
Distlib Documentation

Release 0.3.6

Vinay Sajip

Aug 26, 2022

1	Overview	3
1.1	Distlib evolved out of <code>packaging</code>	3
1.2	What was the problem with <code>packaging</code> ?	3
1.3	How Distlib can help	4
1.4	How you can help	4
1.5	Main features	4
1.6	Python version and platform compatibility	5
1.7	Project status	5
1.8	Change log for <code>distlib</code>	6
1.8.1	0.3.6 (future)	6
1.8.2	0.3.5	6
1.8.3	0.3.4	6
1.8.4	0.3.3	7
1.8.5	0.3.2	7
1.8.6	0.3.1	8
1.8.7	0.3.0	8
1.8.8	0.2.9	9
1.8.9	0.2.8	10
1.8.10	0.2.7	10
1.8.11	0.2.6	11
1.8.12	0.2.5	11
1.8.13	0.2.4	12
1.8.14	0.2.3	12
1.8.15	0.2.2	13
1.8.16	0.2.1	13
1.8.17	0.2.0	14
1.8.18	0.1.9	15
1.8.19	0.1.8	15
1.8.20	0.1.7	16
1.8.21	0.1.6	17
1.8.22	0.1.5	17
1.8.23	0.1.4	17
1.8.24	0.1.3	18
1.8.25	0.1.2	19
1.8.26	0.1.1	20
1.8.27	0.1.0	21

1.9	Next steps	21
2	Tutorial	23
2.1	Installation	23
2.2	Testing	23
2.2.1	PYPI availability	24
2.3	First steps	24
2.3.1	Using the database API	24
2.3.1.1	Distribution paths	24
2.3.1.2	Querying a path for distributions	24
2.3.1.3	Including legacy distributions in the search results	25
2.3.1.4	Distribution properties	25
2.3.1.5	Exporting things from Distributions	26
2.3.2	Distribution dependencies	26
2.3.3	Using the locators API	27
2.3.3.1	Overview	27
2.3.3.2	Under the hood	27
2.3.4	Using the index API	29
2.3.4.1	Overview	29
2.3.4.2	Registering a project	29
2.3.4.3	Uploading a source distribution	29
2.3.4.4	Uploading binary distributions	29
2.3.4.5	Signing a distribution	30
2.3.4.6	Downloading files	30
2.3.4.7	Verifying signatures	31
2.3.4.8	Uploading documentation	31
2.3.4.9	Authentication	31
2.3.4.10	Verifying HTTPS connections	32
2.3.4.11	Saving a default configuration	33
2.3.4.12	Searching PyPI	33
2.3.5	Using the metadata and markers APIs	36
2.3.5.1	Instantiating metadata	36
2.3.5.2	Reading metadata from files and streams	36
2.3.5.3	Writing metadata to paths and streams	37
2.3.5.4	Using markers	37
2.3.6	Using the resource API	37
2.3.6.1	Access to resources in the file system	37
2.3.6.2	Access to resources in the .zip files	38
2.3.6.3	Iterating over resources	39
2.3.7	Using the scripts API	39
2.3.7.1	Specifying scripts to install	39
2.3.7.2	Wrapping callables with scripts	40
2.3.7.3	Specifying a custom executable for shebangs	41
2.3.7.4	Generating variants of a script	41
2.3.7.5	Avoiding overwriting existing scripts	42
2.3.7.6	Generating windowed scripts on Windows	42
2.3.8	Using the version API	42
2.3.8.1	Overview	42
2.3.8.2	Matching versions against constraints	43
2.3.9	Using the wheel API	44
2.3.9.1	Building wheels	44
2.3.9.2	Customising tags during build	44
2.3.9.3	Specifying a wheel's version	45
2.3.9.4	Installing from wheels	45

2.3.9.5	Verifying wheels	46
2.3.9.6	Modifying wheels	46
2.3.9.7	Mounting wheels	47
2.3.9.8	Using vanilla pip to build wheels for existing distributions on PyPI	47
2.3.10	Using the manifest API	53
2.3.10.1	The <code>include</code> directive	54
2.3.10.2	The <code>exclude</code> directive	54
2.3.10.3	The <code>global-include</code> directive	54
2.3.10.4	The <code>global-exclude</code> directive	54
2.3.10.5	The <code>recursive-include</code> directive	55
2.3.10.6	The <code>recursive-exclude</code> directive	55
2.3.10.7	The <code>graft</code> directive	55
2.3.10.8	The <code>prune</code> directive	55
2.4	Next steps	55
3	Distlib's design	57
3.1	The <code>locators</code> API	57
3.1.1	The problem we're trying to solve	57
3.1.2	A minimal solution	58
3.1.2.1	Locating distributions	58
3.1.2.2	Finding dependencies	59
3.2	The <code>index</code> API	60
3.2.1	The problem we're trying to solve	60
3.2.2	A minimal solution	60
3.3	The <code>resources</code> API	61
3.3.1	The problem we're trying to solve	61
3.3.2	A minimal solution	62
3.3.3	Dealing with the requirement for access via file system files	64
3.4	The <code>scripts</code> API	65
3.4.1	The problem we're trying to solve	65
3.4.2	A minimal solution	65
3.4.2.1	Flag formats	66
3.5	The <code>version</code> API	66
3.5.1	The problem we're trying to solve	67
3.5.2	A minimal solution	68
3.5.2.1	Versions	68
3.5.2.2	Matchers	68
3.5.2.3	Version schemes	70
3.6	The <code>wheel</code> API	70
3.6.1	The problem we're trying to solve	71
3.6.2	A minimal solution	71
3.7	Next steps	72
4	API Reference	73
4.1	The <code>distlib.database</code> package	73
4.1.1	Classes	73
4.2	The <code>distlib.resources</code> package	76
4.2.1	Attributes	76
4.2.2	Functions	77
4.2.3	Classes	77
4.3	The <code>distlib.scripts</code> package	79
4.3.1	Classes	79
4.3.2	Functions	81
4.4	The <code>distlib.locators</code> package	81

4.4.1	Classes	81
4.4.2	Functions	84
4.4.3	Variables	85
4.5	The <code>distlib.index</code> package	85
4.5.1	Classes	85
4.6	The <code>distlib.util</code> package	88
4.6.1	Classes	88
4.6.2	Functions	89
4.7	The <code>distlib.wheel</code> package	90
4.7.1	Attributes	90
4.7.2	Classes	90
4.7.3	Functions	93
4.8	Next steps	93
5	Migrating from older APIs	95
5.1	The <code>pkg_resources</code> resource API	95
5.1.1	Basic resource access	95
5.1.2	Resource extraction	95
5.1.3	Provider interface	96
5.2	The <code>pkg_resources</code> entry point API	96
	Index	99

Release:	0.3.6
----------	-------

Welcome to the documentation for `distlib`, **a library of packaging functionality which is intended to be used as the basis for third-party packaging tools**. Using a common layer will improve interoperability and consistency of user experience across those tools which use the library.

Please note: this documentation is a *work in progress*.

Start here for all things `distlib`.

1.1 Distlib evolved out of packaging

`Distlib` is a library which implements low-level functions that relate to packaging and distribution of Python software. It consists in part of the functions in the `packaging` Python package, which was intended to be released as part of Python 3.3, but was removed shortly before Python 3.3 entered beta testing.

Note: The `packaging` package referred to here does not refer to any `packaging` package currently available on PyPI, but to a package that was never released on PyPI but called `packaging` in the Python 3.3 alpha stdlib tree.

1.2 What was the problem with packaging?

The `packaging` software just wasn't ready for inclusion in the Python standard library. The amount of work needed to get it into the desired state was too great, given the number of people able to work on the project, the time they could devote to it, and the Python 3.3 release schedule.

The approach taken by `packaging` was seen to be a good one: to ensure interoperability and consistency between different tools in the packaging space by defining standards for data formats through PEPs, and to do away with the *ad hoc* nature of installation encouraged by the `distutils` approach of using executable Python code in `setup.py`. Where custom code was needed, it could be provided in a standardised way using installation hooks.

While some very good work was done in defining PEPs to codify some of the best practices, `packaging` suffered from some drawbacks, too:

- Not all the PEPs may have been functionally complete, because some important use cases were not considered – for example, built (binary) distributions for Windows.

- It continued the command-based design of `distutils`, which had resulted in `distutils` being difficult to extend in a consistent, easily understood, and maintainable fashion.
- Some important features required by distribution authors were not considered – for example:
 - Access to data files stored in Python packages.
 - Support for plug-in extension points.
 - Support for native script execution on Windows.

These features are supported by third-party tools (like `setuptools` / `Distribute`) using `pkg_resources`, *entry points* and *console scripts*.

- There were a lot of rough edges in the packaging implementation, both in terms of bugs and in terms of incompletely implemented features. This can be seen (with the benefit of hindsight) as due to the goals being set too ambitiously; the project developers bit off more than they could chew.

1.3 How Distlib can help

The idea behind Distlib is expressed in [this python-dev mailing-list post](#), though a different name was suggested for the library. Basically, Distlib contains the implementations of the packaging PEPs and other low-level features which relate to packaging, distribution, and deployment of Python software. If Distlib can be made genuinely useful, then it is possible for third-party packaging tools to transition to using it. Their developers and users then benefit from standardised implementation of low-level functions, time saved by not having to reinvent wheels, and improved interoperability between tools.

1.4 How you can help

If you have some time and the inclination to improve the state of Python packaging, then you can help by trying out Distlib, raising issues where you find problems, contributing feedback and/or patches to the implementation, documentation, and underlying PEPs.

1.5 Main features

Distlib currently offers the following features:

- The package `distlib.database`, which implements a database of installed distributions, as defined by [PEP 376](#), and distribution dependency graph logic. Support is also provided for non-installed distributions (i.e. distributions registered with metadata on an index like PyPI), including the ability to scan for dependencies and building dependency graphs.
- The package `distlib.index`, which implements an interface to perform operations on an index, such as registering a project, uploading a distribution or uploading documentation. Support is included for verifying SSL connections (with domain matching) and signing/verifying packages using GnuPG.

Note: Since this API was developed, a number of features of PyPI have been turned off for various reasons - for example, documentation uploads, XMLRPC search API - and a number of APIs have changed (e.g. PyPI no longer shows GnuPG signatures). For now, the `distlib.index` API should be considered not fully reliable, mostly due to changes in PyPI where there has not been enough time to catch up with them.

- The package `distlib.metadata`, which implements distribution metadata as defined by [PEP 643](#), [PEP 566](#), [PEP 345](#), [PEP 314](#) and [PEP 241](#).

Note: In the past `distlib` has tracked metadata proposals in PEPs even when they were draft, but this has proven to be too time-consuming. The current policy is not to track standards proactively while they’re still being thrashed out, but to look instead at starting to implement them once they’re marked `Final`.

- The package `distlib.markers`, which implements environment markers as defined by [PEP 508](#).
- The package `distlib.manifest`, which implements lists of files used in packaging source distributions.
- The package `distlib.locators`, which allows finding distributions, whether on PyPI (XML-RPC or via the “simple” interface), local directories or some other source.
- The package `distlib.resources`, which allows access to data files stored in Python packages, both in the file system and in `.zip` files.
- The package `distlib.scripts`, which allows installing of scripts with adjustment of shebang lines and support for native Windows executable launchers.
- The package `distlib.version`, which implements version specifiers as defined by [PEP 440](#), but also support for working with “legacy” versions (`setuptools/distribute`) and semantic versions.
- The package `distlib.wheel`, which provides support for building and installing from the Wheel format for binary distributions (see [PEP 427](#)).
- The package `distlib.util`, which contains miscellaneous functions and classes which are useful in packaging, but which do not fit neatly into one of the other packages in `distlib.*` The package implements enhanced globbing functionality such as the ability to use `**` in patterns to specify recursing into subdirectories.

1.6 Python version and platform compatibility

Distlib is intended to be used on any Python version `>= 2.7` and is tested on Python versions 2.7 and 3.6-3.10 on Linux, Windows, and macOS.

1.7 Project status

The project has reached a mature status in its development: there is a test suite and it has been exercised on Windows, Ubuntu and macOS. The project is used by well-known projects such as [pip](#), [virtualenv](#) and [caniusepython3](#).

To work with the project, you can [download a release from PyPI](#), or clone the source repository or download a tarball from it.

The source repository for the project is on GitHub:

<https://github.com/pypa/distlib/>

Coverage results are available at:

<https://coveralls.io/github/vsajip/distlib/>

Continuous integration test results are available at:

<https://github.com/pypa/distlib/actions/>

You can leave feedback by raising a new issue on the [issue tracker](#).

1.8 Change log for distlib

1.8.1 0.3.6 (future)

Released: Not yet.

- scripts
 - Fixed #175: Updated launcher executables to better handle the relationship between launcher and child process in the Job API.

1.8.2 0.3.5

Released: 2022-07-14

- database
 - Fixed #170: Corrected implementation of `get_required_dists()`.
- index
 - Updated coverage pragmas for tests relating to obsolete PyPI APIs.
- locators
 - Changed the default locator configuration.
- metadata
 - Updates in support of PEP 643 / Metadata 2.2.
- scripts
 - Updated launcher executables. Thanks to Michael Bikovitsky for his help with the launcher changes.
 - Fixed #164: Improved support for reproducible builds by allowing a fixed date/time to be inserted into created .exe files. Thanks to Somber Night for the patch.
- util
 - Fixed #161: Updated test case.
- wheel
 - Updated to write archive path of RECORD to RECORD instead of staging path. Thanks to Pieter Pas for the patch.
 - Fixed #169: Removed usage of deprecated `imp` module in favour of `importlib`.
 - Fixed #172: Compute ABI correctly for Python < 3.8.

In addition to the above, `setup.py` was replaced by `setup.cfg` and `pyproject.toml`.

1.8.3 0.3.4

Released: 2021-12-08

- database
 - Fixed #153: Raise warnings in `get_distributions()` if bad metadata seen, but keep going.
- markers

- Fixed #154: Determine Python versions correctly for Python \geq 3.10.
- scripts
 - Updated launcher executables.

Code relating to support for Python 2.6 was also removed.

1.8.4 0.3.3

Released: 2021-09-22

- compat
 - Fixed #152: Removed `splituser()` function which wasn't used and is deprecated.
- markers
 - Fixed #149: Handle version comparisons correctly in environment markers.
- scripts
 - Add ARM-64 launchers and support code to use them. Thanks to Niyas Sait and Adrian Vladu for their contributions.
- util
 - Fixed #148: Handle a single trailing comma following a version. Thanks to Blazej Floch for the report and suggested fix.
- version
 - Fixed #150: Fix incorrect handling of epochs.
- wheel
 - Reverted handling of tags for Python \geq 3.10 (use 310 rather than 3_10). This is because PEP 641 was rejected.
- tests
 - Made changes relating to implementing CI using GitHub Actions.

1.8.5 0.3.2

Released: 2021-05-29

- locators
 - Fixed #141: removed unused regular expression.
- metadata
 - Fixed #140: allowed “Obsoletes” in more scenarios, to better handle faulty metadata already on PyPI.
- resources
 - Fixed #146: added entry for `SourcelessFileLoader` to the finder registry.
- scripts
 - Made the generation of scripts more configurable:

- * the `variant_separator` attribute can be set to determine the separator used between a script basename and its X.Y variant. The default value is '-' and would result in a final script basename like 'foo-X.Y', whereas setting it to '.' would result in a final script basename like 'fooX.Y'.
 - * You can also subclass and override the `get_script_filenames()` method to provide a more customised set of file basenames.
- util
 - Fixed #140: allowed a trailing comma in constraints, to better handle faulty metadata already on PyPI.
 - Moved `get_platform()` logic from `distutils` to here.
 - Fixed #143: removed `normcase()` to avoid some problems on Windows.
 - wheel
 - Dropped any trailing data when computing the Python tag.
 - Added support for manylinux tags.
 - Changed handling of tags for Python ≥ 3.10 (use 3_10 rather than 310).
 - Fixed #147: permission bits are now preserved on POSIX when installing from a wheel.
 - tests
 - Fixed #139: improved handling of errors related to the test PyPI server.

1.8.6 0.3.1

Released: 2020-06-27

The repository has been migrated to Git. References to earlier changesets (commits) in issue comments, etc. will be invalid.

- scripts
 - Fixed #132: Added documentation to help with relative interpreter paths. Thanks to Paul Kienzie for the patch.
 - Fixed #134: Allowed specifying a different target Python version when generating scripts.
 - Fixed #135: Exposed the `enquote_executable` function previously implemented as an internal function.
 - Addressed #138: Improved metadata support for newer metadata versions. Thanks to James Tocknell for the patch.
- wheel
 - Changed the output of flags in entry point definitions. Thanks to `frostming()` for the patch.
 - Stopped writing JSON metadata. Only old-style metadata is written now.

1.8.7 0.3.0

Released: 2019-10-29

- database
 - Issue #102 (partial): `modules` attribute of `InstalledDistribution` was incorrectly computed as a list of bytes, rather than a list of str. This has now been corrected.

- locators
 - Updated Locator._get_digest to check PyPI JSON responses for a “digests” dictionary before trying “algo_digest” keys. Thanks to Jeffery To for the patch.
- scripts
 - Fixed #123: Improved error message if a resource isn’t found.
 - Fixed #124: Stopped norm-casing the executable written into shebangs, as it doesn’t work for some non-ASCII paths.
 - Fixed #125: Updated launchers with versions that correctly report errors containing non-ASCII characters. The updated launchers now also support relative paths (see <http://bit.ly/2JxmOoi> for more information).
 - Changed Python version handling to accommodate versions like e.g. 3.10 (no longer assume a version X.Y where X and Y are single digits).
- util
 - Fixed #127: Allowed hyphens in flags in export specifications.
- wheel
 - Changed Python version handling to accommodate versions like e.g. 3.10 (no longer assume a version X.Y where X and Y are single digits).

1.8.8 0.2.9

Released: 2019-05-14

- index
 - Updated default PyPI URL to <https://pypi.org/pypi>.
- locators
 - Updated default PyPI URL to <https://pypi.org/pypi>.
- metadata
 - Relaxed metadata format checks to ignore ‘Provides’.
- scripts
 - Fixed #33, #34: Simplified script template.
 - Updated Windows launchers.
- util
 - Fixed #116: Corrected parsing of credentials from URLs.
- wheel
 - Fixed #115: Relaxed check for ‘..’ in wheel archive entries by not checking filename parts, only directory segments.
 - Skip entries in archive entries ending with ‘/’ (directories) when verifying or installing.
- docs
 - Updated default PyPI URL to <https://pypi.org/pypi>.
 - Commented out Disqus comment section.
 - Changed theme configuration.

- Updated some out-of-date argument lists.
- tests
 - Updated default PyPI URL to <https://pypi.org/pypi>.
 - Preserved umask on POSIX across a test.

1.8.9 0.2.8

Released: 2018-10-01

- database
 - Fixed #108: Updated metadata scan to look for the METADATA file as well as the JSON formats.
- locators
 - Fixed #112: Handled wheel tags and platform-dependent downloads correctly in SimpleScrapingLocator.
- metadata
 - Fixed #107: Updated documentation on testing to include information on setting PYTHONHASHSEED.
- scripts
 - Fixed #111: Avoided unnecessary newlines in script preambles, which caused problems with detecting encoding declarations. Thanks to Wim Glenn for the report and patch.
- util
 - Fixed #109: Removed existing files (which might have been symlinks) before overwriting.

1.8.10 0.2.7

Released: 2018-04-16

- compat
 - Fixed #105: `cache_from_source` is now imported from `importlib.util` where available.
- database
 - Addressed #102: `InstalledDistributions` now have a `modules` attribute which is a list of top-level modules as read from `top_level.txt`, if that is in the distribution info.
- locators
 - Fixed #103: Thanks to Saulius Žemaitaitis for the patch.
- metadata
 - Added support for PEP 566 / Metadata 1.3.
- scripts
 - Fixed #104: Updated launcher binaries. Thanks to Atsushi Odagiri for the diagnosis and fix.

1.8.11 0.2.6

Released: 2017-10-28

- compat
 - Fixed #99: Updated to handle a case where `sys.getfilesystemencoding()` returns `None`.
- database
 - Fixed #97: Eliminated a crash in `EggInfoDistribution.list_distinfo_files()` which was caused by trying to open a non-existent file.
 - Handled a case where an installed distribution didn't have 'Provides:' metadata.
- locators
 - Fixed #96: `SimpleScrapingLocator` no longer fails prematurely when scraping links due to invalid versions.
- markers
 - Improved error messages issued when interpreting markers
- scripts
 - Improved the shebangs written into installed scripts when the interpreter path is very long or contains spaces (to cater for a limitation in shebang line parsing on Linux)
 - Updated launcher binaries.
- tests
 - Numerous test refinements, not detailed further here.

1.8.12 0.2.5

Released: 2017-05-06

- general
 - Changed regular expressions to be compatible with 3.6 as regards escape sequences. Thanks to Ville Skyttä for the patch.
 - closed some resource leaks related to XML-RPC proxies.
 - Removed Python 2.6 from the support list.
- locators
 - Made downloadability a factor in scoring URLs for preferences.
- markers
 - Replaced the implementation with code which parses requirements in accordance with PEP 508 and evaluates marker expressions according to PEP 508.
- util
 - Changed `_csv_open` to use utf-8 across all platforms on Python 3.x. Thanks to Alastair McCormack for the patch.
- wheel
 - Changed to look for metadata in `metadata.json` as well as `pydist.json`.
- version

- Updated requirement parsing in version matchers to use the new PEP 508-compliant code.
- tests
 - Numerous test refinements, not detailed further here.

1.8.13 0.2.4

Released: 2016-09-30

- compat
 - Updated to not fail on import if SSL is unavailable.
- index
 - Switch from using gpg in preference to gpg2 for signing. This is to avoid gpg2's behaviour of prompting for passwords, which interferes with the tests on some machines.
- locators
 - Changed project name comparisons to follow PEP 503. Thanks to Steven Arcangeli for the patch.
 - Added errors queue to Locator.
- manifest
 - Changed match logic to work under Python 3.6, due to differences in how fnmatch.translate behaves.
- resources
 - Updated finder registry logic to reflect changes in Python 3.6.
- scripts
 - Fixed regular expression in generated script boilerplate.
- util
 - Updated to not fail on import if SSL is unavailable.
 - Added normalize_name for project name comparisons using PEP 503.
- tests
 - Updated to skip certain tests if SSL is unavailable.
 - Numerous other test refinements, not detailed further here.

1.8.14 0.2.3

Released: 2016-04-30

- util
 - Changed get_executable to return Unicode rather than bytes.
 - Fixed #84: Allow + character in output script names.
 - Relaxed too-stringent test looking for application/json in headers.
- wheel
 - sorted the entries in RECORD before writing to file.
- tests

- Numerous test refinements, not detailed further here.

