方法), 610
- `setframerate()` (*aifc.aifc* 方法), 1232
- `setframerate()` (*sunau.AU_write* 方法), 1235
- `setframerate()` (*wave.Wave_write* 方法), 1237
- `setgid()` (在 *os* 模块中), 514
- `setgroups()` (在 *os* 模块中), 514
- `seth()` (在 *turtle* 模块中), 1269
- `setheading()` (在 *turtle* 模块中), 1269
- `sethostname()` (在 *socket* 模块中), 799
- `SetInteger()` (*msilib.Record* 方法), 1692
- `setitem()` (在 *operator* 模块中), 336
- `setitimer()` (在 *signal* 模块中), 922
- `setLevel()` (*logging.Handler* 方法), 610
- `setLevel()` (*logging.Logger* 方法), 606
- `setlocale()` (在 *locale* 模块中), 1256
- `setLocale()` (*xml.sax.xmlreader.XMLReader* 方法), 1072
- `setLoggerClass()` (在 *logging* 模块中), 619
- `setlogmask()` (在 *syslog* 模块中), 1723
- `setLogRecordFactory()` (在 *logging* 模块中), 619
- `setmark()` (*aifc.aifc* 方法), 1232
- `setMaxConns()` (*urllib.request.CacheFTPHandler* 方法), 1116
- `setmode()` (在 *msvcrt* 模块中), 1695
- `setName()` (*threading.Thread* 方法), 712
- `setnchannels()` (*aifc.aifc* 方法), 1232
- `setnchannels()` (*sunau.AU_write* 方法), 1234
- `setnchannels()` (*wave.Wave_write* 方法), 1237
- `setnframes()` (*aifc.aifc* 方法), 1232
- `setnframes()` (*sunau.AU_write* 方法), 1235
- `setnframes()` (*wave.Wave_write* 方法), 1237
- `SetParamEntityParsing()`
(*xml.parsers.expat.xmlparser* 方法), 1077
- `setparameters()` (*ossaudiodev.oss_audio_device* 方法), 1243
- `setparams()` (*aifc.aifc* 方法), 1232
- `setparams()` (*sunau.AU_write* 方法), 1235
- `setparams()` (*wave.Wave_write* 方法), 1237
- `setpassword()` (*zipfile.ZipFile* 方法), 446
- `setpgid()` (在 *os* 模块中), 514
- `setpgrp()` (在 *os* 模块中), 514
- `setpos()` (*aifc.aifc* 方法), 1231
- `setpos()` (*sunau.AU_read* 方法), 1234
- `setpos()` (在 *turtle* 模块中), 1268
- `setpos()` (*wave.Wave_read* 方法), 1236
- `setposition()` (在 *turtle* 模块中), 1268
- `setpriority()` (在 *os* 模块中), 514
- `setprofile()` (在 *sys* 模块中), 1554
- `setprofile()` (在 *threading* 模块中), 710
- `SetProperty()` (*msilib.SummaryInformation* 方法), 1691
- `setProperty()` (*xml.sax.xmlreader.XMLReader* 方法), 1073
- `setPublicId()` (*xml.sax.xmlreader.InputSource* 方法), 1073
- `setquota()` (*imaplib.IMAP4* 方法), 1151
- `setraw()` (在 *tty* 模块中), 1713
- `setrecursionlimit()` (在 *sys* 模块中), 1554
- `setregid()` (在 *os* 模块中), 514
- `setresgid()` (在 *os* 模块中), 514
- `setresuid()` (在 *os* 模块中), 515
- `setreuid()` (在 *os* 模块中), 515
- `setrlimit()` (在 *resource* 模块中), 1719
- `setsampwidth()` (*aifc.aifc* 方法), 1232
- `setsampwidth()` (*sunau.AU_write* 方法), 1234

- `setsampwidth()` (*wave.Wave_write* 方法), 1237
- `setscrreg()` (*curses.window* 方法), 656
- `setsid()` (在 *os* 模块中), 515
- `setsockopt()` (*socket.socket* 方法), 805
- `setstate()` (*codecs.IncrementalDecoder* 方法), 151
- `setstate()` (*codecs.IncrementalEncoder* 方法), 150
- `setstate()` (在 *random* 模块中), 301
- `SetStream()` (*msilib.Record* 方法), 1692
- `SetString()` (*msilib.Record* 方法), 1692
- `setswitchinterval()` (在 *sys* 模块中), 1555
- `setSystemId()` (*xml.sax.xmlreader.InputSource* 方法), 1074
- `setsyx()` (在 *curses* 模块中), 649
- `setTarget()` (*logging.handlers.MemoryHandler* 方法), 641
- `settiltangle()` (在 *turtle* 模块中), 1280
- `settimeout()` (*socket.socket* 方法), 805
- `setTimeout()` (*urllib.request.CacheFTPHandler* 方法), 1116
- `settrace()` (在 *sys* 模块中), 1555
- `settrace()` (在 *threading* 模块中), 710
- `setuid()` (在 *os* 模块中), 515
- `setundobuffer()` (在 *turtle* 模块中), 1283
- `setup()` (*socketserver.BaseRequestHandler* 方法), 1180
- `setUp()` (*unittest.TestCase* 方法), 1400
- `setup()` (在 *turtle* 模块中), 1290
- `SETUP_ANNOTATIONS` (*opcode*), 1675
- `SETUP_ASYNC_WITH` (*opcode*), 1674
- `setup_envIRON()` (*wsgiref.handlers.BaseHandler* 方法), 1102
- `SETUP_EXCEPT` (*opcode*), 1678
- `SETUP_FINALLY` (*opcode*), 1678
- `SETUP_LOOP` (*opcode*), 1678
- `setup_python()` (*venv.EnvBuilder* 方法), 1531
- `setup_scripts()` (*venv.EnvBuilder* 方法), 1531
- `setup_testing_defaults()` (在 *wsgiref.util* 模块中), 1096
- `SETUP_WITH` (*opcode*), 1675
- `setUpClass()` (*unittest.TestCase* 方法), 1400
- `setupterm()` (在 *curses* 模块中), 649
- `SetValue()` (在 *winreg* 模块中), 1700
- `SetValueEx()` (在 *winreg* 模块中), 1700
- `setworldcoordinates()` (在 *turtle* 模块中), 1285
- `setx()` (在 *turtle* 模块中), 1269
- `setxattr()` (在 *os* 模块中), 541
- `sety()` (在 *turtle* 模块中), 1269
- `SF_APPEND` (在 *stat* 模块中), 368
- `SF_ARCHIVED` (在 *stat* 模块中), 368
- `SF_IMMUTABLE` (在 *stat* 模块中), 368
- `SF_MNOWAIT` (在 *os* 模块中), 521
- `SF_NODISKIO` (在 *os* 模块中), 521
- `SF_NOUNLINK` (在 *stat* 模块中), 368
- `SF_SNAPSHOT` (在 *stat* 模块中), 368
- `SF_SYNC` (在 *os* 模块中), 521
- `shape` (*memoryview* 属性), 68
- `Shape` (*turtle* 中的类), 1291
- `shape()` (在 *turtle* 模块中), 1279
- `shapettransform()` (在 *turtle* 模块中), 1280
- `share()` (*socket.socket* 方法), 805
- `shared_ciphers()` (*ssl.SSLSocket* 方法), 824
- `shearfactor()` (在 *turtle* 模块中), 1279
- `Shelf` (*shelve* 中的类), 402
- `shelve`
 - 模块, 403
- `shelve` (模块), 401
- `shield()` (在 *asyncio* 模块中), 877
- `shift()` (*decimal.Context* 方法), 288
- `shift()` (*decimal.Decimal* 方法), 282
- `shift_path_info()` (在 *wsgiref.util* 模块中), 1095
- `shifting`
 - operations, 30
- `shlex` (*shlex* 中的类), 1301
- `shlex` (模块), 1301
- `shortDescription()` (*unittest.TestCase* 方法), 1407
- `shorten()` (在 *textwrap* 模块中), 127
- `shouldFlush()` (*logging.handlers.BufferingHandler* 方法), 641
- `shouldFlush()` (*logging.handlers.MemoryHandler* 方法), 641
- `shouldStop` (*unittest.TestResult* 属性), 1412
- `show()` (*curses.panel.Panel* 方法), 666
- `show_code()` (在 *dis* 模块中), 1669
- `showsyntaxerror()` (*code.InteractiveInterpreter* 方法), 1616
- `showtraceback()` (*code.InteractiveInterpreter* 方法), 1616
- `showturtle()` (在 *turtle* 模块中), 1278
- `showwarning()` (在 *warnings* 模块中), 1567
- `shuffle()` (在 *random* 模块中), 302
- `shutdown()` (*concurrent.futures.Executor* 方法), 762
- `shutdown()` (*imaplib.IMAP4* 方法), 1151
- `shutdown()` (*multiprocessing.managers.BaseManager* 方法), 739
- `shutdown()` (*socketserver.BaseServer* 方法), 1178
- `shutdown()` (*socket.socket* 方法), 805
- `shutdown()` (在 *logging* 模块中), 619
- `shutdown_asyncgens()` (*asyncio.AbstractEventLoop* 方法), 853
- `shutil` (模块), 378
- `side_effect` (*unittest.mock.Mock* 属性), 1425
- `SIG_BLOCK` (在 *signal* 模块中), 921
- `SIG_DFL` (在 *signal* 模块中), 920
- `SIG_IGN` (在 *signal* 模块中), 920
- `SIG_SETMASK` (在 *signal* 模块中), 921
- `SIG_UNBLOCK` (在 *signal* 模块中), 921
- `siginterrupt()` (在 *signal* 模块中), 922
- `signal`

- 模块, 787
- signal (模块), 919
- signal() (在 *signal* 模块中), 923
- Signature (*inspect* 中的类), 1600
- signature (*inspect.BoundArguments* 属性), 1603
- signature() (在 *inspect* 模块中), 1600
- sigpending() (在 *signal* 模块中), 923
- sigtimedwait() (在 *signal* 模块中), 923
- sigwait() (在 *signal* 模块中), 923
- sigwaitinfo() (在 *signal* 模块中), 923
- Simple Mail Transfer Protocol, 1160
- SimpleCookie (*http.cookies* 中的类), 1190
- simplefilter() (在 *warnings* 模块中), 1567
- SimpleHandler (*wsgiref.handlers* 中的类), 1100
- SimpleHTTPRequestHandler (*http.server* 中的类), 1187
- SimpleNamespace (*types* 中的类), 233
- SimpleQueue (*multiprocessing* 中的类), 730
- SimpleXMLRPCRequestHandler (*xmlrpc.server* 中的类), 1209
- SimpleXMLRPCServer (*xmlrpc.server* 中的类), 1209
- sin() (在 *cmath* 模块中), 270
- sin() (在 *math* 模块中), 267
- single dispatch -- 单分派, 1768
- SingleAddressHeader (*email.headerregistry* 中的类), 955
- singledispatch() (在 *functools* 模块中), 330
- sinh() (在 *cmath* 模块中), 271
- sinh() (在 *math* 模块中), 268
- SIO_KEEPAIVE_VALS() (在 *socket* 模块中), 794
- SIO_LOOPBACK_FAST_PATH() (在 *socket* 模块中), 794
- SIO_RCVALL() (在 *socket* 模块中), 794
- site (模块), 1610
- site 命令行选项
 - user-base, 1612
 - user-site, 1612
- sitecustomize
 - 模块, 1611
- site-packages
 - directory, 1610
- sixtofour (*ipaddress.IPv6Address* 属性), 1217
- size (*struct.Struct* 属性), 144
- size (*tarfile.TarInfo* 属性), 456
- size (*tracemalloc.Statistic* 属性), 1523
- size (*tracemalloc.StatisticDiff* 属性), 1523
- size (*tracemalloc.Trace* 属性), 1523
- size() (*ftplib.FTP* 方法), 1144
- size() (*mmap.mmap* 方法), 927
- size_diff (*tracemalloc.StatisticDiff* 属性), 1523
- Sized (*collections.abc* 中的类), 211
- Sized (*typing* 中的类), 1359
- sizeof() (在 *ctypes* 模块中), 702
- skip() (*chunk.Chunk* 方法), 1239
- SKIP() (在 *doctest* 模块中), 1376
- skip() (在 *unittest* 模块中), 1398
- skip_unless_symlink() (在 *test.support* 模块中), 1483
- skipIf() (在 *unittest* 模块中), 1398
- skipinitialspace (*csv.Dialect* 属性), 465
- skipped (*unittest.TestResult* 属性), 1412
- skippedEntity() (*xml.sax.handler.ContentHandler* 方法), 1069
- SkipTest, 1398
- skipTest() (*unittest.TestCase* 方法), 1400
- skipUnless() (在 *unittest* 模块中), 1398
- SLASH() (在 *token* 模块中), 1657
- SLASHEQUAL() (在 *token* 模块中), 1657
- slave() (*nntplib.NNTP* 方法), 1159
- sleep() (在 *asyncio* 模块中), 877
- sleep() (在 *time* 模块中), 568
- slice
 - assignment, 37
 - operation, 35
 - Ⓕ置函数, 1679
- slice -- 切片, 1768
- slice (Ⓕ置类), 20
- SMTP
 - protocol, 1160
- SMTP (*smtplib* 中的类), 1160
- SMTP() (在 *email.policy* 模块中), 950
- smtp_server (*smtpd.SMTPChannel* 属性), 1169
- SMTP_SSL (*smtplib* 中的类), 1161
- smtp_state (*smtpd.SMTPChannel* 属性), 1169
- SMTPAuthenticationError, 1162
- SMTPChannel (*smtpd* 中的类), 1168
- SMTPConnectError, 1162
- smtpd (模块), 1167
- SMTPDataError, 1162
- SMTPException, 1161
- SMTPHandler (*logging.handlers* 中的类), 640
- SMTPHeloError, 1162
- smtplib (模块), 1160
- SMTPNotSupportedError, 1162
- SMTPRecipientsRefused, 1162
- SMTPResponseException, 1161
- SMTPSenderRefused, 1162
- SMTPServer (*smtpd* 中的类), 1167
- SMTPServerDisconnected, 1161
- SMTPUTF8() (在 *email.policy* 模块中), 950
- Snapshot (*tracemalloc* 中的类), 1521
- SND_ALIAS() (在 *winsound* 模块中), 1705
- SND_ASYNC() (在 *winsound* 模块中), 1705
- SND_FILENAME() (在 *winsound* 模块中), 1704
- SND_LOOP() (在 *winsound* 模块中), 1705
- SND_MEMORY() (在 *winsound* 模块中), 1705
- SND_NODEFAULT() (在 *winsound* 模块中), 1705
- SND_NOSTOP() (在 *winsound* 模块中), 1705

