
Python Developer's Guide Documentation

Brett Cannon

Jul 28, 2020

Contents

1	Quick Reference	3
2	Quick Links	5
3	Status of Python branches	7
4	Contributing	9
5	Proposing changes to Python itself	11
6	Other Interpreter Implementations	13
7	Key Resources	15
8	Additional Resources	17
9	Code of Conduct	19
10	Full Table of Contents	21
10.1	Getting Started	21
10.1.1	Install <code>git</code>	21
10.1.2	Get the source code	22
10.1.3	Compile and build	22
10.1.4	Install dependencies	25
10.1.5	Regenerate <code>configure</code>	27
10.1.6	Troubleshoot the build	27
10.1.7	Editors and Tools	27
10.1.8	Directory structure	27
10.2	Where to Get Help	28
10.2.1	Ask <code>#python-dev</code>	28
10.2.2	Zulip	28
10.2.3	Core Mentorship	28
10.2.4	Core Developers Office Hours	29
10.2.5	Mailing Lists	29
10.2.6	File a Bug	29
10.3	Lifecycle of a Pull Request	29
10.3.1	Introduction	29
10.3.2	Quick Guide	29

10.3.3	Step-by-step Guide	30
10.3.4	Making Good PRs	32
10.3.5	patchcheck	32
10.3.6	Making Good Commits	33
10.3.7	Licensing	33
10.3.8	Submitting	34
10.3.9	Converting an Existing Patch from b.p.o to GitHub	34
10.3.10	Reviewing	34
10.3.11	Leaving a Pull Request Review on GitHub	35
10.3.12	Dismissing Review from Another Core Developer	35
10.3.13	Committing/Rejecting	35
10.3.14	Crediting	36
10.4	Running & Writing Tests	36
10.4.1	Running	36
10.4.2	Writing	37
10.4.3	Benchmarks	37
10.5	Increase Test Coverage	38
10.5.1	Common Gotchas	38
10.5.2	Measuring Coverage	38
10.5.3	Filing the Issue	41
10.5.4	Measuring coverage of C code with gcov and lcov	41
10.6	Helping with Documentation	41
10.6.1	Python Documentation	42
10.6.2	Helping with documentation issues	42
10.6.3	Proofreading	42
10.6.4	Helping with the Developer’s Guide	42
10.6.5	Developer’s Guide workflow	43
10.7	Documenting Python	43
10.7.1	Introduction	43
10.7.2	Style guide	44
10.7.3	reStructuredText Primer	47
10.7.4	Additional Markup Constructs	50
10.7.5	Building the documentation	63
10.7.6	Translating	64
10.8	Silence Warnings From the Test Suite	67
10.9	Fixing “easy” Issues (and Beyond)	67
10.10	Issue Tracking	67
10.10.1	Using the Issue Tracker	67
10.10.2	Disagreement With a Resolution on the Issue Tracker	68
10.10.3	Helping Triage Issues	69
10.10.4	Gaining the “Developer” Role on the Issue Tracker	70
10.10.5	The Meta Tracker	70
10.11	Triaging an Issue	71
10.11.1	Python triage team	71
10.11.2	Becoming a member of the Python triage team	71
10.11.3	Fields in the Issue Tracker	73
10.11.4	Generating Special Links in a Comment	78
10.11.5	Checklist for Triaging	78
10.12	Following Python’s Development	79
10.12.1	Mailing Lists	79
10.12.2	Zulip	79
10.12.3	IRC	80
10.12.4	Blogs	80
10.12.5	Standards of behaviour in these communication channels	80

10.12.6	Setting Expectations for Open Source Participation	80
10.12.7	Additional Repositories	80
10.13	Porting Python to a new platform	80
10.14	How to Become a Core Developer	81
10.14.1	What it Takes	81
10.14.2	What it Means	81
10.14.3	Gaining Commit Privileges	82
10.14.4	Responsibilities	83
10.15	Developer Log	83
10.15.1	Procedure for Granting or Dropping Access	87
10.16	Accepting Pull Requests	87
10.16.1	Is the PR ready to be accepted?	87
10.16.2	Handling Others' Code	88
10.16.3	Contributor Licensing Agreements	88
10.16.4	Checking if the CLA has been received	89
10.16.5	What's New and News Entries	89
10.16.6	Working with Git	90
10.17	Development Cycle	91
10.17.1	Branches	91
10.17.2	Stages	93
10.17.3	Repository Administration	94
10.18	Continuous Integration	95
10.18.1	Checking results of automatic builds	96
10.18.2	Stability	96
10.18.3	Flags-dependent failures	96
10.18.4	Ordering-dependent failures	97
10.18.5	Transient failures	98
10.18.6	Custom builders	98
10.19	Adding to the Stdlib	99
10.19.1	Adding to a pre-existing module	99
10.19.2	Adding a new module	99
10.20	Changing the Python Language	100
10.20.1	What Qualifies	101
10.20.2	PEP Process	101
10.20.3	Suggesting new features and language changes	101
10.21	Experts Index	102
10.21.1	Stdlib	102
10.21.2	Tools	107
10.21.3	Platforms	107
10.21.4	Miscellaneous	107
10.21.5	Documentation Translations	108
10.22	gdb Support	108
10.22.1	gdb 7 and later	109
10.22.2	gdb 6 and earlier	115
10.22.3	Updating auto-load-safe-path to allow test_gdb to run	115
10.23	Exploring CPython's Internals	115
10.23.1	CPython Source Code Layout	115
10.23.2	Additional References	116
10.24	Changing CPython's Grammar	117
10.24.1	Abstract	117
10.24.2	Checklist	117
10.25	Design of CPython's Compiler	117
10.25.1	Abstract	117
10.25.2	Parsing	118

10.25.3	Abstract Syntax Trees (AST)	118
10.25.4	Memory Management	119
10.25.5	Parse Tree to AST	120
10.25.6	Control Flow Graphs	120
10.25.7	AST to CFG to Bytecode	120
10.25.8	Introducing New Bytecode	122
10.25.9	Code Objects	123
10.25.10	Important Files	123
10.25.11	Known Compiler-related Experiments	124
10.25.12	References	124
10.26	Design of CPython’s Garbage Collector	124
10.26.1	Abstract	124
10.26.2	Memory layout and object structure	125
10.26.3	Identifying reference cycles	126
10.26.4	Destroying unreachable objects	129
10.26.5	Optimization: generations	129
10.26.6	Optimization: reusing fields to save memory	131
10.26.7	Optimization: delay tracking containers	131
10.27	Updating standard library extension modules	132
10.28	Coverity Scan	132
10.28.1	Access to analysis reports	133
10.28.2	Building and uploading analysis	133
10.28.3	Known limitations	133
10.28.4	Modeling	134
10.28.5	Workflow	134
10.28.6	Contact	134
10.29	Dynamic Analysis with Clang	134
10.29.1	What is Clang?	135
10.29.2	What are Sanitizers?	135
10.29.3	Clang/LLVM Setup	135
10.29.4	Python Build Setup	137
10.30	Running a buildbot worker	140
10.30.1	Preparing for buildbot worker setup	140
10.30.2	Setting up the buildbot worker	140
10.30.3	Buildbot worker operation	143
10.30.4	Required Ports	144
10.30.5	Required Resources	144
10.30.6	Security Considerations	144
10.31	Core Developer Motivations and Affiliations	145
10.31.1	Published entries	145
10.31.2	Goals of this page	148
10.31.3	Limitations on scope	149
10.32	Git Bootcamp and Cheat Sheet	149
10.32.1	Forking CPython GitHub Repository	150
10.32.2	Cloning a Forked CPython Repository	151
10.32.3	Listing the Remote Repositories	151
10.32.4	Setting Up Your Name and Email Address	151
10.32.5	Enabling autocrlf on Windows	151
10.32.6	Creating and Switching Branches	151
10.32.7	Deleting Branches	152
10.32.8	Staging and Committing Files	152
10.32.9	Reverting Changes	153
10.32.10	Stashing Changes	153
10.32.11	Committing Changes	153

10.32.12	Pushing Changes	153
10.32.13	Creating a Pull Request	153
10.32.14	Syncing with Upstream	154
10.32.15	Applying a Patch from Mercurial to Git	154
10.32.16	Downloading Other's Patches	155
10.32.17	Accepting and Merging a Pull Request	155
10.32.18	Backporting Merged Changes	156
10.32.19	Editing a Pull Request Prior to Merging	157
10.33	Appendix: Topics	157
10.33.1	Basics for contributors	157
10.33.2	Core developers	158
10.33.3	Development workflow for contributors	158
10.33.4	Documenting Python and style guide	158
10.33.5	Issue tracking and triaging	158
10.33.6	Language development in depth	158
10.33.7	Testing and continuous integration	159
	Bibliography	161
	Index	163

This guide is a comprehensive resource for *contributing* to Python – for both new and experienced contributors. It is *maintained* by the same community that maintains Python. We welcome your contributions to Python!

Quick Reference

Here are the basic steps needed to get *set up* and contribute a patch. This is meant as a checklist, once you know the basics. For complete instructions please see the *setup guide*.

1. Install and set up *Git* and other dependencies (see the *Git Setup* page for detailed information).
2. Fork the CPython repository to your GitHub account and *get the source code* using:

```
git clone https://github.com/<your_username>/cpython
cd cpython
```

3. Build Python, on UNIX and Mac OS use:

```
./configure --with-pydebug && make -j
```

and on Windows use:

```
PCbuild\build.bat -e -d
```

See also *more detailed instructions, how to install and build dependencies*, and the platform-specific pages for *UNIX, Mac OS, and Windows*.

4. *Run the tests*:

```
./python -m test -j3
```

On *most* Mac OS X systems, replace `./python` with `./python.exe`. On Windows, use `python.bat`. With Python 2.7, replace `test` with `test.regrtest`.

5. Create a new branch where your work for the issue will go, e.g.:

```
git checkout -b fix-issue-12345 master
```

If an issue does not already exist, please *create it*. Trivial issues (e.g. typo fixes) do not require any issue to be created.

6. Once you fixed the issue, run the tests, run `make patchcheck`, and if everything is ok, commit.

7. Push the branch on your fork on GitHub and *create a pull request*. Include the issue number using `bpo-NNNN` in the pull request description. For example:

```
bpo-12345: Fix some bug in spam module
```

Note: First time contributors will need to sign the Contributor Licensing Agreement (CLA) as described in the *Licensing* section of this guide.

CHAPTER 2

Quick Links

Here are some links that you probably will reference frequently while contributing to Python:

- [Issue tracker](#)
- [Buildbot status](#)
- [Where to Get Help](#)
- [PEPs \(Python Enhancement Proposals\)](#)
- [Git Bootcamp and Cheat Sheet](#)

Status of Python branches

Branch	Schedule	Status	First release	End-of-life	Release manager
master	PEP 619	features	<i>2021-10-04</i>	<i>TBD</i>	Pablo Galindo Salgado
3.9	PEP 596	bugfix	<i>2020-10-05</i>	<i>TBD</i>	Łukasz Langa
3.8	PEP 569	bugfix	2019-10-14	<i>2024-10</i>	Łukasz Langa
3.7	PEP 537	security	2018-06-27	<i>2023-06-27</i>	Ned Deily
3.6	PEP 494	security	2016-12-23	<i>2021-12-23</i>	Ned Deily
3.5	PEP 478	security	2015-09-13	<i>2020-09-13</i>	Larry Hastings

The master branch is currently the future Python 3.10, and is the only branch that accepts new features. The latest release for each Python version can be found on the [download page](#).

Status:

features new features, bugfixes, and security fixes are accepted.

prerelease feature fixes, bugfixes, and security fixes are accepted for the upcoming feature release.

bugfix bugfixes and security fixes are accepted, new binaries are still released. (Also called **maintenance** mode or **stable** release)

security only security fixes are accepted and no more binaries are released, but new source-only versions can be released

end-of-life release cycle is frozen; no further changes can be pushed to it.

Dates in *italic* are scheduled and can be adjusted.

By default, the end-of-life is scheduled 5 years after the first release, but can be adjusted by the release manager of each branch. All Python 2 versions have reached end-of-life.

See also the [Development Cycle](#) page for more information about branches.

Contributing

We encourage everyone to contribute to Python and that's why we have put up this developer's guide. If you still have questions after reviewing the material in this guide, then the [Core Python Mentorship](#) group is available to help guide new contributors through the process.

A number of individuals from the Python community have contributed to a series of excellent guides at [Open Source Guides](#).

Core developers and contributors alike will find the following guides useful:

- [How to Contribute to Open Source](#)
- [Building Welcoming Communities](#)

Guide for contributing to Python:

New Contributors	Documentarians	Triagers	Core Developers
<i>Getting Started</i>	<i>Helping with Documentation</i>	<i>Issue Tracking</i>	<i>How to Become a Core Developer</i>
<i>Where to Get Help</i>	<i>Documenting Python</i>	<i>Triaging an Issue</i>	<i>Developer Log</i>
<i>Lifecycle of a Pull Request</i>	<i>Style guide</i>	<i>Helping Triage Issues</i>	<i>Accepting Pull Requests</i>
<i>Running & Writing Tests</i>	<i>reStructuredText Primer</i>	<i>Experts Index</i>	<i>Development Cycle</i>
<i>Fixing "easy" Issues (and Beyond)</i>	<i>Translating</i>		<i>Core Developer Motivations and Affiliations</i>
<i>Following Python's Development</i>			<i>Core Developers Office Hours</i>
<i>Git Bootcamp and Cheat Sheet</i>			

Advanced tasks and topics for once you are comfortable:

- [Silence Warnings From the Test Suite](#)
- [Fixing issues found by the *buildbots*](#)

- *Coverity Scan*
- Helping out with reviewing open pull requests. See *how to review a Pull Request*.
- *Fixing “easy” Issues (and Beyond)*

It is **recommended** that the above documents be read as needed. New contributors will build understanding of the CPython workflow by reading the sections mentioned in this table. You can stop where you feel comfortable and begin contributing immediately without reading and understanding these documents all at once. If you do choose to skip around within the documentation, be aware that it is written assuming preceding documentation has been read so you may find it necessary to backtrack to fill in missing concepts and terminology.

Proposing changes to Python itself

Improving Python’s code, documentation and tests are ongoing tasks that are never going to be “finished”, as Python operates as part of an ever-evolving system of technology. An even more challenging ongoing task than these necessary maintenance activities is finding ways to make Python, in the form of the standard library and the language definition, an even better tool in a developer’s toolkit.

While these kinds of change are much rarer than those described above, they do happen and that process is also described as part of this guide:

- *Adding to the Stdlib*
- *Changing the Python Language*

Other Interpreter Implementations

This guide is specifically for contributing to the Python reference interpreter, also known as CPython (while most of the standard library is written in Python, the interpreter core is written in C and integrates most easily with the C and C++ ecosystems).

There are other Python implementations, each with a different focus. Like CPython, they always have more things they would like to do than they have developers to work on them. Some major examples that may be of interest are:

- **PyPy**: A Python interpreter focused on high speed (JIT-compiled) operation on major platforms
- **Jython**: A Python interpreter focused on good integration with the Java Virtual Machine (JVM) environment
- **IronPython**: A Python interpreter focused on good integration with the Common Language Runtime (CLR) provided by .NET and Mono
- **Stackless**: A Python interpreter focused on providing lightweight microthreads while remaining largely compatible with CPython specific extension modules

Key Resources

- **Coding style guides**
 - **PEP 7** (Style Guide for C Code)
 - **PEP 8** (Style Guide for Python Code)
- **Issue tracker**
 - *Meta tracker* (issue tracker for the issue tracker)
 - *Experts Index*
- Buildbot status
- **Source code**
 - Browse online
 - Snapshot of the **master** branch
 - Daily OS X installer
- PEPs (Python Enhancement Proposals)
- *Where to Get Help*
- *Developer Log*

Additional Resources

- Anyone can clone the sources for this guide. See *Helping with the Developer's Guide*.
- **Help with ...**
 - *Exploring CPython's Internals*
 - *Changing CPython's Grammar*
 - *Design of CPython's Compiler*
 - *Design of CPython's Garbage Collector*
- **Tool support**
 - *gdb Support*
 - *Dynamic Analysis with Clang*
 - Various tools with configuration files as found in the `Misc` directory
 - Information about editors and their configurations can be found in the `wiki`
- `python.org` maintenance
- Search this guide

CHAPTER 9

Code of Conduct

Please note that all interactions on [Python Software Foundation](#)-supported infrastructure is covered by the [PSF Code of Conduct](#), which includes all infrastructure used in the development of Python itself (e.g. mailing lists, issue trackers, GitHub, etc.). In general this means everyone is expected to be open, considerate, and respectful of others no matter what their position is within the project.

10.1 Getting Started

These instructions cover how to get a working copy of the source code and a compiled version of the CPython interpreter (CPython is the version of Python available from <https://www.python.org/>). It also gives an overview of the directory structure of the CPython source code.

Alternatively, if you have [Docker](#) installed you might want to use [our official images](#). These contain the latest releases of several Python versions, along with `git head`, and are provided for development and testing purposes only.

See also:

The [Quick Reference](#) gives brief summary of the process from installing `git` to submitting a pull request.

10.1.1 Install `git`

CPython is developed using `git` for version control. The `git` command line program is named `git`; this is also used to refer to `git` itself. `git` is easily available for all common operating systems.

- **Install**

As the CPython repo is hosted on GitHub, please refer to either the [GitHub setup instructions](#) or the [git project instructions](#) for step-by-step installation directions. You may also want to consider a graphical client such as [TortoiseGit](#) or [GitHub Desktop](#).

- **Configure**

Configure *your name and email* and create an [SSH key](#) as this will allow you to interact with GitHub without typing a username and password each time you execute a command, such as `git pull`, `git push`, or `git fetch`. On Windows, you should also [enable autocrlf](#).

10.1.2 Get the source code

The CPython repo is hosted on GitHub. To get a copy of the source code you should *fork the Python repository on GitHub, create a local clone of your personal fork, and configure the remotes*.

You will only need to execute these steps once:

1. Go to <https://github.com/python/cpython>.
2. Press *Fork* on the top right.
3. When asked where to fork the repository, choose to fork it to your username.
4. Your fork will be created at `https://github.com/<username>/cpython`.
5. Clone your GitHub fork (replace `<username>` with your username):

```
$ git clone git@github.com:<username>/cpython.git
```

(You can use both SSH-based or HTTPS-based URLs.)

6. Configure an upstream remote:

```
$ cd cpython
$ git remote add upstream git@github.com:python/cpython.git
```

7. Verify that your setup is correct:

```
$ git remote -v
origin  git@github.com:<your-username>/cpython.git (fetch)
origin  git@github.com:<your-username>/cpython.git (push)
upstream    git@github.com:python/cpython.git (fetch)
upstream    git@github.com:python/cpython.git (push)
```

If you did everything correctly, you should now have a copy of the code in the `cpython` directory and two remotes that refer to your own GitHub fork (`origin`) and the official CPython repository (`upstream`).

If you want a working copy of an already-released version of Python, i.e., a version in *maintenance mode*, you can checkout a release branch. For instance, to checkout a working copy of Python 3.8, do `git checkout 3.8`.

You will need to re-compile CPython when you do such an update.

Do note that CPython will notice that it is being run from a working copy. This means that if you edit CPython's source code in your working copy, changes to Python code will be picked up by the interpreter for immediate use and testing. (If you change C code, you will need to recompile the affected files as described below.)

Patches for the documentation can be made from the same repository; see *Documenting Python*.

10.1.3 Compile and build

CPython provides several compilation flags which help with debugging various things. While all of the known flags can be found in the `Misc/SpecialBuilds.txt` file, the most critical one is the `Py_DEBUG` flag which creates what is known as a “pydebug” build. This flag turns on various extra sanity checks which help catch common issues. The use of the flag is so common that turning on the flag is a basic compile option.

You should always develop under a pydebug build of CPython (the only instance of when you shouldn't is if you are taking performance measurements). Even when working only on pure Python code the pydebug build provides several useful checks that one should not skip.

UNIX

The core CPython interpreter only needs a C compiler to be built, however, some of the extension modules will need development headers for additional libraries (such as the `zlib` library for compression). Depending on what you intend to work on, you might need to install these additional requirements so that the compiled interpreter supports the desired features.

If you want to install these optional dependencies, consult the *Install dependencies* section below.

If you don't need to install them, the basic steps for building Python for development is to configure it and then compile it.

Configuration is typically:

```
./configure --with-pydebug
```

More flags are available to `configure`, but this is the minimum you should do to get a pydebug build of CPython.

Note: You might need to run `make clean` before or after re-running `configure` in a particular build directory.

Once `configure` is done, you can then compile CPython with:

```
make -s -j2
```

This will build CPython with only warnings and errors being printed to `stderr` and utilize up to 2 CPU cores. If you are using a multi-core machine with more than 2 cores (or a single-core machine), you can adjust the number passed into the `-j` flag to match the number of cores you have (or if your version of Make supports it, you can use `-j` without a number and Make will not limit the number of steps that can run simultaneously.).

At the end of the build you should see a success message, possibly followed by a list of extension modules that haven't been built because their dependencies were missing:

```
Python build finished successfully!
The necessary bits to build these optional modules were not found:
_bz2                _dbm                _gdbm
_lzma               _sqlite3           _ssl
_tkinter           _uuid              readline
zlib
To find the necessary bits, look in setup.py in detect_modules()
for the module's name.
```

If the build failed and you are using a C89 or C99-compliant compiler, please [open a bug report](#).