1.8.15 0.2.2

Released: 2016-01-30

- database
 - Issue #81: Added support for detecting distributions installed by wheel versions ≥ 0.23 (which use meta-data.json rather than pydist.json). Thanks to Te-jé Rodgers for the patch.
- locators
 - Updated default PyPI URL to <https://pypi.python.org/pypi>
- metadata
 - Updated to use different formatting for description field for V1.1 metadata.
 - Corrected “classifier” to “classifiers” in the mapping for V1.0 metadata.
- scripts
 - Improved support for Jython when quoting executables in output scripts.
- util
 - Issue #77: Made the internal URL used for extended metadata fetches configurable via a module attribute.
 - Issue #78: Improved entry point parsing to handle leading spaces in ini-format files.
- docs
 - Numerous documentation updates, not detailed further here.
- tests
 - renamed environment variable SKIP_SLOW to SKIP_ONLINE in tests and applied to some more tests.
 - Numerous other test refinements, not detailed further here.

1.8.16 0.2.1

Released: 2015-07-07

- locators
 - Issue #58: Return a Distribution instance or None from `locate()`.
 - Issue #59: Skipped special keys when looking for versions.
 - Improved behaviour of PyPIJSONLocator to be analogous to that of other locators.
- resource
 - Added resource iterator functionality.
- scripts
 - Issue #71: Updated launchers to decode shebangs using UTF-8. This allows non-ASCII pathnames to be correctly handled.
 - Ensured that the executable written to shebangs is normcased.
 - Changed ScriptMaker to work better under Jython.

- util
 - Changed the mode setting method to work better under Jython.
 - Changed `get_executable()` to return a normcased value.
- wheel
 - Handled multiple-architecture wheel filenames correctly.
- docs
 - Numerous documentation updates, not detailed further here.
- tests
 - Numerous test refinements, not detailed further here.

1.8.17 0.2.0

Released: 2014-12-17

- compat
 - Updated `match_hostname` to use the latest Python implementation.
- database
 - Added *download_urls* and *digests* attributes to `Distribution`.
- locators
 - Issue #48: Fixed the problem of adding a tuple containing a set (unhashable) to a set, by wrapping with `frozenset()`.
 - Issue #55: Return multiple download URLs for distributions, if available.
- manifest
 - Issue #57: Remove unhelpful warnings about pattern matches.
- metadata
 - Updated to reflect changes to PEP 426.
- resources
 - Issue #50: The type of the path needs to be preserved on 2.x.
- scripts
 - Updated (including launchers) to support providing arguments to interpreters in shebang lines.
 - The launcher sources are now included in the repository and the source distribution (they are to be found in the PC directory).
 - Added frames support in IronPython (patch by Pawel Jasinski).
 - Issue #51: encode shebang executable using utf-8 rather than fsencode.
- util
 - Removed reference to `__PYENVV_LAUNCHER__` when determining executable for scripts (relevant only on macOS).
 - Updated to support changes to PEP 426.
- version

- Updated to reflect changes to versioning proposed in PEP 440.
- wheel
 - Updated build() code to respect interpreter arguments in prebuilt scripts.
 - Updated to support changes to PEP 426 / PEP 440.
- docs
 - Numerous documentation updates, not detailed further here.
- tests
 - Numerous test refinements, not detailed further here.

1.8.18 0.1.9

Released: 2014-05-19

- index
 - Added `keystore` keyword argument to signing and verification APIs.
- scripts
 - Issue #47: Updated binary launchers to fix double-quoting bug where script executable paths have spaces.
- docs
 - Numerous documentation updates, not detailed further here.
- tests
 - Numerous test refinements, not detailed further here.

1.8.19 0.1.8

Released: 2014-03-18

- index
 - Improved thread-safety in `SimpleScrapingLocator` (issue #45).
 - Replaced absolute imports with relative ones.
 - Added `search` method to `PackageIndex`.
- locators
 - Improved thread-safety in `SimpleScrapingLocator` (issue #45).
- metadata
 - Fixed bug in `add_requirements` implementation.
- resources
 - The `Cache` class was refactored into `distlib.util.Cache` and `distlib.resources.ResourceCache` classes.
- scripts
 - Implement quoting for executables with spaces in them.
- util

- Gained the `Cache` class, which is also used in `distlib.wheel`.
- version
 - Allowed versions with a single numeric component and a local version component.
 - Adjusted pre-release computation for legacy versions to be the same as the logic in the `setuptools` documentation.
- wheel
 - Added `verify`, `update`, `is_compatible` and `is_mountable` methods to the `Wheel` class.
 - Converted local version separators from ‘-’ to ‘_’ and back.
 - If SOABI not available, used `Py_DEBUG`, `Py_UNICODE_SIZE` and `WITH_PYMALLOC` to derive the ABI.
 - Added “exists” property to `Wheel` instances.
 - Factored out `RECORD` writing and zip building to separate methods.
 - Provided the ability to determine the location where extensions are extracted, by using the `distlib.util.Cache` class.
 - Avoided using `pydist.json` in 1.0 wheels (`bdist_wheel` writes a non-conforming `pydist.json`).
 - Improved computation of compatible tags on macOS, and made `COMPATIBLE_TAGS` a set.
- _backport/sysconfig
 - Replaced an absolute import with a relative one.
- docs
 - Numerous documentation updates, not detailed further here.
- tests
 - Numerous test refinements, not detailed further here.

1.8.20 0.1.7

Released: 2014-01-16

- metadata
 - Added some more fields to the metadata for the index.
- resources
 - Use native literal string in cache path.
 - Issue #40: Now does path adjustments differently for files and zips.
- scripts
 - Improved checking for venvs when generating scripts.
- util
 - Issue #39: Fall back to temporary directory for cache if home directory unavailable.
- wheel
 - Use native literal string in cache path.

1.8.21 0.1.6

Released: 2013-12-31

- scripts
 - Updated binary launchers because the wrong variant was shipped with the previous release.
- version
 - Added support for local component in PEP 440 versions.
- tests
 - Numerous test refinements, not detailed further here.

1.8.22 0.1.5

Released: 2013-12-15

- compat
 - Changed source of import for unescape in Python ≥ 3.4 .
- index
 - Used `dummy_threading` when `threading` isn't available.
 - Used `https` for default index.
- locators
 - Used `dummy_threading` when `threading` isn't available.
- scripts
 - Defaulted to setting script mode bits on POSIX.
 - Use uncompressed executable launchers, since some anti-virus products raise false positive errors.
- util
 - Used `dummy_threading` when `threading` isn't available.
- docs
 - Updated out-of-date links in overview.
- tests
 - Used `dummy_threading` when `threading` isn't available.

1.8.23 0.1.4

Released: 2013-10-31

- scripts
 - Updated the logic for finding the distlib package using a relative, rather than absolute method. This fixes a problem for pip, where distlib is kept in the `pip.vendor.distlib` package.
- `_backport/sysconfig`
 - The analogous change to that made for scripts, described above.

1.8.24 0.1.3

Released: 2013-10-18

- database
 - Added support for PEP 426 JSON metadata (pydist.json).
 - Generalised digests to support e.g. SHA256.
 - Fixed a bug in parsing legacy metadata from .egg directories.
 - Removed duplicated code relating to parsing “provides” fields.
- index
 - Changes relating to support for PEP 426 JSON metadata (pydist.json).
- locators
 - Changes relating to support for PEP 426 JSON metadata (pydist.json).
 - Fixed a bug in scoring download URLs for preference when multiple URLs are available.
 - The legacy scheme is used for the default locator.
 - Made changes relating to parsing “provides” fields.
 - Generalised digests to support e.g. SHA256.
 - If no release version is found for a requirement, prereleases are now considered even if not explicitly requested.
- markers
 - Added support for markers as specified in PEP 426.
- metadata
 - Added support for PEP 426 JSON metadata (pydist.json). The old metadata class is renamed to LegacyMetadata, and the (new) Metadata class wraps the JSON format (and also the legacy format, through LegacyMetadata).
 - Removed code which was only used if docutils was installed. This code implemented validation of .rst descriptions, which is not done in distlib.
- scripts
 - Updated the logic for writing executable files to deal as best we can with files which are already in use and hence cannot be deleted on Windows.
 - Changed the script generation when launchers are used to write a single executable which wraps a script (whether pre-built or generated) and includes a manifest to avoid UAC prompts on Windows.
 - Changed the interface for script generation options: the `make` and `make_multiple` methods of `ScriptMaker` now take an optional `options` dictionary.
- util
 - Added `extract_by_key()` to copy selected keys from one dict to another.
 - Added `parse_name_and_version()` for use in parsing “provides” fields.
 - Made `split_filename` more flexible.
- version
 - Added support for PEP 440 version matching.

- Removed AdaptiveVersion, AdaptiveMatcher etc. as they don't add sufficient value to justify keeping them in.
- wheel
 - Added wheel_version kwarg to Wheel.build API.
 - Changed Wheel.install API (after consultation on distutils-sig).
 - Added support for PEP 426 JSON metadata (pydist.json).
 - Added lib_only flag to install() method.
- docs
 - Numerous documentation updates, not detailed further here.
- tests
 - Numerous test refinements, not detailed further here.

1.8.25 0.1.2

Released: 2013-04-30

- compat
 - Added BaseConfigurator backport for 2.6.
- database
 - Return RECORD path from write_installed_files (or None if dry_run).
 - Explicitly return None from write_shared_locations if dry run.
- metadata
 - Added missing condition in `todict()`.
- scripts
 - Add variants and clobber flag for generation of foo/fooX/foo-X.Y.
 - Added .exe manifests for Windows.
- util
 - Regularised recording of written files.
 - Added Configurator.
- version
 - Tidypups, most suggested by Donald Stufft: Made key functions private, removed _Common class, removed checking for huge version numbers, made UnsupportedVersionError a ValueError.
- wheel
 - Replaced absolute import with relative.
 - Handle None return from write_shared_locations correctly.
 - Fixed bug in Mounter for extension modules not in sub-packages.
 - Made dylib-cache Python version-specific.
- docs
 - Numerous documentation updates, not detailed further here.

- tests
 - Numerous test refinements, not detailed further here.
- other
 - Corrected setup.py to ensure that sysconfig.cfg is included.

1.8.26 0.1.1

Released: 2013-03-22

- database
 - Updated requirements logic to use extras and environment markers.
 - Made it easier to subclass Distribution and EggInfoDistribution.
- locators
 - Added method to clear locator caches.
 - Added the ability to skip pre-releases.
- manifest
 - Fixed bug which caused side-effect when sorting a manifest.
- metadata
 - Updated to handle most 2.0 fields, though PEP 426 is still a draft.
 - Added the option to skip unset fields when writing.
- resources
 - Made separate subclasses ResourceBase, Resource and ResourceContainer from Resource. Thanks to Thomas Kluyver for the suggestion and patch.
- scripts
 - Fixed bug which prevented writing shebang lines correctly on Windows.
- util
 - Made get_cache_base more useful by parameterising the suffix to use.
 - Fixed a bug when reading CSV streams from .zip files under 3.x.
- version
 - Added is_prerelease property to versions.
 - Moved to PEP 426 version formats and sorting.
- wheel
 - Fixed CSV stream reading under 3.x and handled UTF-8 in zip entries correctly.
 - Added metadata and info properties, and updated the install method to return the installed distribution.
 - Added mount/unmount functionality.
 - Removed compatible_tags() function in favour of COMPATIBLE_TAGS attribute.
- docs
 - Numerous documentation updates, not detailed further here.

- tests
 - Numerous test refinements, not detailed further here.

1.8.27 0.1.0

Released: 2013-03-02

- Initial release.

1.9 Next steps

You might find it helpful to look at the *[Tutorial](#)*, or the *[API Reference](#)*.

2.1 Installation

Distlib is a pure-Python library. You should be able to install it using:

```
pip install distlib
```

for installing `distlib` into a virtualenv or other directory where you have write permissions. On Posix platforms, you may need to invoke using `sudo` if you need to install `distlib` in a protected location such as your system Python's `site-packages` directory.

2.2 Testing

A full test suite is included with `distlib`. To run it, you'll need to download the source distribution, unpack it and run `python tests/test_all.py` in the top-level directory of the package.

If running the tests under Python $\geq 3.2.3$, remember to first set the environment variable `PYTHONHASHSEED=0` to disable hash randomisation, which is needed for the tests. (The environment variable also needs to be set if running Python 2.x with `-R`, which is only available in Python 2.6.8 and later.)

Continuous integration test results are available at:

<https://github.com/pypa/distlib/actions>

Coverage results are available at:

<https://app.codecov.io/gh/pypa/distlib>

Note that the actual coverage is higher than that shown, because coverage under Windows is not included in the above coverage figures.

Note that the index tests are configured, by default, to use a local test server, though they can be configured to run against PyPI itself. This local test server is not bundled with `distlib`, but is available from:

<https://raw.githubusercontent.com/vsajip/pypiserver/standalone/pypi-server-standalone.py>

This is a slightly modified version of Ralf Schmitt's `pypiserver`. To use, the script needs to be copied to the `tests` folder of the `distlib` distribution.

If the server script is not available, the tests which use it will be skipped. Naturally, this will also affect the coverage statistics.

2.2.1 PYPI availability

If PyPI is unavailable or slow, then some of the tests can fail or become painfully slow. To skip tests that might be sometimes slow, set the `SKIP_SLOW` environment variable:

```
$ SKIP_SLOW=1 PYTHONHASHSEED=0 python tests/test_all.py
```

on Posix, or:

```
C:\> set SKIP_SLOW=1
C:\> set PYTHONHASHSEED=0
C:\> python tests/test_all.py
```

on Windows.

2.3 First steps

For now, we just list how to use particular parts of the API as they take shape.

2.3.1 Using the database API

You can use the `distlib.database` package to access information about installed distributions. This information is available through the following classes:

- *DistributionPath*, which represents a set of distributions installed on a path.
- *Distribution*, which represents an individual distribution, conforming to recent packaging PEPs ([PEP 643](#), [PEP 566](#), [PEP 508](#), [PEP 440](#), [PEP 386](#), [PEP 376](#), [PEP 345](#), [PEP 314](#) and [PEP 241](#)).
- *EggInfoDistribution*, which represents a legacy distribution in egg format.

2.3.1.1 Distribution paths

The *Distribution* and *EggInfoDistribution* classes are normally not instantiated directly; rather, they are returned by querying *DistributionPath* for distributions. To create a *DistributionPath* instance, you can do

```
>>> from distlib.database import DistributionPath
>>> dist_path = DistributionPath()
```

2.3.1.2 Querying a path for distributions

In this most basic form, `dist_path` will provide access to all non-legacy distributions on `sys.path`. To get these distributions, you invoke the `get_distributions()` method, which returns an iterable. Let's try it:

```
>>> list(dist_path.get_distributions())
[]
```

This may seem surprising if you've just started looking at `distlib`, as you won't *have* any non-legacy distributions.

2.3.1.3 Including legacy distributions in the search results

To include distributions created and installed using `setuptools` or `distribute`, you need to create the `DistributionPath` by specifying an additional keyword argument, like so:

```
>>> dist_path = DistributionPath(include_egg=True)
```

and then you'll get a less surprising result:

```
>>> len(list(dist_path.get_distributions()))
77
```

The exact number returned will be different for you, of course. You can ask for a particular distribution by name, using the `get_distribution()` method:

```
>>> dist_path.get_distribution('setuptools')
<EggInfoDistribution u'setuptools' 0.6c11 at '/usr/lib/python2.7/dist-packages/
↳setuptools.egg-info'>
```

If you want to look at a specific path other than `sys.path`, you specify it as a positional argument to the `DistributionPath` constructor:

```
>>> from pprint import pprint
>>> special_dists = DistributionPath(['tests/fake_dists'], include_egg=True)
>>> pprint([d.name for d in special_dists.get_distributions()])
['babar',
 'choxie',
 'towel-stuff',
 'grammar',
 'truffles',
 'coconuts-aster',
 'nut',
 'bacon',
 'banana',
 'cheese',
 'strawberry']
```

or, if you leave out egg-based distributions:

```
>>> special_dists = DistributionPath(['tests/fake_dists'])
>>> pprint([d.name for d in special_dists.get_distributions()])
['babar',
 'choxie',
 'towel-stuff',
 'grammar']
```

2.3.1.4 Distribution properties

Once you have a `Distribution` instance, you can use it to get more information about the distribution. For example:

- The `metadata` attribute gives access to the distribution's metadata (see [Using the metadata and markers APIs](#) for more information).
- The `name_and_version` attribute shows the name and version in the format `name (X.Y)`.
- The `key` attribute holds the distribution's name in lower-case, as you generally want to search for distributions without regard to case sensitivity.

2.3.1.5 Exporting things from Distributions

Each distribution has a dictionary of *exports*. The exports dictionary is functionally equivalent to “entry points” in `distribute / setuptools`.

The keys to the dictionary are just names in a hierarchical namespace delineated with periods (like Python packages, so we'll refer to them as *pkgnames* in the following discussion). The keys indicate categories of information which the distribution's author wishes to export. In each such category, a distribution may publish one or more entries.

The entries can be used for many purposes, and can point to callable code or data. A common purpose is for publishing callables in the distribution which adhere to a particular protocol.

To give a concrete example, the [Babel](#) library for internationalisation support provides a mechanism for extracting, from a variety of sources, message text to be internationalised. Babel itself provides functionality to extract messages from e.g. Python and JavaScript source code, but helpfully offers a mechanism whereby providers of other sources of message text can provide their own extractors. It does this by providing a category `'babel.extractors'`, under which other software can register extractors for their sources. The [Jinja2](#) template engine, for example, makes use of this to provide a message extractor for Jinja2 templates. Babel itself registers its own extractors under the same category, so that a unified view of all extractors in a given Python environment can be obtained, and Babel's extractors are treated by other parts of Babel in exactly the same way as extractors from third parties.

Any installed distribution can offer up values for any category, and a set of distributions (such as the set of installed distributions on `sys.path`) conceptually has an aggregation of these values.

The values associated with a category are a list of strings with the format:

```
name = prefix [ ":" suffix ] [ "[" flags "]" ]
```

where `name`, `prefix`, and `suffix` are *pkgnames*. `suffix` and `flags` are optional and `flags` follow the description in [Flag formats](#).

Any installed distribution can offer up values for any category, and a set of distributions (such as the set of installed distributions on `sys.path`) conceptually has an aggregation of these values.

For callables, the `prefix` is the package or module name which contains the callable, `suffix` is the path to the callable in the module, and `flags` can be used for any purpose determined by the distribution author (for example, the `extras` feature in `distribute / setuptools`).

This entry format is used in the `distlib.scripts` package for installing scripts based on Python callables.

2.3.2 Distribution dependencies

You can use the `distlib.locators` package to locate the dependencies that a distribution has. The `distlib.database` package has code which allow you to analyse the relationships between a set of distributions:

- `make_graph()`, which generates a dependency graph from a list of distributions.
- `get_dependent_dists()`, which takes a list of distributions and a specific distribution in that list, and returns the distributions that are dependent on that specific distribution.

- `get_required_dists()`, which takes a list of distributions and a specific distribution in that list, and returns the distributions that are required by that specific distribution.

The graph returned by `make_graph()` is an instance of *DependencyGraph*.

2.3.3 Using the locators API

2.3.3.1 Overview

To locate a distribution in an index, we can use the `locate()` function. This returns a potentially downloadable distribution (in the sense that it has a download URL – of course, there are no guarantees that there will actually be a downloadable resource at that URL). The return value is an instance of `distlib.database.Distribution` which can be queried for any distributions it requires, so that they can also be located if desired. Here is a basic example:

```
>>> from distlib.locators import locate
>>> flask = locate('flask')
>>> flask
<Distribution Flask (0.10.1) [https://pypi.org/packages/source/F/Flask/Flask-0.10.1.
↳tar.gz]>
>>> dependencies = [locate(r) for r in flask.run_requires]
>>> from pprint import pprint
>>> pprint(dependencies)
[<Distribution Werkzeug (0.9.1) [https://pypi.org/packages/source/W/Werkzeug/Werkzeug-
↳0.9.1.tar.gz]>,
 <Distribution Jinja2 (2.7) [https://pypi.org/packages/source/J/Jinja2/Jinja2-2.7.tar.
↳gz]>,
 <Distribution itsdangerous (0.21) [https://pypi.org/packages/source/i/itsdangerous/
↳itsdangerous-0.21.tar.gz]>]
>>>
```

The values in the `run_requires` property are just strings. Here's another example, showing a little more detail:

```
>>> authy = locate('authy')
>>> authy.run_requires
set(['httplib2 (>= 0.7, < 0.8)', 'simplejson'])
>>> authy
<Distribution authy (1.0.0) [http://pypi.org/packages/source/a/authy/authy-1.0.0.tar.
↳gz]>
>>> deps = [locate(r) for r in authy.run_requires]
>>> pprint(deps)
[<Distribution httplib2 (0.7.7) [http://pypi.org/packages/source/h/httplib2/httplib2-
↳0.7.7.zip]>,
 <Distribution simplejson (3.3.0) [http://pypi.org/packages/source/s/simplejson/
↳simplejson-3.3.0.tar.gz]>]
>>>
```

Note that the constraints on the dependencies were honoured by `locate()`.

2.3.3.2 Under the hood

Under the hood, `locate()` uses *locators*. Locators are a mechanism for finding distributions from a range of sources. Although the `pypi` subpackage has been copied from `distutils2` to `distlib`, there may be benefits in a higher-level API, and so the `distlib.locators` package has been created as an experiment. Locators are objects which locate distributions. A locator instance's `get_project()` method is called, passing in a project name: The method

returns a dictionary containing information about distribution releases found for that project. The keys of the returned dictionary are versions, and the values are instances of `distlib.database.Distribution`.

The following locators are provided:

- *DirectoryLocator* – this is instantiated with a base directory and will look for archives in the file system tree under that directory. Name and version information is inferred from the filenames of archives, and the amount of information returned about the download is minimal. The locator searches all subdirectories by default, but can be set to only look in the specified directory by setting the `recursive` keyword argument to `False`.
- *PyPIRPCLocator*. – This takes a base URL for the RPC service and will locate packages using PyPI's XML-RPC API. This locator is a little slow (the scraping interface seems to work faster) and case-sensitive. For example, searching for `'flask'` will throw up no results, but you get the expected results when searching from `'Flask'`. This appears to be a limitation of the underlying XML-RPC API. Note that 20 versions of a project necessitate 41 network calls (one to get the versions, and two more for each version – one to get the metadata, and another to get the downloads information).
- *PyPIJSONLocator*. – This takes a base URL for the JSON service and will locate packages using PyPI's JSON API. This locator is case-sensitive. For example, searching for `'flask'` will throw up no results, but you get the expected results when searching from `'Flask'`. This appears to be a limitation of the underlying JSON API. Note that unlike the XML-RPC service, only non-hidden releases will be returned.
- *SimpleScrapingLocator* – this takes a base URL for the site to scrape, and locates packages using a similar approach to the `PackageFinder` class in `pip`, or as documented in the `setuptools` documentation as the approach used by `easy_install`.
- *DistPathLocator* – this takes a *DistributionPath* instance and locates installed distributions. This can be used with *AggregatingLocator* to satisfy requirements from installed distributions before looking elsewhere for them.
- *JSONLocator* – this uses an improved JSON metadata schema and returns data on all versions of a distribution, including dependencies, using a single network request.
- *AggregatingLocator* – this takes a list of other aggregators and delegates finding projects to them. It can either return the first result found (i.e. from the first aggregator in the list provided which returns a non-empty result), or a merged result from all the aggregators in the list.