- SND_NOWAIT() (在 *winsound* 模块中), 1705
 SND_PURGE() (在 *winsound* 模块中), 1705
 sndhdr (模块), 1240
 sniff() (*csv.Sniffer* 方法), 464
 Sniffer (*csv* 中的类), 464
 sock_accept() (*asyncio.AbstractEventLoop* 方法), 859
 SOCK_CLOEXEC() (在 *socket* 模块中), 792
 sock_connect() (*asyncio.AbstractEventLoop* 方法), 859
 SOCK_DGRAM() (在 *socket* 模块中), 792
 SOCK_NONBLOCK() (在 *socket* 模块中), 792
 SOCK_RAW() (在 *socket* 模块中), 792
 SOCK_RDM() (在 *socket* 模块中), 792
 sock_recv() (*asyncio.AbstractEventLoop* 方法), 859
 sock_sendall() (*asyncio.AbstractEventLoop* 方法), 859
 SOCK_SEQPACKET() (在 *socket* 模块中), 792
 SOCK_STREAM() (在 *socket* 模块中), 792
 socket
 对象, 789
 模块, 1085
 socket (*socketserver.BaseServer* 属性), 1178
 socket (模块), 789
 socket() (*imaplib.IMAP4* 方法), 1151
 socket() (in module *socket*), 843
 socket() (在 *socket* 模块中), 795
 socket_type (*socketserver.BaseServer* 属性), 1179
 SocketHandler (*logging.handlers* 中的类), 636
 socketpair() (在 *socket* 模块中), 795
 sockets (*asyncio.Server* 属性), 862
 socketserver (模块), 1176
 SocketType() (在 *socket* 模块中), 796
 SOL_ALG() (在 *socket* 模块中), 794
 SOL_RDS() (在 *socket* 模块中), 794
 SOMAXCONN() (在 *socket* 模块中), 793
 sort() (*imaplib.IMAP4* 方法), 1151
 sort() (*list* 方法), 37
 sort_stats() (*pstats.Stats* 方法), 1504
 sorted() (内置函数), 20
 --sort-keys
 json.tool 命令行选项, 995
 sortTestMethodsUsing (*unittest.TestLoader* 属性), 1411
 source (*doctest.Example* 属性), 1384
 source (*pdb* command), 1498
 source (*shlex.shlex* 属性), 1303
 source_from_cache() (在 *imp* 模块中), 1754
 source_from_cache() (在 *importlib.util* 模块中), 1640
 SOURCE_SUFFIXES() (在 *importlib.machinery* 模块中), 1635
 source_to_code() (*importlib.abc.InspectLoader* 静态方法), 1633
 SourceFileLoader (*importlib.machinery* 中的类), 1637
 sourcehook() (*shlex.shlex* 方法), 1302
 SourcelessFileLoader (*importlib.machinery* 中的类), 1637
 SourceLoader (*importlib.abc* 中的类), 1634
 space
 in printf-style formatting, 48, 61
 in string formatting, 93
 span() (*re.match* 方法), 111
 spawn() (在 *pty* 模块中), 1714
 spawnl() (在 *os* 模块中), 545
 spawnle() (在 *os* 模块中), 545
 spawnlp() (在 *os* 模块中), 545
 spawnlpe() (在 *os* 模块中), 545
 spawnv() (在 *os* 模块中), 545
 spawnve() (在 *os* 模块中), 545
 spawnvp() (在 *os* 模块中), 545
 spawnvpe() (在 *os* 模块中), 545
 spec_from_file_location() (在 *importlib.util* 模块中), 1641
 spec_from_loader() (在 *importlib.util* 模块中), 1641
 special method -- 特殊方法, 1768
 specified_attributes
 (*xml.parsers.expat.xmlparser* 属性), 1077
 speed() (*ossaudiodev.oss_audio_device* 方法), 1243
 speed() (在 *turtle* 模块中), 1272
 split() (*bytearray* 方法), 55
 split() (*bytes* 方法), 55
 split() (*re.regex* 方法), 108
 split() (*str* 方法), 45
 split() (在 *os.path* 模块中), 361
 split() (在 *re* 模块中), 105
 split() (在 *shlex* 模块中), 1301
 splitdrive() (在 *os.path* 模块中), 361
 splitext() (在 *os.path* 模块中), 361
 splitlines() (*bytearray* 方法), 58
 splitlines() (*bytes* 方法), 58
 splitlines() (*str* 方法), 45
 SplitResult (*urllib.parse* 中的类), 1127
 SplitResultBytes (*urllib.parse* 中的类), 1127
 splitunc() (在 *os.path* 模块中), 361
 SpooledTemporaryFile() (在 *tempfile* 模块中), 372
 sprintf-style formatting, 48, 60
 spwd (模块), 1709
 sqlite3 (模块), 408
 sqlite_version() (在 *sqlite3* 模块中), 410
 sqlite_version_info() (在 *sqlite3* 模块中), 410
 sqrt() (*decimal.Context* 方法), 288
 sqrt() (*decimal.Decimal* 方法), 282
 sqrt() (在 *cmath* 模块中), 270
 sqrt() (在 *math* 模块中), 267

- SSL, 810
- ssl (模块), 810
- SSL_CERT_FILE, 841
- SSL_CERT_PATH, 841
- ssl_version (*fiplib.FTP_TLS* 属性), 1144
- SSLContext (*ssl* 中的类), 826
- SSLEOFError, 811
- SSLError, 811
- SSLErrorNumber (*ssl* 中的类), 822
- SSLObject (*ssl* 中的类), 838
- SSLSession (*ssl* 中的类), 839
- SSLSocket (*ssl* 中的类), 822
- SSLSyscallError, 811
- SSLWantReadError, 811
- SSLWantWriteError, 811
- SSLZeroReturnError, 811
- st () (在 *turtle* 模块中), 1278
- st2list () (在 *parser* 模块中), 1647
- st2tuple () (在 *parser* 模块中), 1647
- st_atime (*os.stat_result* 属性), 534
- ST_ETIME () (在 *stat* 模块中), 366
- st_atime_ns (*os.stat_result* 属性), 534
- st_birthtime (*os.stat_result* 属性), 535
- st_blksize (*os.stat_result* 属性), 535
- st_blocks (*os.stat_result* 属性), 535
- st_creator (*os.stat_result* 属性), 535
- st_ctime (*os.stat_result* 属性), 534
- ST_CTIME () (在 *stat* 模块中), 366
- st_ctime_ns (*os.stat_result* 属性), 534
- st_dev (*os.stat_result* 属性), 534
- ST_DEV () (在 *stat* 模块中), 365
- st_file_attributes (*os.stat_result* 属性), 535
- st_flags (*os.stat_result* 属性), 535
- st_gen (*os.stat_result* 属性), 535
- st_gid (*os.stat_result* 属性), 534
- ST_GID () (在 *stat* 模块中), 365
- st_ino (*os.stat_result* 属性), 534
- ST_INO () (在 *stat* 模块中), 365
- st_mode (*os.stat_result* 属性), 534
- ST_MODE () (在 *stat* 模块中), 365
- st_mtime (*os.stat_result* 属性), 534
- ST_MTIME () (在 *stat* 模块中), 366
- st_mtime_ns (*os.stat_result* 属性), 534
- st_nlink (*os.stat_result* 属性), 534
- ST_NLINK () (在 *stat* 模块中), 365
- st_rdev (*os.stat_result* 属性), 535
- st_rsize (*os.stat_result* 属性), 535
- st_size (*os.stat_result* 属性), 534
- ST_SIZE () (在 *stat* 模块中), 366
- st_type (*os.stat_result* 属性), 535
- st_uid (*os.stat_result* 属性), 534
- ST_UID () (在 *stat* 模块中), 365
- stack (*traceback.TracebackException* 属性), 1587
- stack viewer, 1342
- stack () (在 *inspect* 模块中), 1607
- stack_effect () (在 *dis* 模块中), 1671
- stack_size () (在 *_thread* 模块中), 786
- stack_size () (在 *threading* 模块中), 710
- stackable
 - streams, 145
- StackSummary (*traceback* 中的类), 1588
- stamp () (在 *turtle* 模块中), 1271
- standard_b64decode () (在 *base64* 模块中), 1017
- standard_b64encode () (在 *base64* 模块中), 1017
- standarderror (2to3 fixer), 1477
- standend () (*curses.window* 方法), 656
- standout () (*curses.window* 方法), 656
- STAR () (在 *token* 模块中), 1657
- STAREQUAL () (在 *token* 模块中), 1657
- starmap () (*multiprocessing.pool.Pool* 方法), 746
- starmap () (在 *itertools* 模块中), 321
- starmap_async () (*multiprocessing.pool.Pool* 方法), 747
- start (*range* 属性), 39
- start (*UnicodeError* 属性), 84
- start () (*logging.handlers.QueueListener* 方法), 643
- start () (*multiprocessing.managers.BaseManager* 方法), 739
- start () (*multiprocessing.Process* 方法), 726
- start () (*re.match* 方法), 111
- start () (*threading.Thread* 方法), 711
- start () (*tkinter.ttk.Progressbar* 方法), 1325
- start () (在 *tracemalloc* 模块中), 1520
- start () (*xml.etree.ElementTree.TreeBuilder* 方法), 1045
- start_color () (在 *curses* 模块中), 649
- start_component () (*msilib.Directory* 方法), 1693
- start_new_thread () (在 *_thread* 模块中), 786
- start_server () (在 *asyncio* 模块中), 889
- start_unix_server () (在 *asyncio* 模块中), 889
- StartCdataSectionHandler ()
 - (*xml.parsers.expat.xmlparser* 方法), 1079
- StartDoctypeDeclHandler ()
 - (*xml.parsers.expat.xmlparser* 方法), 1078
- startDocument () (*xml.sax.handler.ContentHandler* 方法), 1067
- startElement () (*xml.sax.handler.ContentHandler* 方法), 1068
- StartElementHandler ()
 - (*xml.parsers.expat.xmlparser* 方法), 1078
- startElementNS () (*xml.sax.handler.ContentHandler* 方法), 1068
- STARTF_USESHOWWINDOW () (在 *subprocess* 模块中), 776
- STARTF_USESTDHANDLES () (在 *subprocess* 模块中), 776
- startfile () (在 *os* 模块中), 547