If you decide to *Install dependencies*, you will need to re-run both `configure` and `make`.

Once CPython is done building you will then have a working build that can be run in-place; `./python` on most machines (and what is used in all examples), `./python.exe` wherever a case-insensitive filesystem is used (e.g. on OS X by default), in order to avoid conflicts with the `Python` directory. There is normally no need to install your built copy of Python! The interpreter will realize where it is being run from and thus use the files found in the working copy. If you are worried you might accidentally install your working copy build, you can add `--prefix=/tmp/python` to the configuration step. When running from your working directory, it is best to avoid using the `--enable-shared` flag to `configure`; unless you are very careful, you may accidentally run with code from an older, installed shared Python library rather than from the interpreter you just built.

Clang

If you are using `clang` to build CPython, some flags you might want to set to quiet some standard warnings which are specifically superfluous to CPython are `-Wno-unused-value` `-Wno-empty-body` `-Qunused-arguments`. You can set your `CFLAGS` environment variable to these flags when running `configure`.

If you are using `clang` with `ccache`, turn off the noisy `parentheses-equality` warnings with the `-Wno-parentheses-equality` flag. These warnings are caused by clang not having enough information to detect that extraneous parentheses in expanded macros are valid, because the preprocessing is done separately by `ccache`.

If you are using LLVM 2.8, also use the `-no-integrated-as` flag in order to build the `ctypes` module (without the flag the rest of CPython will still build properly).

Windows

For a quick guide to building you can read [this documentation](#) from Victor Stinner.

Python 3.6 and later can use Microsoft Visual Studio 2017. You can download and use any of the free or paid versions of [Visual Studio 2017](#).

When installing Visual Studio 2017, select the **Python development** workload and the optional **Python native development tools** component to obtain all of the necessary build tools. If you do not already have git installed, you can find git for Windows on the **Individual components** tab of the installer.

Note: If you want to build MSI installers, be aware that the build toolchain for them has a dependency on the Microsoft .NET Framework Version 3.5 (which may not be configured on recent versions of Windows, such as Windows 10). If you are building on a recent Windows version, use the Control Panel (Programs | Programs and Features | Turn Windows Features on or off) and ensure that the entry “.NET Framework 3.5 (includes .NET 2.0 and 3.0)” is enabled.

Your first build should use the command line to ensure any external dependencies are downloaded:

```
PCBuild\build.bat
```

After this build succeeds, you can open the `PCBuild\pcbuid.sln` solution in Visual Studio to continue development.

See the [readme](#) for more details on what other software is necessary and how to build.

Note: **Python 2.7** uses Microsoft Visual Studio 2008, which is most easily obtained through an MSDN subscription. To use the build files in the `PCbuild` directory you will also need Visual Studio 2010, see the [2.7 readme](#) for more details. If you have VS 2008 but not 2010 you can use the build files in the `PC/VS9.0` directory, see the [VS9 readme](#) for details.

Note: If you are using the Windows Subsystem for Linux (WSL), clone the repository from a native Windows terminal program like `cmd.exe` command prompt or PowerShell as well as use a build of git targeted for Windows, e.g., the official one from <https://git-scm.com>. Otherwise, Visual Studio will not be able to find all the project's files and will fail the build.

10.1.4 Install dependencies

This section explains how to install additional extensions (e.g. `zlib`) on *Linux* and *macOS/OS X*. On Windows, extensions are already included and built automatically.

Linux

For UNIX based systems, we try to use system libraries whenever available. This means optional components will only build if the relevant system headers are available. The best way to obtain the appropriate headers will vary by distribution, but the appropriate commands for some popular distributions are below.

On **Fedora, Red Hat Enterprise Linux** and other `yum` based systems:

```
$ sudo yum install yum-utils
$ sudo yum-builddep python3
```

On **Fedora** and other `DNF` based systems:

```
$ sudo dnf install dnf-plugins-core # install this to use 'dnf builddep'
$ sudo dnf builddep python3
```

On **Debian, Ubuntu**, and other `apt` based systems, try to get the dependencies for the Python you're working on by using the `apt` command.

First, make sure you have enabled the source packages in the sources list. You can do this by adding the location of the source packages, including URL, distribution name and component name, to `/etc/apt/sources.list`. Take Ubuntu Bionic for example:

```
deb-src http://archive.ubuntu.com/ubuntu/ bionic main
```

For other distributions, like Debian, change the URL and names to correspond with the specific distribution.

Then you should update the packages index:

```
$ sudo apt-get update
```

Now you can install the build dependencies via `apt`:

```
$ sudo apt-get build-dep python3.6
```

If that package is not available for your system, try reducing the minor version until you find a package that is available.

macOS and OS X

For **macOS systems** (versions 10.12+) and **OS X 10.9 and later**, the Developer Tools can be downloaded and installed automatically; you do not need to download the complete Xcode application.

If necessary, run the following:

```
$ xcode-select --install
```

This will also ensure that the system header files are installed into `/usr/include`.

On **Mac OS X systems** (versions 10.0 - 10.7) and **OS X 10.8**, use the C compiler and other development utilities provided by Apple's Xcode Developer Tools. The Developer Tools are not shipped with Mac OS X.

For these **older releases (versions 10.0 - 10.8)**, you will need to download either the correct version of the Command Line Tools, if available, or install them from the full Xcode app or package for that OS X release. Older versions may be available either as a no-cost download through Apple's App Store or from [the Apple Developer web site](#).

Also note that OS X does not include several libraries used by the Python standard library, including `libzma`, so expect to see some extension module build failures unless you install local copies of them. As of OS X 10.11, Apple no longer provides header files for the deprecated system version of OpenSSL which means that you will not be able to build the `_ssl` extension. One solution is to install these libraries from a third-party package manager, like [Homebrew](#) or [MacPorts](#), and then add the appropriate paths for the header and library files to your `configure` command. For example,

with **Homebrew**:

```
$ brew install openssl xz gdbm
```

and configure Python versions `>= 3.7`:

```
./configure --with-pydebug --with-openssl=$(brew --prefix openssl)
```

or configure Python versions `< 3.7`:

```
$ CPPFLAGS="-I$(brew --prefix openssl)/include" \  
  LDFLAGS="-L$(brew --prefix openssl)/lib" \  
  ./configure --with-pydebug
```

and make:

```
$ make -s -j2
```

or **MacPorts**:

```
$ sudo port install pkgconfig openssl xz gdbm
```

and configure:

```
$ CPPFLAGS="-I/opt/local/include" \  
  LDFLAGS="-L/opt/local/lib" \  
  ./configure --with-pydebug
```

and make:

```
$ make -s -j2
```

There will sometimes be optional modules added for a new release which won't yet be identified in the OS level build dependencies. In those cases, just ask for assistance on the core-mentorship list. If working on bug fixes for Python 2.7, use `python` in place of `python3` in the above commands.

Explaining how to build optional dependencies on a UNIX based system without root access is beyond the scope of this guide.

Note: While you need a C compiler to build CPython, you don't need any knowledge of the C language to contribute! Vast areas of CPython are written completely in Python: as of this writing, CPython contains slightly more Python code than C.

10.1.5 Regenerate `configure`

If a change is made to Python which relies on some POSIX system-specific functionality (such as using a new system call), it is necessary to update the `configure` script to test for availability of the functionality.

Python's `configure` script is generated from `configure.ac` using Autoconf. Instead of editing `configure`, edit `configure.ac` and then run `autoreconf` to regenerate `configure` and a number of other files (such as `pyconfig.h`).

When submitting a patch with changes made to `configure.ac`, you should also include the generated files.

Note that running `autoreconf` is not the same as running `autoconf`. For example, `autoconf` by itself will not regenerate `pyconfig.h.in`. `autoreconf` runs `autoconf` and a number of other tools repeatedly as is appropriate.

Python's `configure.ac` script typically requires a specific version of Autoconf. At the moment, this reads: `AC_PREREQ(2.69)`.

If the system copy of Autoconf does not match this version, you will need to install your own copy of Autoconf.

10.1.6 Troubleshoot the build

This section lists some of the common problems that may arise during the compilation of Python, with proposed solutions.

Avoid recreating auto-generated files

Under some circumstances you may encounter Python errors in scripts like `Parser/asdl_c.py` or `Python/makeopcodetargets.py` while running `make`. Python auto-generates some of its own code, and a full build from scratch needs to run the auto-generation scripts. However, this makes the Python build require an already installed Python interpreter; this can also cause version mismatches when trying to build an old (2.x) Python with a new (3.x) Python installed, or vice versa.

To overcome this problem, auto-generated files are also checked into the Git repository. So if you don't touch the auto-generation scripts, there's no real need to auto-generate anything.

10.1.7 Editors and Tools

Python is used widely enough that practically all code editors have some form of support for writing Python code. Various coding tools also include Python support.

For editors and tools which the core developers have felt some special comment is needed for coding *in* Python, see *Additional Resources*.

10.1.8 Directory structure

There are several top-level directories in the CPython source tree. Knowing what each one is meant to hold will help you find where a certain piece of functionality is implemented. Do realize, though, there are always exceptions to every rule.

Doc The official documentation. This is what <https://docs.python.org/> uses. See also *Building the documentation*.

Grammar Contains the EBNF (Extended Backus-Naur Form) grammar file for Python.

Include Contains all interpreter-wide header files.

Lib The part of the standard library implemented in pure Python.

Mac Mac-specific code (e.g., using IDLE as an OS X application).

Misc Things that do not belong elsewhere. Typically this is varying kinds of developer-specific documentation.

Modules The part of the standard library (plus some other code) that is implemented in C.

Objects Code for all built-in types.

PC Windows-specific code.

PCbuild Build files for the version of MSVC currently used for the Windows installers provided on python.org.

Parser Code related to the parser. The definition of the AST nodes is also kept here.

Programs Source code for C executables, including the main function for the CPython interpreter (in versions prior to Python 3.5, these files are in the Modules directory).

Python The code that makes up the core CPython runtime. This includes the compiler, eval loop and various built-in modules.

Tools Various tools that are (or have been) used to maintain Python.

10.2 Where to Get Help

If you are working on Python it is very possible you will come across an issue where you need some assistance to solve it (this happens to core developers all the time).

Should you require help, there are a *variety of options available* to seek assistance. If the question involves process or tool usage then please check the rest of this guide first as it should answer your question.

10.2.1 Ask #python-dev

If you are comfortable with IRC you can try asking on #python-dev (on the [freenode](#) network). Typically there are a number of experienced developers, ranging from triagers to core developers, who can answer questions about developing for Python. Just remember that #python-dev is for questions involving the development *of* Python whereas #python is for questions concerning development *with* Python.

10.2.2 Zulip

An alternative to IRC is our own [Zulip](#) instance. There are different streams for asking help with core development, as well as core developers' office hour stream. It is preferred that you ask questions here first or schedule an office hour, before posting to python-dev mailing list or filing bugs.

10.2.3 Core Mentorship

If you are interested in improving Python and contributing to its development, but don't yet feel entirely comfortable with the public channels mentioned above, [Python Mentors](#) are here to help you. Python is fortunate to have a community of volunteer core developers willing to mentor anyone wishing to contribute code, work on bug fixes or improve documentation. Everyone is welcomed and encouraged to contribute.

10.2.4 Core Developers Office Hours

Several core developers have set aside time to host mentorship office hours. During the office hour, core developers are available to help contributors with our process, answer questions, and help lower the barrier of contributing and becoming Python core developers.

The PSF's code of conduct applies for interactions with core developers during office hours.

Core Developer	Schedule	Details
Zachary Ware	See details link	Schedule at https://calendly.com/zware
Mariatta Wijaya	Thursdays 7PM - 8PM Pacific (Vancouver, Canada Timezone)	In Python's Zulip Chat , Core > Office Hour stream. A reminder will be posted to both Zulip and Mariatta's twitter account 24 hours before the start.

10.2.5 Mailing Lists

Further options for seeking assistance include the [python-ideas](#) and [python-dev](#) mailing lists. Python-ideas contains discussion of speculative Python language ideas for possible inclusion into the language. If an idea gains traction it can then be discussed and honed to the point of becoming a solid proposal and presented on python-dev. Python-dev contains discussion of current Python design issues, release mechanics, and maintenance of existing releases. As with `#python-dev`, these mailing lists are for questions involving the development *of* Python, **not** for development *with* Python.

10.2.6 File a Bug

If you strongly suspect you have stumbled on a bug (be it in the build process, in the test suite, or in other areas), then open an issue on the [issue tracker](#). As with every bug report it is strongly advised that you detail which conditions triggered it (including the OS name and version, and what you were trying to do), as well as the exact error message you encountered.

10.3 Lifecycle of a Pull Request

10.3.1 Introduction

CPython uses a workflow based on pull requests. What this means is that you create a branch in Git, make your changes, push those changes to your fork on GitHub (`origin`), and then create a pull request against the official CPython repository (`upstream`).

10.3.2 Quick Guide

Clear communication is key to contributing to any project, especially an Open Source project like CPython.

Here is a quick overview of how you can contribute to CPython:

1. [Create an issue](#) that describes your change*⁰

⁰ If an issue is trivial (e.g. typo fixes), or if an issue already exists, you can skip this step.

2. *Create a new branch in Git*
3. Work on changes (e.g. fix a bug or add a new feature)
4. *Run tests* and make `patchcheck`
5. *Commit* and *push* changes to your GitHub fork
6. *Create Pull Request* on GitHub to merge a branch from your fork
7. Review and address *comments on your Pull Request*
8. When your changes are merged, you can *delete the PR branch*
9. Celebrate contributing to CPython! :)

Note: In order to keep the commit history intact, please avoid squashing or amending history and then force-pushing to the PR. Reviewers often want to look at individual commits.

10.3.3 Step-by-step Guide

You should have already *set up your system*, *got the source code*, and *built Python*.

- Update data from your upstream repository:

```
git fetch upstream
```

- Create a new branch in your local clone:

```
git checkout -b <branch-name> upstream/master
```

- Make changes to the code, and use `git status` and `git diff` to see them.

(Learn more about *Making Good PRs*)

- Make sure the changes are fine and don't cause any test failure:

```
make patchcheck  
./python -m test
```

(Learn more about *patchcheck* and about *Running & Writing Tests*)

- Once you are satisfied with the changes, add the files and commit them:

```
git add <filenames>  
git commit -m '<message>'
```

(Learn more about *Making Good Commits*)

- Then push your work to your GitHub fork:

```
git push origin <branch-name>
```

- Finally go on `https://github.com/<your-username>/cpython`: you will see a box with the branch you just pushed and a green button that allows you to create a pull request against the official CPython repository.
- When people start adding review comments, you can address them by switching to your branch, making more changes, committing them, and pushing them to automatically update your PR:

```
git checkout <branch-name>
# make changes and run tests
git add <filenames>
git commit -m '<message>'
git push origin <branch-name>
```

- If a core developer reviewing your PR pushed one or more commits to your PR branch, then after checking out your branch and before editing, run:

```
git pull origin <branch-name> # pull = fetch + merge
```

If you have made local changes that have not been pushed to your fork and there are merge conflicts, git will warn you about this and enter conflict resolution mode. See [Resolving Merge Conflicts](#) below.

- If time passes and there are merge conflicts with the master branch, GitHub will show a warning to this end and you may be asked to address this. Merge the changes from the master branch while resolving the conflicts locally:

```
git checkout <branch-name>
git pull upstream master # pull = fetch + merge
# resolve conflicts: see "Resolving Merge Conflicts" below
git push origin <branch-name>
```

- After your PR has been accepted and merged, you can *delete the branch*:

```
git branch -D <branch-name> # delete local branch
git push origin -d <branch-name> # delete remote branch
```

Note: You can still upload a patch to bugs.python.org, but the GitHub pull request workflow is **strongly** preferred.

Resolving Merge Conflicts

When merging changes from different branches (or variants of a branch on different repos), the two branches may contain incompatible changes to one or more files. These are called “merge conflicts” and need to be manually resolved as follows:

1. Check which files have merge conflicts:

```
git status
```

2. Edit the affected files and bring them to their intended final state. Make sure to remove the special “conflict markers” inserted by git.
3. Commit the affected files:

```
git add <filenames>
git merge --continue
```

When running the final command, git may open an editor for writing a commit message. It is usually okay to leave that as-is and close the editor.

See [the merge command's documentation](#) for a detailed technical explanation.

10.3.4 Making Good PRs

When creating a pull request for submission, there are several things that you should do to help ensure that your pull request is accepted.

First, make sure to follow Python's style guidelines. For Python code you should follow [PEP 8](#), and for C code you should follow [PEP 7](#). If you have one or two discrepancies those can be fixed by the core developer who merges your pull request. But if you have systematic deviations from the style guides your pull request will be put on hold until you fix the formatting issues.

Second, be aware of backwards-compatibility considerations. While the core developer who eventually handles your pull request will make the final call on whether something is acceptable, thinking about backwards-compatibility early will help prevent having your pull request rejected on these grounds. Put yourself in the shoes of someone whose code will be broken by the change(s) introduced by the pull request. It is quite likely that any change made will break someone's code, so you need to have a good reason to make a change as you will be forcing someone to update their code. (This obviously does not apply to new classes or functions; new arguments should be optional and have default values which maintain the existing behavior.) If in doubt, have a look at [PEP 387](#) or *discuss* the issue with experienced developers.

Third, make sure you have proper tests to verify your pull request works as expected. Pull requests will not be accepted without the proper tests!

Fourth, make sure the entire test suite *runs without failure* because of your changes. It is not sufficient to only run whichever test seems impacted by your changes, because there might be interferences unknown to you between your changes and some other part of the interpreter.

Fifth, proper *documentation* additions/changes should be included.

10.3.5 patchcheck

`patchcheck` is a simple automated patch checklist that guides a developer through the common patch generation checks. To run `patchcheck`:

On *UNIX* (including Mac OS X):

```
make patchcheck
```

On *Windows* (after any successful build):

```
python.bat Tools\scripts\patchcheck.py
```

The automated patch checklist runs through:

- Are there any whitespace problems in Python files? (using `Tools/scripts/reindent.py`)
- Are there any whitespace problems in C files?
- Are there any whitespace problems in the documentation? (using `Tools/scripts/reindent-rst.py`)
- Has the documentation been updated?
- Has the test suite been updated?
- Has an entry under `Misc/NEWS.d/next` been added?
- Has `Misc/ACKS` been updated?
- Has `configure` been regenerated, if necessary?
- Has `pyconfig.h.in` been regenerated, if necessary?

The automated patch check doesn't actually *answer* all of these questions. Aside from the whitespace checks, the tool is a memory aid for the various elements that can go into making a complete patch.

10.3.6 Making Good Commits

Each feature or bugfix should be addressed by a single pull request, and for each pull request there may be several commits. In particular:

- Do **not** fix more than one issue in the same commit (except, of course, if one code change fixes all of them).
- Do **not** do cosmetic changes to unrelated code in the same commit as some feature/bugfix.

Commit messages should follow the following structure:

```
bpo-42: Make the spam module more spammy (GH-NNNN)

The spam module sporadically came up short on spam. This change
raises the amount of spam in the module by making it more spammy.
```

The first line or sentence is meant to be a dense, to-the-point explanation of what the purpose of the commit is. The imperative form (used in the example above) is strongly preferred to a descriptive form such as 'the spam module is now more spammy'. Use `git log --oneline` to see existing title lines. Furthermore, the first line should not end in a period.

If this is not enough detail for a commit, a new paragraph(s) can be added to explain in proper depth what has happened (detail should be good enough that a core developer reading the commit message understands the justification for the change).

Check *the git bootcamp* for further instructions on how the commit message should look like when merging a pull request.

Note: [How to Write a Git Commit Message](#) is a nice article that describes how to write a good commit message.

10.3.7 Licensing

To accept your change we must have your formal approval for distributing your work under the [PSF license](#). Therefore, you need to sign a [contributor agreement](#) which allows the [Python Software Foundation](#) to license your code for use with Python (you retain the copyright).

Note: You only have to sign this document once, it will then apply to all your further contributions to Python.

Here are the steps needed in order to sign the CLA:

1. If you don't have an account on bugs.python.org (aka b.p.o), please [register](#) to create one.
2. Make sure your GitHub username is listed in the "Your Details" section at b.p.o.
3. Fill out and sign the [PSF contributor form](#). The "bugs.python.org username" requested by the form is the "Login name" field under "Your Details".

After signing the CLA, please **wait at least one US business day** and then check the status by going to the [check-python-cla](#) website. The check will also be run automatically the next time you push changes to your PR.

10.3.8 Submitting

Once you are satisfied with your work you will want to commit your changes to your branch. In general you can run `git commit -a` and that will commit everything. You can always run `git status` to see what changes are outstanding.

When all of your changes are committed (i.e. `git status` doesn't list anything), you will want to push your branch to your fork:

```
git push origin <branch name>
```

This will get your changes up to GitHub.

Now you want to [create a pull request from your fork](#). If this is a pull request in response to a pre-existing issue on the [issue tracker](#), please make sure to reference the issue number using `bp0-NNNN` in the pull request title or message.

If this is a pull request for an unreported issue (assuming you already performed a search on the issue tracker for a pre-existing issue), create a new issue and reference it in the pull request. Please fill in as much relevant detail as possible to prevent reviewers from having to delay reviewing your pull request because of lack of information.

If this issue is so simple that there's no need for an issue to track any discussion of what the pull request is trying to solve (e.g. fixing a spelling mistake), then the pull request needs to have the "skip issue" label added to it by someone with commit access.