There is a default locator, available at `distlib.locators.default_locator`.

The `locators` package also contains a function, `get_all_distribution_names()`, which retrieves the names of all distributions registered on PyPI:

```
>>> from distlib.locators import get_all_distribution_names
>>> names = get_all_distribution_names()
>>> len(names)
31905
>>>
```

This is implemented using the XML-RPC API.

Apart from *JSONLocator*, none of the locators currently returns enough metadata to allow dependency resolution to be carried out, but that is a result of the fact that metadata relating to dependencies are not indexed, and would require not just downloading the distribution archives and inspection of contained metadata files, but potentially also introspecting `setup.py`! This is the downside of having vital information only available via keyword arguments to the `setup()` call: hopefully, a move to fully declarative metadata will facilitate indexing it and allowing the provision of improved features.

The locators will skip binary distributions other than wheels. (`.egg` files are currently treated as binary distributions).

The PyPI locator classes don't yet support the use of mirrors, but that can be added in due course – once the basic functionality is working satisfactorily.

2.3.4 Using the index API

You can use the `distlib.index` package to perform operations relating to a package index compatible with PyPI. This includes things like registering a project, uploading a distribution or uploading documentation.

2.3.4.1 Overview

You access index functionality through an instance of the `PackageIndex` class. This is instantiated with the URL of the repository (which can be omitted if you want to use PyPI itself):

```
>>> from distlib.index import PackageIndex
>>> index = PackageIndex()
>>> index.url
'http://pypi.org/pypi'
```

To use a local test server, you might do this:

```
>>> index = PackageIndex('http://localhost:8080/')
```

2.3.4.2 Registering a project

Registering a project can be done using a `Metadata` instance which holds the index metadata used for registering. A simple example:

```
>>> from distlib.metadata import Metadata
>>> metadata = Metadata()
>>> metadata.name = 'tatterdemalion'
>>> metadata.version = '0.1'
>>> # other fields omitted
>>> response = index.register(metadata)
```

The `register()` method returns an HTTP response, such as might be returned by a call to `urlopen`. If an error occurs, a `HTTPError` will be raised. Otherwise, the `response.code` should be 200.

2.3.4.3 Uploading a source distribution

To upload a source distribution, you need to do the following as a minimum:

```
>>> metadata = ... # get a populated Metadata instance
>>> response = index.upload_file(metadata, archive_name)
```

The `upload_file()` method returns an HTTP response or, in case of error, raises an `HTTPError`.

2.3.4.4 Uploading binary distributions

When uploading binary distributions, you need to specify the file type and Python version, as in the following example:

```
>>> response = index.upload_file(metadata, archive_name,
...                               filetype='bdist_dumb',
...                               pyversion='2.6')
```

2.3.4.5 Signing a distribution

To sign a distribution, you will typically need GnuPG. The default implementation looks for `gpg` or `gpg2` on the path, but if not available there, you can explicitly specify an absolute path indicating where the signing program is to be found:

```
>>> index.gpg = '/path/to/gpg'
```

Once this is set, you can sign the archive before uploading, as follows:

```
>>> response = index.upload_file(metadata, archive_name,
...                               signer='Test User',
...                               sign_password='secret',
...                               keystore='/path/to/keys')
```

As an alternative to passing the keystore with each call, you can specify that in an instance attribute:

```
>>> index.gpg_home = '/path/to/keys'
```

The `keystore` is a directory which contains the GnuPG key database (files like `pubring.gpg`, `secring.gpg`, and `trustdb.gpg`).

When you sign a distribution, both the distribution and the signature are uploaded to the index.

2.3.4.6 Downloading files

The `PackageIndex` class contains a utility method which allows you to download distributions (and other files, such as signatures):

```
>>> index.download_file(url, destfile, digest=None, reporthook=None)
```

This is similar in function to `urlretrieve()` in the standard library. Provide a `digest` if you want the call to check that the has digest of the downloaded file matches a specific value: if not provided, no matching is done. The value passed can just be a plain string in the case of an MD5 digest or, if you want to specify the hashing algorithm to use, specify a tuple such as `('sha1', '0123456789abcdef...')`. The hashing algorithm must be one that's supported by the `hashlib` module.

Benefits to using this method over plain `urlretrieve()` are:

- It will use the `ssl_verifier`, if set, to ensure that the download is coming from where you think it is (see [Verifying HTTPS connections](#)).
- It will compute the digest as it downloads, saving you from having to read the whole of the downloaded file just to compute its digest.

Note that the url you download from doesn't actually need to be on the index – in theory, it could be from some other site. Note that if you have an `ssl_verifier` set on the index, it will perform its checks according to whichever url you supply – whether it's a resource on the index or not.

2.3.4.7 Verifying signatures

For any archive downloaded from an index, you can retrieve any signature by just appending `.asc` to the path portion of the download URL for the archive, and downloading that. The `index` class offers a `verify_signature()` method for validating a signature. If you have files `'good.bin'`, `'bad.bin'` which are different from each other, and `'good.bin.asc'` has the signature for `'good.bin'`, then you can verify signatures like this:

```
>>> index.verify_signature('good.bin.asc', 'good.bin', '/path/to/keys')
True
>>> index.verify_signature('good.bin.asc', 'bad.bin', '/path/to/keys')
False
```

The last argument, which is optional, specifies a directory which holds the GnuPG keys used for verification – the *keystore*. Instead of specifying the keystore location in each call, you can specify the location in an instance attribute:

```
>>> index.gpg_home = '/path/to/keys'
```

If you do this, you don't need to pass the keystore location.

Note that if you don't have the `gpg` or `gpg2` programs on the path, you may need to specify the location of the verifier program explicitly:

```
>>> index.gpg = '/path/to/gpg'
```

Some caveats about verified signatures

In order to be able to perform signature verification, you'll have to ensure that the public keys of whoever signed those distributions are in your key store. However, having these keys shouldn't give you a false sense of security; unless you can be sure that those keys actually belong to the people or organisations they purport to represent, the signature has no real value, even if it is verified without error. For you to be able to trust a key, it would need to be signed by someone you trust, who vouches for it – and this requires there to be either a signature from a valid certifying authority (e.g. Verisign, Thawte etc.) or a [Web of Trust](#) around the keys that you want to rely on.

An index may itself countersign distributions (so *it* deals with the keys of the distribution publishers, but you need only deal with the public signing key belonging to the index). If you trust the index, you can trust the verified signature if it's signed by the index.

2.3.4.8 Uploading documentation

To upload documentation, you need to specify the metadata and the directory which is the root of the documentation (typically, if you use Sphinx to build your documentation, this will be something like `<project>/docs/_build/html`):

```
>>> response = index.upload_documentation(metadata, doc_dir)
```

The `upload_documentation()` method returns an HTTP response or, in case of error, raises an `HTTPError`. The call will zip up the entire contents of the passed directory `doc_dir` and upload the zip file to the index.

2.3.4.9 Authentication

Operations which update the index (all of the above) will require authenticated requests. You can specify a username and password to use for requests sent to the index:

```
>>> index.username = 'test'
>>> index.password = 'secret'
```

For your convenience, these will be automatically read from any `.pypirc` file which you have; if it contains entries for multiple indexes, a repository key in `.pypirc` must match `index.url` to identify which username and password are to be read from `.pypirc`. Note that to ensure compatibility, `distlib` uses `distutils` code to read the `.pypirc` configuration. Thus, given the `.pypirc` file:

```
[distutils]
index-servers =
    pypi
    test

[pypi]
username: me
password: my_strong_password

[test]
repository: http://localhost:8080/
username: test
password: secret
```

you would see the following:

```
>>> index = PackageIndex()
>>> index.username
'me'
>>> index.password
'my_strong_password'
>>> index = PackageIndex('http://localhost:8080/')
>>> index.username
'test'
>>> index.password
'secret'
```

2.3.4.10 Verifying HTTPS connections

Although Python has full support for SSL, it does not, by default, verify SSL connections to servers. That's because in order to do so, a set of certificates which certify the identity of the server needs to be provided (see [the relevant Python documentation](#) for details).

Support for verifying SSL connections is provided in `distlib` through a handler, `distlib.util.HTTPSHandler`. To use it, set the `ssl_verifier` attribute of the index to a suitably configured instance. For example:

```
>>> from distlib.util import HTTPSHandler
>>> verifier = HTTPSHandler('/path/to/root/certs.pem')
>>> index.ssl_verifier = verifier
```

By default, the handler will attempt to match domains, including wildcard matching. This means that (for example) you access `foo.org` or `www.foo.org` which have a certificate for `*.foo.org`, the domains will match. If the domains don't match, the handler raises a `CertificateError` (a subclass of `ValueError`).

Domain mismatches can, however, happen for valid reasons. Say a hosting server `bar.com` hosts `www.foo.org`, which we are trying to access using SSL. If the server holds a certificate for `www.foo.org`, it will present it to the client, as long as both support [Server Name Indication \(SNI\)](#). While `distlib` supports SNI where Python supports

it, Python 2.x does not include SNI support. For this or some other reason, you may wish to turn domain matching off. To do so, instantiate the verifier like this:

```
>>> verifier = HTTPSHandler('/path/to/root/certs.pem', False)
```

Ensuring that *only* HTTPS connections are made

You may want to ensure that traffic is *only* HTTPS for a particular interaction with a server – for example:

- Deal with a Man-In-The-Middle proxy server which listens on port 443 but talks HTTP rather than HTTPS
- Deal with situations where an index page obtained via HTTPS contains links with a scheme of `http` rather than `https`.

To do this, instead of using `HTTPSHandler` as shown above, use the `HTTPSOnlyHandler` class instead, which disallows any HTTP traffic. It's used in the same way as `HTTPSHandler`:

```
>>> from distlib.util import HTTPSOnlyHandler
>>> verifier = HTTPSOnlyHandler('/path/to/root/certs.pem')
>>> index.ssl_verifier = verifier
```

Note that with this handler, you can't make *any* HTTP connections at all - it will raise `URLLError` if you try.

Getting hold of root certificates

At the time of writing, you can find a file in the appropriate format on the [cURL website](#). Just download the `cacert.pem` file and pass the path to it when instantiating your verifier.

2.3.4.11 Saving a default configuration

If you don't have a `.pyirc` file but want to save one, you can do this by setting the username and password and calling the `save_configuration()` method:

```
>>> index = PackageIndex()
>>> index.username = 'fred'
>>> index.password = 'flintstone'
>>> index.save_configuration()
```

This will use `distutils` code to save a default `.pyirc` file which specifies a single index – PyPI – with the specified username and password.

2.3.4.12 Searching PyPI

You can use the `search()` method of `PackageIndex` to search for distributions on PyPI:

```
>>> index = PackageIndex()
>>> from pprint import pprint
>>> pprint(index.search('tatterdema'))
[{'_pypi_ordering': 0,
  'name': 'tatterdemalion',
  'summary': 'A dummy distribution',
  'version': '0.1.0'}]
```

If a string is specified, just the name is searched for. Alternatively, you can specify a dictionary of attributes to search for, along with values to match. For example:

```
>>> pprint(index.search({'summary': 'dummy'}))
[{'_pypi_ordering': 5,
  'name': 'collective.lorem',
  'summary': 'A package that provides dummy content generation.',
  'version': '0.2.3'},
 {'_pypi_ordering': 7,
  'name': 'collective.loremipsum',
  'summary': 'Creates dummy content with populated Lorem Ipsum.',
  'version': '0.8'},
 {'_pypi_ordering': 1,
  'name': 'cosent.dummyspackage',
  'summary': 'A dummy package for buildtools testing',
  'version': '0.4'},
 {'_pypi_ordering': 0,
  'name': 'django-dummyimage',
  'summary': 'Dynamic Dummy Image Generator For Django!',
  'version': '0.1.1'},
 {'_pypi_ordering': 1,
  'name': 'django-plainpasswordhasher',
  'summary': 'Dummy (plain text) password hashing for Django.',
  'version': '0.2'},
 {'_pypi_ordering': 2,
  'name': 'django-plainpasswordhasher',
  'summary': 'Dummy (plain text) password hashing for Django.',
  'version': '0.3'},
 {'_pypi_ordering': 1,
  'name': 'dummyscache',
  'summary': 'A dummy in-memory cache for development and testing. (Not recommended_
→ for production use.)',
  'version': '0.0.2'},
 {'_pypi_ordering': 0,
  'name': 'dummy-txredis',
  'summary': 'Dummy txRedis client and factory.',
  'version': '0.5'},
 {'_pypi_ordering': 7,
  'name': 'eea.eggmonkeytesttarget',
  'summary': 'A dummy package to test eea.eggmonkey',
  'version': '5.7'},
 {'_pypi_ordering': 8,
  'name': 'invewrapper',
  'summary': 'dummy/transitional package that depends on "pew"',
  'version': '0.1.8'},
 {'_pypi_ordering': 0,
  'name': 'monoprocessing',
  'summary': 'A dummy implementation of multiprocessing.Pool',
  'version': '0.1'},
 {'_pypi_ordering': 0,
  'name': 'myFun',
  'summary': 'This is a dummy function which prints given list data.',
  'version': '1.0.0'},
 {'_pypi_ordering': 0,
  'name': 'ReadableDict-a-dict-without-brackets',
  'summary': 'provides a dummy implementation of a dict without brackets',
  'version': '0.0'},
 {'_pypi_ordering': 4,
```

(continues on next page)

(continued from previous page)

```
'name': 'setuptools_dummy',
'summary': 'Setuptools Dummy Filefinder',
'version': '0.1.0.4'},
{'_pypi_ordering': 0,
 'name': 'tatterdemalion',
 'summary': 'A dummy distribution',
 'version': '0.1.0'}]
```

If you specify multiple attributes, then the search returns the intersection of matches – an `and` operation:

```
>>> pprint(index.search({'summary': 'dummy', 'name': 'ta'}))
[{'_pypi_ordering': 7,
  'name': 'eea.eggmonkeytesttarget',
  'summary': 'A dummy package to test eea.eggmonkey',
  'version': '5.7'},
 {'_pypi_ordering': 0,
  'name': 'tatterdemalion',
  'summary': 'A dummy distribution',
  'version': '0.1.0'}]
```

If you want a union of matches – an `or` operation – specify a second argument to the `PackageIndex.search()` method with the value `'or'`:

```
>>> pprint(index.search({'version': '2013.9', 'name': 'pytzp'}, 'or'))
[{'_pypi_ordering': 65,
  'name': 'pytz',
  'summary': 'World timezone definitions, modern and historical',
  'version': '2013.9'},
 {'_pypi_ordering': 2,
  'name': 'pytzpure',
  'summary': 'A pure-Python version of PYTZ (timezones).',
  'version': '0.2.4'}]
```

The search functionality makes use of PyPI's XML-RPC interface, so it will only work for indexes which supply a compatible implementation. The following search attributes are currently supported:

- name
- version
- stable_version
- author
- author_email
- maintainer
- maintainer_email
- home_page
- license
- summary
- description
- keywords
- platform

- `download_url`
- `classifiers` (list of classifier strings)
- `project_url`
- `docs_url` (URL of the pythonhosted.org docs if they've been supplied)

2.3.5 Using the metadata and markers APIs

The metadata API is exposed through a `Metadata` class. This class can read and write metadata files complying with any of the defined versions: 1.0 ([PEP 241](#)), 1.1 ([PEP 314](#)), 1.2 ([PEP 345](#)), 2.1 ([PEP 566](#)) and 2.2 ([PEP 643](#)). It implements methods to parse and write metadata files.

2.3.5.1 Instantiating metadata

You can simply instantiate a `Metadata` instance and start populating it:

```
>>> from distlib.metadata import Metadata
>>> md = Metadata()
>>> md.name = 'foo'
>>> md.version = '1.0'
```

An instance so created may not be valid unless it has some minimal properties which meet certain constraints, as specified in the [Core metadata specifications](#).

These constraints aren't applicable to legacy metadata. Therefore, when creating `Metadata` instances to deal with such metadata, you can specify the `scheme` keyword when creating the instance:

```
>>> legacy_metadata = Metadata(scheme='legacy')
```

The term 'legacy' refers to the version scheme.

Whether dealing with current or legacy metadata, an instance's `validate()` method can be called to ensure that the metadata has no missing or invalid data. This raises a `DistlibException` (either `MetadataMissingError` or `MetadataInvalidError`) if the metadata isn't valid.

You can initialise an instance with a dictionary using the following form:

```
>>> metadata = Metadata(mapping=a_dictionary)
```

2.3.5.2 Reading metadata from files and streams

The `Metadata` class can be instantiated with the path of the metadata file. Here's an example with legacy metadata:

```
>>> from distlib.metadata import Metadata
>>> metadata = Metadata(path='PKG-INFO')
>>> metadata.name
'CLVault'
>>> metadata.version
'0.5'
>>> metadata.run_requires
['keyring']
```

Instead of using the `path` keyword argument to specify a file location, you can also specify a `fileobj` keyword argument to specify a file-like object which contains the data.

2.3.5.3 Writing metadata to paths and streams

Writing metadata can be done using the `write` method:

```
>>> metadata.write(path='/to/my/pydist.json')
```

You can also specify a file-like object to write to, using the `fileobj` keyword argument.

2.3.5.4 Using markers

Environment markers are implemented in the `distlib.markers` package and accessed via a single function, `interpret()`.

See [PEP 508](#) for more information about environment markers. The `interpret()` function takes a string argument which represents a Boolean expression, and returns either `True` or `False`:

```
>>> from distlib.markers import interpret
>>> interpret('python_version >= "1.0"')
True
```

You can pass in a context dictionary which is checked for values before the environment:

```
>>> interpret('python_version >= "1.0"', {'python_version': '0.5'})
False
```

You won't normally need to work with markers in this way – they are dealt with by the `Metadata` and `Distribution` logic when needed.

2.3.6 Using the resource API

You can use the `distlib.resources` package to access data stored in Python packages, whether in the file system or .zip files. Consider a package which contains data alongside Python code:

```
foofoo
├── bar
│   ├── bar_resource.bin
│   ├── baz.py
│   └── __init__.py
├── foo_resource.bin
├── __init__.py
└── nested
    └── nested_resource.bin
```

2.3.6.1 Access to resources in the file system

You can access these resources like so:

```
>>> from distlib.resources import finder
>>> f = finder('foofoo')
>>> r = f.find('foo_resource.bin')
>>> r.is_container
False
>>> r.size
10
```

(continues on next page)

(continued from previous page)

```

>>> r.bytes
b'more_data\n'
>>> s = r.as_stream()
>>> s.read()
b'more_data\n'
>>> s.close()
>>> r = f.find('nested')
>>> r.is_container
True
>>> r.resources
{'nested_resource.bin'}
>>> r = f.find('nested/nested_resource.bin')
>>> r.size
12
>>> r.bytes
b'nested data\n'
>>> f = finder('foofoo.bar')
>>> r = f.find('bar_resource.bin')
>>> r.is_container
False
>>> r.bytes
b'data\n'

```

2.3.6.2 Access to resources in the .zip files

It works the same way if the package is in a .zip file. Given the zip file `foo.zip`:

```

$ unzip -l foo.zip
Archive:  foo.zip
  Length      Date    Time    Name
-----
      10  2012-09-20  21:34   foo/foo_resource.bin
       8  2012-09-20  21:42   foo/__init__.py
      14  2012-09-20  21:42   foo/bar/baz.py
       8  2012-09-20  21:42   foo/bar/__init__.py
       5  2012-09-20  21:33   foo/bar/bar_resource.bin
-----
      45                          5 files

```

You can access its resources as follows:

```

>>> import sys
>>> sys.path.append('foo.zip')
>>> from distlib.resources import finder
>>> f = finder('foo')
>>> r = f.find('foo_resource.bin')
>>> r.is_container
False
>>> r.size
10
>>> r.bytes
'more_data\n'

```

and so on.

2.3.6.3 Iterating over resources

You can iterate over resources as shown in the following example:

```
>>> from distlib.resources import finder
>>> f = finder('foofoo')
>>> iterator = f.iterator('')
>>> for r in iterator: print('%-20s %s' % (r.name, r.is_container))
...
foo_resource.bin      True
__init__.py          False
bar                   True
bar/bar_resource.bin  False
bar/baz.py            False
bar/__init__.py       False
nested                True
nested/nested_resource.bin False
```

It works with zipped resources, too:

```
>>> import sys
>>> sys.path.append('foo.zip')
>>> from distlib.resources import finder
>>> f = finder('foo')
>>> iterator = f.iterator('')
>>> for r in iterator: print('%-20s %s' % (r.name, r.is_container))
...
foo_resource.bin      True
__init__.py          False
bar                   True
bar/bar_resource.bin  False
bar/baz.py            False
bar/__init__.py       False
```

2.3.7 Using the scripts API

You can use the `distlib.scripts` API to install scripts. Installing scripts is slightly more involved than just copying files:

- You may need to adjust shebang lines in scripts to point to the interpreter to be used to run scripts. This is important in virtual environments (venvs), and also in other situations where you may have multiple Python installations on a single computer.
- On Windows, on systems where the [PEP 397](#) launcher isn't installed, it is not easy to ensure that the correct Python interpreter is used for a script. You may wish to install native Windows executable launchers which run the correct interpreter, based on a shebang line in the script.

2.3.7.1 Specifying scripts to install

To install scripts, create a *ScriptMaker* instance, giving it the source and target directories for scripts:

```
>>> from distlib.scripts import ScriptMaker
>>> maker = ScriptMaker(source_dir, target_dir)
```

You can then install a script `foo.py` like this:

```
>>> maker.make('foo.py')
```

The string passed to `make` can take one of the following forms:

- A filename, relative to the source directory for scripts, such as `foo.py` or `subdir/bar.py`.
- A reference to a callable, given in the form:

```
name = some_package.some_module:some_callable [flags]
```

where the *flags* part is optional.

For more information about flags, see *Flag formats*.

Note that this format is exactly the same as for export entries in a distribution (see *Exporting things from Distributions*).

When this form is passed to the `ScriptMaker.make()` method, a Python stub script is created with the appropriate shebang line and with code to load and call the specified callable with no arguments, returning its value as the return code from the script.