- StartNamespaceDeclHandler()
(*xml.parsers.expat.xmlparser* 方法), 1079
- startPrefixMapping()
(*xml.sax.handler.ContentHandler* 方法), 1068
- startswith() (*bytearray* 方法), 54
- startswith() (*bytes* 方法), 54
- startswith() (*str* 方法), 46
- startTest() (*unittest.TestResult* 方法), 1413
- startTestRun() (*unittest.TestResult* 方法), 1413
- starttls() (*imaplib.IMAP4* 方法), 1152
- starttls() (*nntplib.NNTP* 方法), 1156
- starttls() (*smtpplib.SMTP* 方法), 1164
- STARTUPINFO (*subprocess* 中的类), 775
- stat
模块, 533
- stat (模块), 364
- stat() (*nntplib.NNTP* 方法), 1158
- stat() (*os.DirEntry* 方法), 533
- stat() (*pathlib.Path* 方法), 351
- stat() (*poplib.POP3* 方法), 1146
- stat() (在 *os* 模块中), 533
- stat_float_times() (在 *os* 模块中), 536
- stat_result (*os* 中的类), 534
- state() (*tkinter.ttk.Widget* 方法), 1321
- statement -- 语句, 1768
- staticmethod() (☐置函数), 20
- Statistic (*tracemalloc* 中的类), 1522
- StatisticDiff (*tracemalloc* 中的类), 1523
- statistics (模块), 306
- statistics() (*tracemalloc.Snapshot* 方法), 1522
- StatisticsError, 312
- Stats (*pstats* 中的类), 1503
- status (*http.client.HTTPResponse* 属性), 1138
- status() (*imaplib.IMAP4* 方法), 1152
- statvfs() (在 *os* 模块中), 536
- STD_ERROR_HANDLE() (在 *subprocess* 模块中), 776
- STD_INPUT_HANDLE() (在 *subprocess* 模块中), 776
- STD_OUTPUT_HANDLE() (在 *subprocess* 模块中), 776
- StdButtonBox (*tkinter.tix* 中的类), 1337
- stderr (*asyncio.asyncio.subprocess.Process* 属性), 898
- stderr (*subprocess.CalledProcessError* 属性), 769
- stderr (*subprocess.CompletedProcess* 属性), 768
- stderr (*subprocess.Popen* 属性), 775
- stderr (*subprocess.TimeoutExpired* 属性), 769
- stderr() (在 *sys* 模块中), 1556
- stdev() (在 *statistics* 模块中), 311
- stdin (*asyncio.asyncio.subprocess.Process* 属性), 898
- stdin (*subprocess.Popen* 属性), 774
- stdin() (在 *sys* 模块中), 1556
- stdout (*asyncio.asyncio.subprocess.Process* 属性), 898
- stdout (*subprocess.CalledProcessError* 属性), 769
- stdout (*subprocess.CompletedProcess* 属性), 768
- stdout (*subprocess.Popen* 属性), 774
- stdout (*subprocess.TimeoutExpired* 属性), 769
- STDOUT() (在 *subprocess* 模块中), 768
- stdout() (在 *sys* 模块中), 1556
- step (*pdb command*), 1497
- step (*range* 属性), 39
- step() (*tkinter.ttk.Progressbar* 方法), 1325
- stereocontrols() (*ossaudiodev.oss_mixer_device* 方法), 1245
- stls() (*poplib.POP3* 方法), 1146
- stop (*range* 属性), 39
- stop() (*asyncio.AbstractEventLoop* 方法), 853
- stop() (*logging.handlers.QueueListener* 方法), 643
- stop() (*tkinter.ttk.Progressbar* 方法), 1325
- stop() (*unittest.TestResult* 方法), 1412
- stop() (在 *tracemalloc* 模块中), 1520
- stop_here() (*bdb.Bdb* 方法), 1489
- StopAsyncIteration, 83
- StopIteration, 82
- stopListening() (在 *logging.config* 模块中), 623
- stopTest() (*unittest.TestResult* 方法), 1413
- stopTestRun() (*unittest.TestResult* 方法), 1413
- storbinary() (*ftplib.FTP* 方法), 1142
- store() (*imaplib.IMAP4* 方法), 1152
- STORE_ACTIONS (*optparse.Option* 属性), 1750
- STORE_ANNOTATION (*opcode*), 1678
- STORE_ATTR (*opcode*), 1676
- STORE_DEREF (*opcode*), 1678
- STORE_FAST (*opcode*), 1678
- STORE_GLOBAL (*opcode*), 1676
- STORE_NAME (*opcode*), 1676
- STORE_SUBSCR (*opcode*), 1673
- storlines() (*ftplib.FTP* 方法), 1142
- str (built-in class)
(see also string), 40
- str (☐置类), 40
- str() (在 *locale* 模块中), 1259
- strcoll() (在 *locale* 模块中), 1259
- StreamError, 452
- StreamHandler (*logging* 中的类), 631
- StreamReader (*asyncio* 中的类), 889
- StreamReader (*codecs* 中的类), 152
- streamreader (*codecs.CodecInfo* 属性), 145
- StreamReaderProtocol (*asyncio* 中的类), 891
- StreamReaderWriter (*codecs* 中的类), 153
- StreamRecoder (*codecs* 中的类), 153
- StreamRequestHandler (*socketserver* 中的类), 1180
- streams, 145
stackable, 145
- StreamWriter (*asyncio* 中的类), 890
- StreamWriter (*codecs* 中的类), 151
- streamwriter (*codecs.CodecInfo* 属性), 145
- strerror (*OSError* 属性), 82
- strerror() (在 *os* 模块中), 515
- strftime() (*datetime.date* 方法), 169

- `strftime()` (*datetime.datetime* 方法), 177
- `strftime()` (*datetime.time* 方法), 181
- `strftime()` (在 *time* 模块中), 569
- `strict` (*csv.Dialect* 属性), 465
- `strict()` (在 *email.policy* 模块中), 950
- `strict_domain` (*http.cookiejar.DefaultCookiePolicy* 属性), 1198
- `strict_errors()` (在 *codecs* 模块中), 149
- `strict_ns_domain` (*http.cookiejar.DefaultCookiePolicy* 属性), 1198
- `strict_ns_set_initial_dollar` (*http.cookiejar.DefaultCookiePolicy* 属性), 1198
- `strict_ns_set_path` (*http.cookiejar.DefaultCookiePolicy* 属性), 1199
- `strict_ns_unverifiable` (*http.cookiejar.DefaultCookiePolicy* 属性), 1198
- `strict_rfc2965_unverifiable` (*http.cookiejar.DefaultCookiePolicy* 属性), 1198
- `strides` (*memoryview* 属性), 68
- `string`
 - `format()` (*built-in function*), 11
 - `formatting`, `printf`, 48
 - `interpolation`, `printf`, 48
 - `methods`, 41
 - `str` (*built-in class*), 40
 - `str()` (*built-in function*), 20
 - text sequence type, 40
 - 对象, 40
 - 模块, 1260
- `string` (*re.match* 属性), 111
- `string` (模块), 89
- `STRING()` (在 *token* 模块中), 1657
- `string_at()` (在 *ctypes* 模块中), 702
- `StringIO` (*io* 中的类), 564
- `stringprep` (模块), 132
- `strip()` (*bytearray* 方法), 56
- `strip()` (*bytes* 方法), 56
- `strip()` (*str* 方法), 46
- `strip_dirs()` (*pstats.Stats* 方法), 1503
- `stripspaces` (*curses.textpad.Textbox* 属性), 662
- `strptime()` (*datetime.datetime* 类方法), 172
- `strptime()` (在 *time* 模块中), 570
- `struct`
 - 模块, 805
- `Struct` (*struct* 中的类), 144
- `struct` (模块), 139
- `struct sequence` -- 结构序列, 1768
- `struct_time` (*time* 中的类), 570
- `Structure` (*ctypes* 中的类), 705
- `structures`
 - C, 139
- `strxfrm()` (在 *locale* 模块中), 1259
- `STType()` (在 *parser* 模块中), 1648
- `Style` (*tkinter.ttk* 中的类), 1332
- `sub()` (*re.regex* 方法), 109
- `sub()` (在 *operator* 模块中), 335
- `sub()` (在 *re* 模块中), 106
- `subdirs` (*filecmp.dircmp* 属性), 370
- `SubElement()` (在 *xml.etree.ElementTree* 模块中), 1040
- `submit()` (*concurrent.futures.Executor* 方法), 762
- `submodule_search_locations` (*importlib.machinery.ModuleSpec* 属性), 1639
- `subn()` (*re.regex* 方法), 109
- `subn()` (在 *re* 模块中), 107
- `subnets()` (*ipaddress.IPv4Network* 方法), 1220
- `subnets()` (*ipaddress.IPv6Network* 方法), 1222
- `Subnormal` (*decimal* 中的类), 291
- `suboffsets` (*memoryview* 属性), 69
- `subpad()` (*curses.window* 方法), 656
- `subprocess` (模块), 767
- `subprocess_exec()` (*asyncio.AbstractEventLoop* 方法), 896
- `subprocess_shell()` (*asyncio.AbstractEventLoop* 方法), 896
- `SubprocessError`, 768
- `SubprocessProtocol` (*asyncio* 中的类), 882
- `subscribe()` (*imaplib.IMAP4* 方法), 1152
- `subscript`
 - assignment, 37
 - operation, 35
- `subsequent_indent` (*textwrap.TextWrapper* 属性), 130
- `substitute()` (*string.Template* 方法), 97
- `subTest()` (*unittest.TestCase* 方法), 1401
- `subtract()` (*collections.Counter* 方法), 198
- `subtract()` (*decimal.Context* 方法), 288
- `subtype` (*email.headerregistry.ContentTypeHeader* 属性), 955
- `subwin()` (*curses.window* 方法), 656
- `successful()` (*multiprocessing.pool.AsyncResult* 方法), 747
- `suffix_map` (*mimetypes.MimeTypes* 属性), 1016
- `suffix_map()` (在 *mimetypes* 模块中), 1015
- `suite()` (在 *parser* 模块中), 1646
- `suiteClass` (*unittest.TestLoader* 属性), 1411
- `sum()` (`☐`置函数), 20
- `summarize()` (*doctest.DocTestRunner* 方法), 1387
- `summarize_address_range()` (在 *ipaddress* 模块中), 1225
- `sunau` (模块), 1232
- `super` (*pyclbr.Class* 属性), 1663
- `super()` (`☐`置函数), 21
- `supernet()` (*ipaddress.IPv4Network* 方法), 1221

`supernet()` (*ipaddress.IPv6Network* 方法), 1222
`supports_bytes_environ()` (在 *os* 模块中), 515
`supports_dir_fd()` (在 *os* 模块中), 536
`supports_effective_ids()` (在 *os* 模块中), 537
`supports_fd()` (在 *os* 模块中), 537
`supports_follow_symlinks()` (在 *os* 模块中), 537
`supports_unicode_filenames()` (在 *os.path* 模块中), 361
`SupportsAbs` (*typing* 中的类), 1359
`SupportsBytes` (*typing* 中的类), 1359
`SupportsComplex` (*typing* 中的类), 1359
`SupportsFloat` (*typing* 中的类), 1359
`SupportsInt` (*typing* 中的类), 1359
`SupportsRound` (*typing* 中的类), 1359
`suppress()` (在 *contextlib* 模块中), 1569
`SuppressCrashReport` (*test.support* 中的类), 1485
`SW_HIDE()` (在 *subprocess* 模块中), 776
`swapcase()` (*bytearray* 方法), 59
`swapcase()` (*bytes* 方法), 59
`swapcase()` (*str* 方法), 47
`sym_name()` (在 *symbol* 模块中), 1657
`Symbol` (*symtable* 中的类), 1656
`symbol` (模块), 1657
`SymbolTable` (*symtable* 中的类), 1655
`symlink()` (在 *os* 模块中), 537
`symlink_to()` (*pathlib.Path* 方法), 356
`symmetric_difference()` (*frozenset* 方法), 70
`symmetric_difference_update()` (*frozenset* 方法), 71
`symtable` (模块), 1654
`symtable()` (在 *symtable* 模块中), 1655
`sync()` (*dbm.dumb.dumbdbm* 方法), 408
`sync()` (*dbm.gnu.gdbm* 方法), 407
`sync()` (*ossaudiodev.oss_audio_device* 方法), 1243
`sync()` (*shelve.Shelf* 方法), 401
`sync()` (在 *os* 模块中), 538
`syncdown()` (*curses.window* 方法), 656
`synchronized()` (在 *multiprocessing.sharedctypes* 模块中), 738
`SyncManager` (*multiprocessing.managers* 中的类), 740
`syncok()` (*curses.window* 方法), 656
`syncup()` (*curses.window* 方法), 656
`SyntaxErr`, 1056
`SyntaxError`, 83
`SyntaxWarning`, 86
`sys`
 模块, 17
`sys` (模块), 1543
`sys_exc` (*2to3 fixer*), 1477
`sys_version` (*http.server.BaseHTTPRequestHandler* 属性), 1185
`sysconf()` (在 *os* 模块中), 552
`sysconf_names()` (在 *os* 模块中), 552