Your pull request may involve several commits as a result of addressing code review comments. Please keep the commit history in the pull request intact by not squashing, amending, or anything that would require a force push to GitHub. A detailed commit history allows reviewers to view the diff of one commit to another so they can easily verify whether their comments have been addressed. The commits will be squashed when the pull request is merged.

10.3.9 Converting an Existing Patch from b.p.o to GitHub

When a patch exists in the [issue tracker](#) that should be converted into a GitHub pull request, please first ask the original patch author to prepare their own pull request. If the author does not respond after a week, it is acceptable for another contributor to prepare the pull request based on the existing patch. In this case, both parties should sign the [CLA](#). When creating a pull request based on another person's patch, provide attribution to the original patch author by adding "Co-authored-by: Author Name <email_address> ." to the pull request description and commit message. See [the GitHub article](#) on how to properly add the co-author info.

See also [Applying a Patch from Mercurial to Git](#).

10.3.10 Reviewing

To begin with, please be patient! There are many more people submitting pull requests than there are people capable of reviewing your pull request. Getting your pull request reviewed requires a reviewer to have the spare time and motivation to look at your pull request (we cannot force anyone to review pull requests and no one is employed to look at pull requests). If your pull request has not received any notice from reviewers (i.e., no comment made) after one month, first "ping" the issue on the [issue tracker](#) to remind the nosy list that the pull request needs a review. If you don't get a response within a week after pinging the issue, then you can try emailing python-dev@python.org to ask for someone to review your pull request.

When someone does manage to find the time to look at your pull request they will most likely make comments about how it can be improved (don't worry, even core developers of Python have their pull requests sent back to them for changes). It is then expected that you update your pull request to address these comments, and the review process will thus iterate until a satisfactory solution has emerged.

How to Review a Pull Request

One of the bottlenecks in the Python development process is the lack of code reviews. If you browse the bug tracker, you will see that numerous issues have a fix, but cannot be merged into the main source code repository, because no one has reviewed the proposed solution. Reviewing a pull request can be just as informative as providing a pull request and it will allow you to give constructive comments on another developer's work. This guide provides a checklist for submitting a code review. It is a common misconception that in order to be useful, a code review has to be perfect. This is not the case at all! It is helpful to just test the pull request and/or play around with the code and leave comments in the pull request or issue tracker.

1. If you have not already done so, get a copy of the CPython repository by following the *setup guide*, build it and run the tests.
2. Check the bug tracker to see what steps are necessary to reproduce the issue and confirm that you can reproduce the issue in your version of the Python REPL (the interactive shell prompt), which you can launch by executing `./python` inside the repository.
3. Checkout and apply the pull request (Please refer to the instruction *Downloading Other's Patches*)
4. If the changes affect any C file, run the build again.
5. Launch the Python REPL (the interactive shell prompt) and check if you can reproduce the issue. Now that the pull request has been applied, the issue should be fixed (in theory, but mistakes do happen! A good review aims to catch these before the code is merged into the Python repository). You should also try to see if there are any corner cases in this or related issues that the author of the fix may have missed.
6. If you have time, run the entire test suite. If you are pressed for time, run the tests for the module(s) where changes were applied. However, please be aware that if you are recommending a pull request as 'merge-ready', you should always make sure the entire test suite passes.

10.3.11 Leaving a Pull Request Review on GitHub

When you review a pull request, you should provide additional details and context of your review process.

Instead of simply "approving" the pull request, leave comments. For example:

1. If you tested the PR, report the result and the system and version tested on, such as 'Windows 10', 'Ubuntu 16.4', or 'Mac High Sierra'.
2. If you request changes, try to suggest how.
3. Comment on what is "good" about the pull request, not just the "bad". Doing so will make it easier for the PR author to find the good in your comments.

10.3.12 Dismissing Review from Another Core Developer

A core developer can dismiss another core developer's review if they confirmed that the requested changes have been made. When a core developer has assigned the PR to themselves, then it is a sign that they are actively looking after the PR, and their review should not be dismissed.

10.3.13 Committing/Rejecting

Once your pull request has reached an acceptable state (and thus considered "accepted"), it will either be merged or rejected. If it is rejected, please do not take it personally! Your work is still appreciated regardless of whether your pull request is merged. Balancing what *does* and *does not* go into Python is tricky and we simply cannot accept everyone's contributions.

But if your pull request is merged it will then go into Python's VCS (version control system) to be released with the next major release of Python. It may also be backported to older versions of Python as a bugfix if the core developer doing the merge believes it is warranted.

10.3.14 Crediting

Non-trivial contributions are credited in the `Misc/ACKS` file (and, most often, in a contribution's news entry as well). You may be asked to make these edits on the behalf of the core developer who accepts your pull request.

10.4 Running & Writing Tests

Note: This document assumes you are working from an *in-development* checkout of Python. If you are not then some things presented here may not work as they may depend on new features not available in earlier versions of Python.

10.4.1 Running

The shortest, simplest way of running the test suite is the following command from the root directory of your checkout (after you have *built Python*):

```
./python -m test
```

You may need to change this command as follows throughout this section. On *most* Mac OS X systems, replace `./python` with `./python.exe`. On Windows, use `python.bat`. If using Python 2.7, replace `test` with `test.regrtest`.

If you don't have easy access to a command line, you can run the test suite from a Python or IDLE shell:

```
>>> from test import autotest
```

This will run the majority of tests, but exclude a small portion of them; these excluded tests use special kinds of resources: for example, accessing the Internet, or trying to play a sound or to display a graphical interface on your desktop. They are disabled by default so that running the test suite is not too intrusive. To enable some of these additional tests (and for other flags which can help debug various issues such as reference leaks), read the help text:

```
./python -m test -h
```

If you want to run a single test file, simply specify the test file name (without the extension) as an argument. You also probably want to enable verbose mode (using `-v`), so that individual failures are detailed:

```
./python -m test -v test_abc
```

To run a single test case, use the `unittest` module, providing the import path to the test case:

```
./python -m unittest -v test.test_abc.TestABC
```

If you have a multi-core or multi-CPU machine, you can enable parallel testing using several Python processes so as to speed up things:

```
./python -m test -j0
```

If you are running a version of Python prior to 3.3 you must specify the number of processes to run simultaneously (e.g. `-j2`).

Finally, if you want to run tests under a more strenuous set of settings, you can run `test` as:

```
./python -bb -E -Wd -m test -r -w -uall
```

The various extra flags passed to Python cause it to be much stricter about various things (the `-Wd` flag should be `-W error` at some point, but the test suite has not reached a point where all warnings have been dealt with and so we cannot guarantee that a bug-free Python will properly complete a test run with `-W error`). The `-r` flag to the test runner causes it to run tests in a more random order which helps to check that the various tests do not interfere with each other. The `-w` flag causes failing tests to be run again to see if the failures are transient or consistent. The `-uall` flag allows the use of all available resources so as to not skip tests requiring, e.g., Internet access.

To check for reference leaks (only needed if you modified C code), use the `-R` flag. For example, `-R 3:2` will first run the test 3 times to settle down the reference count, and then run it 2 more times to verify if there are any leaks.

You can also execute the `Tools/scripts/run_tests.py` script as found in a CPython checkout. The script tries to balance speed with thoroughness. But if you want the most thorough tests you should use the strenuous approach shown above.

Unexpected Skips

Sometimes when running the test suite, you will see “unexpected skips” reported. These represent cases where an entire test module has been skipped, but the test suite normally expects the tests in that module to be executed on that platform.

Often, the cause is that an optional module hasn't been built due to missing build dependencies. In these cases, the missing module reported when the test is skipped should match one of the modules reported as failing to build when *Compile and build*.

In other cases, the skip message should provide enough detail to help figure out and resolve the cause of the problem (for example, the default security settings on some platforms will disallow some tests)

10.4.2 Writing

Writing tests for Python is much like writing tests for your own code. Tests need to be thorough, fast, isolated, consistently repeatable, and as simple as possible. We try to have tests both for normal behaviour and for error conditions. Tests live in the `Lib/test` directory, where every file that includes tests has a `test_` prefix.

One difference with ordinary testing is that you are encouraged to rely on the `test.support` module. It contains various helpers that are tailored to Python's test suite and help smooth out common problems such as platform differences, resource consumption and cleanup, or warnings management. That module is not suitable for use outside of the standard library.

When you are adding tests to an existing test file, it is also recommended that you study the other tests in that file; it will teach you which precautions you have to take to make your tests robust and portable.

10.4.3 Benchmarks

Benchmarking is useful to test that a change does not degrade performance.

The [Python Benchmark Suite](#) has a collection of benchmarks for all Python implementations. Documentation about running the benchmarks is in the [README.txt](#) of the repo.

10.5 Increase Test Coverage

Python development follows a practice that all semantic changes and additions to the language and `STDLIB` (standard library) are accompanied by appropriate unit tests. Unfortunately Python was in existence for a long time before the practice came into effect. This has left chunks of the `stdlib` untested which is not a desirable situation to be in.

A good, easy way to become acquainted with Python's code and to help out is to help increase the test coverage for Python's `stdlib`. Ideally we would like to have 100% coverage, but any increase is a good one. Do realize, though, that getting 100% coverage is not always possible. There could be platform-specific code that simply will not execute for you, errors in the output, etc. You can use your judgement as to what should and should not be covered, but being conservative and assuming something should be covered is generally a good rule to follow.

Choosing what module you want to increase test coverage for can be done in a couple of ways. You can simply run the entire test suite yourself with coverage turned on and see what modules need help. This has the drawback of running the entire test suite under coverage measuring which takes some time to complete, but you will have an accurate, up-to-date notion of what modules need the most work.

Another is to follow the examples below and simply see what coverage your favorite module has. This is "stabbing in the dark", though, and so it might take some time to find a module that needs coverage help.

Do make sure, though, that for any module you do decide to work on that you run coverage for just that module. This will make sure you know how good the explicit coverage of the module is from its own set of tests instead of from implicit testing by other code that happens to use the module.

10.5.1 Common Gotchas

Please realize that coverage reports on modules already imported before coverage data starts to be recorded will be wrong. Typically you can tell a module falls into this category by the coverage report saying that global statements that would obviously be executed upon import have gone unexecuted while local statements have been covered. In these instances you can ignore the global statement coverage and simply focus on the local statement coverage.

When writing new tests to increase coverage, do take note of the style of tests already provided for a module (e.g., whitebox, blackbox, etc.). As some modules are primarily maintained by a single core developer they may have a specific preference as to what kind of test is used (e.g., whitebox) and prefer that other types of tests not be used (e.g., blackbox). When in doubt, stick with whitebox testing in order to properly exercise the code.

10.5.2 Measuring Coverage

It should be noted that a quirk of running coverage over Python's own `stdlib` is that certain modules are imported as part of interpreter startup. Those modules required by Python itself will not be viewed as executed by the coverage tools and thus look like they have very poor coverage (e.g., the `stat` module). In these instances the module will appear to not have any coverage of global statements but will have proper coverage of local statements (e.g., function definitions will not be traced, but the function bodies will). Calculating the coverage of modules in this situation will simply require manually looking at what local statements were not executed.

Using `coverage.py`

One of the most popular third-party coverage tools is `coverage.py` which provides very nice HTML output along with advanced features such as *branch coverage*. If you prefer to stay with tools only provided by the `stdlib` then you can use `test.regrtest`.

Install Coverage

By default, pip will not install into the in-development version of Python you just built, and this built version of Python will not see packages installed into your default version of Python. One option is to use a virtual environment to install coverage:

```
./python -m venv ../cpython-venv
source ../cpython-venv/bin/activate
pip install coverage
```

On *most* Mac OS X systems, replace `./python` with `./python.exe`. On Windows, use `python.bat`.

You can now use `python` without the `./` for the rest of these instructions, as long as your `venv` is activated. For more info on `venv` see [Virtual Environment](#) documentation.

If this does not work for you for some reason, you should try using the in-development version of `coverage.py` to see if it has been updated as needed. To do this you should clone/check out the development version of `coverage.py`:

```
git clone https://github.com/nedbat/coveragepy.git
```

You will need to use the full path to the installation.

Another option is to use an installed copy of `coverage.py`, if you already have it. For this, you will again need to use the full path to that installation.

Basic Usage

The following command will tell you if your copy of `coverage` works (substitute `COVERAGEDIR` with the directory where your clone exists, e.g. `../coveragepy`):

```
./python COVERAGEDIR
```

`Coverage.py` will print out a little bit of helper text verifying that everything is working. If you are using an installed copy, you can do the following instead (note this must be installed using the built copy of Python, such as by `venv`):

```
./python -m coverage
```

The rest of the examples on how to use `coverage.py` will assume you are using a cloned copy, but you can substitute the above and all instructions should still be valid.

To run the test suite under `coverage.py`, do the following:

```
./python COVERAGEDIR run --pylib Lib/test/regtest.py
```

To run only a single test, specify the module/package being tested in the `--source` flag (so as to prune the coverage reporting to only the module/package you are interested in) and then append the name of the test you wish to run to the command:

```
./python COVERAGEDIR run --pylib --source=abc Lib/test/regtest.py test_abc
```

To see the results of the coverage run, you can view a text-based report with:

```
./python COVERAGEDIR report
```

You can use the `--show-missing` flag to get a list of lines that were not executed:

```
./python COVERAGEDIR report --show-missing
```

But one of the strengths of `coverage.py` is its HTML-based reports which let you visually see what lines of code were not tested:

```
./python COVERAGEDIR html -i --include=`pwd`/Lib/* --omit="Lib/test/*,Lib/*/tests/*"
```

This will generate an HTML report in a directory named `htmlcov` which ignores any errors that may arise and ignores modules for which test coverage is unimportant (e.g. tests, temp files, etc.). You can then open the `htmlcov/index.html` file in a web browser to view the coverage results along with pages that visibly show what lines of code were or were not executed.

Branch Coverage

For the truly daring, you can use another powerful feature of `coverage.py`: branch coverage. Testing every possible branch path through code, while a great goal to strive for, is a secondary goal to getting 100% line coverage for the entire stdlib (for now).

If you decide you want to try to improve branch coverage, simply add the `--branch` flag to your coverage run:

```
./python COVERAGEDIR run --pylib --branch <arguments to run test(s)>
```

This will lead to the report stating not only what lines were not covered, but also what branch paths were not executed.

Coverage Results For Modules Imported Early On

For the *truly truly* daring, you can use a hack to get `coverage.py` to include coverage for modules that are imported early on during CPython's startup (e.g. the `encodings` module). Do not worry if you can't get this to work or it doesn't make any sense; it's entirely optional and only important for a small number of modules.

If you still choose to try this, the first step is to build `coverage.py`'s C extension code. Assuming that `coverage.py`'s clone is at `COVERAGEDIR` and your clone of CPython is at `CPYTHONDIR`, you execute the following in your `coverage.py` clone:

```
CPPFLAGS="-I CPYTHONDIR -I CPYTHONDIR/Include" CPYTHONDIR/python setup.py build_ext --  
↪ inplace
```

This will build `coverage.py`'s C extension code in-place, allowing the previous instructions on how to gather coverage to continue to work.

To get `coverage.py` to be able to gather the most accurate coverage data on as many modules as possible **with a HORRIBLE HACK that you should NEVER use in your own code**, run the following from your CPython clone:

```
PYTHONPATH=COVERAGEDIR/coverage/fullcoverage ./python COVERAGEDIR run --pylib Lib/  
↪ test/regrtest.py
```

This will give you the most complete coverage possible for CPython's standard library.

Using `test.regrtest`

If you prefer to rely solely on the stdlib to generate coverage data, you can do so by passing the appropriate flags to `test` (along with any other flags you want to):

```
./python -m test --coverage -D `pwd`/coverage_data <test arguments>
```


Do note the argument to `-D`; if you do not specify an absolute path to where you want the coverage data to end up it will go somewhere you don't expect.

Note: If you are running coverage over the entire test suite, make sure to add `-x test_importlib test_runpy test_trace` to exclude those tests as they trigger exceptions during coverage; see <https://bugs.python.org/issue10541> and <https://bugs.python.org/issue10991>.

Once the tests are done you will find the directory you specified contains files for each executed module along with which lines were executed how many times.

10.5.3 Filing the Issue

Once you have increased coverage, you need to create an issue on the [issue tracker](#) and submit a *pull request*. On the issue set the “Components” to “Test” and “Versions” to the version of Python you worked on (i.e., the in-development version).

10.5.4 Measuring coverage of C code with `gcov` and `lcov`

It's also possible to measure the function, line and branch coverage of Python's C code. Right now only GCC with `gcov` is supported. In order to create an instrumented build of Python with `gcov`, run:

```
make coverage
```

Then run some code and gather coverage data with the `gcov` command. In order to create a HTML report you can install `lcov`. The command:

```
make coverage-lcov
```

assembles coverage data, removes 3rd party and system libraries and finally creates a report. You can skip both steps and just run:

```
make coverage-report
```

if you like to generate a coverage report for Python's stdlib tests. It takes about 20 to 30 minutes on a modern computer.

Note: Multiple test jobs may not work properly. C coverage reporting has only been tested with a single test process.

10.6 Helping with Documentation

Python is known for having well-written documentation. Maintaining the documentation's accuracy and keeping a high level of quality takes a lot of effort. Community members, like you, help with writing, editing, and updating content, and these contributions are appreciated and welcomed.

This high-level **Helping with Documentation** section provides:

- an overview of Python's documentation
- how to help with documentation issues
- information on proofreading

- guidance on contributing to this Developer's Guide

The next chapter, *Documenting Python*, gives extensive, detailed information on how to write documentation and submit changes.

10.6.1 Python Documentation

The *Documenting Python* section covers the details of how Python's documentation works. It includes information about the markup language used, specific formats, and style recommendations. Looking at pre-existing documentation source files can be very helpful when getting started. *How to build the documentation* walks you through the steps to create a draft build which lets you see how your changes will look and validates that your new markup is correct.

You can view the documentation built from *in-development* and *maintenance* branches at <https://docs.python.org/dev/>. The in-development and most recent 3.x (as well as 2.x) maintenance branches are rebuilt once per day.

If you would like to be more involved with documentation, consider subscribing to the docs@python.org mailing list. The [issue tracker](#) sends new documentation issues to this mailing list, and, less frequently, the list receives some directly mailed bug reports. The docs-sig@python.org mailing list discusses the documentation toolchain, projects, and standards.

10.6.2 Helping with documentation issues

If you look at [documentation issues](#) on the [issue tracker](#), you will find various documentation problems that may need work. Issues vary from typos to unclear documentation and items lacking documentation.

If you see a documentation issue that you would like to tackle, you can:

- check to see if there is a paperclip or [octocat](#) icon at the end of the issue's title column. If there is, then someone has already created a pull request for the issue.
- leave a comment on the issue saying you are going to try and create a pull request and roughly how long you think you will take to do so (this allows others to take on the issue if you happen to forget or lose interest).
- submit a [pull request](#) for the issue.

By following the steps in the *Quick Guide to Pull Requests*, you will learn the workflow for documentation pull requests.

10.6.3 Proofreading

While an issue filed on the [issue tracker](#) means there is a known issue somewhere, that does not mean there are not other issues lurking about in the documentation. Proofreading a part of the documentation, such as a "How to" or OS specific document, can often uncover problems (e.g., documentation that needs updating for Python 3).

If you decide to proofread, read a section of the documentation from start to finish, filing issues in the issue tracker for each major type of problem you find. Simple typos don't require issues of their own, but, instead, submit a pull request directly. It's best to avoid filing a single issue for an entire section containing multiple problems; instead, file several issues so that it is easier to break the work up for multiple people and more efficient review.

10.6.4 Helping with the Developer's Guide

The Developer's Guide (what you're reading now) uses the same process as the main Python documentation, except for some small differences. The source lives in a [separate repository](#) and bug reports should be submitted to the [devguide GitHub tracker](#).

Our devguide workflow uses continuous integration and deployment so changes to the devguide are normally published when the pull request is merged. Changes to CPython documentation follows the workflow of a CPython release and is published in the release.

10.6.5 Developer's Guide workflow

To submit a *pull request*, you can fork the *devguide repo* to your GitHub account and clone it using:

```
$ git clone https://github.com/<your_username>/devguide
```

For your PR to be accepted, you will also need to sign the *contributor agreement*.

To build the devguide, some additional dependencies are required (most importantly, *Sphinx*), and the standard way to install dependencies in Python projects is to create a virtualenv, and then install dependencies from a `requirements.txt` file. For your convenience, this is all *automated for you*. To build the devguide on a Unix-like system use:

```
$ make html
```

in the checkout directory. On Windows use:

```
> .\make html
```

You will find the generated files in `_build/html`. Note that `make check` runs automatically when you submit a *pull request*. You may wish to run `make check` and `make linkcheck` to make sure that it runs without errors.

10.7 Documenting Python

The Python language has a substantial body of documentation, much of it contributed by various authors. The markup used for the Python documentation is *reStructuredText*, developed by the *docutils* project, amended by custom directives and using a toolset named *Sphinx* to post-process the HTML output.

This document describes the style guide for our documentation as well as the custom *reStructuredText* markup introduced by *Sphinx* to support Python documentation and how it should be used.

The documentation in HTML, PDF or EPUB format is generated from text files written using the *reStructuredText* format and contained in the *CPython Git repository*.

Note: If you're interested in contributing to Python's documentation, there's no need to write *reStructuredText* if you're not so inclined; plain text contributions are more than welcome as well. Send an e-mail to docs@python.org or open an issue on the [tracker](#).