You can pass an optional `options` dictionary to the `make()` method. This is meant to contain options which control script generation. There are two options currently in use:

`gui`: This Boolean value, if `True`, indicates on Windows that a Windows executable launcher (rather than a launcher which is a console application) should be used. (This only applies if `add_launchers` is `true`.)

`interpreter_args`: If provided, this should be a list of strings which are added to the shebang line following the interpreter. If there are values with spaces, you will need to surround them with double quotes.

Note: Use of this feature may affect portability, since POSIX does not standardise how these arguments are passed to the interpreter (see https://en.wikipedia.org/wiki/Shebang_line#Portability for more information).

For example, you can pass `{'gui': True}` to generate a windowed script.

2.3.7.2 Wrapping callables with scripts

Let's see how wrapping a callable works. Consider the following file:

```
$ cat scripts/foo.py
def main():
    print('Hello from foo')

def other_main():
    print('Hello again from foo')

we can try wrapping ``main`` and ``other_main`` as callables::

>>> from distlib.scripts import ScriptMaker
>>> maker = ScriptMaker('scripts', '/tmp/scratch')
>>> maker.make_multiple(('foo = foo:main', 'bar = foo:other_main'))
['/tmp/scratch/foo', '/tmp/scratch/bar']
>>>

we can inspect the resulting scripts. First, ``foo``::
```

(continues on next page)

(continued from previous page)

```

$ ls /tmp/scratch/
bar foo
$ cat /tmp/scratch/foo
#!/usr/bin/python

if __name__ == '__main__':
    import sys, re

    def _resolve(module, func):
        __import__(module)
        mod = sys.modules[module]
        parts = func.split('.')
        result = getattr(mod, parts.pop(0))
        for p in parts:
            result = getattr(result, p)
        return result

    try:
        sys.argv[0] = re.sub('-script.pyw?$', '', sys.argv[0])

        func = _resolve('foo', 'main')
        rc = func() # None interpreted as 0
    except Exception as e: # only supporting Python >= 2.6
        sys.stderr.write('%s\n' % e)
        rc = 1
    sys.exit(rc)

```

The other script, `bar`, is different only in the essentials:

```

$ diff /tmp/scratch/foo /tmp/scratch/bar
16c16
<         func = _resolve('foo', 'main')
---
>         func = _resolve('foo', 'other_main')

```

2.3.7.3 Specifying a custom executable for shebangs

You may need to specify a custom executable for shebang lines. To do this, set the `executable` attribute of a *ScriptMaker* instance to the absolute Unicode path of the executable which you want to be written to the shebang lines of scripts. If not specified, the executable running the *ScriptMaker* code is used. If the value has spaces, you should surround it with double quotes. You can use the `enquote_executable()` function for this.

Changed in version 0.3.1: The `enquote_executable()` function was an internal function `_enquote_executable()` in earlier versions.

For relocatable `.exe` files under Windows, you can specify the location of the python executable relative to the script by putting `<launcher_dir>` as the beginning of the executable path. Since windows places `python.exe` in the root install directory and the application scripts in the *Scripts* subdirectory, setting `maker.executable = r"<launcher_dir>\\.python.exe"` will allow you to move a python installation which is installed together with an application to a different path or a different machine and the `.exe` files will still run.

2.3.7.4 Generating variants of a script

When installing a script `foo`, it is not uncommon to want to install version-specific variants such as `foo3` or `foo-3`. You can control exactly which variants of the script get written through the *ScriptMaker* instance's `variants`

attribute. This defaults to `set(('', 'X.Y'))`, which means that by default a script `foo` would be installed as `foo` and `foo-3.2` under Python 3.2. If the value of the `variants` attribute were set to `set(('', 'X', 'X.Y'))` then the `foo` script would be installed as `foo`, `foo3` and `foo-3.2` when run under Python 3.2.

Note: If you need to generate variants for a different version of Python than the one running the *ScriptMaker* code, set the `version_info` attribute of the *ScriptMaker* instance to a 2-tuple holding the major and minor version numbers of the target Python version.

New in version 0.3.1.

2.3.7.5 Avoiding overwriting existing scripts

In some scenarios, you might overwrite existing scripts when you shouldn't. For example, if you use Python 2.7 to install a distribution with script `foo` in the user site (see [PEP 370](#)), you will write (on POSIX) scripts `~/.local/bin/foo` and `~/.local/bin/foo-2.7`. If you then install the same distribution with Python 3.2, you would write (on POSIX) scripts `~/.local/bin/foo` and `~/.local/bin/foo-3.2`. However, by overwriting the `~/.local/bin/foo` script, you may prevent verification or removal of the 2.7 installation to fail, because the overwritten file may be different (and so have a different hash from what was computed during the 2.7 installation).

To control overwriting of generated scripts this way, you can use the `clobber` attribute of a *ScriptMaker* instance. This is set to `False` by default, which prevents overwriting; to force overwriting, set it to `True`.

2.3.7.6 Generating windowed scripts on Windows

The `make()` and `make_multiple()` methods take an optional second `options` argument, which can be used to control script generation. If specified, this should be a dictionary of options. Currently, only the value for the `gui` key in the dictionary is inspected: if `True`, it generates scripts with `.pyw` extensions (rather than `.py`) and, if `add_launchers` is specified as `True` in the *ScriptMaker* instance, then (on Windows) a windowed native executable launcher is created (otherwise, the native executable launcher will be a console application).

2.3.8 Using the version API

2.3.8.1 Overview

The `NormalizedVersion` class implements a [PEP 440](#) compatible version:

```
>>> from distlib.version import NormalizedVersion
>>> v1 = NormalizedVersion('1.0')
>>> v2 = NormalizedVersion('1.0a1')
>>> v3 = NormalizedVersion('1.0b1')
>>> v4 = NormalizedVersion('1.0c1')
>>> v5 = NormalizedVersion('1.0.post1')
>>>
```

These sort in the expected order:

```
>>> v2 < v3 < v4 < v1 < v5
True
>>>
```

You can't pass any old thing as a version number:


```
>>> NormalizedVersion('foo')
Traceback (most recent call last):
File "<stdin>", line 1, in <module>
File "distlib/version.py", line 49, in __init__
    self._parts = parts = self.parse(s)
File "distlib/version.py", line 254, in parse
    def parse(self, s): return normalized_key(s)
File "distlib/version.py", line 199, in normalized_key
    raise UnsupportedVersionError(s)
distlib.version.UnsupportedVersionError: foo
>>>
```

2.3.8.2 Matching versions against constraints

The `NormalizedMatcher` is used to match version constraints against versions:

```
>>> from distlib.version import NormalizedMatcher
>>> m = NormalizedMatcher('foo (1.0b1)')
>>> m
NormalizedMatcher('foo (1.0b1)')
>>> [m.match(v) for v in v1, v2, v3, v4, v5]
[False, False, True, False, False]
>>>
```

Specifying `'foo (1.0b1)'` is equivalent to specifying `'foo (==1.0b1)'`, i.e. only the exact version is matched. You can also specify inequality constraints:

```
>>> m = NormalizedMatcher('foo (<1.0c1)')
>>> [m.match(v) for v in v1, v2, v3, v4, v5]
[False, True, True, False, False]
>>>
```

and multiple constraints:

```
>>> m = NormalizedMatcher('foo (>= 1.0b1, <1.0.post1)')
>>> [m.match(v) for v in v1, v2, v3, v4, v5]
[True, False, True, True, False]
>>>
```

You can do exactly the same thing as above with `setuptools/` `distribute` version numbering (use `LegacyVersion` and `LegacyMatcher`) or with semantic versioning (use `SemanticVersion` and `SemanticMatcher`). However, you can't mix and match versions of different types:

```
>>> from distlib.version import SemanticVersion, LegacyVersion
>>> nv = NormalizedVersion('1.0.0')
>>> lv = LegacyVersion('1.0.0')
>>> sv = SemanticVersion('1.0.0')
>>> lv == sv
Traceback (most recent call last):
File "<stdin>", line 1, in <module>
File "distlib/version.py", line 61, in __eq__
    self._check_compatible(other)
File "distlib/version.py", line 58, in _check_compatible
    raise TypeError('cannot compare %r and %r' % (self, other))
TypeError: cannot compare LegacyVersion('1.0.0') and SemanticVersion('1.0.0')
```

(continues on next page)

(continued from previous page)

```
>>> nv == sv
Traceback (most recent call last):
File "<stdin>", line 1, in <module>
File "distlib/version.py", line 61, in __eq__
    self._check_compatible(other)
File "distlib/version.py", line 58, in _check_compatible
    raise TypeError('cannot compare %r and %r' % (self, other))
TypeError: cannot compare NormalizedVersion('1.0.0') and SemanticVersion('1.0.0')
>>>
```

2.3.9 Using the wheel API

You can use the `distlib.wheel` package to build and install from files in the Wheel format, defined in [PEP 427](#).

2.3.9.1 Building wheels

Building wheels is straightforward:

```
from distlib.wheel import Wheel

wheel = Wheel()

# Set the distribution's identity
wheel.name = 'name_of_distribution'
wheel.version = '0.1'

# Indicate where the files to go in the wheel are to be found
paths = {
    'prefix': '/path/to/installation/prefix',
    'purelib': '/path/to/purelib', # only one of purelib
    'platlib': '/path/to/platlib', # or platlib should be set
    'scripts': '/path/to/scripts',
    'headers': '/path/to/headers',
    'data': '/path/to/data',
}

wheel.dirname = '/where/you/want/the/wheel/to/go'
# Now build
wheel.build(paths)
```

If the `'data'`, `'headers'` and `'scripts'` keys are absent, or point to paths which don't exist, nothing will be added to the wheel for these categories. The `'prefix'` key and one of `'purelib'` or `'platlib'` *must* be provided, and the paths referenced should exist.

2.3.9.2 Customising tags during build

By default, the `build()` method will use default tags depending on whether or not the build is a pure-Python build:

- For a pure-Python build, the `pyver` will be set to `pyXY` where `XY` is the version of the building Python. The `abi` tag will be `none` and the `arch` tag will be `any`.
- For a build which is not pure-Python (i.e. contains C code), the `pyver` will be set to e.g. `cpXY`, and the `abi` and `arch` tags will be set according to the building Python.

If you want to override these default tags, you can pass a `tags` parameter to the `build()` method which has the tags you want to declare. For example, for a pure build where we know that the code in the wheel will be compatible with the major version of the building Python:

```
from wheel import PYVER
tags = {
    'pyver': [PYVER[:-1], PYVER],
}
wheel.build(paths, tags)
```

This would set the `pyver` tags to be `pyX.pyXY` where `X` and `Y` relate to the building Python. You can similarly pass values using the `abi` and `arch` keys in the `tags` dictionary.

2.3.9.3 Specifying a wheel's version

You can also specify a particular “Wheel-Version” to be written to the wheel metadata of a wheel you’re building. Simply pass a (major, minor) tuple in the `wheel_version` keyword argument to `build()`. If not specified, the most recent version supported is written.

2.3.9.4 Installing from wheels

Installing from wheels is similarly straightforward. You just need to indicate where you want the files in the wheel to be installed:

```
from distlib.wheel import Wheel
from distlib.scripts import ScriptMaker

wheel = Wheel('/path/to/my_dist-0.1-py32-none-any.whl')

# Indicate where the files in the wheel are to be installed to.
# All the keys should point to writable paths.
paths = {
    'prefix': '/path/to/installation/prefix',
    'purelib': '/path/to/purelib',
    'platlib': '/path/to/platlib',
    'scripts': '/path/to/scripts',
    'headers': '/path/to/headers',
    'data': '/path/to/data',
}

maker = ScriptMaker(None, None)
# You can specify a custom executable in script shebang lines, whether
# or not to install native executable launchers, whether to do a dry run
# etc. by setting attributes on the maker, wither when creating it or
# subsequently.

# Now install. The method accepts optional keyword arguments:
#
# - A ``warner`` argument which, if specified, should be a callable that
#   will be called with (software_wheel_version, file_wheel_version) if
#   they differ. They will both be in the form (major_ver, minor_ver).
#
# - A ``lib_only`` argument which indicates that only the library portion
#   of the wheel should be installed - no scripts, header files or
#   non-package data.
```

(continues on next page)

(continued from previous page)

```
wheel.install(paths, maker)
```

Only one of the `purelib` or `platlib` paths will actually be written to (assuming that they are different, which isn't often the case). Which one it is depends on whether the wheel metadata declares that the wheel contains pure Python code.

2.3.9.5 Verifying wheels

You can verify that a wheel's contents match the declared contents in the wheel's `RECORD` entry, by calling the `verify()` method. This will raise a `DistlibException` if a size or digest mismatch is found.

2.3.9.6 Modifying wheels

Note: In an ideal world one would not need to modify wheels, but in the short term there might be a need to do so (for example, to add dependency information which is missing). If you are working with wheels on your own projects, you *shouldn't* use the method described here, as you will have full control of the wheels you build yourself. However, if you are working with third party wheels which you don't build yourself but you need to modify in some way, then the approach described below might be useful.

You can update existing wheels with `distlib` by calling the `update()` method of a wheel. This is called as follows:

```
modified = wheel.update(modifier, dest_dir, **kwargs)
```

where the `modifier` is a callable which you specify, and `kwargs` are options you want to pass to it (currently, the `update()` method passes `kwargs` unchanged to the `modifier`). The `dest_dir` argument indicates where you want any new wheel to be written - it is optional and if not specified, *the existing wheel will be overwritten*.

The `update()` method extracts the entire contents of the wheel to a temporary location, and then calls `modifier` as follows:

```
modified = modifier(path_map, **kwargs)
```

where `path_map` is a dictionary mapping archive paths to the location of the corresponding extracted archive entry, and `kwargs` is whatever was passed to the `update` method. If the `modifier` returns `True`, a new wheel is built from the (possibly updated) contents of `path_map` and its path name. The passed `path_map` will contain all of the wheel's entries other than the `RECORD` entry (which will be recreated if a new wheel is built).

For example, if you wanted to add `numpy` as a dependency in a `scipy` wheel, you might do something like this:

```
def add_numpy_dependency(path_map, **kwargs):
    mdp_path = path_map['scipy-0.11.dist-info/pydist.json']
    md = Metadata(path=mdp_path)
    md.add_requirements(['numpy'])
    md.write(path=mdp_path)
    return True

wheel = Wheel('scipy-0.11-py27-abi3-linux_x86_64.whl')
wheel.update(add_numpy_dependency)
```

In the above example, the modifier doesn't actually use `kwargs`, but you could pass useful information which can be used to control the modifier's operation. For example, you might make the function work with other distributions than `scipy`, or other versions of `scipy`:

```
def add_numpy_dependency(path_map, **kwargs):
    name = kwargs.get('name', 'scipy')
    version = kwargs.get('version', '0.11')
    key = '%s-%s.dist-info/pydist.json' % (name, version)
    mdpath = path_map[key]
    md = Metadata(path=mdpath)
    md.add_requirements(['numpy'])
    md.write(path=mdpath)
    return True
```

2.3.9.7 Mounting wheels

One of Python's perhaps under-used features is `zipimport`, which gives the ability to import Python source from `.zip` files. Since wheels are `.zip` files, they can sometimes be used to provide functionality without needing to be installed. Whereas `.zip` files contain no convention for indicating compatibility with a particular Python, wheels *do* contain this compatibility information. Thus, it is possible to check if a wheel can be directly imported from, and the wheel support in `distlib` allows you to take advantage of this using the `mount()` and `unmount()` methods. When you mount a wheel, its absolute path name is added to `sys.path`, allowing the Python code in it to be imported. (A `DistlibException` is raised if the wheel isn't compatible with the Python which calls the `mount()` method.)

The `mount()` method takes an optional keyword parameter `append` which defaults to `False`, meaning the a mounted wheel's pathname is added to the beginning of `sys.path`. If you pass `True`, the pathname is appended to `sys.path`.

The `mount()` method goes further than just enabling Python imports – any C extensions in the wheel are also made available for import. For this to be possible, the wheel has to be built with additional metadata about extensions – a JSON file called `EXTENSIONS` which serialises an extension mapping dictionary. This maps extension module names to the names in the wheel of the shared libraries which implement those modules.

Running `unmount()` on the wheel removes its absolute pathname from `sys.path` and makes its C extensions, if any, also unavailable for import.

Note: The C extension mounting functionality may not work in all cases, though it should work in a useful subset of cases. Use with care. Note that extension information is currently only available in wheels built using `distil` – for wheels built using e.g. `pip`, this note will not apply, because C extensions will never be available for import.

- There might be subtle differences in binary compatibility between the extension and the running Python, because the compatibility tag framework currently does not capture all the relevant ABI information. This is a situation which can be expected to improve over time.
 - If the extension uses custom dynamically linked libraries which are bundled with the extension, it may not be found by the dynamic loading machinery, for reasons that are platform-dependent. In such cases, you should have a good understanding of how dynamic loading works on your platforms, before taking advantage of this feature.
-

2.3.9.8 Using vanilla pip to build wheels for existing distributions on PyPI

Although work is afoot to add wheel support to `pip`, you don't need this to build wheels for existing PyPI distributions if you use `distlib`. The following script shows how you can use an unpatched, vanilla `pip` to build wheels:

```
#!/usr/bin/env python
# -*- coding: utf-8 -*-
#
# Copyright (C) 2013 Vinay Sajip. License: MIT
#

import logging
import optparse      # for 2.6
import os
import re
import shutil
import subprocess
import sys
import tempfile

logger = logging.getLogger('wheeler')

from distlib.compat import configparser, filter
from distlib.database import DistributionPath, Distribution, make_graph
from distlib.locators import (JSONLocator, SimpleScrapingLocator,
                              AggregatingLocator, DependencyFinder)
from distlib.manifest import Manifest
from distlib.metadata import Metadata
from distlib.util import parse_requirement, get_package_data
from distlib.wheel import Wheel

EGG_INFO_RE = re.compile(r'(-py\d\.\d)?\.egg-info', re.I)

INSTALLED_DISTS = DistributionPath(include_egg=True)

def get_requirements(data):
    lines = []
    for line in data.splitlines():
        line = line.strip()
        if not line or line[0] == '#':
            continue
        lines.append(line)
    reqts = []
    extras = {}
    result = {'install': reqts, 'extras': extras}
    for line in lines:
        if line[0] != '[':
            reqts.append(line)
        else:
            i = line.find(']', 1)
            if i < 0:
                raise ValueError('unrecognised line: %r' % line)
            extra = line[1:i]
            extras[extra] = reqts = []
    return result

def convert_egg_info(libdir, prefix, options):
    files = os.listdir(libdir)
    ei = list(filter(lambda d: d.endswith('.egg-info'), files))[0]
    olddn = os.path.join(libdir, ei)
```

(continues on next page)

(continued from previous page)

```

di = EGG_INFO_RE.sub('.dist-info', ei)
newdn = os.path.join(libdir, di)
os.rename(olddn, newdn)
if options.compatible:
    renames = {}
else:
    renames = {
        'entry_points.txt': 'EXPORTS',
    }
excludes = set([
    'SOURCES.txt',          # of no interest in/post WHEEL
    'installed-files.txt',  # replaced by RECORD, so not needed
    'requires.txt',        # added to METADATA, so not needed
    'PKG-INFO',            # replaced by METADATA
    'not-zip-safe',        # not applicable
])
files = os.listdir(newdn)
metadata = mdname = reqts = None
for oldfn in files:
    pn = os.path.join(newdn, oldfn)
    if oldfn in renames:
        os.rename(pn, os.path.join(newdn, renames[oldfn]))
    else:
        if oldfn == 'requires.txt':
            with open(pn, 'r') as f:
                reqts = get_requirements(f.read())
        elif oldfn == 'PKG-INFO':
            metadata = Metadata(path=pn)
            pd = get_package_data(metadata.name, metadata.version)
            metadata = Metadata(mapping=pd['index-metadata'])
            mdname = os.path.join(newdn, 'pydist.json')
            if oldfn in excludes or not options.compatible:
                os.remove(pn)
if metadata:
    # Use Metadata 1.2 or later
    metadata.provides += ['%s (%s)' % (metadata.name,
                                       metadata.version)]
    # Update if not set up by get_package_data
    if reqts and not metadata.run_requires:
        metadata.dependencies = reqts
    metadata.write(path=mdname)
manifest = Manifest(os.path.dirname(libdir))
manifest.findall()
paths = manifest.allfiles
dp = DistributionPath([libdir])
dist = next(dp.get_distributions())
dist.write_installed_files(paths, prefix)

def install_dist(distname, workdir, options):
    pfx = '--install-option='
    purelib = pfx + '--install-purelib=%s/purelib' % workdir
    platlib = pfx + '--install-platlib=%s/platlib' % workdir
    headers = pfx + '--install-headers=%s/headers' % workdir
    scripts = pfx + '--install-scripts=%s/scripts' % workdir
    data = pfx + '--install-data=%s/data' % workdir
    # Use the pip adjacent to sys.executable, if any (for virtualenvs)

```

(continues on next page)

(continued from previous page)

```

d = os.path.dirname(sys.executable)
files = filter(lambda o: o in ('pip', 'pip.exe'), os.listdir(d))
if not files:
    prog = 'pip'
else:
    prog = os.path.join(d, next(files))
cmd = [prog, 'install',
        '--no-deps', '--quiet',
        '--index-url', 'http://pypi.org/simple/',
        '--timeout', '3', '--default-timeout', '3',
        purelib, platlib, headers, scripts, data, distname]
result = {
    'scripts': os.path.join(workdir, 'scripts'),
    'headers': os.path.join(workdir, 'headers'),
    'data': os.path.join(workdir, 'data'),
}
print('Pipping %s ...' % distname)
p = subprocess.Popen(cmd, shell=False, stdout=sys.stdout,
                     stderr=subprocess.STDOUT)
stdout, _ = p.communicate()
if p.returncode:
    raise ValueError('pip failed to install %s:\n%s' % (distname, stdout))
for dn in ('purelib', 'platlib'):
    libdir = os.path.join(workdir, dn)
    if os.path.isdir(libdir):
        result[dn] = libdir
        break
convert_egg_info(libdir, workdir, options)
dp = DistributionPath([libdir])
dist = next(dp.get_distributions())
md = dist.metadata
result['name'] = md.name
result['version'] = md.version
return result

def build_wheel(distname, options):
    result = None
    r = parse_requirement(distname)
    if not r:
        print('Invalid requirement: %r' % distname)
    else:
        dist = INSTALLED_DISTs.get_distribution(r.name)
        if dist:
            print('Can\'t build a wheel from already-installed '
                  'distribution %s' % dist.name_and_version)
        else:
            workdir = tempfile.mkdtemp()  # where the Wheel input files will live
            try:
                paths = install_dist(distname, workdir, options)
                paths['prefix'] = workdir
                wheel = Wheel()
                wheel.name = paths.pop('name')
                wheel.version = paths.pop('version')
                wheel.dirname = options.destdir
                wheel.build(paths)
                result = wheel
            
```

(continues on next page)