`sysconfig` (模块), 1558
`syslog` (模块), 1723
`syslog()` (在 *syslog* 模块中), 1723
`SysLogHandler` (*logging.handlers* 中的类), 637
`system()` (在 *os* 模块中), 547
`system()` (在 *platform* 模块中), 667
`system_alias()` (在 *platform* 模块中), 667
`SystemError`, 83
`SystemExit`, 83
`systemId` (*xml.dom.DocumentType* 属性), 1051
`SystemRandom` (*random* 中的类), 303
`SystemRandom` (*secrets* 中的类), 505
`SystemRoot`, 772

T

`-t` <tarfile>
 tarfile 命令行选项, 458
`-t` <zipfile>
 zipfile 命令行选项, 450
`-t`, `--time`
 timeit 命令行选项, 1510
`-t`, `--top-level-directory` directory
 unittest-discover 命令行选项, 1394
`-t`, `--trace`
 trace 命令行选项, 1513
`-T`, `--trackcalls`
 trace 命令行选项, 1513
`T_FMT()` (在 *locale* 模块中), 1257
`T_FMT_AMPM()` (在 *locale* 模块中), 1257
`tab()` (*tkinter.ttk.Notebook* 方法), 1324
`TabError`, 83
`tabnanny` (模块), 1662
`tabs()` (*tkinter.ttk.Notebook* 方法), 1324
`tabsize` (*textwrap.TextWrapper* 属性), 129
`tabular`
 data, 461
`tag` (*xml.etree.ElementTree.Element* 属性), 1041
`tag_bind()` (*tkinter.ttk.Treeview* 方法), 1331
`tag_configure()` (*tkinter.ttk.Treeview* 方法), 1331
`tag_has()` (*tkinter.ttk.Treeview* 方法), 1331
`tagName` (*xml.dom.Element* 属性), 1053
`tail` (*xml.etree.ElementTree.Element* 属性), 1041
`take_snapshot()` (在 *tracemalloc* 模块中), 1520
`takewhile()` (在 *itertools* 模块中), 321
`tan()` (在 *cmath* 模块中), 270
`tan()` (在 *math* 模块中), 267
`tanh()` (在 *cmath* 模块中), 271
`tanh()` (在 *math* 模块中), 268
`TarError`, 452
`TarFile` (*tarfile* 中的类), 452, 453
`tarfile` (模块), 451
`tarfile` 命令行选项
 `-c` <tarfile> <source1> ...
 <sourceN>, 458

```

--create <tarfile> <source1> ...
    <sourceN>, 458
-e <tarfile> [<output_dir>], 458
--extract <tarfile> [<output_dir>],
    458
-l <tarfile>, 458
--list <tarfile>, 458
-t <tarfile>, 458
--test <tarfile>, 458
-v, --verbose, 458
target (xml.dom.ProcessingInstruction 属性), 1055
TarInfo (tarfile 中的类), 456
Task (asyncio 中的类), 873
task_done() (asyncio.Queue 方法), 906
task_done() (multiprocessing.JoinableQueue 方法),
    730
task_done() (queue.Queue 方法), 784
tau() (在 cmath 模块中), 272
tau() (在 math 模块中), 268
tb_locals (unittest.TestResult 属性), 1412
tbreak (pdb command), 1497
tcdrain() (在 termios 模块中), 1713
tcflow() (在 termios 模块中), 1713
tcflush() (在 termios 模块中), 1713
tcgetattr() (在 termios 模块中), 1712
tcgetpgrp() (在 os 模块中), 522
Tcl() (在 tkinter 模块中), 1308
TCPServer (socketserver 中的类), 1176
tcsendbreak() (在 termios 模块中), 1713
tcsetattr() (在 termios 模块中), 1712
tcsetpgrp() (在 os 模块中), 522
tearDown() (unittest.TestCase 方法), 1400
tearDownClass() (unittest.TestCase 方法), 1400
tee() (在 itertools 模块中), 322
tell() (aifc.aifc 方法), 1231, 1232
tell() (chunk.Chunk 方法), 1239
tell() (io.IOBase 方法), 558
tell() (io.TextIOBase 方法), 563
tell() (mmap.mmap 方法), 927
tell() (sunau.AU_read 方法), 1234
tell() (sunau.AU_write 方法), 1235
tell() (wave.Wave_read 方法), 1236
tell() (wave.Wave_write 方法), 1237
Telnet (telnetlib 中的类), 1170
telnetlib (模块), 1170
TEMP, 373
temp_cwd() (在 test.support 模块中), 1482
temp_dir() (在 test.support 模块中), 1482
temp_umask() (在 test.support 模块中), 1483
tempdir() (在 tempfile 模块中), 373
tempfile (模块), 371
Template (pipes 中的类), 1717
Template (string 中的类), 97
template (string.Template 属性), 97
temporary
    file, 371
    file name, 371
TemporaryDirectory() (在 tempfile 模块中), 372
TemporaryFile() (在 tempfile 模块中), 371
teredo (ipaddress.IPv6Address 属性), 1217
TERM, 649
termattrs() (在 curses 模块中), 649
terminal_size (os 中的类), 523
terminate() (asyncio.asyncio.subprocess.Process 方
    法), 898
terminate() (asyncio.BaseSubprocessTransport 方法),
    882
terminate() (multiprocessing.pool.Pool 方法), 747
terminate() (multiprocessing.Process 方法), 727
terminate() (subprocess.Popen 方法), 774
termios (模块), 1712
termname() (在 curses 模块中), 649
test (doctest.DocTestFailure 属性), 1390
test (doctest.UnexpectedException 属性), 1390
test (模块), 1477
--test <tarfile>
    tarfile 命令行选项, 458
test() (在 cgi 模块中), 1091
TestCase (unittest 中的类), 1399
TestFailed, 1480
testfile() (在 doctest 模块中), 1379
TESTFN() (在 test.support 模块中), 1480
TestLoader (unittest 中的类), 1409
testMethodPrefix (unittest.TestLoader 属性), 1411
testmod() (在 doctest 模块中), 1380
TestResult (unittest 中的类), 1411
tests() (在 imghdr 模块中), 1240
testsource() (在 doctest 模块中), 1389
testsRun (unittest.TestResult 属性), 1412
TestSuite (unittest 中的类), 1409
test.support (模块), 1480
testzip() (zipfile.ZipFile 方法), 447
text (traceback.TracebackException 属性), 1587
Text (typing 中的类), 1362
text (xml.etree.ElementTree.Element 属性), 1041
text encoding -- 文本编码, 1768
text file -- 文本文件, 1768
text mode, 17
text() (msilib.Dialog 方法), 1694
text() (在 msilib 模块中), 1694
text_factory (sqlite3.Connection 属性), 416
Textbox (curses.textpad 中的类), 661
TextCalendar (calendar 中的类), 193
textdomain() (在 gettext 模块中), 1248
textdomain() (在 locale 模块中), 1261
textinput() (在 turtle 模块中), 1287
TextIO (typing 中的类), 1362
TextIOBase (io 中的类), 562

```

- TextIOWrapper (*io* 中的类), 563
- TextTestResult (*unittest* 中的类), 1414
- TextTestRunner (*unittest* 中的类), 1414
- textwrap (模块), 127
- TextWrapper (*textwrap* 中的类), 129
- theme_create() (*tkinter.ttk.Style* 方法), 1334
- theme_names() (*tkinter.ttk.Style* 方法), 1334
- theme_settings() (*tkinter.ttk.Style* 方法), 1334
- theme_use() (*tkinter.ttk.Style* 方法), 1334
- THOUSEP() (在 *locale* 模块中), 1258
- Thread (*threading* 中的类), 711
- thread() (*imaplib.IMAP4* 方法), 1152
- thread_info() (在 *sys* 模块中), 1557
- threading (模块), 709
- ThreadingMixIn (*socketserver* 中的类), 1177
- ThreadingTCPServer (*socketserver* 中的类), 1177
- ThreadingUDPServer (*socketserver* 中的类), 1177
- ThreadPoolExecutor (*concurrent.futures* 中的类), 763
- threads
- POSIX, 786
- throw (2to3 fixer), 1477
- ticket_lifetime_hint (*ssl.SSLSession* 属性), 839
- tigetflag() (在 *curses* 模块中), 649
- tigetnum() (在 *curses* 模块中), 649
- tigetstr() (在 *curses* 模块中), 650
- TILDE() (在 *token* 模块中), 1657
- tilt() (在 *turtle* 模块中), 1280
- tiltangle() (在 *turtle* 模块中), 1280
- time (*datetime* 中的类), 179
- time (*ssl.SSLSession* 属性), 839
- time (模块), 566
- time() (*asyncio.AbstractEventLoop* 方法), 854
- time() (*datetime.datetime* 方法), 174
- time() (在 *time* 模块中), 570
- Time2Internaldate() (在 *imaplib* 模块中), 1148
- timedelta (*datetime* 中的类), 165
- TimedRotatingFileHandler (*logging.handlers* 中的类), 634
- timegm() (在 *calendar* 模块中), 194
- timeit (模块), 1507
- timeit 命令行选项
- c, --clock, 1510
 - h, --help, 1510
 - n N, --number=N, 1510
 - p, --process, 1510
 - r N, --repeat=N, 1510
 - s S, --setup=S, 1510
 - t, --time, 1510
 - u, --unit=U, 1510
 - v, --verbose, 1510
- timeit() (*timeit.Timer* 方法), 1509
- timeit() (在 *timeit* 模块中), 1508
- timeout, 792
- timeout (*socketserver.BaseServer* 属性), 1179
- timeout (*ssl.SSLSession* 属性), 839
- timeout (*subprocess.TimeoutExpired* 属性), 768
- timeout() (*curses.window* 方法), 656
- TIMEOUT_MAX() (在 *_thread* 模块中), 786
- TIMEOUT_MAX() (在 *threading* 模块中), 710
- TimeoutError, 85, 728, 766, 871
- TimeoutExpired, 768
- Timer (*threading* 中的类), 718
- Timer (*timeit* 中的类), 1508
- times() (在 *os* 模块中), 547
- timestamp() (*datetime.datetime* 方法), 175
- timetuple() (*datetime.date* 方法), 169
- timetuple() (*datetime.datetime* 方法), 175
- timetz() (*datetime.datetime* 方法), 174
- timezone (*datetime* 中的类), 188
- timezone() (在 *time* 模块中), 573
- title() (*bytearray* 方法), 59
- title() (*bytes* 方法), 59
- title() (*str* 方法), 47
- title() (在 *turtle* 模块中), 1290
- Tix, 1335
- tix_addbitmapdir() (*tkinter.tix.tixCommand* 方法), 1339
- tix_cget() (*tkinter.tix.tixCommand* 方法), 1339
- tix_configure() (*tkinter.tix.tixCommand* 方法), 1339
- tix_filedialog() (*tkinter.tix.tixCommand* 方法), 1339
- tix_getbitmap() (*tkinter.tix.tixCommand* 方法), 1339
- tix_getimage() (*tkinter.tix.tixCommand* 方法), 1339
- TIX_LIBRARY, 1336
- tix_option_get() (*tkinter.tix.tixCommand* 方法), 1339
- tix_resetoptions() (*tkinter.tix.tixCommand* 方法), 1339
- tixCommand (*tkinter.tix* 中的类), 1338
- Tk, 1307
- Tk (*tkinter* 中的类), 1308
- Tk (*tkinter.tix* 中的类), 1335
- Tk Option Data Types, 1315
- Tkinter, 1307
- tkinter (模块), 1307
- tkinter.scrolledtext (模块), 1340
- tkinter.tix (模块), 1335
- tkinter.ttk (模块), 1318
- TList (*tkinter.tix* 中的类), 1338
- TLS, 810
- TMP, 373
- TMPDIR, 373
- to_bytes() (*int* 方法), 31
- to_eng_string() (*decimal.Context* 方法), 289
- to_eng_string() (*decimal.Decimal* 方法), 282