10.7.1 Introduction

Python's documentation has long been considered to be good for a free programming language. There are a number of reasons for this, the most important being the early commitment of Python's creator, Guido van Rossum, to providing documentation on the language and its libraries, and the continuing involvement of the user community in providing assistance for creating and maintaining documentation.

The involvement of the community takes many forms, from authoring to bug reports to just plain complaining when the documentation could be more complete or easier to use.

This document is aimed at authors and potential authors of documentation for Python. More specifically, it is for people contributing to the standard documentation and developing additional documents using the same tools as the standard documents. This guide will be less useful for authors using the Python documentation tools for topics other than Python, and less useful still for authors not using the tools at all.

If your interest is in contributing to the Python documentation, but you don't have the time or inclination to learn reStructuredText and the markup structures documented here, there's a welcoming place for you among the Python contributors as well. Any time you feel that you can clarify existing documentation or provide documentation that's missing, the existing documentation team will gladly work with you to integrate your text, dealing with the markup for you. Please don't let the material in this document stand between the documentation and your desire to help out!

10.7.2 Style guide

Use of whitespace

All reST files use an indentation of 3 spaces; no tabs are allowed. The maximum line length is 80 characters for normal text, but tables, deeply indented code samples and long links may extend beyond that. Code example bodies should use normal Python 4-space indentation.

Make generous use of blank lines where applicable; they help group things together.

A sentence-ending period may be followed by one or two spaces; while reST ignores the second space, it is customarily put in by some users, for example to aid Emacs' auto-fill mode.

Footnotes

Footnotes are generally discouraged, though they may be used when they are the best way to present specific information. When a footnote reference is added at the end of the sentence, it should follow the sentence-ending punctuation. The reST markup should appear something like this:

```
This sentence has a footnote reference. [#]_ This is the next sentence.
```

Footnotes should be gathered at the end of a file, or if the file is very long, at the end of a section. The docutils will automatically create backlinks to the footnote reference.

Footnotes may appear in the middle of sentences where appropriate.

Capitalization

Sentence case

Sentence case is a set of capitalization rules used in English sentences: the first word is always capitalized and other words are only capitalized if there is a specific rule requiring it.

In the Python documentation, the use of sentence case in section titles is preferable, but consistency within a unit is more important than following this rule. If you add a section to a chapter where most sections are in title case, you can either convert all titles to sentence case or use the dominant style in the new section title.

Sentences that start with a word for which specific rules require starting it with a lower case letter should be avoided.

Note: Sections that describe a library module often have titles in the form of “`modulename` — Short description of the module.” In this case, the description should be capitalized as a stand-alone sentence.

Many special names are used in the Python documentation, including the names of operating systems, programming languages, standards bodies, and the like. Most of these entities are not assigned any special markup, but the preferred spellings are given here to aid authors in maintaining the consistency of presentation in the Python documentation.

Other terms and words deserve special mention as well; these conventions should be used to ensure consistency throughout the documentation:

CPU For “central processing unit.” Many style guides say this should be spelled out on the first use (and if you must use it, do so!). For the Python documentation, this abbreviation should be avoided since there’s no reasonable way to predict which occurrence will be the first seen by the reader. It is better to use the word “processor” instead.

POSIX The name assigned to a particular group of standards. This is always uppercase.

Python The name of our favorite programming language is always capitalized.

reST For “reStructuredText,” an easy to read, plaintext markup syntax used to produce Python documentation. When spelled out, it is always one word and both forms start with a lower case ‘r’.

Unicode The name of a character coding system. This is always written capitalized.

Unix The name of the operating system developed at AT&T Bell Labs in the early 1970s.

Affirmative Tone

The documentation focuses on affirmatively stating what the language does and how to use it effectively.

Except for certain security or segfault risks, the docs should avoid wording along the lines of “feature x is dangerous” or “experts only”. These kinds of value judgments belong in external blogs and wikis, not in the core documentation.

Bad example (creating worry in the mind of a reader):

Warning: failing to explicitly close a file could result in lost data or excessive resource consumption.
Never rely on reference counting to automatically close a file.

Good example (establishing confident knowledge in the effective use of the language):

A best practice for using files is use a try/finally pair to explicitly close a file after it is used. Alternatively, using a with-statement can achieve the same effect. This assures that files are flushed and file descriptor resources are released in a timely manner.

Economy of Expression

More documentation is not necessarily better documentation. Err on the side of being succinct.

It is an unfortunate fact that making documentation longer can be an impediment to understanding and can result in even more ways to misread or misinterpret the text. Long descriptions full of corner cases and caveats can create the impression that a function is more complex or harder to use than it actually is.

Security Considerations (and Other Concerns)

Some modules provided with Python are inherently exposed to security issues (e.g. shell injection vulnerabilities) due to the purpose of the module (e.g. `ssl`). Littering the documentation of these modules with red warning boxes for

problems that are due to the task at hand, rather than specifically to Python's support for that task, doesn't make for a good reading experience.

Instead, these security concerns should be gathered into a dedicated "Security Considerations" section within the module's documentation, and cross-referenced from the documentation of affected interfaces with a note similar to "Please refer to the `:ref:`security-considerations`` section for important information on how to avoid common mistakes."

Similarly, if there is a common error that affects many interfaces in a module (e.g. OS level pipe buffers filling up and stalling child processes), these can be documented in a "Common Errors" section and cross-referenced rather than repeated for every affected interface.

Code Examples

Short code examples can be a useful adjunct to understanding. Readers can often grasp a simple example more quickly than they can digest a formal description in prose.

People learn faster with concrete, motivating examples that match the context of a typical use case. For instance, the `str.rpartition()` method is better demonstrated with an example splitting the domain from a URL than it would be with an example of removing the last word from a line of Monty Python dialog.

The ellipsis for the `sys.ps2` secondary interpreter prompt should only be used sparingly, where it is necessary to clearly differentiate between input lines and output lines. Besides contributing visual clutter, it makes it difficult for readers to cut-and-paste examples so they can experiment with variations.

Code Equivalents

Giving pure Python code equivalents (or approximate equivalents) can be a useful adjunct to a prose description. A documenter should carefully weigh whether the code equivalent adds value.

A good example is the code equivalent for `all()`. The short 4-line code equivalent is easily digested; it re-emphasizes the early-out behavior; and it clarifies the handling of the corner-case where the iterable is empty. In addition, it serves as a model for people wanting to implement a commonly requested alternative where `all()` would return the specific object evaluating to False whenever the function terminates early.

A more questionable example is the code for `itertools.groupby()`. Its code equivalent borders on being too complex to be a quick aid to understanding. Despite its complexity, the code equivalent was kept because it serves as a model to alternative implementations and because the operation of the "grouper" is more easily shown in code than in English prose.

An example of when not to use a code equivalent is for the `oct()` function. The exact steps in converting a number to octal doesn't add value for a user trying to learn what the function does.

Audience

The tone of the tutorial (and all the docs) needs to be respectful of the reader's intelligence. Don't presume that the readers are stupid. Lay out the relevant information, show motivating use cases, provide glossary links, and do your best to connect-the-dots, but don't talk down to them or waste their time.

The tutorial is meant for newcomers, many of whom will be using the tutorial to evaluate the language as a whole. The experience needs to be positive and not leave the reader with worries that something bad will happen if they make a misstep. The tutorial serves as guide for intelligent and curious readers, saving details for the how-to guides and other sources.

Be careful accepting requests for documentation changes from the rare but vocal category of reader who is looking for vindication for one of their programming errors ("I made a mistake, therefore the docs must be wrong..."). Typically,

the documentation wasn't consulted until after the error was made. It is unfortunate, but typically no documentation edit would have saved the user from making false assumptions about the language ("I was surprised by ...").

10.7.3 reStructuredText Primer

This section is a brief introduction to reStructuredText (reST) concepts and syntax, intended to provide authors with enough information to author documents productively. Since reST was designed to be a simple, unobtrusive markup language, this will not take too long.

See also:

The authoritative [reStructuredText User Documentation](#).

Paragraphs

The paragraph is the most basic block in a reST document. Paragraphs are simply chunks of text separated by one or more blank lines. As in Python, indentation is significant in reST, so all lines of the same paragraph must be left-aligned to the same level of indentation.

Inline markup

The standard reST inline markup is quite simple: use

- one asterisk: `*text*` for emphasis (italics),
- two asterisks: `**text**` for strong emphasis (boldface), and
- backquotes: ``text`` for code samples.

If asterisks or backquotes appear in running text and could be confused with inline markup delimiters, they have to be escaped with a backslash.

Be aware of some restrictions of this markup:

- it may not be nested,
- content may not start or end with whitespace: `* text*` is wrong,
- it must be separated from surrounding text by non-word characters. Use a backslash escaped space to work around that: `this is\ *one*\ word`.

These restrictions may be lifted in future versions of the docutils.

reST also allows for custom "interpreted text roles", which signify that the enclosed text should be interpreted in a specific way. Sphinx uses this to provide semantic markup and cross-referencing of identifiers, as described in the appropriate section. The general syntax is `:rolename:`content``.

Lists and Quotes

List markup is natural: just place an asterisk at the start of a paragraph and indent properly. The same goes for numbered lists; they can also be automatically numbered using a # sign:

```
* This is a bulleted list.
* It has two items, the second
  item uses two lines.

1. This is a numbered list.
```

(continues on next page)

(continued from previous page)

```
2. It has two items too.  
  
#. This is a numbered list.  
#. It has two items too.
```

Nested lists are possible, but be aware that they must be separated from the parent list items by blank lines:

```
* this is  
* a list  
  
  * with a nested list  
  * and some subitems  
  
* and here the parent list continues
```

Definition lists are created as follows:

```
term (up to a line of text)  
  Definition of the term, which must be indented  
  
  and can even consist of multiple paragraphs  
  
next term  
  Description.
```

Paragraphs are quoted by just indenting them more than the surrounding paragraphs.

Source Code

Literal code blocks are introduced by ending a paragraph with the special marker `::`. The literal block must be indented:

```
This is a normal text paragraph. The next paragraph is a code sample::  
  
  It is not processed in any way, except  
  that the indentation is removed.  
  
  It can span multiple lines.  
  
This is a normal text paragraph again.
```

The handling of the `::` marker is smart:

- If it occurs as a paragraph of its own, that paragraph is completely left out of the document.
- If it is preceded by whitespace, the marker is removed.
- If it is preceded by non-whitespace, the marker is replaced by a single colon.

That way, the second sentence in the above example's first paragraph would be rendered as "The next paragraph is a code sample:".

Hyperlinks

External links

Use ``Link text <http://target>`_` for inline web links. If the link text should be the web address, you don't need special markup at all, the parser finds links and mail addresses in ordinary text.

Internal links

Internal linking is done via a special reST role, see the section on specific markup, *Cross-linking markup*.

Sections

Section headers are created by underlining (and optionally overlining) the section title with a punctuation character, at least as long as the text:

```
=====
This is a heading
=====
```

Normally, there are no heading levels assigned to certain characters as the structure is determined from the succession of headings. However, for the Python documentation, here is a suggested convention:

- # with overline, for parts
- * with overline, for chapters
- =, for sections
- -, for subsections
- ^, for subsubsections
- ", for paragraphs

Explicit Markup

“Explicit markup” is used in reST for most constructs that need special handling, such as footnotes, specially-highlighted paragraphs, comments, and generic directives.

An explicit markup block begins with a line starting with `..` followed by whitespace and is terminated by the next paragraph at the same level of indentation. (There needs to be a blank line between explicit markup and normal paragraphs. This may all sound a bit complicated, but it is intuitive enough when you write it.)

Directives

A directive is a generic block of explicit markup. Besides roles, it is one of the extension mechanisms of reST, and Sphinx makes heavy use of it.

Basically, a directive consists of a name, arguments, options and content. (Keep this terminology in mind, it is used in the next chapter describing custom directives.) Looking at this example,

```
.. function:: foo(x)
                foo(y, z)
:bar: no

Return a line of text input from the user.
```

`function` is the directive name. It is given two arguments here, the remainder of the first line and the second line, as well as one option `bar` (as you can see, options are given in the lines immediately following the arguments and indicated by the colons).

The directive content follows after a blank line and is indented relative to the directive start.

Footnotes

For footnotes, use `[#]_` to mark the footnote location, and add the footnote body at the bottom of the document after a “Footnotes” rubric heading, like so:

```
Lorem ipsum [#]_ dolor sit amet ... [#]_

.. rubric:: Footnotes

.. [#] Text of the first footnote.
.. [#] Text of the second footnote.
```

You can also explicitly number the footnotes for better context.

Comments

Every explicit markup block which isn't a valid markup construct (like the footnotes above) is regarded as a comment.

Source encoding

Since the easiest way to include special characters like em dashes or copyright signs in reST is to directly write them as Unicode characters, one has to specify an encoding:

All Python documentation source files must be in UTF-8 encoding, and the HTML documents written from them will be in that encoding as well.

Gotchas

There are some problems one commonly runs into while authoring reST documents:

- **Separation of inline markup:** As said above, inline markup spans must be separated from the surrounding text by non-word characters, you have to use an escaped space to get around that.

10.7.4 Additional Markup Constructs

Sphinx adds a lot of new directives and interpreted text roles to standard reST markup. This section contains the reference material for these facilities. Documentation for “standard” reST constructs is not included here, though they are used in the Python documentation.

Note: This is just an overview of Sphinx' extended markup capabilities; full coverage can be found in [its own documentation](#).

Meta-information markup

sectionauthor

Identifies the author of the current section. The argument should include the author's name such that it can be used for presentation (though it isn't) and email address. The domain name portion of the address should be lower case. Example:

```
.. sectionauthor:: Guido van Rossum <guido@python.org>
```

Currently, this markup isn't reflected in the output in any way, but it helps keep track of contributions.

Module-specific markup

The markup described in this section is used to provide information about a module being documented. Each module should be documented in its own file. Normally this markup appears after the title heading of that file; a typical file might start like this:

```
:mod:`parrot` -- Dead parrot access
=====

.. module:: parrot
   :platform: Unix, Windows
   :synopsis: Analyze and reanimate dead parrots.
.. moduleauthor:: Eric Cleese <eric@python.invalid>
.. moduleauthor:: John Idle <john@python.invalid>
```

As you can see, the module-specific markup consists of two directives, the `module` directive and the `moduleauthor` directive.

module

This directive marks the beginning of the description of a module, package, or submodule. The name should be fully qualified (i.e. including the package name for submodules).

The `platform` option, if present, is a comma-separated list of the platforms on which the module is available (if it is available on all platforms, the option should be omitted). The keys are short identifiers; examples that are in use include "IRIX", "Mac", "Windows", and "Unix". It is important to use a key which has already been used when applicable.

The `synopsis` option should consist of one sentence describing the module's purpose – it is currently only used in the Global Module Index.

The `deprecated` option can be given (with no value) to mark a module as deprecated; it will be designated as such in various locations then.

moduleauthor

The `moduleauthor` directive, which can appear multiple times, names the authors of the module code, just like `sectionauthor` names the author(s) of a piece of documentation. It too does not result in any output currently.

Note: It is important to make the section title of a module-describing file meaningful since that value will be inserted in the table-of-contents trees in overview files.

Information units

There are a number of directives used to describe specific features provided by modules. Each directive requires one or more signatures to provide basic information about what is being described, and the content should be the description. The basic version makes entries in the general index; if no index entry is desired, you can give the directive option flag `:noindex:`. The following example shows all of the features of this directive type:

```
.. function:: spam(eggs)
                ham(eggs)
:noindex:

    Spam or ham the foo.
```

The signatures of object methods or data attributes should not include the class name, but be nested in a class directive. The generated files will reflect this nesting, and the target identifiers (for HTML output) will use both the class and method name, to enable consistent cross-references. If you describe methods belonging to an abstract protocol such as context managers, use a class directive with a (pseudo-)type name too to make the index entries more informative.

The directives are:

c:function

Describes a C function. The signature should be given as in C, e.g.:

```
.. c:function:: PyObject* PyType_GenericAlloc(PyTypeObject *type, Py_ssize_t_
↪nitems)
```

This is also used to describe function-like preprocessor macros. The names of the arguments should be given so they may be used in the description.

Note that you don't have to backslash-escape asterisks in the signature, as it is not parsed by the reST inliner.

c:member

Describes a C struct member. Example signature:

```
.. c:member:: PyObject* PyTypeObject.tp_bases
```

The text of the description should include the range of values allowed, how the value should be interpreted, and whether the value can be changed. References to structure members in text should use the `member` role.

c:macro

Describes a “simple” C macro. Simple macros are macros which are used for code expansion, but which do not take arguments so cannot be described as functions. This is not to be used for simple constant definitions. Examples of its use in the Python documentation include `PyObject_HEAD` and `Py_BEGIN_ALLOW_THREADS`.

c:type

Describes a C type. The signature should just be the type name.

c:var

Describes a global C variable. The signature should include the type, such as:

```
.. c:var:: PyObject* PyClass_Type
```

data

Describes global data in a module, including both variables and values used as “defined constants.” Class and object attributes are not documented using this directive.

exception

Describes an exception class. The signature can, but need not include parentheses with constructor arguments.

function

Describes a module-level function. The signature should include the parameters, enclosing optional parameters in brackets. Default values can be given if it enhances clarity. For example:

```
.. function:: repeat([repeat=3[, number=1000000]])
```

Object methods are not documented using this directive. Bound object methods placed in the module namespace as part of the public interface of the module are documented using this, as they are equivalent to normal functions for most purposes.

The description should include information about the parameters required and how they are used (especially whether mutable objects passed as parameters are modified), side effects, and possible exceptions. A small example may be provided.

coroutinefunction

Describes a module-level coroutine. The description should include similar information to that described for `function`.

decorator

Describes a decorator function. The signature should *not* represent the signature of the actual function, but the usage as a decorator. For example, given the functions

```
def removename(func):
    func.__name__ = ''
    return func

def setnewname(name):
    def decorator(func):
        func.__name__ = name
        return func
    return decorator
```

the descriptions should look like this:

```
.. decorator:: removename

    Remove name of the decorated function.

.. decorator:: setnewname(name)

    Set name of the decorated function to *name*.
```

There is no `deco` role to link to a decorator that is marked up with this directive; rather, use the `:func:` role.

class

Describes a class. The signature can include parentheses with parameters which will be shown as the constructor arguments.

attribute

Describes an object data attribute. The description should include information about the type of the data to be expected and whether it may be changed directly. This directive should be nested in a class directive, like in this example:

```
.. class:: Spam

    Description of the class.

    .. attribute:: ham
```

(continues on next page)

(continued from previous page)

```
Description of the attribute.
```

If it is also possible to document an attribute outside of a class directive, for example if the documentation for different attributes and methods is split in multiple sections. The class name should then be included explicitly:

```
.. attribute:: Spam.eggs
```

method

Describes an object method. The parameters should not include the `self` parameter. The description should include similar information to that described for `function`. This directive should be nested in a `class` directive, like in the example above.

coroutinemethod

Describes an object coroutine method. The parameters should not include the `self` parameter. The description should include similar information to that described for `function`. This directive should be nested in a `class` directive.

decoratormethod

Same as `decorator`, but for decorators that are methods.

Refer to a decorator method using the `:meth:` role.

staticmethod

Describes an object static method. The description should include similar information to that described for `function`. This directive should be nested in a `class` directive.

classmethod

Describes an object class method. The parameters should not include the `cls` parameter. The description should include similar information to that described for `function`. This directive should be nested in a `class` directive.

abstractmethod

Describes an object abstract method. The description should include similar information to that described for `function`. This directive should be nested in a `class` directive.

opcode

Describes a Python `bytecode` instruction.

cmdoption

Describes a Python command line option or switch. Option argument names should be enclosed in angle brackets. Example:

```
.. cmdoption:: -m <module>  
  
Run a module as a script.
```

envvar

Describes an environment variable that Python uses or defines.

There is also a generic version of these directives:

describe

This directive produces the same formatting as the specific ones explained above but does not create index entries or cross-referencing targets. It is used, for example, to describe the directives in this document. Example:

```
.. describe:: opcode  
  
Describes a Python bytecode instruction.
```

Showing code examples

Examples of Python source code or interactive sessions are represented using standard reST literal blocks. They are started by a `:` at the end of the preceding paragraph and delimited by indentation.

Representing an interactive session requires including the prompts and output along with the Python code. No special markup is required for interactive sessions. After the last line of input or output presented, there should not be an “unused” primary prompt; this is an example of what *not* to do:

```
>>> 1 + 1
2
>>>
```

Syntax highlighting is handled in a smart way:

- There is a “highlighting language” for each source file. By default, this is 'python' as the majority of files will have to highlight Python snippets.
- Within Python highlighting mode, interactive sessions are recognized automatically and highlighted appropriately.
- The highlighting language can be changed using the `highlight` directive, used as follows:

```
.. highlight:: c
```

This language is used until the next `highlight` directive is encountered.

- The `code-block` directive can be used to specify the highlight language of a single code block, e.g.:

```
.. code-block:: c

#include <stdio.h>

void main() {
    printf("Hello world!\n");
}
```

- The values normally used for the highlighting language are:
 - python (the default)
 - c
 - rest
 - none (no highlighting)
- If highlighting with the current language fails, the block is not highlighted in any way.

Longer displays of verbatim text may be included by storing the example text in an external file containing only plain text. The file may be included using the `literalinclude` directive.¹ For example, to include the Python source file `example.py`, use:

```
.. literalinclude:: example.py
```

The file name is relative to the current file’s path. Documentation-specific include files should be placed in the `Doc/includes` subdirectory.