(continued from previous page)

```

        finally:
            shutil.rmtree(workdir)
    return result

def main(args=None):
    parser = optparse.OptionParser(usage='%prog [options] requirement [requirement ...
→]')
    parser.add_option('-d', '--dest', dest='destdir', metavar='DESTDIR',
                      default=os.getcwd(), help='Where you want the wheels '
                      'to be put.')
    parser.add_option('-n', '--no-deps', dest='deps', default=True,
                      action='store_false',
                      help='Don\'t build dependent wheels.')
    options, args = parser.parse_args(args)
    options.compatible = True    # may add flag to turn off later
    if not args:
        parser.print_usage()
    else:
        # Check if pip is available; no point in continuing, otherwise
        try:
            with open(os.devnull, 'w') as f:
                p = subprocess.call(['pip', '--version'], stdout=f, stderr=subprocess.
→STDOUT)
        except Exception:
            p = 1
        if p:
            print('pip appears not to be available. Wheeler needs pip to '
                  'build wheels.')
            return 1
        if options.deps:
            # collect all the requirements, including dependencies
            u = 'http://pypi.org/simple/'
            locator = AggregatingLocator(JSONLocator(),
                                         SimpleScrapingLocator(u, timeout=3.0),
                                         scheme='legacy')
            finder = DependencyFinder(locator)
            wanted = set()
            for arg in args:
                r = parse_requirement(arg)
                if not r.constraints:
                    dname = r.name
                else:
                    dname = '%s (%s)' % (r.name, ', '.join(r.constraints))
                print('Finding the dependencies of %s ...' % arg)
                dists, problems = finder.find(dname)
                if problems:
                    print('There were some problems resolving dependencies '
                          'for %r.' % arg)
                    for _, info in problems:
                        print('  Unsatisfied requirement %r' % info)
                wanted |= dists
            want_ordered = True    # set to False to skip ordering
            if not want_ordered:
                wanted = list(wanted)
            else:
                graph = make_graph(wanted, scheme=locator.scheme)

```

(continues on next page)

(continued from previous page)

```

slist, cycle = graph.topological_sort()
if cycle:
    # Now sort the remainder on dependency count.
    cycle = sorted(cycle, reverse=True,
                   key=lambda d: len(graph.reverse_list[d]))
wanted = slist + cycle

# get rid of any installed distributions from the list
for w in list(wanted):
    dist = INSTALLED_DISTS.get_distribution(w.name)
    if dist or w.name in ('setuptools', 'distribute'):
        wanted.remove(w)
        s = w.name_and_version
        print('Skipped already-installed distribution %s' % s)

# converted wanted list to pip-style requirements
args = ['%s==%s' % (dist.name, dist.version) for dist in wanted]

# Now go build
built = []
for arg in args:
    wheel = build_wheel(arg, options)
    if wheel:
        built.append(wheel)
if built:
    if options.dest_dir == os.getcwd():
        dest = ''
    else:
        dest = ' in %s' % options.dest_dir
    print('The following wheels were built%s:' % dest)
    for wheel in built:
        print(' %s' % wheel.filename)

if __name__ == '__main__':
    logging.basicConfig(format='%(levelname)-8s %(name)s %(message)s',
                        filename='wheeler.log', filemode='w')
    try:
        rc = main()
    except Exception as e:
        print('Failed - sorry! Reason: %s\nPlease check the log.' % e)
        logger.exception('Failed.')
        rc = 1
    sys.exit(rc)

```

This script, `wheeler.py`, is also available [here](#). Note that by default, it downloads dependencies of any distribution you specify and builds separate wheels for each distribution. It's smart about not repeating work if dependencies are common across multiple distributions you specify:

```

$ python wheeler.py sphinx flask
Finding the dependencies of sphinx ...
Finding the dependencies of flask ...
Pipping Jinja2==2.6 ...
Pipping docutils==0.10 ...
Pipping Pygments==1.6 ...
Pipping Werkzeug==0.8.3 ...
Pipping Sphinx==1.1.3 ...

```

(continues on next page)

(continued from previous page)

```
Pipping Flask==0.9 ...
The following wheels were built:
  Jinja2-2.6-py27-none-any.whl
  docutils-0.10-py27-none-any.whl
  Pygments-1.6-py27-none-any.whl
  Werkzeug-0.8.3-py27-none-any.whl
  Sphinx-1.1.3-py27-none-any.whl
  Flask-0.9-py27-none-any.whl
```

Note that the common dependency – Jinja2 – was only built once.

You can opt to not build dependent wheels by specifying `--no-deps` on the command line.

Note that the script also currently uses an `http:` URL for PyPI – this may need to change to an `https:` URL in the future.

Note: It can't be used to build wheels from existing distributions, as `pip` will either refuse to install to custom locations (because it views a distribution as already installed), or will try to upgrade and thus uninstall the existing distribution, even though installation is requested to a custom location (and uninstallation is not desirable). For best results, run it in a fresh `venv`:

```
$ my_env/bin/python wheeler.py some_dist
```

It should use the `venv`'s `pip`, if one is found.

2.3.10 Using the manifest API

You can use the `distlib.manifest` API to construct lists of files when creating distributions. This functionality is an improved version of the equivalent functionality in `distutils`, where it was not a public API.

You can create instances of the `Manifest` class to work with a set of files rooted in a particular directory:

```
>>> from distlib.manifest import Manifest
>>> manifest = Manifest('/path/to/my/sources')
```

This sets the `base` attribute to the passed in root directory. You can add one or multiple files using names relative to the base directory:

```
>>> manifest.add('abc')
>>> manifest.add_many(['def', 'ghi'])
```

As a result of the above two statements, the manifest will consist of `'/path/to/my/sources/abc'`, `'/path/to/my/sources/def'` and `'/path/to/my/sources/ghi'`. No check is made regarding the existence of these files.

You can get all the files below the base directory of the manifest:

```
>>> manifest.findall()
```

This will populate the `allfiles` attribute of `manifest` with a list of all files in the directory tree rooted at the base. However, the manifest is still empty:

```
>>> manifest.files
>>> set()
```

You can populate the manifest – the `files` attribute – by running a number of *directives*, using the `process_directive()` method. Each directive will either add files from `allfiles` to `files`, or remove files from `allfiles` if they were added by a previous directive. A directive is a string which must have a specific syntax: malformed lines will result in a `DistlibException` being raised. The following directives are available: they are compatible with the syntax of `MANIFEST.in` files processed by `distutils`.

Consider the following directory tree:

```
testsrc/
├── keep
│   └── keep.txt
├── LICENSE
├── README.txt
├── subdir
│   ├── lose
│   │   └── lose.txt
│   ├── somedata.txt
│   └── subsubdir
│       └── somedata.bin
```

This will be used to illustrate how the directives work, in the following sections.

2.3.10.1 The `include` directive

This takes the form of the word `include` (case-sensitive) followed by a number of file-name patterns (as used in `MANIFEST.in` in `distutils`). All files in `allfiles` matching the patterns (considered relative to the base directory) are added to `files`. For example:

```
>>> manifest.process_directive('include R*.txt LIC* keep/*.txt')
```

This will add `README.txt`, `LICENSE` and `keep/keep.txt` to the manifest.

2.3.10.2 The `exclude` directive

This takes the form of the word `exclude` (case-sensitive) followed by a number of file-name patterns (as used in `MANIFEST.in` in `distutils`). All files in `files` matching the patterns (considered relative to the base directory) are removed from `files`. For example:

```
>>> manifest.process_directive('exclude LIC*')
```

This will remove `'LICENSE'` from the manifest, as it was added in the section above.

2.3.10.3 The `global-include` directive

This works just like `include`, but will add matching files at all levels of the directory tree:

```
>>> manifest.process_directive('global-include *.txt')
```

This will add `subdir/somedata.txt` and `subdir/lose/lose.txt` from the manifest.

2.3.10.4 The `global-exclude` directive

This works just like `exclude`, but will remove matching files at all levels of the directory tree:

```
>>> manifest.process_directive('global-exclude l*.txt')
```

This will remove `subdir/lose/lose.txt` from the manifest.

2.3.10.5 The `recursive-include` directive

This directive takes a directory name (relative to the base) and a set of patterns. The patterns are used as in `global-include`, but only for files under the specified directory:

```
>>> manifest.process_directive('recursive-include subdir l*.txt')
```

This will add `subdir/lose/lose.txt` back to the manifest.

2.3.10.6 The `recursive-exclude` directive

This works like `recursive-include`, but excludes matching files under the specified directory if they were already added by a previous directive:

```
>>> manifest.process_directive('recursive-exclude subdir lose*')
```

This will remove `subdir/lose/lose.txt` from the manifest again.

2.3.10.7 The `graft` directive

This directive takes the name of a directory (relative to the base) and copies all the names under it from `allfiles` to `files`.

2.3.10.8 The `prune` directive

This directive takes the name of a directory (relative to the base) and removes all the names under it from `files`.

2.4 Next steps

You might find it helpful to look at information about *Distlib's design* – or peruse the *API Reference*.

This is the section containing some discussion of how `distlib`'s design was arrived at, as and when time permits.

3.1 The `locators` API

This section describes the design of the `distlib` API relating to accessing distribution metadata, whether stored locally or in indexes like PyPI.

3.1.1 The problem we're trying to solve

People who use distributions need to locate, download and install them. Distributions can be found in a number of places, such as:

- An Internet index such as [The Python Packages Index \(PyPI\)](#), or a mirror thereof.
- Other Internet resources, such as the developer's website, or a source code repository such as GitHub, BitBucket, Google Code or similar.
- File systems, whether local to one computer or shared between several.
- Distributions which have already been installed, and are available in the `sys.path` of a running Python interpreter.

When we're looking for distributions, we don't always know exactly what we want: often, we just want the latest version, but it's not uncommon to want a specific older version, or perhaps the most recent version that meets some constraints on the version. Since we need to be concerned with matching versions, we need to consider the version schemes in use (see [The version API](#)).

It's useful to separate the notion of a *project* from a distribution: The project is the version-independent part of the distribution, i.e. it's described by the *name* of the distribution and encompasses all released distributions which use that name.

We often don't just want a single distribution, either: a common requirement, when installing a distribution, is to locate all distributions that it relies on, which aren't already installed. So we need a *dependency finder*, which itself needs

to locate depended-upon distributions, and recursively search for dependencies until all that are available have been found.

We may need to distinguish between different types of dependencies:

- Post-installation dependencies. These are needed by the distribution after it has been installed, and is in use.
- Build dependencies. These are needed for building and/or installing the distribution, but are not needed by the distribution itself after installation.
- Test dependencies. These are only needed for testing the distribution, but are not needed by the distribution itself after installation.

When testing a distribution, we need all three types of dependencies. When installing a distribution, we need the first two, but not the third.

3.1.2 A minimal solution

3.1.2.1 Locating distributions

It seems that the simplest API to locate a distribution would look like `locate(requirement)`, where `requirement` is a string giving the distribution name and optional version constraints. Given that we know that distributions can be found in different places, it's best to consider a `Locator` class which has a `locate()` method with a corresponding signature, with subclasses for each of the different types of location that distributions inhabit. It's also reasonable to provide a default locator in a module attribute `default_locator`, and a module-level `locate()` function which calls the `locate()` method on the default locator.

Since we'll often need to locate all the versions of a project before picking one, we can imagine that a locator would need a `get_project()` method for fetching all versions of a project; and since we will be likely to want to use caching, we can assume there will be a `_get_project()` method to do the actual work of fetching the version data, which the higher-level `get_project()` will call (and probably cache). So our locator base class will look something like this:

```
class Locator(object):
    """
    Locate distributions.
    """

    def __init__(self, scheme='default'):
        """
        Initialise a locator with the specified version scheme.
        """

    def locate(self, requirement):
        """
        Locate the highest-version distribution which satisfies
        the constraints in ``requirement``, and return a
        ``Distribution`` instance if found, or else ``None``.
        """

    def get_project(self, name):
        """
        Return all known distributions for a project named ``name``,
        returning a dictionary mapping version to ``Distribution``
        instance, or an empty dictionary if nothing was found.
        Use _get_project to do the actual work, and cache the results for
        future use.
```

(continues on next page)

(continued from previous page)

```

"""

def _get_project(self, name):
    """
    Return all known distributions for a project named ``name``,
    returning a dictionary mapping version to ``Distribution``
    instance, or an empty dictionary if nothing was found.
    """

```

When attempting to `locate()`, it would be useful to pass requirement information to `get_project()` / `_get_project()`. This can be done in a `matcher` attribute which is normally `None` but set to a `distlib.version.Matcher` instance when a `locate()` call is in progress.

Note that in order to work with legacy version numbers (those not complying with PEP 440), you need to pass `scheme='legacy'` to the initializer for a locator.

3.1.2.2 Finding dependencies

A dependency finder will depend on a locator to locate dependencies. A simple approach will be to consider a `DependencyFinder` class which takes a locator as a constructor argument. It might look something like this:

```

class DependencyFinder(object):
    """
    Locate dependencies for distributions.
    """

    def __init__(self, locator):
        """
        Initialise an instance, using the specified locator
        to locate distributions.
        """

    def find(self, requirement, meta_extras=None, prereleases=False):
        """
        Find a distribution matching requirement and all distributions
        it depends on. Use the ``meta_extras`` argument to determine
        whether distributions used only for build, test etc. should be
        included in the results. Allow ``requirement`` to be either a
        :class:`Distribution` instance or a string expressing a
        requirement. If ``prereleases`` is True, treat pre-releases as
        normal releases; otherwise only return pre-releases if they're
        all that's available.

        Return a set of :class:`Distribution` instances and a set of
        problems.

        The distributions returned should be such that they have the
        :attr:`required` attribute set to ``True`` if they were
        from the ``requirement`` passed to ``find()`` , and they have the
        :attr:`build_time_dependency` attribute set to ``True`` unless they
        are post-installation dependencies of the ``requirement``.

        The problems should be a tuple consisting of the string
        ``'unsatisfied'`` and the requirement which couldn't be satisfied
        by any distribution known to the locator.
        """

```

3.2 The index API

This section describes the design of the `distlib` API relating to performing certain operations on Python package indexes like PyPI. Note that this API does not support *finding* distributions - the `locators` API is used for that.

3.2.1 The problem we're trying to solve

Operations on a package index that are commonly performed by distribution developers are:

- Register projects on the index.
- Upload distributions relating to projects on the index, with support for signed distributions.
- Upload documentation relating to projects.

Less common operations are:

- Find a list of hosts which mirror the index.
- Save a default `.pypirc` file with default username and password to use.

3.2.2 A minimal solution

The distutils approach was to have several separate command classes called `register`, `upload` and `upload_doc`, where really all that was needed was some methods. That's the approach `distlib` takes, by implementing a `PackageIndex` class with `register()`, `upload_file()` and `upload_documentation()` methods. The `PackageIndex` class contains no user interface code whatsoever: that's assumed to be the domain of the packaging tool. The packaging tool is expected to get the required information from a user using whatever means the developers of that tool deem to be the most appropriate; the required attributes are then set on the `PackageIndex` instance. (Examples of this kind of information: user name, password, whether the user wants to save a default configuration, where the signing program and its keys live.)

The minimal interface to provide the required functionality thus looks like this:

```
class PackageIndex(object):
    def __init__(self, url=None, mirror_host=None):
        """
        Initialise an instance using a specific index URL, and
        a DNS name for a mirror host which can be used to
        determine available mirror hosts for the index.
        """

    def save_configuration(self):
        """
        Save the username and password attributes of this
        instance in a default .pypirc file.
        """

    def register(self, metadata):
        """
        Register a project on the index, using the specified metadata.
        """

    def upload_file(self, metadata, filename, signer=None,
                    sign_password=None, filetype='sdist',
                    pyversion='source'):
        """
```

(continues on next page)

(continued from previous page)

```

Upload a distribution file to the index using the
specified metadata to identify it, with options
for signing and for binary distributions which are
specific to Python versions.
"""

def upload_documentation(self, metadata, doc_dir):
    """
    Upload documentation files in a specified directory
    using the specified metadata to identify it, after
    archiving the directory contents into a .zip file.
    """

```

The following additional attributes can be identified on PackageIndex instances:

- username - the username to use for authentication.
- password - the password to use for authentication.
- mirrors (read-only) - a list of hostnames of known mirrors.

3.3 The resources API

This section describes the design of the distlib API relating to accessing ‘resources’, which is a convenient label for data files associated with Python packages.

3.3.1 The problem we’re trying to solve

Developers often have a need to co-locate data files with their Python packages. Examples of these might be:

- Templates, commonly used in web applications
- Translated messages used in internationalisation/localisation

The stdlib does not provide a uniform API to access these resources. A common approach is to use `__file__` like this:

```

base = os.path.dirname(__file__)
data_filename = os.path.join(base, 'data.bin')
with open(data_filename, 'rb') as f:
    # read the data from f

```

However, this approach fails if the package is deployed in a .zip file.

To consider how to provide a minimal uniform API to access resources in Python packages, we’ll assume that the requirements are as follows:

- All resources are regarded as binary. The consuming application is expected to know how to convert resources to text, where appropriate.
- All resources are read-only.
- It should be possible to access resources either as streams, or as their entire data as a byte-string.
- Resources will have a unique, identifying name which is text. Resources will be hierarchical and named using filesystem-like paths using ‘/’ as a separator. The library will be responsible for converting resource names to the names of the underlying representations (e.g. encoding of file names corresponding to resource names).

- Some resources are containers of other resources, some are not. For example, a resource `nested/nested_resource.bin` in a package would not contain other resources, but implies the existence of a resource `nested`, which contains `nested_resource.bin`.
- Resources can only be associated with packages, not with modules. That's because with peer modules `a.py` and `b.py`, there's no obvious location for data associated only with `a`: both `a` and `b` are in the same directory. With a package, there's no ambiguity, as a package is associated with a specific directory, and no other package can be associated with that directory.
- Support should be provided for access to data deployed in the file system or in packages contained in `.zip` files, and third parties should be able to extend the facilities to work with other storage formats which support import of Python packages.
- It should be possible to access the contents of any resource through a file on the file system. This is to cater for any external APIs which need to access the resource data as files (examples would be a shared library for linking using `dlopen()` on POSIX, or any APIs which need access to resource data via OS-level file handles rather than Python streams).

3.3.2 A minimal solution

We know that we will have to deal with resources, so it seems natural that there would be a `Resource` class in the solution. From the requirements, we can see that a `Resource` would have the following:

- A `name` property identifying the resource.
- A `as_stream` method allowing access to the resource data as a binary stream. This is not a property, because a new stream is returned each time this method is called. The returned stream should be closed by the caller.
- A `bytes` property returning the entire contents of the resource as a byte string.
- A `size` property indicating the size of the resource (in bytes).
- An `is_container` property indicating whether the resource is a container of other resources.
- A `resources` property returning the names of resources contained within the resource.

The `Resource` class would be the logical place to perform sanity checks which relate to all resources. For example:

- It doesn't make sense to ask for the `bytes` or `size` properties or call the `as_stream` method of a container resource.
- It doesn't make sense to ask for the `resources` property of a resource which is *not* a container.

It seems reasonable to raise exceptions for incorrect property or method accesses.

We know that we need to support resource access in the file system as well as `.zip` files, and to support other sources of storage which might be used to import Python packages. Since import and loading of Python packages happens through [PEP 302](#) importers and loaders, we can deduce that the mechanism used to find resources in a package will be closely tied to the loader for that package.

We could consider an API for finding resources in a package like this:

```
def find_resource(pkg, resource_name):  
    # return a Resource instance for the resource
```

and then use it like this:

```
r1 = find_resource(pkg, 'foo')  
r2 = find_resource(pkg, 'bar')
```

However, we'll often have situations where we will want to get multiple resources from a package, and in certain applications we might want to implement caching or other processing of resources before returning them. The above API doesn't facilitate this, so let's consider delegating the finding of resources in a package to a *finder* for that package. Once we get a finder, we can hang on to it and ask it to find multiple resources. Finders can be extended to provide whatever caching and preprocessing an application might need.

To get a finder for a package, let's assume there's a `finder` function:

```
def finder(pkg):
    # return a finder for the specified package
```

We can use it like this:

```
f = finder(pkg)
r1 = f.find('foo')
r2 = f.find('bar')
```

The `finder` function knows what kind of finder to return for a particular package through the use of a registry. Given a package, `finder` can determine the loader for that package, and based on the type of loader, it can instantiate the right kind of finder. The registry maps loader types to callables that return finders. The callable is called with a single argument – the Python module object for the package.

Given that we have finders in the design, we can identify `ResourceFinder` and `ZipResourceFinder` classes for the two import systems we're going to support. We'll make `ResourceFinder` a concrete class rather than an interface - it'll implement getting resources from packages stored in the file system. `ZipResourceFinder` will be a subclass of `ResourceFinder`.

Since there is no loader for file system packages when the C-based import system is used, the registry will come with the following mappings:

- `type(None) -> ResourceFinder`
- `_frozen_importlib.SourceFileLoader -> ResourceFinder`
- `zipimport.zipimporter -> ZipResourceFinder`

Users of the API can add new or override existing mappings using the following function:

```
def register_finder(loader, finder_maker):
    # register `finder_maker` to make finders for packages with a loader
    # of the same type as `loader`.
```

Typically, the `finder_maker` will be a class like `ResourceFinder` or `ZipResourceFinder`, but it can be any callable which takes the Python module object for a package and returns a finder.

Let's consider in more detail what finders look like and how they interact with the `Resource` class. We'll keep the `Resource` class minimal; API users never instantiate `Resource` directly, but call a finder's `find` method to return a `Resource` instance. A finder could return an instance of a `Resource` subclass if really needed, though it shouldn't be necessary in most cases. If a finder can't find a resource, it should return `None`.

The `Resource` constructor will look like this:

```
def __init__(self, finder, name):
    self.finder = finder
    self.name = name
    # other initialisation, not specified
```

and delegate as much work as possible to its finder. That way, new import loader types can be supported just by implementing a suitable `XXXResourceFinder` for that loader type.

What a finder needs to do can be exemplified by the following skeleton for `ResourceFinder`:

```
class ResourceFinder(object):
    def __init__(self, module):
        "Initialise finder for the specified package"

    def find(self, resource_name):
        "Find and return a ``Resource`` instance or ``None``"

    def is_container(self, resource):
        "Return whether resource is a container"

    def get_bytes(self, resource):
        "Return the resource's data as bytes"

    def get_size(self, resource):
        "Return the size of the resource's data in bytes"

    def get_stream(self, resource):
        "Return the resource's data as a binary stream"

    def get_resources(self, resource):
        """
        Return the resources contained in this resource as a set of
        (relative) resource names
        """
```

3.3.3 Dealing with the requirement for access via file system files

To cater for the requirement that the contents of some resources be made available via a file on the file system, we'll assume a simple caching solution that saves any such resources to a local file system cache, and returns the filename of the resource in the cache. We need to divide the work between the finder and the cache. We'll deliver the cache function through a `Cache` class, which will have the following methods:

- A constructor which takes an optional base directory for the cache. If none is provided, we'll construct a base directory of the form:

```
<rootdir>/distlib/resource-cache
```

where `<rootdir>` is the user's home directory. On Windows, if the environment specifies a variable named `LOCALAPPDATA`, its value will be used as `<rootdir>` – otherwise, the user's home directory will be used.