- `to_integral()` (*decimal.Decimal* 方法), 282
- `to_integral_exact()` (*decimal.Context* 方法), 289
- `to_integral_exact()` (*decimal.Decimal* 方法), 283
- `to_integral_value()` (*decimal.Decimal* 方法), 283
- `to_sci_string()` (*decimal.Context* 方法), 289
- `ToASCII()` (在 *encodings.idna* 模块中), 161
- `tobuf()` (*tarfile.TarInfo* 方法), 456
- `tobytes()` (*array.array* 方法), 222
- `tobytes()` (*memoryview* 方法), 65
- `today()` (*datetime.date* 类方法), 167
- `today()` (*datetime.datetime* 类方法), 171
- `tofile()` (*array.array* 方法), 222
- `tok_name()` (在 *token* 模块中), 1657
- `token` (*shlex.shlex* 属性), 1304
- `token` (模块), 1657
- `token_bytes()` (在 *secrets* 模块中), 505
- `token_hex()` (在 *secrets* 模块中), 506
- `token_urlsafe()` (在 *secrets* 模块中), 506
- `TokenError`, 1660
- `tokenize` (模块), 1659
- `tokenize` 命令行选项
 - `-e`, `--exact`, 1661
 - `-h`, `--help`, 1660
- `tokenize()` (在 *tokenize* 模块中), 1659
- `tolist()` (*array.array* 方法), 222
- `tolist()` (*memoryview* 方法), 65
- `tolist()` (*parser.ST* 方法), 1648
- `tomono()` (在 *audioop* 模块中), 1229
- `toordinal()` (*datetime.date* 方法), 169
- `toordinal()` (*datetime.datetime* 方法), 175
- `top()` (*curses.panel.Panel* 方法), 666
- `top()` (*poplib.POP3* 方法), 1146
- `top_panel()` (在 *curses.panel* 模块中), 665
- `toprettyxml()` (*xml.dom.minidom.Node* 方法), 1059
- `tostereo()` (在 *audioop* 模块中), 1229
- `tostring()` (*array.array* 方法), 222
- `tostring()` (在 *xml.etree.ElementTree* 模块中), 1040
- `tostringlist()` (在 *xml.etree.ElementTree* 模块中), 1040
- `total_changes` (*sqlite3.Connection* 属性), 416
- `total_ordering()` (在 *functools* 模块中), 328
- `total_seconds()` (*datetime.timedelta* 方法), 167
- `totuple()` (*parser.ST* 方法), 1648
- `touch()` (*pathlib.Path* 方法), 356
- `touchline()` (*curses.window* 方法), 656
- `touchwin()` (*curses.window* 方法), 657
- `tounicode()` (*array.array* 方法), 222
- `ToUnicode()` (在 *encodings.idna* 模块中), 161
- `towards()` (在 *turtle* 模块中), 1272
- `toxml()` (*xml.dom.minidom.Node* 方法), 1059
- `tparm()` (在 *curses* 模块中), 650
- `Trace` (*trace* 中的类), 1514
- `Trace` (*tracemalloc* 中的类), 1523
- `trace` (模块), 1512
- `trace function`, 710, 1549, 1555
- `trace` 命令行选项
 - `-c`, `--count`, 1513
 - `-C`, `--coverdir=<dir>`, 1513
 - `-f`, `--file=<file>`, 1513
 - `-g`, `--timing`, 1513
 - `--help`, 1512
 - `--ignore-dir=<dir>`, 1513
 - `--ignore-module=<mod>`, 1513
 - `-l`, `--listfuncs`, 1513
 - `-m`, `--missing`, 1513
 - `-R`, `--no-report`, 1513
 - `-r`, `--report`, 1513
 - `-s`, `--summary`, 1513
 - `-t`, `--trace`, 1513
 - `-T`, `--trackcalls`, 1513
 - `--version`, 1512
- `trace()` (在 *inspect* 模块中), 1607
- `trace_dispatch()` (*bdb.Bdb* 方法), 1488
- `traceback`
 - 对象, 1545, 1585
- `Traceback` (*tracemalloc* 中的类), 1523
- `traceback` (*tracemalloc.Statistic* 属性), 1523
- `traceback` (*tracemalloc.StatisticDiff* 属性), 1523
- `traceback` (*tracemalloc.Trace* 属性), 1523
- `traceback` (模块), 1585
- `traceback_limit` (*tracemalloc.Snapshot* 属性), 1522
- `traceback_limit` (*wsgiref.handlers.BaseHandler* 属性), 1102
- `TracebackException` (*traceback* 中的类), 1587
- `tracebacklimit()` (在 *sys* 模块中), 1557
- `tracebacks`
 - in CGI scripts, 1094
- `TracebackType()` (在 *types* 模块中), 232
- `tracemalloc` (模块), 1515
- `tracer()` (在 *turtle* 模块中), 1285
- `traces` (*tracemalloc.Snapshot* 属性), 1522
- `transfercmd()` (*ftplib.FTP* 方法), 1142
- `TransientResource` (*test.support* 中的类), 1485
- `translate()` (*bytearray* 方法), 54
- `translate()` (*bytes* 方法), 54
- `translate()` (*str* 方法), 47
- `translate()` (在 *fnmatch* 模块中), 377
- `translation()` (在 *gettext* 模块中), 1249
- `transport` (*asyncio.StreamWriter* 属性), 891
- Transport Layer Security, 810
- `Tree` (*tkinter.tix* 中的类), 1337
- `TreeBuilder` (*xml.etree.ElementTree* 中的类), 1044
- `Treeview` (*tkinter.ttk* 中的类), 1328
- `triangular()` (在 *random* 模块中), 303
- triple-quoted string -- 三引号字符串, 1768
- `True`, 27, 77
- `true`, 27
- `True` (赋值变量), 25

- truediv() (在 *operator* 模块中), 335
 - trunc() (in module *math*), 29
 - trunc() (在 *math* 模块中), 266
 - truncate() (*io.IOBase* 方法), 558
 - truncate() (在 *os* 模块中), 538
 - truth
 - value, 27
 - truth() (在 *operator* 模块中), 334
 - try
 - 语句, 79
 - ttk, 1318
 - tty
 - I/O control, 1712
 - tty (模块), 1713
 - ttyname() (在 *os* 模块中), 522
 - tuple
 - 对象, 36, 38
 - tuple (☐置类), 38
 - Tuple() (在 *typing* 模块中), 1365
 - tuple2st() (在 *parser* 模块中), 1647
 - tuple_params (2to3 fixer), 1477
 - turnoff_sigfpe() (在 *fpectl* 模块中), 1613
 - turnon_sigfpe() (在 *fpectl* 模块中), 1613
 - Turtle (*turtle* 中的类), 1290
 - turtle (模块), 1263
 - turtledemo (模块), 1294
 - turtles() (在 *turtle* 模块中), 1289
 - TurtleScreen (*turtle* 中的类), 1290
 - turtlesize() (在 *turtle* 模块中), 1279
 - type
 - Boolean, 6
 - operations on dictionary, 71
 - operations on list, 37
 - ☐置函数, 76
 - 对象, 21
 - type -- 类型, 1769
 - type (*optparse.Option* 属性), 1738
 - type (*socket.socket* 属性), 806
 - type (*tarfile.TarInfo* 属性), 456
 - Type (*typing* 中的类), 1358
 - type (*urllib.request.Request* 属性), 1109
 - type (☐置类), 21
 - type alias -- 类型别名, 1769
 - type hint -- 类型提示, 1769
 - TYPE_CHECKER (*optparse.Option* 属性), 1749
 - TYPE_CHECKING() (在 *typing* 模块中), 1366
 - typeahead() (在 *curses* 模块中), 650
 - typecode (*array.array* 属性), 221
 - typecodes() (在 *array* 模块中), 221
 - TYPED_ACTIONS (*optparse.Option* 属性), 1750
 - typed_subpart_iterator() (在 *email.iterators* 模块中), 985
 - TypeError, 83
 - types
 - built-in, 27
 - immutable sequence, 36
 - mutable sequence, 36
 - operations on integer, 30
 - operations on mapping, 71
 - operations on numeric, 29
 - operations on sequence, 35, 37
 - 模块, 76
 - types (2to3 fixer), 1477
 - TYPES (*optparse.Option* 属性), 1749
 - types (模块), 230
 - types_map (*mimetypes.MimeTypes* 属性), 1016
 - types_map() (在 *mimetypes* 模块中), 1015
 - types_map_inv (*mimetypes.MimeTypes* 属性), 1016
 - TypeVar (*typing* 中的类), 1357
 - typing (模块), 1351
 - TZ, 571, 572
 - tzinfo (*datetime* 中的类), 182
 - tzinfo (*datetime.datetime* 属性), 173
 - tzinfo (*datetime.time* 属性), 180
 - tzname() (*datetime.datetime* 方法), 175
 - tzname() (*datetime.time* 方法), 181
 - tzname() (*datetime.timezone* 方法), 189
 - tzname() (*datetime.tzinfo* 方法), 183
 - tzname() (在 *time* 模块中), 573
 - tzset() (在 *time* 模块中), 571
- ## U
- u, --unit=U
 - timeit 命令行选项, 1510
 - ucd_3_2_0() (在 *unicodedata* 模块中), 132
 - udata (*select.kevent* 属性), 848
 - UDPServer (*socketserver* 中的类), 1176
 - UF_APPEND() (在 *stat* 模块中), 368
 - UF_COMPRESSED() (在 *stat* 模块中), 368
 - UF_HIDDEN() (在 *stat* 模块中), 368
 - UF_IMMUTABLE() (在 *stat* 模块中), 368
 - UF_NODUMP() (在 *stat* 模块中), 368
 - UF_NOUNLINK() (在 *stat* 模块中), 368
 - UF_OPAQUE() (在 *stat* 模块中), 368
 - uid (*tarfile.TarInfo* 属性), 456
 - uid() (*imaplib.IMAP4* 方法), 1153
 - uidl() (*poplib.POP3* 方法), 1146
 - u-LAW, 1227, 1232, 1240
 - ulaw2lin() (在 *audioop* 模块中), 1229
 - umask() (在 *os* 模块中), 515
 - unalias (*pdb* command), 1499
 - uname (*tarfile.TarInfo* 属性), 457
 - uname() (在 *os* 模块中), 515
 - uname() (在 *platform* 模块中), 667
 - UNARY_INVERT (*opcode*), 1672
 - UNARY_NEGATIVE (*opcode*), 1672
 - UNARY_NOT (*opcode*), 1672
 - UNARY_POSITIVE (*opcode*), 1672