¹ There is a standard `include` directive, but it raises errors if the file is not found. This one only emits a warning.

Inline markup

As said before, Sphinx uses interpreted text roles to insert semantic markup in documents.

Names of local variables, such as function/method arguments, are an exception, they should be marked simply with `*var*`.

For all other roles, you have to write `:rolename: `content``.

There are some additional facilities that make cross-referencing roles more versatile:

- You may supply an explicit title and reference target, like in reST direct hyperlinks: `:role: `title` <target>`` will refer to *target*, but the link text will be *title*.
- If you prefix the content with `!`, no reference/hyperlink will be created.
- For the Python object roles, if you prefix the content with `~`, the link text will only be the last component of the target. For example, `:meth: `~Queue.Queue.get`` will refer to `Queue.Queue.get` but only display `get` as the link text.

In HTML output, the link's `title` attribute (that is e.g. shown as a tool-tip on mouse-hover) will always be the full target name.

The following roles refer to objects in modules and are possibly hyperlinked if a matching identifier is found:

mod

The name of a module; a dotted name may be used. This should also be used for package names.

func

The name of a Python function; dotted names may be used. The role text should not include trailing parentheses to enhance readability. The parentheses are stripped when searching for identifiers.

data

The name of a module-level variable or constant.

const

The name of a “defined” constant. This may be a C-language `#define` or a Python variable that is not intended to be changed.

class

A class name; a dotted name may be used.

meth

The name of a method of an object. The role text should include the type name and the method name. A dotted name may be used.

attr

The name of a data attribute of an object.

exc

The name of an exception. A dotted name may be used.

The name enclosed in this markup can include a module name and/or a class name. For example, `:func: `filter`` could refer to a function named `filter` in the current module, or the built-in function of that name. In contrast, `:func: `foo.filter`` clearly refers to the `filter` function in the `foo` module.

Normally, names in these roles are searched first without any further qualification, then with the current module name prepended, then with the current module and class name (if any) prepended. If you prefix the name with a dot, this order is reversed. For example, in the documentation of the `codecs` module, `:func: `open`` always refers to the built-in function, while `:func: ` .open`` refers to `codecs.open()`.

A similar heuristic is used to determine whether the name is an attribute of the currently documented class.

The following roles create cross-references to C-language constructs if they are defined in the API documentation:

c:data

The name of a C-language variable.

c:func

The name of a C-language function. Should include trailing parentheses.

c:macro

The name of a “simple” C macro, as defined above.

c:type

The name of a C-language type.

c:member

The name of a C type member, as defined above.

The following roles do not refer to objects, but can create cross-references or internal links:

envvar

An environment variable. Index entries are generated.

keyword

The name of a Python keyword. Using this role will generate a link to the documentation of the keyword. `True`, `False` and `None` do not use this role, but simple code markup (```True```), given that they're fundamental to the language and should be known to any programmer.

option

A command-line option of Python. The leading hyphen(s) must be included. If a matching `cmdoption` directive exists, it is linked to. For options of other programs or scripts, use simple ```code``` markup.

token

The name of a grammar token (used in the reference manual to create links between production displays).

The following role creates a cross-reference to the term in the glossary:

term

Reference to a term in the glossary. The glossary is created using the `glossary` directive containing a definition list with terms and definitions. It does not have to be in the same file as the `term` markup, in fact, by default the Python docs have one global glossary in the `glossary.rst` file.

If you use a term that's not explained in a glossary, you'll get a warning during build.

The following roles don't do anything special except formatting the text in a different style:

command

The name of an OS-level command, such as `rm`.

dfn

Mark the defining instance of a term in the text. (No index entries are generated.)

file

The name of a file or directory. Within the contents, you can use curly braces to indicate a “variable” part, for example:

```
``spam`` is installed in :file:`/usr/lib/python2.{x}/site-packages` ...
```

In the built documentation, the `x` will be displayed differently to indicate that it is to be replaced by the Python minor version.

guilabel

Labels presented as part of an interactive user interface should be marked using `guilabel`. This includes labels from text-based interfaces such as those created using `curses` or other text-based libraries. Any label used in the interface should be marked with this role, including button labels, window titles, field names, menu and menu selection names, and even values in selection lists.

kbd

Mark a sequence of keystrokes. What form the key sequence takes may depend on platform- or application-specific conventions. When there are no relevant conventions, the names of modifier keys should be spelled out, to improve accessibility for new users and non-native speakers. For example, an *xemacs* key sequence may be marked like `:kbd:`C-x C-f``, but without reference to a specific application or platform, the same sequence should be marked as `:kbd:`Control-x Control-f``.

mailheader

The name of an RFC 822-style mail header. This markup does not imply that the header is being used in an email message, but can be used to refer to any header of the same “style.” This is also used for headers defined by the various MIME specifications. The header name should be entered in the same way it would normally be found in practice, with the camel-casing conventions being preferred where there is more than one common usage. For example: `:mailheader:`Content-Type``.

makevar

The name of a **make** variable.

manpage

A reference to a Unix manual page including the section, e.g. `:manpage:`ls(1)``.

menuselection

Menu selections should be marked using the `menuselection` role. This is used to mark a complete sequence of menu selections, including selecting submenus and choosing a specific operation, or any subsequence of such a sequence. The names of individual selections should be separated by `-->`.

For example, to mark the selection “Start > Programs”, use this markup:

```
:menuselection:`Start --> Programs`
```

When including a selection that includes some trailing indicator, such as the ellipsis some operating systems use to indicate that the command opens a dialog, the indicator should be omitted from the selection name.

mimetype

The name of a MIME type, or a component of a MIME type (the major or minor portion, taken alone).

newsgroup

The name of a Usenet newsgroup.

program

The name of an executable program. This may differ from the file name for the executable for some platforms. In particular, the `.exe` (or other) extension should be omitted for Windows programs.

regexp

A regular expression. Quotes should not be included.

samp

A piece of literal text, such as code. Within the contents, you can use curly braces to indicate a “variable” part, as in `:file:.`

If you don't need the “variable part” indication, use the standard ```code``` instead.

The following roles generate external links:

pep

A reference to a Python Enhancement Proposal. This generates appropriate index entries. The text “PEP *number*” is generated; in the HTML output, this text is a hyperlink to an online copy of the specified PEP. Such hyperlinks should not be a substitute for properly documenting the language in the manuals.

rfc

A reference to an Internet Request for Comments. This generates appropriate index entries. The text “RFC *number*” is generated; in the HTML output, this text is a hyperlink to an online copy of the specified RFC.

Note that there are no special roles for including hyperlinks as you can use the standard reST markup for that purpose.

Cross-linking markup

To support cross-referencing to arbitrary sections in the documentation, the standard reST labels are “abused” a bit: Every label must precede a section title; and every label name must be unique throughout the entire documentation source.

You can then reference to these sections using the `:ref:`label-name`` role.

Example:

```
.. _my-reference-label:
Section to cross-reference
-----

This is the text of the section.

It refers to the section itself, see :ref:`my-reference-label`.
```

The `:ref:` invocation is replaced with the section title.

Alternatively, you can reference any label (not just section titles) if you provide the link text `:ref:`link text <reference-label>``.

Paragraph-level markup

These directives create short paragraphs and can be used inside information units as well as normal text:

note

An especially important bit of information about an API that a user should be aware of when using whatever bit of API the note pertains to. The content of the directive should be written in complete sentences and include all appropriate punctuation.

Example:

```
.. note::

    This function is not suitable for sending spam e-mails.
```

warning

An important bit of information about an API that a user should be aware of when using whatever bit of API the warning pertains to. The content of the directive should be written in complete sentences and include all appropriate punctuation. In the interest of not scaring users away from pages filled with warnings, this directive should only be chosen over `note` for information regarding the possibility of crashes, data loss, or security implications.

versionadded

This directive documents the version of Python which added the described feature, or a part of it, to the library or C API. When this applies to an entire module, it should be placed at the top of the module section before any prose.

The first argument must be given and is the version in question. The second argument is optional and can be used to describe the details of the feature.

Example:

```
.. versionadded:: 3.5
```

versionchanged

Similar to `versionadded`, but describes when and what changed in the named feature in some way (new parameters, changed side effects, platform support, etc.). This one *must* have the second argument (explanation of the change).

Example:

```
.. versionchanged:: 3.1
   The *spam* parameter was added.
```

Note that there should be no blank line between the directive head and the explanation; this is to make these blocks visually continuous in the markup.

deprecated

Indicates the version from which the described feature is deprecated.

There is one required argument: the version from which the feature is deprecated.

Example:

```
.. deprecated:: 3.8
```

deprecated-removed

Like `deprecated`, but it also indicates in which version the feature is removed.

There are two required arguments: the version from which the feature is deprecated, and the version in which the feature is removed.

Example:

```
.. deprecated-removed:: 3.8 4.0
```

impl-detail

This directive is used to mark CPython-specific information. Use either with a block content or a single sentence as an argument, i.e. either

```
.. impl-detail::
   This describes some implementation detail.
   More explanation.
```

or

```
.. impl-detail:: This shortly mentions an implementation detail.
```

“CPython implementation detail:” is automatically prepended to the content.

seealso

Many sections include a list of references to module documentation or external documents. These lists are created using the `seealso` directive.

The `seealso` directive is typically placed in a section just before any sub-sections. For the HTML output, it is shown boxed off from the main flow of the text.

The content of the `seealso` directive should be a reST definition list. Example:

```
.. seealso::

   Module :mod:`zipfile`
       Documentation of the :mod:`zipfile` standard module.

   `GNU tar manual, Basic Tar Format <http://link>`_
       Documentation for tar archive files, including GNU tar extensions.
```

rubric

This directive creates a paragraph heading that is not used to create a table of contents node. It is currently used for the “Footnotes” caption.

centered

This directive creates a centered boldfaced paragraph. Use it as follows:

```
.. centered::

   Paragraph contents.
```

Table-of-contents markup

Since reST does not have facilities to interconnect several documents, or split documents into multiple output files, Sphinx uses a custom directive to add relations between the single files the documentation is made of, as well as tables of contents. The `toctree` directive is the central element.

toctree

This directive inserts a “TOC tree” at the current location, using the individual TOCs (including “sub-TOC trees”) of the files given in the directive body. A numeric `maxdepth` option may be given to indicate the depth of the tree; by default, all levels are included.

Consider this example (taken from the library reference index):

```
.. toctree::
   :maxdepth: 2

   intro
   strings
   datatypes
   numeric
   (many more files listed here)
```

This accomplishes two things:

- Tables of contents from all those files are inserted, with a maximum depth of two, that means one nested heading. `toctree` directives in those files are also taken into account.
- Sphinx knows that the relative order of the files `intro`, `strings` and so forth, and it knows that they are children of the shown file, the library index. From this information it generates “next chapter”, “previous chapter” and “parent chapter” links.

In the end, all files included in the build process must occur in one `toctree` directive; Sphinx will emit a warning if it finds a file that is not included, because that means that this file will not be reachable through standard navigation.

The special file `contents.rst` at the root of the source directory is the “root” of the TOC tree hierarchy; from it the “Contents” page is generated.

Index-generating markup

Sphinx automatically creates index entries from all information units (like functions, classes or attributes) like discussed before.

However, there is also an explicit directive available, to make the index more comprehensive and enable index entries in documents where information is not mainly contained in information units, such as the language reference.

The directive is `index` and contains one or more index entries. Each entry consists of a type and a value, separated by a colon.

For example:

```
.. index::
   single: execution; context
   module: __main__
   module: sys
   triple: module; search; path
```

This directive contains five entries, which will be converted to entries in the generated index which link to the exact location of the index statement (or, in case of offline media, the corresponding page number).

The possible entry types are:

single Creates a single index entry. Can be made a subentry by separating the subentry text with a semicolon (this notation is also used below to describe what entries are created).

pair `pair: loop; statement` is a shortcut that creates two index entries, namely `loop; statement` and `statement; loop`.

triple Likewise, `triple: module; search; path` is a shortcut that creates three index entries, which are `module; search path`, `search; path, module` and `path; module search`.

module, keyword, operator, object, exception, statement, builtin These all create two index entries. For example, `module: hashlib` creates the entries `module; hashlib` and `hashlib; module`. The **builtin** entry type is slightly different in that “built-in function” is used in place of “builtin” when creating the two entries.

For index directives containing only “single” entries, there is a shorthand notation:

```
.. index:: BNF, grammar, syntax, notation
```

This creates four index entries.

Grammar production displays

Special markup is available for displaying the productions of a formal grammar. The markup is simple and does not attempt to model all aspects of BNF (or any derived forms), but provides enough to allow context-free grammars to be displayed in a way that causes uses of a symbol to be rendered as hyperlinks to the definition of the symbol. There is this directive:

productionlist

This directive is used to enclose a group of productions. Each production is given on a single line and consists of a name, separated by a colon from the following definition. If the definition spans multiple lines, each continuation line must begin with a colon placed at the same column as in the first line.

Blank lines are not allowed within `productionlist` directive arguments.

The definition can contain token names which are marked as interpreted text (e.g. `unaryneg ::= "-"`integer``) – this generates cross-references to the productions of these tokens.

Note that no further reST parsing is done in the production, so that you don't have to escape `*` or `|` characters.

The following is an example taken from the Python Reference Manual:

```
.. productionlist::
   try_stmt: try1_stmt | try2_stmt
   try1_stmt: "try" ":" `suite`
             : ("except" [ `expression` ["," `target`] ] ":" `suite`)+
             : ["else" ":" `suite`]
             : ["finally" ":" `suite`]
   try2_stmt: "try" ":" `suite`
             : "finally" ":" `suite`
```

Substitutions

The documentation system provides three substitutions that are defined by default. They are set in the build configuration file `conf.py`.

|release|

Replaced by the Python release the documentation refers to. This is the full version string including alpha/beta/release candidate tags, e.g. `2.5.2b3`.

|version|

Replaced by the Python version the documentation refers to. This consists only of the major and minor version parts, e.g. `2.5`, even for version `2.5.1`.

|today|

Replaced by either today's date, or the date set in the build configuration file. Normally has the format `April 14, 2007`.

10.7.5 Building the documentation

The toolset used to build the docs is written in Python and is called `Sphinx`. `Sphinx` is maintained separately and is not included in this tree. Also needed are `blurb`, a tool to create `Misc/NEWS` on demand; and `python-docs-theme`, the `Sphinx` theme for the Python documentation.

To build the documentation, follow the instructions from one of the sections below. You can view the documentation after building the HTML by pointing a browser at the file `Doc/build/html/index.html`.

You are expected to have installed the latest stable version of `Sphinx` and `blurb` on your system or in a `virtualenv` (which can be created using `make venv`), so that the Makefile can find the `sphinx-build` command. You can also specify the location of `sphinx-build` with the `SPHINXBUILD` **make** variable.

Using make / make.bat

On Unix, run the following from the root of your *repository clone* to build the output as HTML:

```
cd Doc
make venv
make html
```

or alternatively `make -C Doc/ venv html`.

You can also use `make help` to see a list of targets supported by **make**. Note that `make check` is automatically run when you submit a *pull request*, so you should make sure that it runs without errors.

On Windows, a `make.bat` batchfile tries to emulate **make** as closely as possible, but the `venv` target is not implemented, so you will probably want to make sure you are working in a virtual environment before proceeding, otherwise all dependencies will be automatically installed on your system.

When ready, run the following from the root of your *repository clone* to build the output as HTML:

```
cd Doc
make html
```

You can also use `make help` to see a list of targets supported by `make.bat`.

See also `Doc/README.rst` for more information.

Using sphinx-build

Sometimes we directly want to execute the `sphinx-build` tool instead of through `make` (although the latter is still the preferred way). In this case, you can use the following command line from the `Doc` directory (make sure to install `Sphinx`, `blurb` and `python-docs-theme` packages from PyPI):

```
sphinx-build -b<builder> . build/<builder>
```

where `<builder>` is one of `html`, `text`, `latex`, or `htmlhelp` (for explanations see the `make` targets above).

10.7.6 Translating

Python documentation translations are governed by the **PEP 545**, they are built by `docsbuild-scripts` and hosted on `docs.python.org`. There are several documentation translations already in production, other are work in progress:

Language	Contact	Links
Arabic (ar)	Abdur-Rahmaan Janhangeer	GitHub
Bengali as spoken in India (bn_IN)	Kushal Das	GitHub
French (fr)	Julien Palard (mdk)	GitHub
Hindi as spoken in india (hi_IN)		GitHub
Hungarian (hu)	Tamás Bajusz (gbtami)	GitHub Mailing List
Indonesian (id)	Oon Arfiandwi	GitHub
Italian (it)		mail
Japanese (ja)	Kinebuchi Tomohiko (cocoatomo)	GitHub Doc
Korean (ko)		GitHub Doc
Lithuanian (lt)		mail
Polish (pl)		mail Translations
Portuguese (pt)	Gustavo Toffo	
Portuguese as spoken in Brasil (pt-br)	Marco Rougeth	GitHub Wiki Telegram article
Russian (ru)		mail
Simplified Chinese (zh-cn)	Shengjing Zhu	Transifex GitHub Doc
Spanish (es)	Raúl Cumplido	GitHub
Traditional Chinese (zh-tw)	Adrian Liaw	GitHub Transifex Doc
Turkish (tr)		GitHub

Starting a new translation

First subscribe to the [doc-sig](#) mailing list, and introduce yourself and the translation you're starting.

Then you can bootstrap your new translation by using our [cookiecutter](#).

The important steps looks like this:

- Create the github repo (anywhere), with the right hierarchy (using the cookiecutter).
- Gather people to help you translating, you can't do it alone.
- You can use any tool to translate, as long as you can synchronize with git. Some are using Transifex, some are using only github, or you can choose another way, it's up to you.
- Ensure we updated this page to reflect your work and progress, either via a PR, or by asking on the doc-sig mailing list.
- When `tutorial/`, `bugs.py` and `library/functions` are complete, ask on doc-sig for your language to be added in the language picker on docs.python.org.

PEP 545 summary:

Here are the essential points of [PEP 545](#):

- Each translation is assigned an appropriate lowercased language tag, with an optional region subtag, joined with a dash, like `pt-br` or `fr`.
- Each translation is under CC0 and marked as so in the README (as in the cookiecutter).
- Translations files are hosted on `https://github.com/python/python-docs-{LANGUAGE_TAG}` (not mandatory to start a translation, but mandatory to land on `docs.python.org`).
- Translations having completed `tutorial/`, `library/stdtypes` and `library/functions` are hosted on `https://docs.python.org/{LANGUAGE_TAG}/{VERSION_TAG}/`.

How to get help

Discussions about translations occur on the [doc-sig](#) mailing list, and there's a freenode IRC channel, [#python-doc](#).

Translation FAQ

Which version of Python documentation should be translated?

Consensus is to work on current stable, you can then propagate your translation from a branch to another using [pomerge](#).

Are there some tools to help in managing the repo?

Here's what's we're using:

- [pomerge](#) to propagate translation from a files to others.
- [pospell](#) to check for typo in po files.
- [powrap](#) to rewrap the `.po` files before committing, this helps keeping git diffs short.
- [potodo](#) to list what needs to be translated.

How a coordinator is elected?

There is no election, each country have to sort this out. Here are some suggestions.

- Coordinator requests are to be public on [doc-sig](#) mailing list.
- If the given language have a native core dev, the core dev have its word on the choice.
- Anyone who wants to become coordinator for its native language, and shows motivation by translating and building a community, will be named coordinator.
- In case of concurrency between two persons, no one will sort this out for you. It is up to you two to organize a local election or whatever needed to sort this out.
- In case a coordinator become inactive or unreachable for a long period of time, someone else can ask for a takeover on [doc-sig](#).

The entry for my translation is missing/not up to date on this page

Ask on [doc-sig](#), or better, make a PR on the [devguide](#).

I have a translation, but not on git, what should I do?

Just ask for help on the [doc-sig](#) mailing list and our [python-fu](#), [git-fu](#) and [bash-fu](#) combined will help you create an appropriate repository. If you use a tool like [transifex](#) don't worry. Keeping them in sync is not that hard.

My git hierarchy does not match yours, can I keep it?

No, inside the `github.com/python` organization we'll all have the exact same hierarchy so bots will be able to build all of our translations. So you may have to convert from one hierarchy to another. Ask for help on the doc-sig mailing list if you're not sure on how to do it.

What hierarchy should I use in my github repository?

As for every project, we have a *branch* per version. We store `po` files in the root of the repository using the `gettext_compact=0` style.

10.8 Silence Warnings From the Test Suite

When running Python's test suite, no warnings should result when you run it under *strenuous testing conditions* (you can ignore the extra flags passed to `test` that cause randomness and parallel execution if you want). Unfortunately new warnings are added to Python on occasion which take some time to eliminate (e.g., `ResourceWarning`). Typically the easy warnings are dealt with quickly, but the more difficult ones that require some thought and work do not get fixed immediately.

If you decide to tackle a warning you have found, open an issue on the [issue tracker](#) (if one has not already been opened) and say you are going to try and tackle the issue, and then proceed to fix the issue.

10.9 Fixing “easy” Issues (and Beyond)

When you feel comfortable enough to want to help tackle issues by trying to create a patch to fix an issue, you can start by looking at the “easy” issues. These issues *should* be ones where it should take no longer than a day or weekend to fix. But because the “easy” classification is typically done at triage time it can turn out to be inaccurate, so do feel free to leave a comment if you think the classification no longer applies.