- A `get()` method which takes a `Resource` and returns a file system filename, such that the contents of that named file will be the contents of the resource.
- An `is_stale()` method which takes a `Resource` and its corresponding file system filename, and returns whether the file system file is stale when compared with the resource. Knowing that cache invalidation is hard, the default implementation just returns `True`.
- A `prefix_to_dir()` method which converts a prefix to a directory name. We'll assume that for the cache, a resource path can be divided into two parts: the *prefix* and the *subpath*. For resources in a .zip file, the prefix would be the pathname of the archive, while the subpath would be the path inside the archive. For a file system resource, since it is already in the file system, the prefix would be `None` and the subpath would be the absolute path name of the resource. The `prefix_to_dir()` method's job is to convert a prefix (if not `None`) to a subdirectory in the cache that holds the cached files for all resources with that prefix. We'll delegate the determination of a resource's prefix and subpath to its finder, using a `get_cache_info()` method on finders, which takes a `Resource` and returns a (prefix, subpath) tuple.

The default implementation will use `os.splitdrive()` to see if there's a Windows drive, if present, and convert its `' : '` to `'---'`. The rest of the prefix will be converted by replacing `' / '` by `'--'`, and appending `' .cache '` to the result.

The cache will be activated when the `file_path` property of a `Resource` is accessed. This will be a cached property, and will call the cache's `get()` method to obtain the file system path.

3.4 The scripts API

This section describes the design of the `distlib` API relating to installing scripts.

3.4.1 The problem we're trying to solve

Installing scripts is slightly more involved than simply copying files from source to target, for the following reasons:

- On POSIX systems, scripts need to be made executable. To cater for scenarios where there are multiple Python versions installed on a computer, scripts need to have their shebang lines adjusted to point to the correct interpreter. This requirement is commonly found when virtual environments (venvs) are in use, but also in other multiple-interpreter scenarios.
- On Windows systems, which don't support shebang lines natively, some alternate means of finding the correct interpreter need to be provided. Following the acceptance and implementation of PEP 397, a shebang- interpreting launcher will be available in Python 3.3 and later and a standalone version of it for use with earlier Python versions is also available. However, where this can't be used, an alternative approach using executable launchers installed with the scripts may be necessary. (That is the approach taken by `setuptools`.) Windows also has two types of launchers - console applications and Windows applications. The appropriate launcher needs to be used for scripts.
- Some scripts are effectively just callable code in a Python package, with boilerplate for importing that code, calling it and returning its return value as the script's return code. It would be useful to have the boilerplate standardised, so that developers need just specify which callables to expose as scripts, and under what names, using e.g. `a name = callable` syntax. (This is the approach taken by `setuptools` for the popular `console_scripts` feature).

3.4.2 A minimal solution

Script handling in `distutils` and `setuptools` is done in two phases: 'build' and 'install'. Whether a particular packaging tool chooses to do the 'heavy lifting' of script creation (i.e. the things referred to above, beyond simple copying) in 'build' or 'install' phases, the job is the same. To abstract out just the functionality relating to scripts, in an extensible way, we can just delegate the work to a class, unimaginatively called `ScriptMaker`. Given the above requirements, together with the more basic requirement of being able to do 'dry-run' installation, we need to provide a `ScriptMaker` with the following items of information:

- Where source scripts are to be found.
- Where scripts are to be installed.
- Whether, on Windows, executable launchers should be added.
- Whether a dry-run mode is in effect.

These dictate the form that `ScriptMaker.__init__()` will take.

In addition, other methods suggest themselves for `ScriptMaker`:

- A `make()` method, which takes a *specification*, which is either a filename or a ‘wrap me a callable’ indicator which looks like this:

```
name = some_package.some_module:some_callable [ flag(=value) ... ]
```

The name would need to be a valid filename for a script, and the `some_package.some_module` part would indicate the module where the callable resides. The `some_callable` part identifies the callable, and optionally you can have flags, which the *ScriptMaker* instance must know how to interpret. One flag would be `'gui'`, indicating that the launcher should be a Windows application rather than a console application, for GUI-based scripts which shouldn’t show a console window.

The above specification is used by `setuptools` for the ‘console_scripts’ feature. See *Flag formats* for more information about flags.

It seems sensible for this method to return a list of absolute paths of files that were installed (or would have been installed, but for the dry-run mode being in effect).

- A `make_multiple()` method, which takes an iterable of specifications and just runs calls `make()` on each item iterated over, aggregating the results to return a list of absolute paths of all files that were installed (or would have been installed, but for the dry-run mode being in effect).

One advantage of having this method is that you can override it in a subclass for post-processing, e.g. to run a tool like `2to3`, or an analysis tool, over all the installed files.

- The details of the callable specification can be encapsulated in a utility function, `get_exports_entry()`. This would take a specification and return `None`, if the specification didn’t match the callable format, or an instance of `ExportEntry` if it did match.

In addition, the following attributes on a *ScriptMaker* could be further used to refine its behaviour:

- `force` to indicate when scripts should be copied from source to target even when timestamps show the target is up to date.
- `set_mode` to indicate whether, on Posix, the execute mode bits should be set on the target script.

3.4.2.1 Flag formats

Flags, if present, are enclosed by square brackets. Each flag can have the format of just an alphanumeric string, optionally followed by an ‘=’ and a value (with no intervening spaces). Multiple flags can be separated by ‘,’ and whitespace. The following would be valid flag sections:

```
[a,b,c]
[a, b, c]
[a=b, c=d, e, f=g, 9=8]
```

whereas the following would be invalid:

```
[]
[\]
[a,]
[a,,b]
[a=,b,c]
```

3.5 The version API

This section describes the design of the `distlib` API relating to versions.

3.5.1 The problem we’re trying to solve

Distribution releases are named by versions and versions have two principal uses:

- Identifying a particular release and determining whether or not it is earlier or later than some other release.
- When specifying other distributions that a distribution release depends on, specifying constraints governing the releases of those distributions that are depended upon.

In addition, qualitative information may be given by the version format about the quality of the release: e.g. alpha versions, beta versions, stable releases, hot-fixes following a stable release. The following excerpt from [PEP 386](#) defines the requirements for versions:

- It should be possible to express more than one versioning level (usually this is expressed as major and minor revision and, sometimes, also a micro revision).
- A significant number of projects need special meaning versions for “pre-releases” (such as “alpha”, “beta”, “rc”), and these have widely used aliases (“a” stands for “alpha”, “b” for “beta” and “c” for “rc”). And these pre-release versions make it impossible to use a simple alphanumerical ordering of the version string components. (e.g. 3.1a1 < 3.1)
- Some projects also need “post-releases” of regular versions, mainly for maintenance purposes, which can’t be clearly expressed otherwise.
- Development versions allow packagers of unreleased work to avoid version clashes with later stable releases.

There are a number of version schemes in use. The ones of most interest in the Python ecosystem are:

- Loose versioning in `distutils`. *Any* version number is allowed, with lexicographical ordering. No support exists for pre- and post-releases, and lexicographical ordering can be unintuitive (e.g. ‘1.10’ < ‘1.2.1’)
- Strict versioning in `distutils`, which supports slightly more structure. It allows for up to three dot-separated numeric components, and support for multiple alpha and beta releases. However, there is no support for release candidates, nor for post-release versions.
- Versioning in `setuptools/distribute`. This is described in [PEP 386](#) in [this section](#) – it’s perhaps the most widely used Python version scheme, but since it tries to be very flexible and work with a wide range of conventions, it ends up allowing a very chaotic mess of version conventions in the Python community as a whole.
- The proposed versioning scheme described in [PEP 440](#).
- [Semantic versioning](#), which is rational, simple and well-regarded in the software community in general.

Although the new versioning scheme mentioned in PEP 386 was implemented in `distutils2` and that code has been copied over to `distlib`, there are many projects on PyPI which do not conform to it, but rather to the “legacy” versioning schemes in `distutils/setuptools/distribute`. These schemes are deserving of some support not because of their intrinsic qualities, but due to their ubiquity in projects registered on PyPI. Below are some results from testing actual projects on PyPI:

```
Packages processed: 24891
Packages with no versions: 217
Packages with versions: 24674
Number of packages clean for all schemes: 19010 (77%)
Number of packages clean for PEP 386: 21072 (85%)
Number of packages clean for PEP 386 + suggestion: 23685 (96%)
Number of packages clean for legacy: 24674 (100%, by you would expect)
Number of packages clean for semantic: 19278 (78%)
```

where “+ suggestion” refers to using the suggested version algorithm to derive a version from a version which would otherwise be incompatible with [PEP 386](#).

3.5.2 A minimal solution

Since `distlib` is a low-level library which might be used by tools which work with existing projects, the internal implementation of versions has changed slightly from `distutils2` to allow better support for legacy version numbering. Since the re-implementation facilitated adding semantic version support at minimal cost, this has also been provided.

3.5.2.1 Versions

The basic scheme is as follows. The differences between versioning schemes is catered for by having a single function for each scheme which converts a string version to an appropriate tuple which acts as a key for sorting and comparison of versions. We have a base class, `Version`, which defines any common code. Then we can have subclasses `NormalizedVersion` (PEP-386), `LegacyVersion` (`distribute/setuptools`) and `SemanticVersion`.

To compare versions, we just check type compatibility and then compare the corresponding tuples.

3.5.2.2 Matchers

Matchers take a name followed by a set of constraints in parentheses. Each constraint is an operation together with a version string which needs to be converted to the corresponding version instance.

In summary, the following attributes can be identified for `Version` and `Matcher`:

Version:

- version string passed in to constructor (stripped)
- parser to convert string to tuple
- compare functions to compare with other versions of same type

Matcher:

- version string passed in to constructor (stripped)
- name of distribution
- list of constraints
- parser to convert string to name and set of constraints, using the same function as for `Version` to convert the version strings in the constraints to version instances
- method to match version to constraints and return `True/False`

Given the above, it appears that all the functionality *could* be provided with a single class per versioning scheme, with the *only* difference between them being the function to convert from version string to tuple. Any instance would act as either version or predicate, would display itself differently according to which it is, and raise exceptions if the wrong type of operation is performed on it (matching only allowed for predicate instances; `<=`, `<`, `>=`, `>` comparisons only allowed for version instances; and `==` and `!=` allowed for either).

However, the use of the same class to implement versions and predicates leads to ambiguity, because of the very loose project naming and versioning schemes allowed by PyPI. For example, “Hello 2.0” could be a valid project name, and “5” is a project name actually registered on PyPI. If distribution names can look like versions, it’s hard to discern the developer’s intent when creating an instance with the string “5”. So, we make separate classes for `Version` and `Matcher`.

For ease of testing, the module will define, for each of the supported schemes, a function to do the parsing (as no information is needed other than the string), and the `parse` method of the class will call that function:

```

def normalized_key(s):
    "parse using PEP-386 logic"

def legacy_key(s):
    "parse using distribute/setuptools logic"

def semantic_key(s):
    "parse using semantic versioning logic"

class Version:
    # defines all common code

    def parse(self, s):
        raise NotImplementedError('Please implement in a subclass')

```

and then:

```

class NormalizedVersion(Version):
    def parse(self, s): return normalized_key(s)

class LegacyVersion(Version):
    def parse(self, s): return legacy_key(s)

class SemanticVersion(Version):
    def parse(self, s): return semantic_key(s)

```

And a custom versioning scheme can be devised to work in the same way:

```

def custom_key(s):
    """
    convert s to tuple using custom logic, raise UnsupportedVersionError
    on problems
    """

class CustomVersion(Version):
    def parse(self, s): return custom_key(s)

```

The matcher classes are pretty minimal, too:

```

class Matcher(object):
    version_class = None

    def match(self, string_or_version):
        """
        If passed a string, convert to version using version_class,
        then do matching in a way independent of version scheme in use
        """

```

and then:

```

class NormalizedMatcher(Matcher):
    version_class = NormalizedVersion

class LegacyMatcher(Matcher):
    version_class = LegacyVersion

class SemanticMatcher(Matcher):
    version_class = SemanticVersion

```

3.5.2.3 Version schemes

Ideally one would want to work with the PEP 386 scheme, but there might be times when one needs to work with the legacy scheme (for example, when investigating dependency graphs of existing PyPI projects). Hence, the important aspects of each scheme are bundled into a simple `VersionScheme` class:

```
class VersionScheme(object):
    def __init__(self, key, matcher):
        self.key = key          # version string -> tuple converter
        self.matcher = matcher  # Matcher subclass for the scheme
```

Of course, the version class is also available through the matcher's `version_class` attribute.

`VersionScheme` makes it easier to work with alternative version schemes. For example, say we decide to experiment with an “adaptive” version scheme, which is based on the PEP 386 scheme, but when handed a non-conforming version, automatically tries to convert it to a normalized version using `suggest_normalized_version()`. Then, code which has to deal with version schemes just has to pick the appropriate scheme by name.

Creating the adaptive scheme is easy:

```
def adaptive_key(s):
    try:
        result = normalized_key(s, False)
    except UnsupportedVersionError:
        s = suggest_normalized_version(s)
        if s is None:
            raise
        result = normalized_key(s, False)
    return result

class AdaptiveVersion(NormalizedVersion):
    def parse(self, s): return adaptive_key(s)

class AdaptiveMatcher(Matcher):
    version_class = AdaptiveVersion
```

The appropriate scheme can be fetched by using the `get_scheme()` function, which is defined thus:

```
def get_scheme(scheme_name):
    "Get a VersionScheme for the given scheme_name."
```

Allowed names are 'normalized', 'legacy', 'semantic', 'adaptive' and 'default' (which points to the same as 'adaptive'). If an unrecognised name is passed in, a `ValueError` is raised.

The reimplemented `distlib.version` module is shorter than the corresponding module in `distutils2`, but the entire test suite passes and there is support for working with three versioning schemes as opposed to just one. However, the concept of “final” versions, which is not in the PEP but which was in the `distutils2` implementation, has been removed because it appears of little value (there's no way to determine the “final” status of versions for many of the project releases registered on PyPI).

3.6 The wheel API

This section describes the design of the `wheel` API which facilitates building and installing from *wheels*, the new binary distribution format for Python described in [PEP 427](#).

3.6.1 The problem we're trying to solve

There are basically two operations which need to be performed on wheels:

- Building a wheel from a source distribution.
- Installing a distribution which has been packaged as a wheel.

3.6.2 A minimal solution

Since we're talking about wheels, it seems likely that a `Wheel` class would be part of the design. This allows for extensibility over a purely function-based API. The `Wheel` would be expected to have methods that support the required operations:

```
class Wheel(object):
    def __init__(self, spec):
        """
        Initialise an instance from a specification. This can either be a
        valid filename for a wheel (for when you want to work with an
        existing wheel), or just the ``name-version-buildver`` portion of
        a wheel's filename (for when you're going to build a wheel for a
        known version and build of a named project).
        """

    def build(self, paths, tags=None):
        """
        Build a wheel. The ``name``, ``version`` and ``buildver`` should
        already have been set correctly. The ``paths`` should be a
        dictionary with keys 'prefix', 'scripts', 'headers', 'data' and one
        of 'purelib' and 'platlib'. These must point to valid paths if
        they are to be included in the wheel. The optional ``tags``
        argument should, if specified, be a dictionary with optional keys
        'pyver', 'abi' and 'arch' indicating lists of tags which
        indicate environments with which the wheel is compatible.
        """

    def install(self, paths, maker, **kwargs):
        """
        Install from a wheel. The ``paths`` should be a dictionary with
        keys 'prefix', 'scripts', 'headers', 'data', 'purelib' and
        'platlib'. These must point to valid paths to which files may
        be written if they are in the wheel. Only one of the 'purelib'
        and 'platlib' paths will be used (in the case where they are
        different), depending on whether the wheel is for a pure-
        Python distribution.

        The ``maker`` argument should be a suitably configured
        :class:`ScriptMaker` instance. The ``source_dir`` and
        ``target_dir`` arguments can be set to ``None`` when creating the
        instance - these will be set to appropriate values inside this
        method.

        The following keyword arguments are recognised:

        * ``warner``, if specified, should be a callable
          that will be called with (software_wheel_ver, file_wheel_ver)
          if they differ. They will both be in the form of tuples

```

(continues on next page)

(continued from previous page)

```
(major_ver, minor_ver). The ``warner`` defaults to ``None``.

* It's conceivable that one might want to install only the library
  portion of a package -- not installing scripts, headers data and
  so on. If ``lib_only`` is specified as ``True``, only the
  ``site-packages`` contents will be installed. The default value
  is ``False`` (meaning everything will be installed).
"""
```

In addition to the above, the following attributes can be identified for a `Wheel` instance:

- `name` – the name of the distribution
- `version` – the version of the distribution
- `buildver` – the build tag for the distribution
- `pyver` – a list of Python versions with which the wheel is compatible
- `abi` – a list of application binary interfaces (ABIs) with which the wheel is compatible
- `arch` – a list of architectures with which the wheel is compatible
- `dirname` – The directory in which a wheel file is found/to be created
- `filename` – The filename of the wheel (computed from the other attributes)

3.7 Next steps

You might find it helpful to look at the [API Reference](#).

This is the place where the functions and classes in `distlib`'s public API are described.

4.1 The `distlib.database` package

4.1.1 Classes

class `DistributionPath`

This class represents a set of distributions which are installed on a Python path (like `PYTHONPATH` / `sys.path`). Both new-style (`distlib`) and legacy (`egg`) distributions are catered for.

Methods:

`__init__` (*path=None, include_egg=False*)
Initialise the instance using a particular path.

Parameters

- **`path`** (*list of str*) – The path to use when looking for distributions. If `None` is specified, `sys.path` is used.
- **`include_egg`** – If `True`, legacy distributions (eggs) are included in the search; otherwise, they aren't.

`enable_cache` ()

Enables a cache, so that metadata information doesn't have to be fetched from disk. The cache is per instance of the `DistributionPath` instance and is enabled by default. It can be disabled using `disable_cache()` and cleared using `clear_cache()` (disabling won't automatically clear it).

`disable_cache` ()

Disables the cache, but doesn't clear it.

`clear_cache` ()

Clears the cache, but doesn't change its enabled/disabled status. If enabled, the cache will be re-populated when querying for distributions.

get_distributions()

The main querying method if you want to look at all the distributions. It returns an iterator which returns *Distribution* and, if `include_egg` was specified as `True` for the instance, also instances of any *EggInfoDistribution* for any legacy distributions found.

get_distribution(name)

Looks for a distribution by name. It returns the first one found with that name (there should only be one distribution with a given name on a given search path). Returns `None` if no distribution was found, or else an instance of *Distribution* (or, if `include_egg` was specified as `True` for the instance, an instance of *EggInfoDistribution* if a legacy distribution was found with that name).

Parameters `name` (*str*) – The name of the distribution to search for.

get_exported_entries(category, name=None)

Returns an iterator for entries exported by distributions on the path.

Parameters

- **category** (*str*) – The export category to look in.
- **name** (*str*) – A specific name to search for. If not specified, all entries in the category are returned.

Returns An iterator which iterates over exported entries (instances of `ExportEntry`).

class Distribution

A class representing a distribution, typically one which hasn't been installed (most likely, one which has been obtained from an index like PyPI).

Properties:

name

The name of the distribution.

version

The version of the distribution.

metadata

The metadata for the distribution. This is a `distlib.metadata.Metadata` instance.

download_url

The download URL for the distribution. If there are multiple URLs, this will be one of the values in *download_urls*.

download_urls

A set of known download URLs for the distribution.

New in version 0.2.0: The `download_urls` attribute was added.

digest

The digest for the source distribution. This is either `None` or a 2-tuple consisting of the hashing algorithm and the digest using that algorithm, e.g. ('sha256', '01234...').

digests

A dictionary mapping download URLs to digests, if and when digests are available.

New in version 0.2.0: The `digests` attribute was added.

locator

The locator for an instance which has been retrieved through a locator. This is `None` for an installed distribution.

class InstalledDistribution (*Distribution*)

A class representing an installed distribution. This class is not instantiated directly, except by packaging tools. Instances of it are returned from querying a *DistributionPath*.

Properties:

requested

Whether the distribution was installed by user request (if not, it may have been installed as a dependency of some other distribution).

exports

The distribution's exports, as described in *Exporting things from Distributions*. This is a cached property.

Methods:

list_installed_files ()

Returns an iterator over all of the individual files installed as part of the distribution, including metadata files. The iterator returns tuples of the form (path, hash, size). The list of files is written by the installer to the RECORD metadata file.

list_distinfo_files ()

Similar to *list_installed_files* (), but only returns metadata files.

check_installed_files ()

Runs over all the installed files to check that the size and checksum are unchanged from the values in the RECORD file, written when the distribution was installed. It returns a list of mismatches. If the files in the distribution haven't been corrupted, an empty list will be returned; otherwise, a list of mismatches will be returned.

Returns

A list which, if non-empty, will contain tuples with the following elements:

- The path in RECORD which failed to match.
- One of the strings 'exists', 'size' or 'hash' according to what didn't match (existence is checked first, then size, then hash).
- The expected value of what didn't match (as obtained from RECORD).
- The actual value of what didn't match (as obtained from the file system).

read_exports (*filename=None*)

Read exports information from a file.

Normal access to a distribution's exports should be through its *exports* attribute. This method is called from there as needed. If no filename is specified, the EXPORTS file in the *.dist-info* directory is read (it is expected to be present).

Parameters *filename* (*str*) – The filename to read from, or *None* to read from the default location.

Returns The exports read from the file.

Return type *dict*

write_exports (*exports, filename=None*)

Write exports information to a file.

If no filename is specified, the EXPORTS file in the *.dist-info* directory is written.

Parameters

- **exports** (*dict*) – A dictionary whose keys are categories and whose values are dictionaries which contain *ExportEntry* instances keyed on their name.

- **filename** (*str*) – The filename to read from, or *None* to read from the default location.

class EggInfoDistribution

Analogous to *Distribution*, but covering legacy distributions. This class is not instantiated directly. Instances of it are returned from querying a *DistributionPath*.

Properties:

name

The name of the distribution.

version

The version of the distribution.

metadata

The metadata for the distribution. This is a `distlib.metadata.Metadata` instance.