- UnboundLocalError, 83
- unbuffered I/O, 17
- UNC paths
 - and `os.makedirs()`, 528
- `unconsumed_tail` (`zlib.Decompress` 属性), 431
- `unctrl()` (在 `curses` 模块中), 650
- `unctrl()` (在 `curses.ascii` 模块中), 664
- Underflow (`decimal` 中的类), 291
- `undisplay` (`pdb` command), 1499
- `undo()` (在 `turtle` 模块中), 1271
- `undobufferentries()` (在 `turtle` 模块中), 1283
- `undoc_header` (`cmd.Cmd` 属性), 1298
- `unescape()` (在 `html` 模块中), 1025
- `unescape()` (在 `xml.sax.saxutils` 模块中), 1070
- UnexpectedException, 1390
- `unexpectedSuccesses` (`unittest.TestResult` 属性), 1412
- `unget_wch()` (在 `curses` 模块中), 650
- `ungetch()` (在 `curses` 模块中), 650
- `ungetch()` (在 `msvcrt` 模块中), 1696
- `ungetmouse()` (在 `curses` 模块中), 650
- `ungetwch()` (在 `msvcrt` 模块中), 1696
- `unhexlify()` (在 `binascii` 模块中), 1021
- Unicode, 130, 145
 - database, 130
- `unicode (2to3 fixer)`, 1477
- `unicodedata` (模块), 130
- `UnicodeDecodeError`, 84
- `UnicodeEncodeError`, 84
- `UnicodeError`, 83
- `UnicodeTranslateError`, 84
- `UnicodeWarning`, 86
- `unidata_version()` (在 `unicodedata` 模块中), 132
- `unified_diff()` (在 `difflib` 模块中), 120
- `uniform()` (在 `random` 模块中), 303
- `UnimplementedFileMode`, 1134
- `Union` (`ctypes` 中的类), 705
- `union()` (`frozenset` 方法), 70
- `Union()` (在 `typing` 模块中), 1364
- `unique()` (在 `enum` 模块中), 242, 245
- `unittest` (模块), 1391
- `unittest` 命令行选项
 - `-b`, `--buffer`, 1393
 - `-c`, `--catch`, 1393
 - `-f`, `--failfast`, 1394
 - `--locals`, 1394
- `unittest-discover` 命令行选项
 - `-p`, `--pattern` pattern, 1394
 - `-s`, `--start-directory` directory, 1394
 - `-t`, `--top-level-directory` directory, 1394
 - `-v`, `--verbose`, 1394
- `unittest.mock` (模块), 1418
- universal newlines
 - `bytearray.splitlines` method, 58
 - `bytes.splitlines` method, 58
 - `csv.reader` function, 461
 - `importlib.abc.InspectLoader.get_source` method, 1632
 - `io.IncrementalNewlineDecoder` class, 565
 - `io.TextIOWrapper` class, 564
 - `open()` built-in function, 16
 - `str.splitlines` method, 45
 - `subprocess` module, 769
- universal newlines -- 通用换行, 1769
- UNIX
 - file control, 1715
 - I/O control, 1715
- `unix_dialect` (`csv` 中的类), 464
- `UnixDatagramServer` (`socketserver` 中的类), 1176
- `UnixStreamServer` (`socketserver` 中的类), 1176
- `unknown_decl()` (`html.parser.HTMLParser` 方法), 1028
- `unknown_open()` (`urllib.request.BaseHandler` 方法), 1112
- `unknown_open()` (`urllib.request.UnknownHandler` 方法), 1116
- `UnknownHandler` (`urllib.request` 中的类), 1109
- `UnknownProtocol`, 1134
- `UnknownTransferEncoding`, 1134
- `unlink()` (`pathlib.Path` 方法), 356
- `unlink()` (在 `os` 模块中), 538
- `unlink()` (`xml.dom.minidom.Node` 方法), 1059
- `unlock()` (`mailbox.Babyl` 方法), 1004
- `unlock()` (`mailbox.Mailbox` 方法), 999
- `unlock()` (`mailbox.Maildir` 方法), 1001
- `unlock()` (`mailbox.mbox` 方法), 1001
- `unlock()` (`mailbox.MH` 方法), 1003
- `unlock()` (`mailbox.MMDF` 方法), 1004
- `unpack()` (`struct.Struct` 方法), 144
- `unpack()` (在 `struct` 模块中), 140
- `unpack_archive()` (在 `shutil` 模块中), 384
- `unpack_array()` (`xdrlib.Unpacker` 方法), 488
- `unpack_bytes()` (`xdrlib.Unpacker` 方法), 488
- `unpack_double()` (`xdrlib.Unpacker` 方法), 488
- `UNPACK_EX` (`opcode`), 1676
- `unpack_farray()` (`xdrlib.Unpacker` 方法), 488
- `unpack_float()` (`xdrlib.Unpacker` 方法), 488
- `unpack_fopaque()` (`xdrlib.Unpacker` 方法), 488
- `unpack_from()` (`struct.Struct` 方法), 144
- `unpack_from()` (在 `struct` 模块中), 140
- `unpack_fstring()` (`xdrlib.Unpacker` 方法), 488
- `unpack_list()` (`xdrlib.Unpacker` 方法), 488
- `unpack_opaque()` (`xdrlib.Unpacker` 方法), 488
- `UNPACK_SEQUENCE` (`opcode`), 1676
- `unpack_string()` (`xdrlib.Unpacker` 方法), 488

- Unpacker (*xdrlib* 中的类), 486
- unparsedEntityDecl()
 - (*xml.sax.handler.DTDHandler* 方法), 1069
- UnparsedEntityDeclHandler()
 - (*xml.parsers.expat.xmlparser* 方法), 1079
- Unpickler (*pickle* 中的类), 391
- UnpicklingError, 390
- unquote() (在 *email.utils* 模块中), 983
- unquote() (在 *urllib.parse* 模块中), 1128
- unquote_plus() (在 *urllib.parse* 模块中), 1128
- unquote_to_bytes() (在 *urllib.parse* 模块中), 1128
- unregister() (*select.devpoll* 方法), 844
- unregister() (*select.epoll* 方法), 845
- unregister() (*selectors.BaseSelector* 方法), 850
- unregister() (*select.poll* 方法), 846
- unregister() (在 *atexit* 模块中), 1584
- unregister() (在 *faulthandler* 模块中), 1493
- unregister_archive_format() (在 *shutil* 模块中), 384
- unregister_dialect() (在 *csv* 模块中), 462
- unregister_unpack_format() (在 *shutil* 模块中), 384
- unset() (*test.support.EnvironmentVarGuard* 方法), 1485
- unsetenv() (在 *os* 模块中), 515
- UnstructuredHeader (*email.headerregistry* 中的类), 954
- unsubscribe() (*imaplib.IMAP4* 方法), 1153
- UnsupportedOperation, 555
- until (*pdb* command), 1498
- untokenize() (在 *tokenize* 模块中), 1660
- untouchwin() (*curses.window* 方法), 657
- unused_data (*bz2.BZ2Decompressor* 属性), 437
- unused_data (*lzma.LZMADecompressor* 属性), 441
- unused_data (*zlib.Decompress* 属性), 431
- unverifiable (*urllib.request.Request* 属性), 1110
- unwrap() (*ssl.SSLSocket* 方法), 825
- unwrap() (在 *inspect* 模块中), 1605
- up (*pdb* command), 1496
- up() (在 *turtle* 模块中), 1274
- update() (*collections.Counter* 方法), 198
- update() (*dict* 方法), 73
- update() (*frozenset* 方法), 70
- update() (*hashlib.hash* 方法), 495
- update() (*hmac.HMAC* 方法), 504
- update() (*http.cookies.Morsel* 方法), 1191
- update() (*mailbox.Mailbox* 方法), 999
- update() (*mailbox.Maildir* 方法), 1000
- update() (*trace.CoverageResults* 方法), 1514
- update() (在 *turtle* 模块中), 1286
- update_authenticated() (*urllib.request.HTTPPasswordMgrWithPriorAuth* 方法), 1114
- update_lines_cols() (在 *curses* 模块中), 650
- update_panels() (在 *curses.panel* 模块中), 665
- update_visible() (*mailbox.BabylMessage* 方法), 1010
- update_wrapper() (在 *functools* 模块中), 332
- upper() (*bytearray* 方法), 60
- upper() (*bytes* 方法), 60
- upper() (*str* 方法), 47
- urandom() (在 *os* 模块中), 553
- URL, 1087, 1122, 1130, 1184
 - parsing, 1122
 - relative, 1122
- url (*xmlrpc.client.ProtocolError* 属性), 1206
- url2pathname() (在 *urllib.request* 模块中), 1106
- urllibcleanup() (在 *urllib.request* 模块中), 1120
- urldefrag() (在 *urllib.parse* 模块中), 1126
- urlencode() (在 *urllib.parse* 模块中), 1128
- URLError, 1129
- urljoin() (在 *urllib.parse* 模块中), 1125
- urllib (*2to3* fixer), 1477
- urllib (模块), 1104
- urllib.error (模块), 1129
- urllib.parse (模块), 1122
- urllib.request
 - 模块, 1133
- urllib.request (模块), 1104
- urllib.response (模块), 1122
- urllib.robotparser (模块), 1130
- urlopen() (在 *urllib.request* 模块中), 1104
- URLopener (*urllib.request* 中的类), 1120
- urlparse() (在 *urllib.parse* 模块中), 1122
- urlretrieve() (在 *urllib.request* 模块中), 1119
- urlsafe_b64decode() (在 *base64* 模块中), 1017
- urlsafe_b64encode() (在 *base64* 模块中), 1017
- urlsplit() (在 *urllib.parse* 模块中), 1124
- urlunparse() (在 *urllib.parse* 模块中), 1124
- urlunsplit() (在 *urllib.parse* 模块中), 1125
- urn (*uuid.UUID* 属性), 1174
- use_default_colors() (在 *curses* 模块中), 650
- use_env() (在 *curses* 模块中), 650
- use_rawinput (*cmd.Cmd* 属性), 1298
- UseForeignDTD() (*xml.parsers.expat.xmlparser* 方法), 1077
- USER, 644
- user
 - effective id, 512
 - id, 513
 - id, setting, 515
- user() (*poplib.POP3* 方法), 1146
- USER_BASE() (在 *site* 模块中), 1612
- user_call() (*bdb.Bdb* 方法), 1489
- user_exception() (*bdb.Bdb* 方法), 1489
- user_line() (*bdb.Bdb* 方法), 1489
- user_return() (*bdb.Bdb* 方法), 1489
- USER_SITE() (在 *site* 模块中), 1611

--user-base
 site 命令行选项, 1612
 usercustomize
 模块, 1611
 UserDict (*collections* 中的类), 209
 UserList (*collections* 中的类), 209
 USERNAME, 512, 644
 username (*email.headerregistry.Address* 属性), 957
 USERPROFILE, 358
 userptr () (*curses.panel.Panel* 方法), 666
 --user-site
 site 命令行选项, 1612
 UserString (*collections* 中的类), 209
 UserWarning, 85
 USTAR_FORMAT () (在 *tarfile* 模块中), 453
 UTC, 566
 utc (*datetime.timezone* 属性), 189
 utcfromtimestamp () (*datetime.datetime* 类方法), 172
 utcnow () (*datetime.datetime* 类方法), 171
 utcoffset () (*datetime.datetime* 方法), 175
 utcoffset () (*datetime.time* 方法), 181
 utcoffset () (*datetime.timezone* 方法), 189
 utcoffset () (*datetime.tzinfo* 方法), 182
 utctimetuple () (*datetime.datetime* 方法), 175
 utf8 (*email.policy.EmailPolicy* 属性), 949
 utf8 () (*poplib.POP3* 方法), 1146
 utf8_enabled (*imaplib.IMAP4* 属性), 1153
 utime () (在 *os* 模块中), 538
 uu
 模块, 1020
 uu (模块), 1022
 UUID (*uuid* 中的类), 1173
 uuid (模块), 1173
 uuid1, 1174
 uuid1 () (在 *uuid* 模块中), 1174
 uuid3, 1174
 uuid3 () (在 *uuid* 模块中), 1174
 uuid4, 1174
 uuid4 () (在 *uuid* 模块中), 1174
 uuid5, 1174
 uuid5 () (在 *uuid* 模块中), 1174
 UuidCreate () (在 *msilib* 模块中), 1689

V

-v, --verbose
 tarfile 命令行选项, 458
 timeit 命令行选项, 1510
 unittest-discover 命令行选项, 1394
 v4_int_to_packed () (在 *ipaddress* 模块中), 1225
 v6_int_to_packed () (在 *ipaddress* 模块中), 1225
 validator () (在 *wsgiref.validate* 模块中), 1099
 value
 truth, 27

value (*ctypes._SimpleCData* 属性), 703
 value (*http.cookiejar.Cookie* 属性), 1199
 value (*http.cookies.Morsel* 属性), 1191
 value (*xml.dom.Attr* 属性), 1054
 Value () (*multiprocessing.managers.SyncManager* 方法), 741
 Value () (在 *multiprocessing* 模块中), 736
 Value () (在 *multiprocessing.sharedctypes* 模块中), 737
 value_decode () (*http.cookies.BaseCookie* 方法), 1190
 value_encode () (*http.cookies.BaseCookie* 方法), 1190
 ValueError, 84
 valuerefs () (*weakref.WeakValueDictionary* 方法), 225
 values
 Boolean, 77
 values () (*dict* 方法), 73
 values () (*email.message.EmailMessage* 方法), 933
 values () (*email.message.Message* 方法), 970
 values () (*mailbox.Mailbox* 方法), 998
 values () (*types.MappingProxyType* 方法), 233
 ValuesView (*collections.abc* 中的类), 212
 ValuesView (*typing* 中的类), 1360
 variable annotation -- 变量标注, 1769
 variance () (在 *statistics* 模块中), 311
 variant (*uuid.UUID* 属性), 1174
 vars () (设置函数), 22
 vbar (*tkinter.scrolledtext.ScrolledText* 属性), 1340
 VBAR () (在 *token* 模块中), 1657
 VBAREQUAL () (在 *token* 模块中), 1657
 设置函数
 compile, 76, 231, 1647
 complex, 29
 eval, 76, 236, 237, 1647
 exec, 10, 76, 1647
 float, 29
 hash, 36
 int, 29
 len, 35, 71
 max, 35
 min, 35
 slice, 1679
 type, 76
 Vec2D (*turtle* 中的类), 1291
 venv (模块), 1527
 VERBOSE () (在 *re* 模块中), 104
 verbose () (在 *tabnanny* 模块中), 1663
 verbose () (在 *test.support* 模块中), 1480
 verify () (*smtpplib.SMTP* 方法), 1163
 verify_client_post_handshake ()
 (*ssl.SSLSocket* 方法), 825
 VERIFY_CRL_CHECK_CHAIN () (在 *ssl* 模块中), 817
 VERIFY_CRL_CHECK_LEAF () (在 *ssl* 模块中), 817