For the truly adventurous looking for a challenge, you can look for issues that are not considered easy and try to fix those. It must be warned, though, that it is quite possible that a bug that has been left open has been left into that state because of the difficulty compared to the benefit of the fix. It could also still be open because no consensus has been reached on how to fix the issue (although having a patch that proposes a fix can turn the tides of the discussion to help bring it to a close). Regardless of why the issue is open, you can also always provide useful comments if you do attempt a fix, successful or not.

10.10 Issue Tracking

10.10.1 Using the Issue Tracker

If you think you have found a bug in Python, you can report it to the [issue tracker](#). The [issue tracker](#) is also commonly referred to as *bugs.python.org* and *bpo*. Documentation bugs can also be reported there.

You can report bugs with the issue tracker itself to the [meta tracker](#).

If you would like to file an issue about this devguide, please do so at the [devguide repo](#).

Checking if a bug already exists

The first step before filing an issue report is to see whether the problem has already been reported. Checking if the problem is an existing issue will:

- help you see if the problem has already been resolved or has been fixed for the next release
- save time for you and the developers
- help you learn what needs to be done to fix it
- determine if additional information, such as how to replicate the issue, is needed

To do see if the issue already exists, search the bug database using the search box on the top of the issue tracker page. An [advanced search](#) is also available by clicking on “Search” in the sidebar.

Reporting an issue

If the problem you're reporting is not already in the [issue tracker](#), you need to log in by entering your user and password in the form on the left. If you don't already have a tracker account, select the “Register” link or, if you use [OpenID](#), one of the OpenID provider logos in the sidebar.

It is not possible to submit a bug report anonymously.

Once logged in, you can submit a bug by clicking on the “Create New” link in the sidebar.

The submission form has a number of fields, and they are described in detail in the [Triaging an Issue](#) page. This is a short summary:

- in the **Title** field, enter a *very* short description of the problem; less than ten words is good;
- in the **Type** field, select the type of your problem (usually behavior);
- if you know which **Components** and **Versions** are affected by the issue, you can select these too; otherwise, leave them blank;
- last but not least, you have to describe the problem in detail, including what you expected to happen, what did happen, and how to replicate the problem in the **Comment** field. Be sure to include whether any extension modules were involved, and what hardware and software platform you were using (including version information as appropriate).

Understanding the issue's progress and status

The triaging team will take care of setting other fields, and possibly assign the issue to a specific developer. You will automatically receive an update each time an action is taken on the bug.

10.10.2 Disagreement With a Resolution on the Issue Tracker

As humans, we will have differences of opinions from time to time. First and foremost, please be respectful that care, thought, and volunteer time went into the resolution.

With this in mind, take some time to consider any comments made in association with the resolution of the issue. On reflection, the resolution steps may seem more reasonable than you initially thought.

If you still feel the resolution is incorrect, then raise a thoughtful question on [python-dev](#). Further argument and disrespectful discourse on [python-dev](#) after a consensus has been reached amongst the core developers is unlikely to win any converts.

As a reminder, issues closed by a core developer have already been carefully considered. Please do not reopen a closed issue.

10.10.3 Helping Triage Issues

Once you know your way around how Python's source files are structured and you are comfortable working with patches, a great way to contribute is to help triage issues. Do realize, though, that experience working on Python is needed in order to effectively help triage.

Around the clock, new issues are being opened on the [issue tracker](#) and existing issues are being updated. Every issue needs to be triaged to make sure various things are in proper order. Even without special privileges you can help with this process.

Classifying Reports

For bugs, an issue needs to:

- clearly explain the bug so it can be reproduced
- include all relevant platform details
- state what version(s) of Python are affected by the bug.

These are things you can help with once you have experience developing for Python:

- try reproducing the bug: For instance, if a bug is not clearly explained enough for you to reproduce it then there is a good chance a core developer won't be able to either.
- see if the issue happens on a different Python version: It is always helpful to know if a bug not only affects the in-development version of Python, but whether it also affects other versions in maintenance mode.
- write a unit test: If the bug lacks a unit test that should end up in Python's test suite, having that written can be very helpful.

This is all helpful as it allows triagers (i.e., *people with the Developer role on the issue tracker*) to properly classify an issue so it can be handled by the right core developers in a timely fashion.

Reviewing Patches

If an issue has a patch attached that has not been reviewed, you can help by making sure the patch:

- follows the style guides
- applies cleanly to an up-to-date clone
- is a good solution to the problem it is trying to solve
- includes proper tests
- includes proper documentation changes
- submitter is listed in `Misc/ACKS`, either already or the patch adds them

Doing all of this allows core developers and *triagers* to more quickly look for subtle issues that only people with extensive experience working on Python's code base will notice.

Finding an Issue You Can Help With

If you want to help triage issues, you might also want to search for issues in modules which you have a working knowledge. Search for the name of a module in the issue tracker or use the [advanced search](#) to search for specific components (e.g. “Windows” if you are a Windows developer, “Extension Modules” if you are familiar with C, etc.). Finally you can use the “Random issue” link in the sidebar to pick random issues until you find an issue that you like. You may find old issues that can be closed, either because they are no longer valid or they have a patch that is ready to be committed, but no one has had the time to do so.

In the sidebar you can also find links to summaries for easy issues and issues with a patch.

10.10.4 Gaining the “Developer” Role on the Issue Tracker

When you have consistently shown the ability to properly help triage issues without guidance, you may request that you be given the “Developer” role on the [issue tracker](#). You can make the request of any person who already has the Developer role. If they decide you are ready to gain the extra privileges on the tracker they will then act as a mentor to you until you are ready to do things entirely on your own. There is no set rule as to how many issues you need to have helped with before or how long you have been participating. The key requirements are that you show the desire to help, you are able to work well with others (especially those already with the Developer role), and that have a firm grasp of how to do things on the issue tracker properly on your own.

Gaining the Developer role will allow you to set any value on any issue in the tracker, releasing you from the burden of having to ask others to set values on an issue for you in order to properly triage something. This will not only help speed up and simplify your work in helping out, but also help lessen the workload for everyone by gaining your help.

10.10.5 The Meta Tracker

If you find an issue with the [issue tracker](#), you can report it to the [meta tracker](#). The meta tracker is where you file issues against anything you come across when working with the issue tracker itself (e.g you can't attach a file, the layout is broken on your browser, Rietveld gave you an error, etc.).

If you want to contribute to the tracker you can get a checkout of the source and install a local instance where to experiment. You can find detailed instructions on the [Tracker Development](#) page.

See also:

Issues with Python and documentation

The Python issue tracker Where to report issues about Python.

The New-bugs-announce mailing list Where all the new issues created on the tracker are reported.

The Python-bugs-list mailing list Where all the changes to issues are reported.

The meta tracker and its development

The meta tracker Where to report issues about the tracker itself.

The Tracker development wiki page Instructions about setting up a local instance of the bug tracker.

The Tracker-discuss mailing list Discussions about the bug tracker.

10.11 Triageing an Issue

This section of the devguide documents the [issue tracker](#) for users and developers.

Contributors with the Developer role on the [issue tracker](#) can triage issues directly without any assistance.

Additionally, this section provides an overview of the Python triage team.

10.11.1 Python triage team

The Python triage team is a group dedicated towards improving workflow efficiency through thoughtful review and triage of open issues and pull requests. This helps contributors receive timely feedback and enables core developers to focus on reviewed items which reduces their workload. The expectations of this role expand upon the “Developer” role on the [issue tracker](#). The responsibilities listed below are primarily centered around the Python GitHub repositories. This extends beyond CPython, and, as needed, to other repos such as devguide and core-workflow.

Responsibilities include:

- **PR/issue management**
 - Renaming PRs
 - Reviewing PRs
 - Assisting contributors
 - Notifying appropriate core developers
- **Applying appropriate labels to PRs/Issues**
 - Skip news
 - Skip issue
 - Good first issue
 - Other categorizations

As triagers gain experience, they may have some intuition of when a PR should be closed. Triagers can recommend closing a PR, but the final decision must be made by a core developer. By having triagers and core developers work together, the author receives a careful consideration of their PR. This encourages future contributions, regardless of whether their PR is accepted or closed.

Triagers can make use of the `invalid` and `stale` labels to suggest that a PR may be suitable for closure. For more information, see the [GitHub PR labels](#) section.

It is also of paramount importance to treat every contributor to the Python project kindly and with respect. Regardless of whether they're entirely new or a veteran core developer, they're actively choosing to voluntarily donate their time towards the improvement of Python. As is the case with any member of the Python Software Foundation, always follow the [PSF Code of Conduct](#).

10.11.2 Becoming a member of the Python triage team

Any Python core developers are welcome to invite a Python contributor to the Python triage team. Do note that the responsibilities of a Python triager is more elevated than a developer on bpo. For example, the Python triager has access to more repositories than just CPython. Triagers will be responsible to handle not just issues, but also pull requests, and even managing backports.

Any existing developers on b.p.o can transition into becoming a Python triager. They can request this to any core developer, and the core developer can pass the request to the [Python organization admin](#) on GitHub. The request

can be made confidentially via a DM in Zulip or Discourse, or publicly by opening an [issue in the core-workflow repository](#).

Any contributor who is not already a developer on b.p.o can also self-nominate to be a member of Python triage team. They can request this to any core developer, confidentially via DM in Zulip or Discourse, or publicly by opening an issue in core-workflow. If a core developer agrees and is willing to vouch for them, the core developer can pass the request to the GitHub administrator. They should also be added as developer on bpo.

For every new triager, it would be great to announce them in the python-committers mailing list and core-workflow category in Discourse. [Example announcement](#).

GitHub Labels for PRs

An important component of triaging PRs for the CPython repo involves appropriately categorizing them through the usage of labels.

Labels for PRs include:

DO-NOT-MERGE Used on PRs to prevent miss-islington from being able to automatically merge to pull request. This label is appropriate when a PR has a non-trivial conflict with the branch it is being merged into.

expert-asyncio Used for PRs which involve changes to the asyncio module or other asynchronous frameworks that utilize it.

invalid Used manually for PRs that do not meet basic requirements and automatically added by bedevere when PR authors attempt to merge maintenance branches into the master branch. During events such as the October Hacktoberfest, this label will prevent the PR from counting toward the author's contributions.

needs backport to X.Y Used for PRs which are appropriate to backport to branches prior to master. Generally, backports to the maintenance branches are primarily bugfixes and documentation clarifications. Backports to the security branches are strictly reserved for PRs involving security fixes, such as crashes, privilege escalation, and DoS. The use of this label will cause miss-islington to attempt to automatically merge the PR into the branches specified.

OS-X Used for PRs involving changes which only have an effect upon a specific operating system. Current variations of the label include OS-Windows and OS-Mac.

skip issue Used for PRs which involve trivial changes, such as typo fixes, comment changes, and section rephrases. The majority of PRs require an issue to be attached to, but if there are no code changes and the section being modified retains the same meaning, this label might be appropriate.

skip news Similar to the skip issue label, this label is used for PRs which involve trivial changes, backports, or already have a relevant news entry in another PR. Any potentially impactful changes should have a corresponding news entry, but for trivial changes it's commonly at the discretion of the PR author if they wish to opt-out of making one.

sprint Used for PRs authored during an in-person sprint, such as at PyCon, EuroPython, or other official Python events. The label is used to prioritize the review of those PRs during the sprint.

stale Used for PRs that include changes which are no longer relevant or when the author hasn't responded to feedback in a long period of time. This label helps core developers quickly identify PRs that are candidates for closure or require a ping to the author.

type-bugfix Used for PRs that address unintentional behavior, but do not pose significant security concerns. Generally, bugfixes will be attached to a specific issue where the unintended behavior was first reported.

type-documentation Used for PRs that exclusively involve changes to the documentation. Documentation includes *.rst files, docstrings, and code comments.

type-enhancement Used for PRs that provide additional functionality or capabilities beyond the existing specifications.

type-performance Used for PRs that provide performance optimizations.

type-security Used for PRs that involve critical security issues. Less severe security concerns can instead use the `type-bugfix` label.

type-tests Used for PRs that exclusively involve changes to the tests.

10.11.3 Fields in the Issue Tracker

The major elements found in an issue report include:

- Classification (including *Title*) - These fields categorize the issue. The fields include *Title*, *Type*, *Stage*, *Components*, and *Version*.
- Process - These fields indicate the state of the issue and its progress toward resolution. The fields are *Status*, *Resolution*, *Dependencies*, *Supersedes*, *Assigned To*, *Nosy List*, *Priority*, *Keywords*, *Comment*, *File*, *File Description*, *Remote hg repo*, *GitHub PR*.
- Messages
- History

Title

A brief description of the issue. Review whether the title is too generic or specifies an incorrect term or library.

(Optional) Add a prefix at the start of the title to indicate the module, e.g. IDLE, doc, or asyncio.

Type

Describes the type of issue. If an issue does not fit within any specific type, please do not set a type.

Type	Description
be-hav-ior	Unexpected behavior, result, or exception. Most bugs will have this type.
com-pile error	Errors reported by the compiler while compiling Python.
crash	Hard crashes of the Python interpreter – possibly with a core dump or a Windows error box.
en-hance-ment	Issues that propose the addition of new functionality, such as new functions, classes, modules, or even new arguments for existing functions. Also used for improvements in the documentation, test suite and other refactorings. A good place to discuss enhancements prior to filing an issue is python-ideas mailing list.
per-for-mance	Situations where too much time is necessary to complete the task. For example, a common task now takes significantly longer to complete.
re-source us-age	Situations where too many resources (e.g. memory) are used.
se-cu-rity	Issues that might have security implications. Report security vulnerabilities using the procedure found in the Reporting security issues in Python page on the python.org website.

Stage

A needed next action to advance the issue. The *stage* needn't be set until it is clear that the issue has been initially triaged and determined work will be needed.

Stage	Description
test needed	The steps which are needed to reproduce the issue. The bug reporter should post a script, instructions, or example to help someone test or reproduce the issue.
needs patch	A patch or pull request is needed to solve the problem (i.e. fixing the bug or adding the requested improvement).
patch review	A patch or pull request exists, but it needs review. Any triager or core developer may do the review.
commit review	A triager performed a patch review and it looks good. This signals to core developers the patch or pull request needs a quick once-over to make sure nothing was overlooked before committing it.
re-solved	The issue is considered closed and addressed (e.g. patch or pull request committed; expected behavior).

Components

The area or Python library affected by the issue. This is a multi-select field.

Choosing certain components, such as *Documentation*, may cause the issue to be auto-assigned, i.e. the issue tracker may automatically fill in the *Assigned To* field after you press `Submit` changes.

One or more components may be selected for an issue:

Component	Description
2to3 (2.x to 3 conversion tool)	The 2to3 conversion tool in Lib/lib2to3 .
Build	The build process.
ctypes	The ctypes package in Lib/ctypes .
Demos and Tools	The files in Tools and Tools/demo .
Distutils	The distutils package in Lib/distutils .
Documentation	The documentation in Doc (source used to build HTML docs for https://docs.python.org/).
email	The email package and related modules.
Extension Modules	C modules in Modules .
IDLE	The Lib/idlelib package.
Installation	The installation process.
Interpreter Core	The interpreter core. The built-in objects in Objects , the Python , Grammar and Parser dirs.
IO	The I/O system, Lib/io.py and Modules/_io .
Library (Lib)	Python modules in Lib .
Macintosh	The Mac OS X operating system.
Regular Expressions	The Lib/re.py and Modules/_sre.c modules.
Tests	The unittest framework in Lib/unittest The doctest framework Lib/doctest.py . The CPython tests in Lib/test . The test runner in Lib/test/regtest.py . The test support utilities in Lib/test/support .
Tkinter	The Lib/tkinter package.
Unicode	Unicode, codecs, str vs bytes, Objects/unicodeobject.c .
Windows	The Windows operating system.
XML	The Lib/xml package.

Versions

The known versions of Python that the issue affects and should be fixed for.

Thus if an issue for a new feature is assigned for e.g., Python 3.8 but is not applied before Python 3.8.0 is released, this field should be updated to say Python 3.9 as the version and drop Python 3.8.

Priority

What is the severity and urgency?

Priority	Description
low	This is for low-impact bugs.
normal	The default value for most issues filed.
high	Try to fix the issue before the next final release.
critical	Should definitely be fixed for next final release.
deferred blocker	The issue will not hold up the next release, <i>n</i> . It will be promoted to a <i>release blocker</i> for the following release, <i>n+1</i> .
release blocker	The issue must be fixed before <i>any</i> release is made, e.g., will block the next release even if it is an alpha release.

As a guideline, *critical* and above are usually reserved for crashes, serious regressions or breakage of very important APIs. Whether a bug is a *release blocker* for the current *release schedule* is decided by the release manager. Triagers may recommend this priority and should add the release manager to the *nosy list*. If needed, consult the *release schedule* and the release's associated PEP for the release manager's name.

Keywords

Various informational flags about the issue. Multiple values are possible.

Key-word	Description
buildbot	A buildbot triggered the issue being reported.
easy	Fixing the issue should not take longer than a day for someone new to contributing to Python to solve.
easy (C)	Fixing the issue should not take longer than a day for someone new contributing to Python, focused on C.
security_issue	This is a security issue or is related to one. The main difference from the “security” issue type is that this is a definite security problem that has to be dealt with.
PEP 3121	The issue is related to PEP PEP 3121 . Extension Module Initialization and Finalization.
new-comer friendly	Issue suitable for newcomer/first time contributors. Not suitable for experienced contributors. Typically it is straightforward, well-defined, low-risk, and optionally someone is able to mentor the new contributor.
gsoc	The issue would fit as, or is related to, a GSoc project.
needs re-view	The patch or pull request attached to the issue is in need of a review.
patch	There is a patch or pull request attached to the issue.
3.3re-gression	The issue is a regression in 3.3.

Nosy List

A list of people who may be interested in an issue.

It is acceptable to add someone to the nosy list if you think the issue should be brought to their attention. Use the *Experts Index* to know who wants to be added to the nosy list for issues targeting specific areas.

If you are logged in and have JavaScript enabled, you can use the `[+]` button to add yourself to the nosy list (remember to click on “Submit Changes” afterwards). Note that you are added to the nosy automatically when you submit a message.

The nosy list also has an autocomplete that lets you search from the lists of developers and *Experts Index*. The search is case-insensitive and works for real names, modules, interest areas, etc., and only adds the username(s) to the nosy once an entry is selected.

Assigned To

Who is expected to take the next step in resolving the issue.

It is acceptable to assign an issue to someone if the issue cannot move forward without their help, e.g., they need to make a technical decision to allow the issue to move forward. Also consult the *Experts Index* as certain stdlib modules should always be assigned to a specific person.

Note that in order to assign an issue to someone, that person **must** have the *Developer role* on the issue tracker.

Dependencies

The issue requires the listed issue(s) to be resolved first before it can move forward.

Superseder

The issue is a duplicate of the listed issue(s).

Status

Status	Description
open	Issue is not resolved.
pending	The issue is blocked until someone (often the OP (original poster)) provides some critical information; the issue will be closed after a set amount time if no reply comes in. Useful when someone opens an issue that lacks enough information to reproduce the bug reported. Requesting additional information and setting status to <i>pending</i> indicates that the issue should be closed if the necessary information is not provided in a timely manner (i.e. one month).
closed	The issue has been resolved (somehow).

Resolution

Why the issue is in its current state. This is not usually used for issues with the “open” status.

Resolution	Description
open	Issue is not resolved.
duplicate	Duplicate of another issue; should have the <i>Superseder</i> field filled out.
fixed	A fix for the issue was committed.
later	Issue is to be worked on in a later release cycle.
not a bug	For some reason the issue is invalid (e.g. the perceived problem is not a bug in Python).
out of date	The issue has already been fixed, or the problem doesn't exist anymore for other reasons.
postponed	Issue will not be worked on at the moment but in a future minor release version.
rejected	Issue was rejected (especially for feature requests).
remind	The issue is acting as a reminder for someone.
wont fix	Issue will not be fixed, typically because it would cause a backwards-compatibility problem.
works for me	Bug cannot be reproduced.

Mercurial Repository

HTTP link to a Mercurial repository that contains a patch for the issue. A *Create Patch* button will appear that computes a diff for the head revision of the remote branch and attaches it to the issue. The button supports only CPython patches.

If you don't indicate a remote branch, `default` is used. You can indicate a remote branch by adding `#BRANCH` to the end of the URL.

10.11.4 Generating Special Links in a Comment

Using the following abbreviations in a comment will automatically generate a link to relevant web pages.

Comment abbreviation	Description
#<number>, issue<number>, or issue <number>	Links to the tracker issue <number>.
msg<number>	Links to the tracker message <number>.
PR <number>, PR<number>, or pull request <number>	Links to GitHub pull requests .
a 10-, 11-, 12-, or 40-digit hex <number>	Indicates a Git or Mercurial changeset identifier and generates a link to changeset <number> on GitHub or https://hg.python.org/ . The <code>git</code> and <code>hg</code> prefixes can also be used to disambiguate, and must precede the number without spaces.
r<number>, rev<number>, or revision <number>	Indicates a legacy Subversion revision number, a reference to a changeset that was checked in prior to 2011-03-05 when the official Python source code repositories were migrated from the SVN (Subversion) VCS to Mercurial. The issue tracker automatically translates the legacy svn revision <number> to its corresponding Mercurial changeset identifier.
Dir/file.ext or Dir/ file.ext:NNN	Links to files in the Python source code repositories , possibly linking to the line number specified after the <code>:</code> . <code>3.6/Dir/file.ext</code> will generate a link with <code>3.6</code> as branch.
PEP <number> or PEP<number>	Link to the PEP (Python Enhancement Proposal) <number>.
devguide, devguide/ triaging, or devguide/ triaging#generating-special-links-in-a-comment	Links to the Devguide, this page, and this section respectively.