Methods:

list_installed_files()

Returns a list all of the individual files installed as part of the distribution.

class DependencyGraph

This class represents a dependency graph between releases. The nodes are distribution instances; the edges model dependencies. An edge from *a* to *b* means that *a* depends on *b*.

add_distribution(distribution)

Add *distribution* to the graph.

add_edge(x, y, label=None)

Add an edge from distribution *x* to distribution *y* with the given *label* (string).

add_missing(distribution, requirement)

Add a missing *requirement* (string) for the given *distribution*.

repr_node(dist, level=1)

Print a subgraph starting from *dist*. *level* gives the depth of the subgraph.

Direct access to the graph nodes and edges is provided through these attributes:

adjacency_list

Dictionary mapping distributions to a list of (*other*, *label*) tuples where *other* is a distribution and the edge is labelled with *label* (i.e. the version specifier, if such was provided).

reverse_list

Dictionary mapping distributions to a list of predecessors. This allows efficient traversal.

missing

Dictionary mapping distributions to a list of requirements that were not provided by any distribution.

4.2 The `distlib.resources` package

4.2.1 Attributes

cache

An instance of *ResourceCache*. This can be set after module import, but before calling any functionality which uses it, to ensure that the cache location is entirely under your control.

If you access the `file_path` property of `Resource` instance, the cache will be needed, and if not set by you, an instance with a default location will be created. See `distlib.util.get_cache_base()` for more information.

4.2.2 Functions

finder (*package*)

Get a finder for the specified package.

If the package hasn't been imported yet, an attempt will be made to import it. If importing fails, an `ImportError` will be raised.

Parameters `package` (*str*) – The name of the package for which a finder is desired.

Returns A finder for the package.

register_finder (*loader, finder_maker*)

Register a callable which makes finders for a particular type of [PEP 302](#) loader.

Parameters

- **loader** – The loader for which a finder is to be returned.
- **finder_maker** – A callable to be registered, which is called when a loader of the specified type is used to load a package. The callable is called with a single argument – the Python module object corresponding to the package – and must return a finder for that package.

4.2.3 Classes

class Resource

A class representing resources. It is never instantiated directly, but always through calling a finder's `find` method.

Properties:

is_container

Whether this instance is a container of other resources.

bytes

All of the resource data as a byte string. Raises an exception if accessed on a container resource.

size

The size of the resource data in bytes. Raises an exception if accessed on a container resource.

resources

The relative names of all the contents of this resource. Raises an exception if accessed on a resource which is *not* a container.

path

This attribute is set by the resource's finder. It is a textual representation of the path, such that if a PEP 302 loader's `get_data()` method is called with the path, the resource's bytes are returned by the loader. This attribute is analogous to the `resource_filename` API in `setuptools`. Note that for resources in zip files, the path will be a pointer to the resource in the zip file, and not directly usable as a filename. While `setuptools` deals with this by extracting zip entries to cache and returning filenames from the cache, this does not seem an appropriate thing to do in this package, as a resource is already made available to callers either as a stream or a string of bytes.

file_path

This attribute is the same as the path for file-based resource. For resources in a .zip file, the relevant resource is extracted to a file in a cache in the file system, and the name of the cached file is returned. This is for use with APIs that need file names, or need to be able to access data through OS-level file handles. See the `Cache` documentation for more information about the cache.

Methods:

as_stream()

A binary stream of the resource's data. This must be closed by the caller when it's finished with.

Raises an exception if called on a container resource.

class ResourceFinder

A base class for resource finders, which finds resources for packages stored in the file system.

__init__(module)

Initialise the finder for the package specified by `module`.

Parameters `module` – The Python module object representing a package.

find(resource_name)

Find a resource with the name specified by `resource_name` and return a `Resource` instance which represents it.

Parameters `resource_name` – A fully qualified resource name, with hierarchical components separated by `'/'`.

Returns A `Resource` instance, or `None` if a resource with that name wasn't found.

iterator(resource_name)

Return a generator which walks through the resources available through `resource_name`.

Parameters `resource_name` – A fully qualified resource name, with hierarchical components separated by `'/'`. You can use `'` to mean the 'root' resource. If the resource name refers to a non-container resource, only that resource is returned. Otherwise, the named resource is returned, followed by its children, recursively. If there is no resource named `resource_name`, `None` is returned.

Returns A generator to iterate over resources, or `None`.

is_container(resource)

Return whether a resource is a container of other resources.

Parameters `resource` (a `Resource` instance) – The resource whose status as container is wanted.

Returns `True` or `False`.

get_stream(resource)

Return a binary stream for the specified resource.

Parameters `resource` (a `Resource` instance) – The resource for which a stream is wanted.

Returns A binary stream for the resource.

get_bytes(resource)

Return the contents of the specified resource as a byte string.

Parameters `resource` (a `Resource` instance) – The resource for which the bytes are wanted.

Returns The data in the resource as a byte string.

get_size(resource)

Return the size of the specified resource in bytes.

Parameters **resource** (a *Resource* instance) – The resource for which the size is wanted.

Returns The size of the resource in bytes.

class ZipResourceFinder

This has the same interface as *ResourceFinder*.

class ResourceCache

This class implements a cache for resources which must be accessible as files in the file system. It is based on *Cache*, and adds resource-specific methods.

__init__ (*base=None*)

Initialise a cache instance with a specific directory which holds the cache. If *base* is not specified, the value *resource-cache* in the directory returned by *get_cache_base()* is used.

get (*resource*)

Ensures that the resource is available as a file in the file system, and returns the name of that file. This method calls the resource's finder's *get_cache_info()* method.

is_stale (*resource, path*)

Returns whether the data in the resource which is cached in the file system is stale compared to the resource's current data. The default implementation returns *True*, causing the resource's data to be re-written to the file every time.

4.3 The `distlib.scripts` package

4.3.1 Classes

class ScriptMaker

A class used to install scripts based on specifications.

Attributes

source_dir

The directory where script sources are to be found.

target_dir

The directory where scripts are to be created.

add_launchers

Whether to create native executable launchers on Windows.

force

Whether to overwrite scripts even when timestamps show they're up to date.

set_mode

Whether, on Posix, the scripts should have their execute mode set.

script_template

The text of a template which should contain `%(shebang)s`, `%(module)s` and `%(func)s` in the appropriate places.

The attribute is defined at class level. You can override it at the instance level to customise your scripts.

version_info

A two-tuple of the Python version to be used when generating scripts, where version-specific variants such as *foo3* or *foo-3.8* are created. This defaults to `sys.version_info`. The provided tuple can have more elements, but only the first two are used.

New in version 0.3.1.

variant_separator

A string value placed between the root basename and the version information in a variant-specific filename. This defaults to `'-'`, which means that a script with root basename `foo` and a variant `X.Y` will have a base filename of `foo-3.8` for target Python version 3.8. If you wanted to write `foo3.8` instead of `foo-3.8`, this attribute could be set to `''`. If you need more control over filename generation, you can subclass `ScriptMaker` and override the `get_script_filenames()` method.

New in version 0.3.2.

Methods

__init__ (*source_directory*, *target_directory*, *add_launchers=True*, *dry_run=False*)

Initialise the instance with options that control its behaviour.

Parameters

- **source_directory** (*str*) – Where to find scripts to install.
- **target_directory** (*str*) – Where to install scripts to.
- **add_launchers** (*bool*) – If true, create executable launchers on Windows. The executables are currently generated from the following project:
https://bitbucket.org/vinay.sajip/simple_launcher/
- **dry_run** – If true, don't actually install scripts - just pretend to.

make (*specification*, *options=None*)

Make a script in the target directory.

Parameters

- **specification** (*str*) – A specification, which can take one of the following forms:
 - A filename, relative to `source_directory`, such as `foo.py` or `subdir/bar.py`.
 - A reference to a callable, given in the form:

`name = some_package.some_module:some_callable [flags]`

where the *flags* part is optional.

When this form is passed, a Python stub script is created with the appropriate shebang line and with code to load and call the specified callable with no arguments, returning its value as the return code from the script.

For more information about flags, see [Flag formats](#).

- **options** (*dict*) – If specified, a dictionary of options used to control script creation. Currently, the following keys are checked:

gui: This should be a `bool` which, if `True`, indicates that the script is a windowed application. This distinction is only drawn on Windows if `add_launchers` is `True`, and results in a windowed native launcher application if `options['gui']` is `True` (otherwise, the native executable launcher is a console application).

interpreter_args: If specified, this should be a list of strings which are appended to the interpreter executable in the shebang line. If there are values with spaces, you will need to surround them with double quotes.

Note: Linux does not handle passing arguments to interpreters particularly well – multiple arguments are bundled up into one when passing to the interpreter – see

https://en.wikipedia.org/wiki/Shebang_line#Portability for more information. This may also affect other POSIX platforms – consult the OS documentation for your system if necessary. On Windows, the `distlib` native executable launchers *do* parse multiple arguments and pass them to the interpreter.

Returns A list of absolute pathnames of files installed (or which would have been installed, but for `dry_run` being true).

make_multiple (*specifications, options*)
Make multiple scripts from an iterable.

This method just calls `make()` once for each value returned by the iterable, but it might be convenient to override this method in some scenarios to do post-processing of the installed files (for example, running `2to3` on them).

Parameters

- **specifications** – an iterable giving the specifications to follow.
- **options** – As for the `make()` method.

Returns A list of absolute pathnames of files installed (or which would have been installed, but for `dry_run` being true).

get_script_filenames (*name*)

Get the names of scripts to be written for the specified base name, based on the `variants` and `version_info` for this instance. You can override this if you need to customise the filenames to be written.

Parameters **name** (*str*) – the basename of the script to be written.

Returns A set of filenames of files to be written as scripts, based on what variants are specified. For example, if the name is `foo` and the variants are `{'X', 'X.Y'}` and the `version_info` is `(3, 8)`, then the result would be `{'foo3', 'foo-3.8'}`.

New in version 0.3.2.

4.3.2 Functions

enquote_executable (*path*)

Cover an executable path in quotes. This only applies quotes if the passed path contains any spaces. It's at least a little careful when doing the quoting - for example, producing e.g. `/usr/bin/env "/dir with spaces/bin/jython"` instead of `"/usr/bin/env /dir with spaces/bin/jython"`

Changed in version 0.3.1: This was an internal function `_enquote_executable()` in earlier versions.

4.4 The `distlib.locators` package

4.4.1 Classes

class Locator

The base class for locators. Implements logic common to multiple locators.

__init__ (*scheme='default'*)

Initialise an instance of the locator.

Parameters **scheme** (*str*) – The version scheme to use.

get_project (*name*)

This method should be implemented in subclasses. It returns a (potentially empty) dictionary whose keys are the versions located for the project named by *name*, and whose values are instances of `distlib.util.Distribution`.

convert_url_to_download_info (*url*, *project_name*)

Extract information from a URL about the name and version of a distribution.

Parameters

- **url** (*str*) – The URL potentially of an archive (though it needn't be).
- **project_name** (*str*) – This must match the project name determined from the archive (case-insensitive matching is used).

Returns

`None` if the URL does not appear to be that of a distribution archive for the named project. Otherwise, a dictionary is returned with the following keys at a minimum:

- **url** – the URL passed in, minus any fragment portion.
- **filename** – a suitable filename to use for the archive locally.

Optional keys returned are:

- **md5_digest** – the MD5 hash of the archive, for verification after downloading. This is extracted from the fragment portion, if any, of the passed-in URL.
- **sha256_digest** – the SHA256 hash of the archive, for verification after downloading. This is extracted from the fragment portion, if any, of the passed-in URL.

Return type `dict`

get_distribution_names ()

Get the names of all distributions known to this locator.

The base class raises `NotImplementedError`; this method should be implemented in a subclass.

Returns All distributions known to this locator.

Return type `set`

locate (*requirement*, *prereleases=False*)

This tries to locate the latest version of a potentially downloadable distribution which matches a requirement (name and version constraints). If a potentially downloadable distribution (i.e. one with a download URL) is not found, `None` is returned – otherwise, an instance of `Distribution` is returned. The returned instance will have, at a minimum, **name**, **version** and **source_url** populated.

Parameters

- **requirement** (*str*) – The name and optional version constraints of the distribution to locate, e.g. `'Flask'` or `'Flask (>= 0.7, < 0.9)'`.
- **prereleases** (*bool*) – If `True`, prereleases are treated like normal releases. The default behaviour is to not return any prereleases unless they are the only ones which match the requirement.

Returns A matching instance of `Distribution`, or `None`.

get_errors ()

This returns a (possibly empty) list of error messages relating to a recent `get_project()` or `locate()` call. Fetching the errors clears the error list.

New in version 0.2.4.

class DirectoryLocator (*Locator*)

This locator scans the file system under a base directory, looking for distribution archives. The locator scans all subdirectories recursively, unless the `recursive` flag is set to `False`.

`__init__` (*base_dir*, ***kwargs*)

Parameters

- **base_dir** (*str*) – The base directory to scan for distribution archives.
- **kwargs** – Passed to base class constructor, apart from the following keyword arguments:
 - `recursive` (defaults to `True`) – if `False`, no recursion into subdirectories occurs.

class PyPIRPCLocator (*Locator*)

This locator uses the PyPI XML-RPC interface to locate distribution archives and other data about downloads.

`__init__` (*url*, ***kwargs*)

param url The base URL to use for the XML-RPC service.

type url *str*

param kwargs Passed to base class constructor.

`get_project` (*name*)

See *Locator.get_project()*.

class PyPIJSONLocator (*Locator*)

This locator uses the PyPI JSON interface to locate distribution archives and other data about downloads. It gets the metadata and URL information in a single call, so it should perform better than the XML-RPC locator.

`__init__` (*url*, ***kwargs*)

param url The base URL to use for the JSON service.

type url *str*

param kwargs Passed to base class constructor.

`get_project` (*name*)

See *Locator.get_project()*.

class SimpleScrapingLocator

This locator uses the PyPI ‘simple’ interface – a Web scraping interface – to locate distribution archives.

`__init__` (*url*, *timeout=None*, *num_workers=10*, ***kwargs*)

Parameters

- **url** (*str*) – The base URL to use for the simple service HTML pages.
- **timeout** (*float*) – How long (in seconds) to wait before giving up on a remote resource.
- **num_workers** (*int*) – The number of worker threads created to perform scraping activities.
- **kwargs** – Passed to base class constructor.

class DistPathLocator

This locator uses a *DistributionPath* instance to locate installed distributions.

`__init__` (*url*, *distpath*, ***kwargs*)

Parameters

- **distpath** (*DistributionPath*) – The distribution path to use.

- **kwargs** – Passed to base class constructor.

class AggregatingLocator (*Locator*)

This locator uses a list of other aggregators and delegates finding projects to them. It can either return the first result found (i.e. from the first aggregator in the list provided which returns a non-empty result), or a merged result from all the aggregators in the list.

__init__ (**locators*, ***kwargs*)

Parameters

- **locators** (*sequence of locators*) – A list of aggregators to delegate finding projects to.
- **merge** (*bool*) – If this *kwarg* is `True`, each aggregator in the list is asked to provide results, which are aggregated into a results dictionary. If `False`, the first non-empty return value from the list of aggregators is returned. The locators are consulted in the order in which they're passed in.

class DependencyFinder

This class allows you to recursively find all the distributions which a particular distribution depends on.

__init__ (*locator*)

Initialise an instance with the locator to be used for locating distributions.

find (*requirement*, *metas_extras=None*, *prereleases=False*)

Find all the distributions needed to fulfill *requirement*.

Parameters

- **requirement** – A string of the form *name (version)* where *version* can include an inequality constraint, or an instance of *Distribution* (e.g. representing a distribution on the local hard disk).
- **meta_extras** – A list of meta extras such as `:test:`, `:build:` and so on, to be included in the dependencies.
- **prereleases** – If `True`, allow pre-release versions to be returned - otherwise, don't return prereleases unless they're all that's available.

Returns

A 2-tuple. The first element is a set of *Distribution* instances. The second element is a set of problems encountered during dependency resolution. Currently, if this set is non-empty, it will contain 2-tuples whose first element is the string 'unsatisfied' and whose second element is a requirement which couldn't be satisfied.

In the set of *Distribution* instances returned, some attributes will be set:

- The instance representing the passed-in *requirement* will have the *requested* attribute set to `True`.
- All requirements which are not installation requirements (in other words, are needed only for build and test) will have the *build_time_dependency* attribute set to `True`.

4.4.2 Functions

get_all_distribution_names (*url=None*)

Retrieve the names of all distributions registered on an index.

Parameters *url* (*str*) – The XML-RPC service URL of the node to query. If not specified, The main PyPI index is queried.

Returns A list of the names of distributions registered on the index. Note that some of the names may be Unicode.

Return type `list`

locate (*requirement*, *prereleases=False*)

This convenience function returns the latest version of a potentially downloadable distribution which matches a requirement (name and version constraints). If a potentially downloadable distribution (i.e. one with a download URL) is not found, `None` is returned – otherwise, an instance of `Distribution` is returned. The returned instance will have, at a minimum, `name`, `version`, `download_url` and `download_urls`.

Parameters

- **requirement** (*str*) – The name and optional version constraints of the distribution to locate, e.g. `'Flask'` or `'Flask (>= 0.7, < 0.9)'`.
- **prereleases** (*bool*) – If `True`, prereleases are treated like normal releases. The default behaviour is to not return any prereleases unless they are the only ones which match the requirement.

Returns A matching instance of `Distribution`, or `None`.

4.4.3 Variables

default_locator

This attribute holds a locator which is used by `locate()` to locate distributions.

4.5 The `distlib.index` package

4.5.1 Classes

class PackageIndex

This class represents a package index which is compatible with PyPI, the Python Package Index. It allows you to register projects, upload source and binary distributions (with support for digital signatures), upload documentation, verify signatures and get a list of hosts which are mirrors for the index.

Methods:

__init__ (*url=None*, *mirror_host=None*)

Initialise an instance, setting instance attributes named from the keyword arguments.

Parameters

- **url** – The root URL for the index. If not specified, the URL for PyPI is used (`'http://pypi.org/pypi'`).
- **mirror_host** – The DNS name for a host which can be used to determine available mirror hosts for the index. If not specified, the value `'last.pypi.python.org'` is used.

register (*metadata*)

Register a project with the index.

Parameters metadata – A `Metadata` instance. This should have at least the `Name` and `Version` fields set, and ideally as much metadata as possible about this distribution. Though it might seem odd to have to specify a version when you are initially registering a project, this is required by PyPI. You can see this in PyPI's Web UI when you click the "Package submission" link in the left-hand side menu.

Returns An `urllib` HTTP response returned by the index. If an error occurs, an `HTTPError` exception will be raised.

upload_file(metadata, filename, signer=None, sign_password=None, filetype='sdist', pyversion='source', keystore=None)

Upload a distribution to the index.

Parameters

- **metadata** – A `Metadata` instance. This should have at least the `Name` and `Version` fields set, and ideally as much metadata as possible about this distribution.
- **file_name** – The path to the file which is to be uploaded.
- **signer** – If specified, this needs to be a string identifying the GnuPG private key which is to be used for signing the distribution.
- **sign_password** – The passphrase which allows access to the private key used for the signature.
- **filetype** – The type of the file being uploaded. This would have values such as `sdist` (for a source distribution), `bdist_wininst` for a Windows installer, and so on. Consult the `distutils` documentation for the full set of possible values.
- **pyversion** – The Python version this distribution is compatible with. If it's a pure-Python distribution, the value to use would be `source` - for distributions which are for specific Python versions, you would use the Python version in the form `X.Y`.
- **keystore** – The path to a directory which contains the keys used in signing. If not specified, the instance's `gpg_home` attribute is used instead. This parameter is not used unless a signer is specified.

Returns An `urllib` HTTP response returned by the index. If an error occurs, an `HTTPError` exception will be raised.

Changed in version 0.1.9: The `keystore` argument was added.

upload_documentation(metadata, doc_dir)

Upload HTML documentation to the index. The contents of the specified directory tree will be packed into a `.zip` file which is then uploaded to the index.

Parameters

- **metadata** – A `Metadata` instance. This should have at least the `Name` and `Version` fields set.
- **doc_dir** – The path to the root directory for the HTML documentation. This directory should be the one that contains `index.html`.

Returns An `urllib` HTTP response returned by the index. If an error occurs, an `HTTPError` exception will be raised.

verify_signature(self, signature_filename, data_filename, keystore=None)

Verify a digital signature against a downloaded distribution.

Parameters

- **signature_filename** – The path to the file which contains the digital signature.
- **data_filename** – The path to the file which was supposedly signed to obtain the signature in `signature_filename`.

- **keystore** – The path to a directory which contains the keys used in verification. If not specified, the instance's `gpg_home` attribute is used instead.

Returns `True` if the signature can be verified, else `False`. If an error occurs (e.g. unable to locate the public key used to verify the signature), a `ValueError` is raised.

Changed in version 0.1.9: The `keystore` argument was added.

search (*query*, *operation=None*)

Search the index for distributions matching a search query.

Parameters

- **query** – The query, either as a string or a dictionary. If a string `'foo'` is passed, it will be treated equivalently to passing the dictionary `{ 'name': 'foo' }`. The dictionary can have the following keys:
 - `name`
 - `version`
 - `stable_version`
 - `author`
 - `author_email`
 - `maintainer`
 - `maintainer_email`
 - `home_page`
 - `license`
 - `summary`
 - `description`
 - `keywords`
 - `platform`
 - `download_url`
 - `classifiers` (list of classifier strings)
 - `project_url`
 - `docs_url` (URL of the pythonhosted.org docs if they've been supplied)
- **operation** – If specified, it should be either `'and'` or `'or'`. If not specified, `'and'` is assumed. This is only used if a passed dictionary has multiple keys. It determines whether the intersection or the union of matches is returned.

Returns

A (possibly empty) list of the distributions matching the query. Each entry in the list will be a dictionary with the following keys:

- `_pypi_ordering` – the internal ordering value (an integer)
- `name` – The name of the distribution
- `version` – the version of the distribution
- `summary` – the summary for that version

New in version 0.1.8.

Additional attributes:

username

The username to use when authenticating with the index.

password

The password to use when authenticating with the index.

gpg

The path to the signing and verification program.

gpg_home

The location of the key database for the signing and verification program.

mirrors

The list of hosts which are mirrors for this index.

boundary

The boundary value to use when MIME-encoding requests to be sent to the index. This should be a byte-string.

4.6 The `distlib.util` package

4.6.1 Classes

class Cache

This base class implements common operations for `distlib` caches.

__init__ (*base*)

Initialise a cache instance with a specific directory which holds the cache.

Warning: If `base` is specified and exists, it should exist and its permissions (relevant on POSIX only) should be set to 0700 - i.e. only the user of the running process has any rights over the directory. If this is not done, the application using this functionality may be vulnerable to security breaches as a result of other processes being able to interfere with the cache.

prefix_to_dir (*prefix*)

Converts a prefix (e.g. the name of a resource's containing `.zip`, or a wheel pathname) into a directory name in the cache. This implementation delegates the work to `path_to_cache_dir()`.

class ExportEntry

Attributes:

A class holding information about a exports entry.

name

The name of the entry.

prefix

The prefix part of the entry. For a callable or data item in a module, this is the name of the package or module containing the item.

suffix

The suffix part of the entry. For a callable or data item in a module, this is a dotted path which points to the item in the module.

flags

A list of flags. See *Flag formats* for more information.

value

The actual value of the entry (a callable or data item in a module, or perhaps just a module). This is a cached property of the instance, and is determined by calling *resolve()* with the prefix and suffix properties.

dist

The distribution which exports this entry. This is normally an instance of *InstalledDistribution*.