VERIFY_DEFAULT() (在 *ssl* 模块中), 817
 verify_flags (*ssl.SSLContext* 属性), 832
 verify_mode (*ssl.SSLContext* 属性), 832
 verify_request() (*socketserver.BaseServer* 方法), 1179
 VERIFY_X509_STRICT() (在 *ssl* 模块中), 817
 VERIFY_X509_TRUSTED_FIRST() (在 *ssl* 模块中), 817
 VerifyFlags (*ssl* 中的类), 818
 VerifyMode (*ssl* 中的类), 817
 --version
 trace 命令行选项, 1512
 version (*email.headerregistry.MIMEVersionHeader* 属性), 955
 version (*http.client.HTTPResponse* 属性), 1138
 version (*http.cookiejar.Cookie* 属性), 1199
 version (*ipaddress.IPv4Address* 属性), 1215
 version (*ipaddress.IPv4Network* 属性), 1219
 version (*ipaddress.IPv6Address* 属性), 1217
 version (*ipaddress.IPv6Network* 属性), 1222
 version (*urllib.request.URLopener* 属性), 1120
 version (*uuid.UUID* 属性), 1174
 version() (*ssl.SSLSocket* 方法), 825
 version() (在 *curses* 模块中), 657
 version() (在 *ensurepip* 模块中), 1527
 version() (在 *marshal* 模块中), 404
 version() (在 *platform* 模块中), 667
 version() (在 *sqlite3* 模块中), 410
 version() (在 *sys* 模块中), 1558
 version_info() (在 *sqlite3* 模块中), 410
 version_info() (在 *sys* 模块中), 1558
 version_string() (*http.server.BaseHTTPRequestHandler* 方法), 1187
 vformat() (*string.Formatter* 方法), 90
 virtual
 Environments, 1527
 virtual environment -- 虚拟环境, 1769
 virtual machine -- 虚拟机, 1769
 visit() (*ast.NodeVisitor* 方法), 1653
 对象
 Boolean, 29
 bytearray, 36, 49, 51
 bytes, 49
 complex number, 29
 dictionary, 71
 floating point, 29
 integer, 29
 io.StringIO, 40
 list, 36, 37
 mapping, 71
 memoryview, 49
 method, 76
 numeric, 29
 range, 38

 sequence, 35
 set, 69
 socket, 789
 string, 40
 traceback, 1545, 1585
 tuple, 36, 38
 type, 21
 vline() (*curses.window* 方法), 657
 voidcmd() (*ftplib.FTP* 方法), 1142
 volume (*zipfile.ZipInfo* 属性), 450
 vonmisesvariate() (在 *random* 模块中), 303

W

W_OK() (在 *os* 模块中), 524
 wait() (*asyncio.asyncio.subprocess.Process* 方法), 897
 wait() (*asyncio.Condition* 方法), 903
 wait() (*asyncio.Event* 方法), 902
 wait() (*multiprocessing.pool.AsyncResult* 方法), 747
 wait() (*subprocess.Popen* 方法), 773
 wait() (*threading.Barrier* 方法), 719
 wait() (*threading.Condition* 方法), 715
 wait() (*threading.Event* 方法), 717
 wait() (在 *asyncio* 模块中), 877
 wait() (在 *concurrent.futures* 模块中), 766
 wait() (在 *multiprocessing.connection* 模块中), 749
 wait() (在 *os* 模块中), 548
 wait3() (在 *os* 模块中), 549
 wait4() (在 *os* 模块中), 549
 wait_closed() (*asyncio.Server* 方法), 862
 wait_for() (*asyncio.Condition* 方法), 903
 wait_for() (*threading.Condition* 方法), 715
 wait_for() (在 *asyncio* 模块中), 878
 waitid() (在 *os* 模块中), 548
 waitpid() (在 *os* 模块中), 548
 walk() (*email.message.EmailMessage* 方法), 935
 walk() (*email.message.Message* 方法), 973
 walk() (在 *ast* 模块中), 1653
 walk() (在 *os* 模块中), 539
 walk_packages() (在 *pkgutil* 模块中), 1623
 walk_stack() (在 *traceback* 模块中), 1587
 walk_tb() (在 *traceback* 模块中), 1587
 want (*doctest.Example* 属性), 1384
 warn() (在 *warnings* 模块中), 1566
 warn_explicit() (在 *warnings* 模块中), 1567
 Warning, 85, 421
 warning() (*logging.Logger* 方法), 608
 warning() (在 *logging* 模块中), 617
 warning() (*xml.sax.handler.ErrorHandler* 方法), 1070
 warnings, 1563
 warnings (模块), 1563
 WarningsRecorder (*test.support* 中的类), 1485
 warnoptions() (在 *sys* 模块中), 1558
 wasSuccessful() (*unittest.TestResult* 方法), 1412

- WatchedFileHandler (*logging.handlers* 中的类), 632
 wave (模块), 1235
 WCONTINUED () (在 *os* 模块中), 549
 WCOREDUMP () (在 *os* 模块中), 549
 WeakKeyDictionary (*weakref* 中的类), 224
 WeakMethod (*weakref* 中的类), 225
 weakref (模块), 223
 WeakSet (*weakref* 中的类), 225
 WeakValueDictionary (*weakref* 中的类), 224
 webbrowser (模块), 1085
 weekday () (*datetime.date* 方法), 169
 weekday () (*datetime.datetime* 方法), 176
 weekday () (在 *calendar* 模块中), 194
 weekheader () (在 *calendar* 模块中), 194
 weibullvariate () (在 *random* 模块中), 303
 WEXITED () (在 *os* 模块中), 548
 WEXITSTATUS () (在 *os* 模块中), 550
 wfile (*http.server.BaseHTTPRequestHandler* 属性), 1185
 what () (在 *imghdr* 模块中), 1240
 what () (在 *sndhdr* 模块中), 1241
 whathdr () (在 *sndhdr* 模块中), 1241
 whatis (*pdb* command), 1498
 where (*pdb* command), 1496
 which () (在 *shutil* 模块中), 381
 whichdb () (在 *dbm* 模块中), 404
 while
 语句, 27
 whitespace (*shlex.shlex* 属性), 1303
 whitespace () (在 *string* 模块中), 90
 whitespace_split (*shlex.shlex* 属性), 1303
 Widget (*tkinter.ttk* 中的类), 1321
 width (*textwrap.TextWrapper* 属性), 129
 width () (在 *turtle* 模块中), 1274
 WIFCONTINUED () (在 *os* 模块中), 549
 WIFEXITED () (在 *os* 模块中), 549
 WIFSIGNALED () (在 *os* 模块中), 549
 WIFSTOPPED () (在 *os* 模块中), 549
 模块
 __main__, 1625, 1626
 _locale, 1255
 array, 49
 base64, 1020
 bdb, 1494
 binhex, 1020
 cmd, 1494
 copy, 400
 crypt, 1708
 dbm.gnu, 402
 dbm.ndbm, 402
 errno, 81
 glob, 376
 imp, 22
 math, 29, 272
 os, 1707
 pickle, 234, 400, 401, 403
 pty, 519
 pwd, 358
 pyexpat, 1075
 re, 41, 376
 shelve, 403
 signal, 787
 sitecustomize, 1611
 socket, 1085
 stat, 533
 string, 1260
 struct, 805
 sys, 17
 types, 76
 urllib.request, 1133
 usercustomize, 1611
 uu, 1020
 win32_ver () (在 *platform* 模块中), 668
 WinDLL (*ctypes* 中的类), 695
 window manager (*widgets*), 1314
 window () (*curses.panel.Panel* 方法), 666
 window_height () (在 *turtle* 模块中), 1289
 window_width () (在 *turtle* 模块中), 1289
 Windows ini file, 468
 WindowsError, 84
 WindowsPath (*pathlib* 中的类), 351
 WindowsRegistryFinder (*importlib.machinery* 中的类), 1636
 winerror (*OSError* 属性), 82
 WinError () (在 *ctypes* 模块中), 702
 WINFUNCTYPE () (在 *ctypes* 模块中), 698
 winreg (模块), 1696
 WinSock, 843
 winsound (模块), 1704
 winver () (在 *sys* 模块中), 1558
 WITH_CLEANUP_FINISH (*opcode*), 1675
 WITH_CLEANUP_START (*opcode*), 1675
 with_hostmask (*ipaddress.IPv4Interface* 属性), 1224
 with_hostmask (*ipaddress.IPv4Network* 属性), 1220
 with_hostmask (*ipaddress.IPv6Interface* 属性), 1224
 with_hostmask (*ipaddress.IPv6Network* 属性), 1222
 with_name () (*pathlib.PurePath* 方法), 350
 with_netmask (*ipaddress.IPv4Interface* 属性), 1224
 with_netmask (*ipaddress.IPv4Network* 属性), 1220
 with_netmask (*ipaddress.IPv6Interface* 属性), 1224
 with_netmask (*ipaddress.IPv6Network* 属性), 1222
 with_prefixlen (*ipaddress.IPv4Interface* 属性), 1224
 with_prefixlen (*ipaddress.IPv4Network* 属性), 1220
 with_prefixlen (*ipaddress.IPv6Interface* 属性), 1224
 with_prefixlen (*ipaddress.IPv6Network* 属性), 1222
 with_suffix () (*pathlib.PurePath* 方法), 350
 with_traceback () (*BaseException* 方法), 80