10.11.5 Checklist for Triaging

- Read the issue comment(s).
- **Review and set classification fields**
 - Title: should be concise with specifics which are helpful to someone scanning a list of issue titles. (Optional, if possible) Add a prefix at the start of the title to indicate the module, e.g. IDLE, doc, or async.
 - Type
 - Stage
 - Components: multiple items may be set
 - Versions: set if known, leave blank if unsure. Multiple items may be set.
- **Review and set process fields**
 - Status
 - Resolution
 - Superseder
 - Assigned To

- Nosy List
 - Priority
 - Keywords
- (Optional) Leave a brief comment about the proposed next action needed. If there is a long message list, a summary can be very helpful.

10.12 Following Python's Development

Python's development is communicated through a myriad of ways, mostly through mailing lists, but also other forms.

10.12.1 Mailing Lists

[python-dev](#) is the primary mailing list for discussions about Python's development. The list is open to the public and is subscribed to by all core developers plus many people simply interested in following Python's development. Discussion is focused on issues related to Python's development, such as how to handle a specific issue, a PEP, etc.

- Ideas about new functionality should **not** start here and instead should be sent to [python-ideas](#).
- Technical support questions should also not be asked here and instead should go to [python-list](#) or [python-help](#).

[Python-ideas](#) is a mailing list open to the public to discuss ideas on changing Python. If a new idea does not start here (or [python-list](#), discussed below), it will get redirected here.

Sometimes people post new ideas to [python-list](#) to gather community opinion before heading to [python-ideas](#). The list is also sometimes known as `comp.lang.python`, the name of the newsgroup it mirrors (it is also known by the abbreviation `c.l.py`).

The [python-committers](#) mailing list is a private mailing list for core developers (the archives are publicly available). If something only affects core developers (e.g., the tree is frozen for commits, etc.), it is discussed here instead of [python-dev](#) to keep traffic down on the latter.

[Python-checkins](#) sends out an email for every commit to Python's various repositories from <https://github.com/python/cpython>. All core developers subscribe to this list and are known to reply to these emails to make comments about various issues they catch in the commit. Replies get redirected to [python-dev](#).

There are two mailing lists related to issues on the [issue tracker](#). If you only want an email for when a new issue is open, subscribe to [new-bugs-announce](#). If you would rather receive an email for all changes made to any issue, subscribe to [python-bugs-list](#).

General Python questions should go to [python-list](#) or [tutor](#) or similar resources, such as [StackOverflow](#) or the `#python` IRC channel on [Freenode](#).

[Core-Workflow](#) mailing list is the place to discuss and work on improvements to the CPython core development workflow.

A complete list of Python mailing lists can be found at <https://mail.python.org>. Most lists are also mirrored at [GMANE](#) and can be read and posted to in various ways, including via web browsers, NNTP newsreaders, and RSS feed readers.

10.12.2 Zulip

We have our own [zulipchat](#) instance. This should be used to discuss the development of Python only.

10.12.3 IRC

Some core developers enjoy spending time on IRC discussing various issues regarding Python's development in the `#python-dev` channel on `irc.freenode.net`. This is not a place to ask for help with Python, but to discuss issues related to Python's own development. You can use Freenode's [Web interface](#) if you don't have an IRC client.

10.12.4 Blogs

Several core developers are active bloggers and discuss Python's development that way. You can find their blogs (and various other developers who use Python) at <http://planetpython.org/>.

10.12.5 Standards of behaviour in these communication channels

We try to foster environments of mutual respect, tolerance and encouragement, as described in the PSF's [Diversity Statement](#). Abiding by the guidelines in this document and asking questions or posting suggestions in the appropriate channels are an excellent way to get started on the mutual respect part, greatly increasing the chances of receiving tolerance and encouragement in return.

10.12.6 Setting Expectations for Open Source Participation

Burn-out is common in open source due to a misunderstanding of what users, contributors, and maintainers should expect from each other. Brett Cannon gave a [talk](#) about this topic that sets out to help everyone set reasonable expectations of each other in order to make open source pleasant for everyone involved.

10.12.7 Additional Repositories

[Python Core Workflow](#) hosts the codebase for tools such as `cherry_picker` and `blurb`.

[Python Performance Benchmark](#) project is intended to be an authoritative source of benchmarks for all Python implementations.

10.13 Porting Python to a new platform

The first step is to familiarize yourself with the development toolchain on the platform in question, notably the C compiler. Make sure you can compile and run a hello-world program using the target compiler.

Next, learn how to compile and run the Python interpreter on a platform to which it has already been ported; preferably Unix, but Windows will do, too. The build process for Python, in particular the `Makefile` in the source distribution, will give you a hint on which files to compile for Python. Not all source files are relevant: some are platform specific, others are only used in emergencies (e.g. `getopt.c`).

It is not recommended to start porting Python without at least medium-level understanding of your target platform; i.e. how it is generally used, how to write platform specific apps, etc. Also, some Python knowledge is required, or you will be unable to verify that your port is working correctly.

You will need a `pyconfig.h` file tailored for your platform. You can start with `pyconfig.h.in`, read the comments, and turn on definitions that apply to your platform. Also, you will need a `config.c` file, which lists the built-in modules you support. Again, starting with `Modules/config.c.in` is recommended.

Finally, you will run into some things that are not supported on your target platform. Forget about the `posix` module in the beginning. You can simply comment it out of the `config.c` file.

Keep working on it until you get a `>>>` prompt. You may have to disable the importing of `site.py` by passing the `-S` option. When you have a prompt, bang on it until it executes very simple Python statements.

At some point you will want to use the `os` module; this is the time to start thinking about what to do with the `posix` module. It is okay to simply comment out functions in the `posix` module that cause problems; the remaining ones will be quite useful.

Before you are done, it is highly recommended to run the Python regression test suite, as described in *Running & Writing Tests*.

10.14 How to Become a Core Developer

10.14.1 What it Takes

When you have consistently contributed patches which meet quality standards without requiring extensive rewrites prior to being committed, you may qualify for commit privileges and become a core developer of Python. You must also work well with other core developers (and people in general) as you become an ambassador for the Python project.

Typically a core developer will offer you the chance to gain commit privilege. The person making the offer will become your mentor and watch your commits for a while to make sure you understand the development process. If other core developers agree that you should gain commit privileges you are then extended an official offer. How core developers come to that agreement are outlined in [PEP 13](#).

10.14.2 What it Means

As contributors to the CPython project, our shared responsibility is to collaborate constructively with other contributors, including core developers. This responsibility covers all forms of contribution, whether that's submitting patches to the implementation or documentation, reviewing other peoples' patches, triaging issues on the issue tracker, or discussing design and development ideas on the core mailing lists.

Core developers accept key additional responsibilities around the ongoing management of the project:

- core developers bear the additional responsibility of handling the consequences of accepting a change into the code base or documentation. That includes reverting or fixing it if it causes problems in the Buildbot fleet or someone spots a problem in post-commit review, as well as helping out the release manager in resolving any problems found during the pre-release testing cycle. While all contributors are free to help out with this part of the process, and it is most welcome when they do, the actual responsibility rests with the core developer that merged the change
- core developers also bear the primary responsibility for deciding when changes proposed on the issue tracker should be escalated to `python-ideas` or `python-dev` for wider discussion, as well as suggesting the use of the Python Enhancement Proposal process to manage the design and justification of complex changes, or changes with a potentially significant impact on end users

As a result of the additional responsibilities they accept, core developers gain the privilege of being able to approve proposed changes, as well as being able to reject them as inappropriate. Core developers are also able to request that even already merged changes be escalated to `python-dev` for further discussion, and potentially even reverted prior to release.

Becoming a core developer isn't a binary "all-or-nothing" status - CPython is a large project, and different core developers accept responsibility for making design and development decisions in different areas (as documented in the *Experts Index* and *Developer Log*).

10.14.3 Gaining Commit Privileges

The steps to gaining commit privileges are:

1. A core developer starts a poll at <https://discuss.python.org/c/committers/>
 - Open for 7 days
 - Results shown upon close
2. The poll is announced on python-committers
3. Wait for the poll to close and see if the results confirm your membership as per the voting results required by PEP 13
4. The person who nominated you emails the steering council with your email address and a request that the council either accept or reject the proposed membership
5. Assuming the steering council does not object, a member of the council will email you asking for:
 - Account details as required by <https://github.com/python/voters/>
 - Your preferred email address to subscribe to python-committers with
 - A reminder about the Code of Conduct and to report issues to the PSF Conduct WG
6. Once you have provided the pertinent details, your various new privileges will be turned on
7. Your details will be added to <https://github.com/python/voters/>
8. They will update the devguide to publicly list your team membership at *Developer Log*
9. An announcement email by the steering council member handling your new membership will be sent to python-committers

Mailing Lists

You are expected to subscribe to python-committers, python-dev, python-checkins, and one of new-bugs-announce or python-bugs-list. See *Following Python's Development* for links to these mailing lists.

Sign a Contributor Agreement

Submitting a [contributor form for Python](#) licenses any code you contribute to the Python Software Foundation. While you retain the copyright, giving the PSF the ability to license your code means it can be put under the PSF license so it can be legally distributed with Python.

This is a very important step! Hopefully you have already submitted a contributor agreement if you have been submitting patches. But if you have not done this yet, it is best to do this ASAP, probably before you even do your first commit so as to not forget. Also do not forget to enter your GitHub username into your details on the issue tracker.

Pull Request merging

Once you have your commit privileges on GitHub you will be able to accept pull requests on GitHub. You should plan to continue to submit your own changes through pull requests as if you weren't a core developer to benefit from various things such as automatic integration testing, but you can accept your own pull requests if you feel comfortable doing so.

10.14.4 Responsibilities

As a core developer, there are certain things that are expected of you.

First and foremost, be a good person. This might sound melodramatic, but you are now a member of the Python project and thus represent the project and your fellow core developers whenever you discuss Python with anyone. We have a reputation for being a very nice group of people and we would like to keep it that way. Core developers responsibilities include following the [PSF Code of Conduct](#).

Second, please be prompt in responding to questions. Many contributors to Python are volunteers so what little free time they can dedicate to Python should be spent being productive. If you have been asked to respond to an issue or answer a question and you put it off it ends up stalling other people's work. It is completely acceptable to say you are too busy, but you need to say that instead of leaving people waiting for an answer. This also applies to anything you do on the issue tracker.

Third, please list what areas you want to be considered an expert in the [Experts Index](#). This allows triagers to direct issues to you which involve an area you are an expert in. But, as stated in the second point above, if you do not have the time to answer questions promptly then please remove yourself as needed from the file so that you will not be bothered in the future. Once again, we all understand how life gets in the way, so no one will be insulted if you remove yourself from the list.

Fourth, please consider whether or not you wish to add your name to the [Core Developer Motivations and Affiliations](#) list. Core contributor participation in the list helps the wider Python community to better appreciate the perspectives currently represented amongst the core development team, the Python Software Foundation to better assess the sustainability of current contributions to CPython core development, and also serves as a referral list for organisations seeking commercial Python support from the core development community.

And finally, enjoy yourself! Contributing to open source software should be fun (overall). If you find yourself no longer enjoying the work then either take a break or figure out what you need to do to make it enjoyable again.

10.15 Developer Log

This page lists the historical members of the Python development team. (The master list is kept in a private repository due to containing sensitive contact information.)

Name	GitHub username	Joined
Lysandros Nikolaou	lysnikolaou	2020-06-29
Kyle Stanley	aeros	2020-04-14
Dong-hee Na	corona10	2020-04-08
Karthikeyan Singaravelan	tirkarthi	2019-12-31
Joannah Nanjekye	nanjekyejoannah	2019-09-23
Abhilash Raj	maxking	2019-08-06
Paul Ganssle	pganssle	2019-06-15
Stéphane Wirtel	matrixise	2019-04-08
Stefan Behnel	scoder	2019-04-08
Cheryl Sabella	csabella	2019-02-19
Lisa Roach	lisroach	2018-09-14
Emily Morehouse	emilyemorehouse	2018-09-14
Pablo Galindo	pablogsal	2018-06-06
Mark Shannon	markshannon	2018-05-15
Petr Viktorin	encukou	2018-04-16
Nathaniel J. Smith	njsmith	2018-01-25
Julien Palard	JulienPalard	2017-12-08

Continued on next page

Table 1 – continued from previous page

Name	GitHub username	Joined
Ivan Levkivskiy	ilevkivskiy	2017-12-06
Carol Willing	willingc	2017-05-24
Mariatta	Mariatta	2017-01-27
Xiang Zhang	zhangyangyu	2016-11-21
Inada Naoki	methane	2016-09-26
Xavier de Gaye	xdegaye	2016-06-03
Davin Potts	applio	2016-03-06
Martin Panter	vadmium	2015-08-10
Paul Moore	pfmoore	2015-03-15
Robert Collins	rBTCollins	2014-10-16
Berker Peksag	berkerpeksag	2014-06-26
Steve Dower	zooba	2014-05-10
Kushal Das	kushaldas	2014-04-14
Steven D'Aprano	stevendaprano	2014-02-08
Yury Selivanov	1stl	2014-01-23
Zachary Ware	zware	2013-11-02
Donald Stufft	dstufft	2013-08-14
Ethan Furman	ethanfurman	2013-05-11
Serhiy Storchaka	serhiy-storchaka	2012-12-26
Chris Jerdonek	cjerdonek	2012-09-24
Eric Snow	ericssnowcurrently	2012-09-05
Peter Moody		2012-05-20
Hynek Schlawack	hynek	2012-05-14
Richard Oudkerk		2012-04-29
Andrew Svetlov	asvetlov	2012-03-13
Petri Lehtinen	akheron	2011-10-22
Meador Inge	meadori	2011-09-19
Jeremy Kloth	jkloth	2011-09-12
Sandro Tosi	sandrotosi	2011-08-01
Alex Gaynor	alex	2011-07-18
Charles-François Natali		2011-05-19
Nadeem Vawda		2011-04-10
Jason R. Coombs	jaraco	2011-03-14
Ross Lagerwall		2011-03-13
Eli Bendersky	eliben	2011-01-11
Ned Deily	ned-deily	2011-01-09
David Malcolm	davidmalcolm	2010-10-27
Tal Einat	taleinat	2010-10-04
Łukasz Langa	ambv	2010-09-08
Daniel Stutzbach		2010-08-22
Éric Araujo	merwok	2010-08-10
Brian Quinlan	brianquinlan	2010-07-26
Alexander Belopolsky	abalkin	2010-05-25
Tim Golden	tjguk	2010-04-21
Giampaolo Rodolà	giampaolo	2010-04-17
Jean-Paul Calderone		2010-04-06
Brian Curtin	briancurtin	2010-03-24
Florent Xicluna		2010-02-25
Dino Viehland	DinoV	2010-02-23

Continued on next page

Table 1 – continued from previous page

Name	GitHub username	Joined
Larry Hastings	larryhastings	2010-02-22
Victor Stinner	vstinner	2010-01-30
Stefan Kraah	skraah	2010-01-05
Doug Hellmann	dhellmann	2009-09-20
Frank Wierzbicki		2009-08-02
Ezio Melotti	ezio-melotti	2009-06-07
Philip Jenvey	pjenvey	2009-05-07
Michael Foord	voidspace	2009-04-01
R. David Murray	bitdancer	2009-03-30
Chris Withers	cjw296	2009-03-08
Tarek Ziadé		2008-12-21
Hirokazu Yamamoto		2008-08-12
Armin Ronacher	mitsuhiko	2008-07-23
Antoine Pitrou	pitrou	2008-07-16
Senthil Kumaran	orsenthil	2008-06-16
Jesse Noller		2008-06-16
Jesús Cea	jcea	2008-05-13
Guilherme Polo		2008-04-24
Jeroen Ruigrok van der Werven		2008-04-12
Benjamin Peterson	benjaminp	2008-03-25
David Wolever	wolever	2008-03-17
Trent Nelson	tpn	2008-03-17
Mark Dickinson	mdickinson	2008-01-06
Amaury Forgeot d'Arc	amauryfa	2007-11-09
Christian Heimes	tiran	2007-10-31
Bill Janssen		2007-08-28
Jeffrey Yasskin		2007-08-09
Mark Summerfield		2007-08-01
Alexandre Vassalotti	avassalotti	2007-05-21
Travis E. Oliphant		2007-04-17
Eric V. Smith	ericvsmith	2007-02-28
Josiah Carlson		2007-01-06
Collin Winter		2007-01-05
Richard Jones		2006-05-23
Kristján Valur Jónsson		2006-05-17
Jack Diederich	jackdied	2006-05-17
Steven Bethard		2006-04-27
Gerhard Häring		2006-04-23
George Yoshida		2006-04-17
Ronald Oussoren	ronaldoussoren	2006-03-03
Nick Coghlan	ncoghlan	2005-10-16
Georg Brandl	birkenfeld	2005-05-28
Terry Jan Reedy	terryreedy	2005-04-07
Bob Ippolito		2005-03-02
Peter Astrand		2004-10-21
Facundo Batista	facundobatista	2004-10-16
Sean Reifschneider		2004-09-17

Continued on next page

Table 1 – continued from previous page

Name	GitHub username	Joined
Johannes Gijsbers		2004-08-14
Matthias Klose	doko42	2004-08-04
PJ Eby	pjeby	2004-03-24
Vinay Sajip	vsajip	2004-02-20
Hye-Shik Chang		2003-12-10
Armin Rigo		2003-10-24
Andrew McNamara		2003-06-09
Samuele Pedroni		2003-05-16
Alex Martelli	alexait	2003-04-22
Brett Cannon	brettcannon	2003-04-18
David Goodger		2003-01-02
Gustavo Niemeyer		2002-11-05
Tony Lownds		2002-09-22
Steve Holden		2002-06-14
Christian Tismer	ctismer	2002-05-17
Jason Tishler		2002-05-15
Walter Dörwald	doerwalter	2002-03-21
Andrew MacIntyre		2002-02-17
Gregory P. Smith	gpshead	2002-01-08
Anthony Baxter		2001-12-21
Neal Norwitz		2001-12-19
Raymond Hettinger	rhettinger	2001-12-10
Chui Tey		2001-10-31
Michael W. Hudson		2001-08-27
Finn Bock		2001-08-23
Piers Lauder		2001-07-20
Kurt B. Kaiser	kbkaiser	2001-07-03
Steven M. Gava		2001-06-25
Steve Purcell		2001-03-22
Jim Fulton		2000-10-06
Ka-Ping Yee		2000-10-03
Lars Gustäbel	gustaebel	2000-09-21
Neil Schemenauer	nascheme	2000-09-15
Martin v. Löwis		2000-09-08
Thomas Heller	theller	2000-09-07
Moshe Zadka		2000-07-29
Thomas Wouters	Yhg1s	2000-07-14
Peter Schneider-Kamp		2000-07-10
Paul Prescod		2000-07-01
Tim Peters	tim-one	2000-06-30
Skip Montanaro	smontanaro	2000-06-30
Fredrik Lundh		2000-06-29
Mark Hammond	mhammond	2000-06-09
Marc-André Lemburg	malemberg	2000-06-07
Trent Mick		2000-06-06
Eric S. Raymond		2000-06-02
Greg Stein		1999-11-07
Just van Rossum		1999-01-22
Greg Ward		1998-12-18

Continued on next page

Table 1 – continued from previous page

Name	GitHub username	Joined
Andrew Kuchling	akuchling	1998-04-09
Ken Manheimer		1998-03-03
Jeremy Hylton	jeremyhylton	1997-08-13
Roger E. Masse		1996-12-09
Fred Drake	freddrake	1996-07-23
Barry Warsaw	warsaw	1994-07-25
Jack Jansen	jackjansen	1992-08-13
Sjoerd Mullender	sjoerdmullender	1992-08-04
Guido van Rossum	gvanrossum	1989-12-25

10.15.1 Procedure for Granting or Dropping Access

To be granted the ability to manage who is a committer, you must be a team maintainer of the [Python core team](#) on GitHub. Once you have that privilege you can add people to the team. They will be asked to accept the membership which they can do by visiting <https://github.com/python> and clicking on the appropriate button that will be displayed to them in the upper part of the page.

10.16 Accepting Pull Requests

This page is aimed to core developers, and covers the steps required to accept, merge, and possibly backport a pull request on the main repository.

10.16.1 Is the PR ready to be accepted?

Before a PR is accepted, you must make sure it is ready to enter the public source tree. Use the following as a checklist of what to check for before merging (details of various steps can be found later in this document):

1. Has the submitter signed the CLA? (delineated by a label on the pull request)
2. Did the test suite pass? (delineated by a pull request check)
3. Did code coverage increase or stay the same? (delineated by a comment on the pull request)
4. Are the changes acceptable?
5. Was `configure` regenerated (if necessary)?
6. Was `pyconfig.h.in` regenerated (if necessary)?
7. Was the submitter added to `Misc/ACKS` (as appropriate)?
8. Was an entry added under `Misc/NEWS.d/next` (as appropriate)?
9. Was “What’s New” updated (as appropriate)?
10. Were appropriate labels added to signify necessary backporting of the pull request?
11. Was the pull request first made against the `master` branch?