4.6.2 Functions

get_cache_base()

Return the base directory which will hold distlib caches. If the directory does not exist, it is created.

On Windows, if `LOCALAPPDATA` is defined in the environment, then it is assumed to be a directory, and will be the parent directory of the result. On POSIX, and on Windows if `LOCALAPPDATA` is not defined, the user's home directory – as determined using `os.expanduser('~')` – will be the parent directory of the result.

The result is just the directory `'.distlib'` in the parent directory as determined above.

If a home directory is unavailable (no such directory, or if it's write-protected), a parent directory for the cache is determined using `tempfile.mkdtemp()`. This returns a directory to which only the running process has access (permission mask 0700 on POSIX), meaning that the cache should be isolated from possible malicious interference by other processes.

Note: This cache is used for the following purposes:

- As a place to cache package resources which need to be in the file system, because they are used by APIs which either expect filesystem paths, or to be able to use OS-level file handles. An example of the former is the `SSLContext.load_verify_locations()` method in Python's `ssl` module. The subdirectory `resource-cache` is used for this purpose.
- As a place to cache shared libraries which are extracted as a result of calling the `mount()` method of the `Wheel` class. The subdirectory `dylib-cache` is used for this purpose.

The application using this cache functionality, whether through the above mechanisms or through using the value returned from here directly, is responsible for any cache cleanup that is desired. Note that on Windows, you may not be able to do cache cleanup if any of the cached files are open (this will generally be the case with shared libraries, i.e. DLLs). The best way to do cache cleanup in this scenario may be on application startup, before any resources have been cached or wheels mounted.

path_to_cache_dir(path)

Converts a path (e.g. the name of an archive) into a directory name suitable for use in a cache. The following algorithm is used:

1. On Windows, any `:` in the drive is replaced with `---`.
2. Any occurrence of `os.sep` is replaced with `--`.
3. `'.cache'` is appended.

get_export_entry(specification)

Return a export entry from a specification, if it matches the expected format, or else `None`.

Parameters *specification* (*str*) – A specification, as documented for the *distlib.scripts.ScriptMaker.make()* method.

Returns `None` if the specification didn't match the expected form for an entry, or else an instance of `ExportEntry` holding information about the entry.

resolve (*module_name*, *dotted_path*)

Given a `module_name` and a `dotted_path` representing an object in that module, resolve the passed parameters to an object and return that object.

If the module has not already been imported, this function attempts to import it, then access the object represented by `dotted_path` in the module's namespace. If `dotted_path` is `None`, the module is returned. If import or attribute access fails, an `ImportError` or `AttributeError` will be raised.

Parameters

- **module_name** (*str*) – The name of a Python module or package, e.g. `os` or `os.path`.
- **dotted_path** (*str*) – The path of an object expected to be in the module's namespace, e.g. `'environ'`, `'sep'` or `'path.supports_unicode_filenames'`.

4.7 The `distlib.wheel` package

This package has functionality which allows you to work with wheels (see [PEP 427](#)).

4.7.1 Attributes

cache

An instance of `distlib.util.Cache`. This can be set after module import, but before calling any functionality which uses it, to ensure that the cache location is entirely under your control.

If you call the `mount` method of a `Wheel` instance, and the wheel is successfully mounted and contains C extensions, the cache will be needed, and if not set by you, an instance with a default location will be created. See `distlib.util.get_cache_base()` for more information.

COMPATIBLE_TAGS

A set of (`pyver`, `abi`, `arch`) tags which are compatible with this Python implementation.

4.7.2 Classes

class Wheel

This class represents wheels – either existing wheels, or wheels to be built.

__init__ (*spec*)

Initialise an instance from a specification.

Parameters **spec** (*str*) – This can either be a valid filename for a wheel (for when you want to work with an existing wheel), or just the `name-version-buildver` portion of a wheel's filename (for when you're going to build a wheel for a known version and build of a named project).

build (*paths*, *tags=None*, *wheel_version=None*)

Build a wheel. The `name`, `version` and `buildver` should already have been set correctly.

Parameters

- **paths** – This should be a dictionary with keys `'prefix'`, `'scripts'`, `'headers'`, `'data'` and one of `'purelib'` or `'platlib'`. These must point to valid paths if they are to be included in the wheel.

- **tags** – If specified, this should be a dictionary with optional keys 'pyver', 'abi' and 'arch' indicating lists of tags which indicate environments with which the wheel is compatible.
- **wheel_version** – If specified, this is written to the wheel's "Wheel-Version" metadata. If not specified, the implementation's latest supported wheel version is used.

install (*self*, *paths*, *maker*, ***kwargs*)

Install from a wheel.

Parameters

- **paths** – This should be a dictionary with keys 'prefix', 'scripts', 'headers', 'data', 'purelib' and 'platlib'. These must point to valid paths to which files may be written if they are in the wheel. Only one of the 'purelib' and 'platlib' paths will be used (in the case where they are different), depending on whether the wheel is for a pure-Python distribution.
- **maker** – This should be set to a suitably configured instance of ScriptMaker. The `source_dir` and `target_dir` arguments can be set to `None` when creating the instance - these will be set to appropriate values inside this method.
- **warner** – If specified, should be a callable that will be called with (`software_wheel_ver`, `file_wheel_ver`) if they differ. They will both be in the form of tuples (`major_ver`, `minor_ver`).
- **lib_only** – It's conceivable that one might want to install only the library portion of a package – not installing scripts, headers, data and so on. If `lib_only` is specified as `True`, only the `site-packages` contents will be installed.

is_compatible ()

Determine whether this wheel instance is compatible with the running Python.

Returns `True` if compatible, else `False`.

is_montable ()

Determine whether this wheel instance is indicated suitable for mounting in its metadata.

Returns `True` if mountable, else `False`.

mount (*append=False*)

Mount the wheel so that its contents can be imported directly, without the need to install the wheel. If the wheel contains C extensions and has metadata about these extensions, the extensions are also available for import.

If the wheel tags indicate it is not compatible with the running Python, a `DistlibException` is raised. (The `is_compatible()` method is used to determine compatibility.)

If the wheel is indicated as not suitable for mounting, a `DistlibException` is raised. (The `is_montable()` method is used to determine mountability.)

param append If `True`, the wheel's pathname is added to the end of `sys.path`. By default, it is added to the beginning.

Note: Wheels may state in their metadata that they are not intended to be mountable, in which case this method will raise a `DistlibException` with a suitable message. If C extensions are extracted, the location for extraction will be under the directory `dylib-cache` in the directory returned by `get_cache_base()`.

Wheels may be marked by their publisher as unmountable to indicate that running directly from a zip is not supported by the packaged software.

unmount ()

Unmount the wheel so that its contents can no longer be imported directly. If the wheel contains C extensions and has metadata about these extensions, the extensions are also made unavailable for import.

Note: Unmounting does not automatically clean up any extracted C extensions, as that may not be desired (and not possible, on Windows, because the files will be open). See the `get_cache_base()` documentation for suggested cleanup scenarios.

verify ()

Verify sizes and hashes of the wheel's contents against the sizes and hashes declared in the wheel's RECORD. Raise a `DistlibException` if a size or digest mismatch is detected.

New in version 0.1.8.

update (modifier, dest_dir=None, **kwargs)

Allows a user-defined callable access to the contents of a wheel. The callable can modify the contents of the wheel, add new entries or remove entries. The method first extracts the wheel's contents to a temporary location, and then calls the modifier like this:

```
modified = modifier(path_map, **kwargs)
```

where `path_map` is a dictionary mapping archive paths to the location of the corresponding extracted archive entry, and `kwargs` is whatever was passed to the `update` method. If the modifier returns `True`, a new wheel is built from the (possibly updated) contents of `path_map` and, as a final step, copied to the location of the original wheel (hence effectively modifying it in-place). The passed `path_map` will contain all of the wheel's entries other than the RECORD entry (which will be recreated if a new wheel is built).

New in version 0.1.8.

name

The name of the distribution.

version

The version of the distribution

buildver

The build tag for the distribution.

pyver

A list of Python versions with which the wheel is compatible. See [PEP 427](#) and [PEP 425](#) for details.

abi

A list of application binary interfaces (ABIs) with which the wheel is compatible. See [PEP 427](#) and [PEP 425](#) for details.

arch

A list of architectures with which the wheel is compatible. See [PEP 427](#) and [PEP 425](#) for details.

dirname

The directory in which a wheel file is found/to be created.

filename

The filename of the wheel (computed from the other attributes)

metadata

The metadata for the distribution in the wheel, as a `Metadata` instance.

info

The wheel metadata (contents of the `WHEEL` metadata file) as a dictionary.

exists

Whether the wheel file exists.

New in version 0.1.8.

4.7.3 Functions

is_compatible (*wheel*, *tags=None*)

Indicate if a wheel is compatible with a set of tags. If any combination of the tags of `wheel` is found in `tags`, then the wheel is considered to be compatible.

Parameters

- **wheel** – A `Wheel` instance or the filename of a wheel.
- **tags** – A set of tags to check for compatibility. If not specified, it defaults to the set of tags which are compatible with this Python implementation.

Returns `True` if compatible, else `False`.

4.8 Next steps

You might find it helpful to look at the [mailing list](#).

Migrating from older APIs

This section has information on migrating from older APIs.

5.1 The `pkg_resources` resource API

5.1.1 Basic resource access

```
resource_exists(package, resource_name) finder(package).find(resource_name) is  
not None  
resource_stream(package, resource_name) finder(package).find(resource_name).  
as_stream()  
resource_string(package, resource_name) finder(package).find(resource_name).  
bytes  
resource_isdir(package, resource_name) finder(package).find(resource_name).  
is_container  
resource_listdir(package, resource_name) finder(package).find(resource_name).  
resources
```

5.1.2 Resource extraction

```
resource_filename(package, resource_name) finder(package).find(resource_name).  
file_path  
set_extraction_path(extraction_path) This has no direct analogue, but you can achieve equivalent  
results by doing something like the following:
```

```
from distlib import resources  
  
resources.cache = resources.Cache(extraction_path)
```

before accessing the `file_path` property of any `Resource`. Note that if you have accessed the `file_path` property for a resource *before* doing this, the cache may already have extracted files.

`cleanup_resources(force=False)` This is not actually implemented in `pkg_resources` – it’s a no-op. You could achieve the analogous result using:

```
from distlib import resources

not_removed = resources.cache.clear()
```

5.1.3 Provider interface

You can provide an `XXXResourceFinder` class which finds resources in custom storage containers, and works like `ResourceFinder`. Although it shouldn’t be necessary, you could also return a subclass of `Resource` from your finders, to deal with custom requirements which aren’t catered for.

`get_cache_path(archive_name, names=())` There’s no analogue for this, as you shouldn’t need to care about whether particular resources are implemented in archives or not. If you need this API, please give feedback with more information about your use cases.

`extraction_error()` There’s no analogue for this. The `Cache.get()` method, which writes a resource’s bytes to a file in the cache, will raise any exception caused by underlying I/O. If you need to handle this in the cache layer, you can subclass `Cache` and override `get()`. If that doesn’t work for you, please give feedback with more information about your use cases.

`postprocess(tempname, filename)` There’s no analogue for this. The `Cache.get()` method, which writes a resource’s bytes to a file in the cache, can be overridden to perform any custom post-processing. If that doesn’t work for you, please give feedback with more information about your use cases.

5.2 The `pkg_resources` entry point API

Entry points in `pkg_resources` are equivalent to per-distribution exports dictionary (see *Exporting things from Distributions*). The keys to the dictionary are just names in a hierarchical namespace delineated with periods (like Python packages). These keys are called *groups* in `pkg_resources` documentation, though that term is a little ambiguous. In Eclipse, for example, they are called *extension point IDs*, which is a little closer to the intended usage, but a bit of a mouthful. In `distlib`, we’ll use the term `category` or `export category`.

In `distlib`, the implementation of exports is slightly different from entry points of `pkg_resources`. A *Distribution* instance has an `exports` attribute, which is a dictionary keyed by category and whose values are dictionaries that map names to `ExportEntry` instances.

Below are the `pkg_resources` functions and how to achieve the equivalent in `distlib`. In cases where the `pkg_resources` functions take distribution names, in `distlib` you get the corresponding *Distribution* instance, using:

```
dist = dist_path.get_distribution(distname)
```

and then ask that instance (or the `dist_path` instance) for the things you need.

`load_entry_point(distname, groupname, name)` `dist.exports[groupname][name].value`

`get_entry_info(distname, groupname, name)` `dist.exports[groupname][name]`

`get_entry_map(distname, groupname=None)` `dist.exports` or `dist.exports[groupname]`

```
iter_entry_points(groupname, name=None) dist_path.get_exported_entries(groupname,  
                             name=None)
```


Symbols

[__init__\(\)](#) (*AggregatingLocator* method), 84
[__init__\(\)](#) (*Cache* method), 88
[__init__\(\)](#) (*DependencyFinder* method), 84
[__init__\(\)](#) (*DirectoryLocator* method), 83
[__init__\(\)](#) (*DistPathLocator* method), 83
[__init__\(\)](#) (*DistributionPath* method), 73
[__init__\(\)](#) (*Locator* method), 81
[__init__\(\)](#) (*PackageIndex* method), 85
[__init__\(\)](#) (*PyPIJSONLocator* method), 83
[__init__\(\)](#) (*PyPIRPCLocator* method), 83
[__init__\(\)](#) (*ResourceCache* method), 79
[__init__\(\)](#) (*ResourceFinder* method), 78
[__init__\(\)](#) (*ScriptMaker* method), 80
[__init__\(\)](#) (*SimpleScrapingLocator* method), 83
[__init__\(\)](#) (*Wheel* method), 90

A

[abi](#) (*Wheel* attribute), 92
[add_distribution\(\)](#) (*DependencyGraph* method), 76
[add_edge\(\)](#) (*DependencyGraph* method), 76
[add_launchers](#) (*ScriptMaker* attribute), 79
[add_missing\(\)](#) (*DependencyGraph* method), 76
[adjacency_list](#) (*DependencyGraph* attribute), 76
[AggregatingLocator](#) (*class in distlib.locators*), 84
 APIs

- [database](#), 24
- [locators](#), 27
- [manifest](#), 53
- [metadata](#), 36
- [package resources](#), 37
- [PyPI](#), 29
- [script installation](#), 39
- [version](#), 42
- [wheel](#), 44

[arch](#) (*Wheel* attribute), 92
[as_stream\(\)](#) (*Resource* method), 78

B

[boundary](#) (*PackageIndex* attribute), 88
[build\(\)](#) (*Wheel* method), 90
[building](#)

- [wheels](#), 44

[buildver](#) (*Wheel* attribute), 92
[bytes](#) (*Resource* attribute), 77

C

[Cache](#) (*class in distlib.util*), 88
[cache](#) (*in module distlib.resources*), 76
[cache](#) (*in module distlib.util*), 90
[check_installed_files\(\)](#) (*InstalledDistribution* method), 75
[clear_cache\(\)](#) (*DistributionPath* method), 73
[COMPATIBLE_TAGS](#) (*in module distlib.util*), 90
[Continuous integration status](#)

- [distlib](#), 23

[convert_url_to_download_info\(\)](#) (*Locator* method), 82
[Coverage status](#)

- [distlib](#), 23

D

[default_locator](#) (*in module distlib.locators*), 85
[DependencyFinder](#) (*class in distlib.locators*), 84
[DependencyGraph](#) (*built-in class*), 76
[digest](#) (*Distribution* attribute), 74
[digests](#) (*Distribution* attribute), 74
[DirectoryLocator](#) (*class in distlib.locators*), 82
[dirname](#) (*Wheel* attribute), 92
[disable_cache\(\)](#) (*DistributionPath* method), 73
[dist](#) (*ExportEntry* attribute), 89
[DistPathLocator](#) (*class in distlib.locators*), 83
[Distribution](#) (*built-in class*), 74
[DistributionPath](#) (*built-in class*), 73
[download_url](#) (*Distribution* attribute), 74
[download_urls](#) (*Distribution* attribute), 74

E

EggInfoDistribution (*built-in class*), 76
enable_cache() (*DistributionPath method*), 73
enquote_executable() (*in module distlib.scripts*), 81
Environment markers
 evaluating, 37
 overriding, 37
exists (*Wheel attribute*), 93
ExportEntry (*class in distlib.util*), 88
exports (*InstalledDistribution attribute*), 75

F

file_path (*Resource attribute*), 77
filename (*Wheel attribute*), 92
find() (*DependencyFinder method*), 84
find() (*ResourceFinder method*), 78
finder() (*in module distlib.resources*), 77
flags (*ExportEntry attribute*), 88
force (*ScriptMaker attribute*), 79

G

get() (*ResourceCache method*), 79
get_all_distribution_names() (*in module distlib.locators*), 84
get_bytes() (*ResourceFinder method*), 78
get_cache_base() (*in module distlib.util*), 89
get_distribution() (*DistributionPath method*), 74
get_distribution_names() (*Locator method*), 82
get_distributions() (*DistributionPath method*), 73
get_errors() (*Locator method*), 82
get_export_entry() (*in module distlib.util*), 89
get_exported_entries() (*DistributionPath method*), 74
get_project() (*Locator method*), 81
get_project() (*PyPIJSONLocator method*), 83
get_project() (*PyPIRPCLocator method*), 83
get_script_filenames() (*ScriptMaker method*), 81
get_size() (*ResourceFinder method*), 78
get_stream() (*ResourceFinder method*), 78
gpg (*PackageIndex attribute*), 88
gpg_home (*PackageIndex attribute*), 88

I

info (*Wheel attribute*), 93
install() (*Wheel method*), 91
Installation
 distlib, 23
InstalledDistribution (*built-in class*), 74
installing

wheels, 45

is_compatible() (*in module distlib.util*), 93
is_compatible() (*Wheel method*), 91
is_container (*Resource attribute*), 77
is_container() (*ResourceFinder method*), 78
is_mountable() (*Wheel method*), 91
is_stale() (*ResourceCache method*), 79
iterator() (*ResourceFinder method*), 78

L

list_distinfo_files() (*InstalledDistribution method*), 75
list_installed_files() (*EggInfoDistribution method*), 76
list_installed_files() (*InstalledDistribution method*), 75
locate() (*in module distlib.locators*), 85
locate() (*Locator method*), 82
Locator (*class in distlib.locators*), 81
locator (*Distribution attribute*), 74

M

make() (*ScriptMaker method*), 80
make_multiple() (*ScriptMaker method*), 81
Manifest
 excluding files, 54
 excluding files globally, 54
 excluding files recursively, 55
 grafting directories, 55
 including files, 54
 including files globally, 54
 including files recursively, 55
 pruning directories, 55

Markers

 evaluating, 37
 overriding, 37

metadata (*Distribution attribute*), 74
metadata (*EggInfoDistribution attribute*), 76
metadata (*Wheel attribute*), 92
mirrors (*PackageIndex attribute*), 88
missing (*DependencyGraph attribute*), 76
modifying
 wheels, 46
mount() (*Wheel method*), 91

N

name (*Distribution attribute*), 74
name (*EggInfoDistribution attribute*), 76
name (*ExportEntry attribute*), 88
name (*Wheel attribute*), 92

P

Package resources
 in .zip files, 38

in the file system, 37
 PackageIndex (class in *distlib.index*), 85
 password (*PackageIndex* attribute), 88
 path (*Resource* attribute), 77
 path_to_cache_dir() (in module *distlib.util*), 89
 prefix (*ExportEntry* attribute), 88
 prefix_to_dir() (*Cache* method), 88
 PyPIJSONLocator (class in *distlib.locators*), 83
 PyPIRPCLocator (class in *distlib.locators*), 83
 Python Enhancement Proposals
 PEP 241, 5, 24, 36
 PEP 302, 62, 77
 PEP 314, 5, 24, 36
 PEP 345, 5, 24, 36
 PEP 370, 42
 PEP 376, 4, 24
 PEP 386, 24, 67
 PEP 397, 39
 PEP 425, 92
 PEP 427, 5, 44, 70, 90, 92
 PEP 440, 5, 24, 42, 67
 PEP 508, 5, 24
 PEP 566, 5, 24, 36
 PEP 643, 5, 24, 36
 pyver (*Wheel* attribute), 92

R

read_exports() (*InstalledDistribution* method), 75
 register() (*PackageIndex* method), 85
 register_finder() (in module *distlib.resources*), 77
 repr_node() (*DependencyGraph* method), 76
 requested (*InstalledDistribution* attribute), 75
 resolve() (in module *distlib.util*), 90
 Resource (class in *distlib.resources*), 77
 ResourceCache (class in *distlib.resources*), 79
 ResourceFinder (class in *distlib.resources*), 78
 Resources
 in .zip files, 38
 in the file system, 37
 resources (*Resource* attribute), 77
 reverse_list (*DependencyGraph* attribute), 76

S

script_template (*ScriptMaker* attribute), 79
 ScriptMaker (class in *distlib.scripts*), 79
 Scripts
 avoid overwriting, 42
 generating variants, 41
 installing, 39
 specifying custom executables, 41
 specifying what to install, 39
 windowed, 42
 wrapping callables, 40

search() (*PackageIndex* method), 87
 set_mode (*ScriptMaker* attribute), 79
 SimpleScrapingLocator (class in *distlib.locators*), 83
 size (*Resource* attribute), 77
 source_dir (*ScriptMaker* attribute), 79
 suffix (*ExportEntry* attribute), 88

T

target_dir (*ScriptMaker* attribute), 79
 Testing
 distlib, 23
 Tests
 speeding up, 24
 Tutorial, 23

U

unmount() (*Wheel* method), 92
 update() (*Wheel* method), 92
 upload_documentation() (*PackageIndex* method), 86
 username (*PackageIndex* attribute), 88

V

value (*ExportEntry* attribute), 89
 variant_separator (*ScriptMaker* attribute), 79
 verify() (*Wheel* method), 92
 verifying
 wheels, 46
 version (*Distribution* attribute), 74
 version (*EggInfoDistribution* attribute), 76
 version (*Wheel* attribute), 92
 version_info (*ScriptMaker* attribute), 79
 Versions
 matching, 43
 overview, 42

W

Wheel (class in *distlib.util*), 90
 wheels
 building, 44
 custom tags when building, 44
 installing, 45
 modifying, 46
 mounting, 47
 verifying, 46
 write_exports() (*InstalledDistribution* method), 75

Z

ZipResourceFinder (class in *distlib.resources*), 79