- WNOHANG() (在 *os* 模块中), 549
- WNOWAIT() (在 *os* 模块中), 548
- wordchars (*shlex.shlex* 属性), 1303
- World Wide Web, 1085, 1122, 1130
- wrap() (*textwrap.TextWrapper* 方法), 130
- wrap() (在 *textwrap* 模块中), 127
- wrap_bio() (*ssl.SSLContext* 方法), 830
- wrap_future() (在 *asyncio* 模块中), 876
- wrap_socket() (*ssl.SSLContext* 方法), 830
- wrap_socket() (在 *ssl* 模块中), 812
- wrapper() (在 *curses* 模块中), 650
- wraps() (在 *functools* 模块中), 333
- writable() (*asyncore.dispatcher* 方法), 914
- writable() (*io.IOBase* 方法), 558
- WRITABLE() (在 *tkinter* 模块中), 1318
- write() (*asyncio.StreamWriter* 方法), 891
- write() (*asyncio.WriteTransport* 方法), 881
- write() (*codecs.StreamWriter* 方法), 151
- write() (*code.InteractiveInterpreter* 方法), 1616
- write() (*configparser.ConfigParser* 方法), 483
- write() (*email.generator.BytesGenerator* 方法), 943
- write() (*email.generator.Generator* 方法), 944
- write() (*io.BufferedIOBase* 方法), 560
- write() (*io.BufferedWriter* 方法), 562
- write() (*io.RawIOBase* 方法), 558
- write() (*io.TextIOBase* 方法), 563
- write() (*mmap.mmap* 方法), 927
- write() (*ossaudiodev.oss_audio_device* 方法), 1242
- write() (*ssl.MemoryBIO* 方法), 839
- write() (*ssl.SSLSocket* 方法), 823
- write() (*telnetlib.Telnet* 方法), 1172
- write() (在 *os* 模块中), 522
- write() (在 *turtle* 模块中), 1278
- write() (*xml.etree.ElementTree.ElementTree* 方法), 1043
- write() (*zipfile.ZipFile* 方法), 447
- write_byte() (*mmap.mmap* 方法), 927
- write_bytes() (*pathlib.Path* 方法), 356
- write_docstringdict() (在 *turtle* 模块中), 1293
- write_eof() (*asyncio.StreamWriter* 方法), 891
- write_eof() (*asyncio.WriteTransport* 方法), 881
- write_eof() (*ssl.MemoryBIO* 方法), 839
- write_history_file() (在 *readline* 模块中), 134
- write_results() (*trace.CoverageResults* 方法), 1514
- write_text() (*pathlib.Path* 方法), 356
- writeall() (*ossaudiodev.oss_audio_device* 方法), 1242
- writeframes() (*aifc.aifc* 方法), 1232
- writeframes() (*sunau.AU_write* 方法), 1235
- writeframes() (*wave.Wave_write* 方法), 1237
- writeframesraw() (*aifc.aifc* 方法), 1232
- writeframesraw() (*sunau.AU_write* 方法), 1235
- writeframesraw() (*wave.Wave_write* 方法), 1237
- writeheader() (*csv.DictWriter* 方法), 466
- writelines() (*asyncio.StreamWriter* 方法), 891
- writelines() (*asyncio.WriteTransport* 方法), 881
- writelines() (*codecs.StreamWriter* 方法), 151
- writelines() (*io.IOBase* 方法), 558
- writePlist() (在 *plistlib* 模块中), 490
- writePlistToBytes() (在 *plistlib* 模块中), 490
- writepy() (*zipfile.PyZipFile* 方法), 448
- writer (*formatter.formatter* 属性), 1684
- writer() (在 *csv* 模块中), 462
- writerow() (*csv.csvwriter* 方法), 466
- writerows() (*csv.csvwriter* 方法), 466
- writestr() (*zipfile.ZipFile* 方法), 447
- WriteTransport (*asyncio* 中的类), 880
- wrdev() (在 *os* 模块中), 522
- writexml() (*xml.dom.minidom.Node* 方法), 1059
- WrongDocumentErr, 1056
- ws_comma (2to3 fixer), 1477
- wsgi_file_wrapper (*wsgiref.handlers.BaseHandler* 属性), 1103
- wsgi_multiprocess (*wsgiref.handlers.BaseHandler* 属性), 1101
- wsgi_multithread (*wsgiref.handlers.BaseHandler* 属性), 1101
- wsgi_run_once (*wsgiref.handlers.BaseHandler* 属性), 1101
- wsgiref (模块), 1095
- wsgiref.handlers (模块), 1100
- wsgiref.headers (模块), 1097
- wsgiref.simple_server (模块), 1098
- wsgiref.util (模块), 1095
- wsgiref.validate (模块), 1099
- WSGIRequestHandler (*wsgiref.simple_server* 中的类), 1098
- WSGIServer (*wsgiref.simple_server* 中的类), 1098
- wShowWindow (*subprocess.STARTUPINFO* 属性), 775
- WSTOPPED() (在 *os* 模块中), 548
- WSTOPSIG() (在 *os* 模块中), 550
- wstring_at() (在 *ctypes* 模块中), 702
- WTERMSIG() (在 *os* 模块中), 550
- WUNTRACED() (在 *os* 模块中), 549
- WWW, 1085, 1122, 1130
- server, 1087, 1184
- ## X
- x regex
- compileall 命令行选项, 1666
- X() (在 *re* 模块中), 104
- X509 certificate, 832
- X_OK() (在 *os* 模块中), 524
- xatom() (*imaplib.IMAP4* 方法), 1153
- XATTR_CREATE() (在 *os* 模块中), 542
- XATTR_REPLACE() (在 *os* 模块中), 542
- XATTR_SIZE_MAX() (在 *os* 模块中), 541
- xcor() (在 *turtle* 模块中), 1272

- XDR, 486
- xdrllib (模块), 486
- xhdr() (*nnplib.NNTP* 方法), 1159
- XHTML, 1026
- XHTML_NAMESPACE() (在 *xml.dom* 模块中), 1048
- xml (模块), 1030
- XML() (在 *xml.etree.ElementTree* 模块中), 1040
- XML_ERROR_ABORTED() (在 *xml.parsers.expat.errors* 模块中), 1083
- XML_ERROR_ASYNC_ENTITY() (在 *xml.parsers.expat.errors* 模块中), 1082
- XML_ERROR_ATTRIBUTE_EXTERNAL_ENTITY_REF() (在 *xml.parsers.expat.errors* 模块中), 1082
- XML_ERROR_BAD_CHAR_REF() (在 *xml.parsers.expat.errors* 模块中), 1082
- XML_ERROR_BINARY_ENTITY_REF() (在 *xml.parsers.expat.errors* 模块中), 1082
- XML_ERROR_CANT_CHANGE_FEATURE_ONCE_PARSING() (在 *xml.parsers.expat.errors* 模块中), 1083
- XML_ERROR_DUPLICATE_ATTRIBUTE() (在 *xml.parsers.expat.errors* 模块中), 1082
- XML_ERROR_ENTITY_DECLARED_IN_PE() (在 *xml.parsers.expat.errors* 模块中), 1083
- XML_ERROR_EXTERNAL_ENTITY_HANDLING() (在 *xml.parsers.expat.errors* 模块中), 1083
- XML_ERROR_FEATURE_REQUIRES_XML_DTD() (在 *xml.parsers.expat.errors* 模块中), 1083
- XML_ERROR_FINISHED() (在 *xml.parsers.expat.errors* 模块中), 1084
- XML_ERROR_INCOMPLETE_PE() (在 *xml.parsers.expat.errors* 模块中), 1083
- XML_ERROR_INCORRECT_ENCODING() (在 *xml.parsers.expat.errors* 模块中), 1082
- XML_ERROR_INVALID_TOKEN() (在 *xml.parsers.expat.errors* 模块中), 1082
- XML_ERROR_JUNK_AFTER_DOC_ELEMENT() (在 *xml.parsers.expat.errors* 模块中), 1082
- XML_ERROR_MISPLACED_XML_PI() (在 *xml.parsers.expat.errors* 模块中), 1082
- XML_ERROR_NO_ELEMENTS() (在 *xml.parsers.expat.errors* 模块中), 1082
- XML_ERROR_NO_MEMORY() (在 *xml.parsers.expat.errors* 模块中), 1082
- XML_ERROR_NOT_STANDALONE() (在 *xml.parsers.expat.errors* 模块中), 1083
- XML_ERROR_NOT_SUSPENDED() (在 *xml.parsers.expat.errors* 模块中), 1083
- XML_ERROR_PARAM_ENTITY_REF() (在 *xml.parsers.expat.errors* 模块中), 1082
- XML_ERROR_PARTIAL_CHAR() (在 *xml.parsers.expat.errors* 模块中), 1082
- XML_ERROR_PUBLICID() (在 *xml.parsers.expat.errors* 模块中), 1083
- XML_ERROR_RECURSIVE_ENTITY_REF() (在 *xml.parsers.expat.errors* 模块中), 1082
- XML_ERROR_SUSPEND_PE() (在 *xml.parsers.expat.errors* 模块中), 1084
- XML_ERROR_SUSPENDED() (在 *xml.parsers.expat.errors* 模块中), 1083
- XML_ERROR_SYNTAX() (在 *xml.parsers.expat.errors* 模块中), 1083
- XML_ERROR_TAG_MISMATCH() (在 *xml.parsers.expat.errors* 模块中), 1083
- XML_ERROR_TEXT_DECL() (在 *xml.parsers.expat.errors* 模块中), 1083
- XML_ERROR_UNBOUND_PREFIX() (在 *xml.parsers.expat.errors* 模块中), 1083
- XML_ERROR_UNCLOSED_CDATA_SECTION() (在 *xml.parsers.expat.errors* 模块中), 1083
- XML_ERROR_UNCLOSED_TOKEN() (在 *xml.parsers.expat.errors* 模块中), 1083
- XML_ERROR_UNDECLARING_PREFIX() (在 *xml.parsers.expat.errors* 模块中), 1083
- XML_ERROR_UNDEFINED_ENTITY() (在 *xml.parsers.expat.errors* 模块中), 1083
- XML_ERROR_UNEXPECTED_STATE() (在 *xml.parsers.expat.errors* 模块中), 1083
- XML_ERROR_UNKNOWN_ENCODING() (在 *xml.parsers.expat.errors* 模块中), 1083
- XML_ERROR_XML_DECL() (在 *xml.parsers.expat.errors* 模块中), 1083
- XML_NAMESPACE() (在 *xml.dom* 模块中), 1048
- xmlcharrefreplace_errors() (在 *codecs* 模块中), 149
- XmlDeclHandler() (*xml.parsers.expat.xmlparser* 方法), 1078
- xml.dom (模块), 1047
- xml.dom.minidom (模块), 1057
- xml.dom.pulldom (模块), 1061
- xml.etree.ElementTree (模块), 1032
- XMLFilterBase (*xml.sax.saxutils* 中的类), 1071
- XMLGenerator (*xml.sax.saxutils* 中的类), 1070
- XMLID() (在 *xml.etree.ElementTree* 模块中), 1040
- XMLNS_NAMESPACE() (在 *xml.dom* 模块中), 1048
- XMLParser (*xml.etree.ElementTree* 中的类), 1045
- xml.parsers.expat (模块), 1075
- xml.parsers.expat.errors (模块), 1082
- xml.parsers.expat.model (模块), 1081
- XMLParserType() (在 *xml.parsers.expat* 模块中), 1075
- XMLPullParser (*xml.etree.ElementTree* 中的类), 1046
- XMLReader (*xml.sax.xmlreader* 中的类), 1071
- xmlrpc.client (模块), 1201
- xmlrpc.server (模块), 1209
- xml.sax (模块), 1063
- xml.sax.handler (模块), 1065
- xml.sax.saxutils (模块), 1070
- xml.sax.xmlreader (模块), 1071

xor() (在 *operator* 模块中), 335
 xover() (*nntplib.NNTP* 方法), 1160
 xpath() (*nntplib.NNTP* 方法), 1160
 xrange (2to3 fixer), 1477
 xreadlines (2to3 fixer), 1477
 xview() (*tkinter.ttk.Treeview* 方法), 1331

Y

Y2K, 566
 ycor() (在 *turtle* 模块中), 1273
 year (*datetime.date* 属性), 168
 year (*datetime.datetime* 属性), 173
 Year 2000, 566
 Year 2038, 566
 yeardatescalendar() (*calendar.Calendar* 方法), 192
 yeardays2calendar() (*calendar.Calendar* 方法), 192
 yeardayscalendar() (*calendar.Calendar* 方法), 192
 YESEXPR() (在 *locale* 模块中), 1258
 YIELD_FROM (*opcode*), 1674
 YIELD_VALUE (*opcode*), 1674
 yiq_to_rgb() (在 *colorsys* 模块中), 1239
 yview() (*tkinter.ttk.Treeview* 方法), 1331

Z

Zen of Python -- Python 之禅, 1769
 ZeroDivisionError, 84
 zfill() (*bytearray* 方法), 60
 zfill() (*bytes* 方法), 60
 zfill() (*str* 方法), 47
 zip (2to3 fixer), 1477
 zip() (置函数), 22
 ZIP_BZIP2() (在 *zipfile* 模块中), 444
 ZIP_DEFLATED() (在 *zipfile* 模块中), 444
 zip_longest() (在 *itertools* 模块中), 322
 ZIP_LZMA() (在 *zipfile* 模块中), 444
 ZIP_STORED() (在 *zipfile* 模块中), 444
 zipapp (模块), 1535
 zipapp 命令行选项
 -h, --help, 1536
 --info, 1536
 -m <mainfn>, --main=<mainfn>, 1536
 -o <output>, --output=<output>, 1536
 -p <interpreter>,
 --python=<interpreter>, 1536
 zipfile (模块), 443
 ZipFile (*zipfile* 中的类), 444
 zipfile 命令行选项
 -c <zipfile> <source1> ...
 <sourceN>, 450
 -e <zipfile> <output_dir>, 450
 -l <zipfile>, 450

 -t <zipfile>, 450
 zipimport (模块), 1619
 zipimporter (*zipimport* 中的类), 1620
 ZipImportError, 1619
 ZipInfo (*zipfile* 中的类), 444
 zlib (模块), 429
 ZLIB_RUNTIME_VERSION() (在 *zlib* 模块中), 432
 ZLIB_VERSION() (在 *zlib* 模块中), 432