Note: If you want to share your work-in-progress code on a feature or bugfix, either open a WIP-prefixed PR, publish patches on the [issue tracker](#) or create a public fork of the repository.

Does the test suite still pass?

You must *run the whole test suite* to ensure that it passes before merging any code changes.

Note: You really need to run the **entire** test suite. Running a single test is not enough as the changes may have unforeseen effects on other tests or library modules.

Running the entire test suite doesn't guarantee that the changes will pass the *continuous integration* tests, as those will exercise more possibilities still (such as different platforms or build options). But it will at least catch non-build specific, non-platform specific errors, therefore minimizing the chance for breakage.

Patch checklist

You should also *run patchcheck* to perform a quick sanity check on the changes.

10.16.2 Handling Others' Code

As a core developer you will occasionally want to commit a patch created by someone else. When doing so you will want to make sure of some things.

First, make sure the patch is in a good state. Both *Lifecycle of a Pull Request* and *Helping Triage Issues* explain what is to be expected of a patch. Typically patches that get cleared by triagers are good to go except maybe lacking `Misc/ACKS` and `Misc/NEWS.d` entries (which a core developer should make sure are updated appropriately).

Second, make sure the patch does not break backwards-compatibility without a good reason. This means *running the entire test suite* to make sure everything still passes. It also means that if semantics do change there must be a good reason for the breakage of code the change will cause (and it **will** break someone's code). If you are unsure if the breakage is worth it, ask on python-dev.

Third, ensure the patch is attributed correctly with the contributor's name in `Misc/ACKS` if they aren't already there (and didn't add themselves in their patch) and by mentioning "Patch by <x>" in the `Misc/NEWS.d` entry and the check-in message. If the patch has been heavily modified then "Initial patch by <x>" is an appropriate alternate wording. GitHub now supports [multiple authors](#) in a commit. Add `Co-authored-by: name <name@example.com>` at the end of the commit message.

If you omit correct attribution in the initial check-in, then update `ACKS` and `NEWS.d` in a subsequent check-in (don't worry about trying to fix the original check-in message in that case).

Finally, make sure that the submitter of the patch has a CLA in place (indicated by an asterisk following their username in the [issue tracker](#) or by the "CLA Signed" label on the pull request). If the submitter lacks a signed CLA and the patch is non-trivial, direct them to the electronic [Contributor Licensing Agreement](#) to ensure the PSF has the appropriate authorizations in place to relicense and redistribute their code.

10.16.3 Contributor Licensing Agreements

Always get a [Contributor Licensing Agreement](#) (CLA) signed unless the change has no possible intellectual property associated with it (e.g. fixing a spelling mistake in documentation). Otherwise it is simply safer from a legal standpoint to require the contributor to sign the CLA.

These days, the CLA can be signed electronically through the form linked above, and this process is strongly preferred to the old mechanism that involved sending a scanned copy of the signed paper form.

As discussed on the PSF [Contribution](#) page, it is the CLA itself that gives the PSF the necessary relicensing rights to redistribute contributions under the Python license stack. This is an additional permission granted above and beyond the normal permissions provided by the chosen open source license.

Some developers may object to the relicensing permissions granted to the PSF by the CLA. They're entirely within their rights to refuse to sign the CLA on that basis, but that refusal *does* mean we **can't accept their patches** for inclusion.

10.16.4 Checking if the CLA has been received

To check if a contributor's CLA has been received, go to the following website:

```
https://check-python-cla.herokuapp.com/
```

and put in their GitHub username.

For further questions about the CLA process, write to: contributors@python.org.

10.16.5 What's New and News Entries

Almost all changes made to the code base deserve an entry in `Misc/NEWS.d`. If the change is particularly interesting for end users (e.g. new features, significant improvements, or backwards-incompatible changes), an entry in the `What's New in Python` document (in `Doc/whatsnew/`) should be added as well. Changes that affect documentation only generally do not require a news entry.

There are two notable exceptions to this general principle, and they both relate to changes that *already* have a news entry, and have not yet been included in any formal release (including alpha and beta releases). These exceptions are:

- If a change is reverted prior to release, then the corresponding entry is simply removed. Otherwise, a new entry must be added noting that the change has been reverted (e.g. when a feature is released in an alpha and then cut prior to the first beta).
- If a change is a fix (or other adjustment) to an earlier unreleased change and the original news entry remains valid, then no additional entry is needed.

Needing a What's New entry almost always means that a change is *not* suitable for inclusion in a maintenance release. A small number of exceptions have been made for Python 2.7 due to the long support period - when implemented, these changes *must* be noted in the "New Additions in Python 2.7 Maintenance Releases" section of the Python 2.7 What's New document.

News entries go into the `Misc/NEWS.d` directory as individual files. The news entry can be created by using `blurb-it`, or the `blurb` tool and its `blurb add` command.

If you are unable to use the tool you can create the news entry file manually. The `Misc/NEWS.d` directory contains a sub-directory named `next` which itself contains various sub-directories representing classifications for what was affected (e.g. `Misc/NEWS.d/next/Library` for changes relating to the standard library). The file name itself should be of the format `<datetime>.bpo-<issue-number>.<nonce>.rst`:

- `<datetime>` is today's date joined with a `-` to the current time, in `YYYY-MM-DD-hh-mm-ss` format, e.g. `2017-05-27-16-46-23`
- `<issue-number>` is the issue number the change is for, e.g. `12345` for `bpo-12345`
- `<nonce>` is some "unique" string to guarantee the file name is unique across branches, e.g. `Y14gI2` (typically six characters, but it can be any length of letters and numbers, and its uniqueness can be satisfied by typing random characters on your keyboard)

So a file name may be `Misc/NEWS.d/next/Library/2017-05-27-16-46-23.bpo-12345.Y14gI2.rst`.

The contents of a news file should be valid reStructuredText. An 80 character column width should be used. There is no indentation or leading marker in the file (e.g. `-`). There is also no need to start the entry with the issue number as it's part of the file name itself. You can use *inline markups* too. Example news entry:

```
Fix warning message when :func:`os.chdir` fails inside
:func:`test.support.temp_cwd`. Patch by Chris Jerdonek.
```

The inline Sphinx roles like `:func:` can be used to assist readers in finding more information. You can build html and verify that the link target is appropriate by using *make html*.

While Sphinx roles can be beneficial to readers, they are not required. Inline ```code blocks``` can be used instead.

10.16.6 Working with Git

See also:

Git Bootcamp and Cheat Sheet

As a core developer, the ability to push changes to the official Python repositories means you have to be more careful with your workflow:

- You should not push new branches to the main repository. You can still use them in your fork that you use for development of patches; you can also push these branches to a **separate** public repository that will be dedicated to maintenance of the work before the work gets integrated in the main repository.
- You should not commit directly into the `master` branch, or any of the maintenance branches (currently 3.9 and 3.8). You should commit against your own feature branch, and create a pull request.
- For a small change, you can make a quick edit through the GitHub web UI. If you choose to use the web UI, be aware that GitHub will create a new branch in the **main** CPython repo (not your fork). Please delete this newly created branch after it has been merged into the `master` branch or any of the maintenance branches. To keep the CPython repo tidy, please try to limit the existence of the new branch to, at most, a few days.

It is recommended to keep a fork of the main repository around, as it allows simple reversion of all local changes (even “committed” ones) if your local clone gets into a state you aren’t happy with.

Active branches

If you do `git branch` you will see a *list of branches*. `master` is the in-development branch, and is the only branch that receives new features. The other branches only receive bug fixes or security fixes.

Backporting Changes to an Older Version

When it is determined that a pull request needs to be backported into one or more of the maintenance branches, a core developer can apply the labels `needs backport to X.Y` to the pull request.

After the pull request has been merged, `miss-islington` (bot) will first try to do the backport automatically. In case that `miss-islington` is unable to do it, then the pull request author or the core developer who merged it should look into backporting it themselves, using the backport generated by `cherry_picker.py` as a starting point.

The commit hash can be obtained from the original pull request, or by using `git log` on the `master` branch. To display the 10 most recent commit hashes and their first line of the commit message:

```
git log -10 --oneline
```

Prefix the backport pull request with the branch, and reference the pull request number from `master`, for example:

```
[3.9] bpo-12345: Fix the Spam Module (GH-NNNN)
```

Note that `cherry_picker.py` adds the branch prefix automatically.

Once the backport pull request has been created, remove the `needs backport to X.Y` label from the original pull request. (Only Core Developers and members of the [Python Triage Team](#) can apply labels to GitHub pull requests).

Reverting a Merged Pull Request

To revert a merged pull request, press the `Revert` button at the bottom of the pull request. It will bring up the page to create a new pull request where the commit can be reverted. It also creates a new branch on the main CPython repository. Delete the branch once the pull request has been merged.

Always include the reason for reverting the commit to help others understand why it was done. The reason should be included as part of the commit message, for example:

```
Revert bpo-NNNN: Fix Spam Module (GH-111)

Reverts python/cpython#111.
Reason: This commit broke the buildbot.
```

10.17 Development Cycle

The responsibilities of a core developer shift based on what kind of branch of Python a developer is working on and what stage the branch is in.

To clarify terminology, Python uses a `major.minor.micro` nomenclature for production-ready releases. So for Python 3.1.2 final, that is a *major version* of 3, a *minor version* of 1, and a *micro version* of 2.

- new *major versions* are exceptional; they only come when strongly incompatible changes are deemed necessary, and are planned very long in advance;
- new *minor versions* are feature releases; they get released annually, from the current *in-development* branch;
- new *micro versions* are bugfix releases; they get released roughly every 2 months; they are prepared in *maintenance* branches.

We also publish non-final versions which get an additional qualifier: *Alpha*, *Beta*, *release candidate*. These versions are aimed at testing by advanced users, not production use.

Each release of Python is tagged in the source repo with a tag of the form `vX.Y.ZTN`, where `X` is the major version, `Y` is the minor version, `Z` is the micro version, `T` is the release level (`a` for alpha releases, `b` for beta, `rc` release candidate, and `null` for final releases), and `N` is the release serial number. Some examples of release tags: `v3.7.0a1`, `v3.6.3`, `v2.7.14rc1`.

10.17.1 Branches

There is a branch for each *feature version*, whether released or not (e.g. 3.7, 3.8).

In-development (main) branch

The `master` branch is the branch for the next feature release; it is under active development for all kinds of changes: new features, semantic changes, performance improvements, bug fixes.

At some point during the life-cycle of a release, a new *maintenance branch* is created to host all bug fixing activity for further micro versions in a feature version (3.8.1, 3.8.2, etc.).

For versions 3.4 and before, this was conventionally done when the final release was cut (for example, 3.4.0 final).

Starting with the 3.5 release, we create the release maintenance branch (e.g. 3.5) at the time we enter beta (3.5.0 beta 1). This allows feature development for the release 3.n+1 to occur within the master branch alongside the beta and release candidate stabilization periods for release 3.n.

Maintenance branches

A branch for a previous feature release, currently being maintained for bug fixes. There are usually two maintenance branches at any given time for Python 3.x. Only during the beta/rc phase of a new minor/feature release will there be three active maintenance branches, e.g. during the beta phase for Python 3.8 there are master, 3.8, 3.7, and 3.6 branches open. Releases produced from a maintenance branch are called **maintenance** or **bugfix** releases; the terms are used interchangeably. These releases have a **micro version** number greater than zero.

The only changes allowed to occur in a maintenance branch without debate are bug fixes. Also, a general rule for maintenance branches is that compatibility must not be broken at any point between sibling minor releases (3.5.1, 3.5.2, etc.). For both rules, only rare exceptions are accepted and **must** be discussed first.

Sometime after a new maintenance branch is created (after a new *minor version* is released), the old maintenance branch on that major version will go into *security mode*, usually after one last maintenance release at the discretion of the release manager. For example, the 3.4 maintenance branch was put into *security mode* after the 3.4.4 final maintenance release following the release of 3.5.1.

Security branches

A branch less than 5 years old but no longer in maintenance mode is a security branch.

The only changes made to a security branch are those fixing issues exploitable by attackers such as crashes, privilege escalation and, optionally, other issues such as denial of service attacks. Any other changes are **not** considered a security risk and thus not backported to a security branch. You should also consider fixing hard-failing tests in open security branches since it is important to be able to run the tests successfully before releasing.

Commits to security branches are to be coordinated with the release manager for the corresponding feature version, as listed in the *Status of Python branches*. Any release made from a security branch is source-only and done only when actual security patches have been applied to the branch. These releases have a **micro version** number greater than the last **bugfix** release.

End-of-life branches

The code base for a release cycle which has reached end-of-life status is frozen and no longer has a branch in the repo. The final state of the end-of-lived branch is recorded as a tag with the same name as the former branch, e.g. 3.3 or 2.6.

For reference, here are the Python versions that most recently reached their end-of-life:

Branch	Schedule	First release	End-of-life	Release manager
3.4	PEP 429	2014-03-16	2019-03-18	Larry Hastings
3.3	PEP 398	2012-09-29	2017-09-29	Georg Brandl, Ned Deily (3.3.7+)
3.2	PEP 392	2011-02-20	2016-02-20	Georg Brandl
3.1	PEP 375	2009-06-27	2012-04-09	Benjamin Peterson
3.0	PEP 361	2008-12-03	2009-06-27	Barry Warsaw
2.6	PEP 361	2008-10-01	2013-10-29	Barry Warsaw
2.7	PEP 373	2010-07-03	2020-01-01	Benjamin Peterson

The latest release for each Python version can be found on the [download page](#).

10.17.2 Stages

Based on what stage the *in-development* version of Python is in, the responsibilities of a core developer change in regards to commits to the VCS.

Pre-alpha

The branch is in this stage when no official release has been done since the latest final release. There are no special restrictions placed on commits, although the usual advice applies (getting patches reviewed, avoiding breaking the buildbots).

Alpha

Alpha releases typically serve as a reminder to core developers that they need to start getting in changes that change semantics or add something to Python as such things should not be added during a *Beta*. Otherwise no new restrictions are in place while in alpha.

Beta

After a first beta release is published, no new features are accepted. Only bug fixes can now be committed. This is when core developers should concentrate on the task of fixing regressions and other new issues filed by users who have downloaded the alpha and beta releases.

Being in beta can be viewed much like being in *RC* but without the extra overhead of needing commit reviews.

Please see the note in the *In-development (main) branch* section above for new information about the creation of the 3.5 maintenance branch during beta.

Release Candidate (RC)

A branch preparing for an RC release can only have bugfixes applied that have been reviewed by other core developers. Generally, these issues must be severe enough (e.g. crashes) that they deserve fixing before the final release. All other issues should be deferred to the next development cycle, since stability is the strongest concern at this point.

You **cannot** skip the peer review during an RC, no matter how small! Even if it is a simple copy-and-paste change, **everything** requires peer review from a core developer.

Final

When a final release is being cut, only the release manager (RM) can make changes to the branch. After the final release is published, the full *development cycle* starts again for the next minor version.

10.17.3 Repository Administration

The source code is currently hosted on [GitHub](#) in the [Python organization](#).

Organization Repository Policy

Within the [Python organization](#), repositories are expected to fall within these general categories:

1. The reference implementation of Python and related repositories (i.e. [CPython](#))
2. Reference implementations of PEPs (e.g. [mypy](#))
3. Tooling and support around CPython and the language (e.g. [python.org repository](#))
4. PSF-related repositories (e.g. the [Code of Conduct](#))
5. PSF Infrastructure repositories (e.g. the [PSF Infrastructure Salt configurations](#))

For any repository which does not explicitly and clearly fall under one of these categories, permission should be sought from the [Python steering council](#).

Organization Owner Policy

The GitHub Organization Owner role allows for full management of all aspects of the Python organization. Allowing for visibility and management of all aspects at all levels including organization membership, team membership, access control, and merge privileges on all repositories. For full details of the permission levels see [GitHub's documentation on Organization permission levels](#). This role is paramount to the security of the Python Language, Community, and Infrastructure.

The Executive Director of the Python Software Foundation delegates authority on GitHub Organization Owner Status to Ernest W. Durbin III - Python Software Foundation Director of Infrastructure. Common reasons for this role are: Infrastructure Staff Membership, Python Software Foundation General Counsel, and Python Software Foundation Staff as fallback.

Inactive or unreachable members may be removed with or without notice. Members who no longer necessitate this level of access will be removed with notice.

Multi-Factor Authentication must be enabled by the user in order to remain an Owner of the Python Organization.

Current Owners

Name	Role	GitHub Username
Benjamin Peterson	Infrastructure Staff	benjaminp
Noah Kantrowitz	Infrastructure Staff	coderranger
Donald Stufft	Infrastructure Staff	dstufft
Ewa Jodlowska	PSF Executive Director	ejodlowska
Ernest W. Durbin III	PSF Director of Infrastructure	ewdurbin
Van Lindberg	PSF General Counsel	VanL

Repository Administrator Role Policy

The Administrator role on the repository allows for managing all aspects including collaborators, access control, integrations, webhooks, and branch protection. For full details of the permission levels see [GitHub's documentation on Repository permission levels](#). Common reasons for this role are: maintenance of Core Developer Workflow tooling, Release Managers for all *in-development*, *maintenance*, and *security mode* releases, and additional Python Core Developers as necessary for redundancy. Occasional temporary administrator access is acceptable as necessary for Core Developer workflow projects.

Inactive or unreachable members may be removed with or without notice. Members who no longer necessitate this level of access will be removed with notice.

Multi-Factor Authentication must be enabled by the user in order to remain an Administrator of the repository.

Current Administrators

Name	Role	GitHub Username
Łukasz Langa	Python 3.8 and 3.9 Release Manager	ambv
Ned Deily	Python 3.7 Release Manager	ned-deily
Lary Hastings	Python 3.5 Release Manager	larryhastings
Berker Peksag	Maintainer of bpo-linkify and cpython-emailer-webhook	berkerpeksag
Brett Cannon	Maintainer of bedevere and the-knights-who-say-ni	brettcannon
Ezio Melotti	Maintainer of bugs.python.org GitHub webhook integration	ezio-melotti
Mariatta Wijaya	Maintainer of blurb_it and miss-islington	Mariatta
Pablo Galindo	Maintainer of buildbot.python.org	pablogsal

Repository Release Manager Role Policy

Release Managers for *in-development*, *maintenance*, and *security mode* Python releases are granted Administrator privileges on the repository. Once a release branch has entered *end-of-life*, the Release Manager for that branch is removed as an Administrator and granted sole privileges (out side of repository administrators) to merge changes to that branch.

Multi-Factor Authentication must be enabled by the user in order to retain access as a Release Manager of the branch.

10.18 Continuous Integration

To assert that there are no regressions in the *development and maintenance branches*, Python has a set of dedicated machines (called *buildbots* or *build workers*) used for continuous integration. They span a number of hardware/operating system combinations. Furthermore, each machine hosts several *builders*, one per active branch: when a new change is pushed to this branch on the public [GitHub repository](#), all corresponding builders will schedule a new build to be run as soon as possible.

The build steps run by the buildbots are the following:

- Checkout of the source tree for the changeset which triggered the build
- Compiling Python
- Running the test suite using *strenuous settings*
- Cleaning up the build tree

It is your responsibility, as a core developer, to check the automatic build results after you push a change to the repository. It is therefore important that you get acquainted with the way these results are presented, and how various kinds of failures can be explained and diagnosed.

10.18.1 Checking results of automatic builds

There are three ways of visualizing recent build results:

- The Web interface for each branch at <https://www.python.org/dev/buildbot/>, where the so-called “waterfall” view presents a vertical rundown of recent builds for each builder. When interested in one build, you’ll have to click on it to know which changesets it corresponds to. Note that the buildbot web pages are often slow to load, be patient.
- The command-line `bbreport.py` client, which you can get from <https://code.google.com/archive/p/bbreport>. Installing it is trivial: just add the directory containing `bbreport.py` to your system path so that you can run it from any filesystem location. For example, if you want to display the latest build results on the development (“master”) branch, type:

```
bbreport.py -q 3.x
```

- The buildbot “console” interface at <http://buildbot.python.org/all/#/console> This works best on a wide, high resolution monitor. Clicking on the colored circles will allow you to open a new page containing whatever information about that particular build is of interest to you. You can also access builder information by clicking on the builder status bubbles in the top line.

If you like IRC, having an IRC client open to the `#python-dev` channel on `irc.freenode.net` is useful. Any time a builder changes state (last build passed and this one didn’t, or vice versa), a message is posted to the channel. Keeping an eye on the channel after pushing a changeset is a simple way to get notified that there is something you should look in to.

Some buildbots are much faster than others. Over time, you will learn which ones produce the quickest results after a build, and which ones take the longest time.

Also, when several changesets are pushed in a quick succession in the same branch, it often happens that a single build is scheduled for all these changesets.

10.18.2 Stability

A subset of the buildbots are marked “stable”. They are taken into account when making a new release. The rule is that all stable builders must be free of persistent failures when the release is cut. It is absolutely **vital** that core developers fix any issue they introduce on the stable buildbots, as soon as possible.

This does not mean that other builders’ test results can be taken lightly, either. Some of them are known for having platform-specific issues that prevent some tests from succeeding (or even terminating at all), but introducing additional failures should generally not be an option.

10.18.3 Flags-dependent failures

Sometimes, while you have run the *whole test suite* before committing, you may witness unexpected failures on the buildbots. One source of such discrepancies is if different flags have been passed to the test runner or to Python itself. To reproduce, make sure you use the same flags as the buildbots: they can be found out simply by clicking the **stdio** link for the failing build’s tests. For example:

```
./python.exe -Wd -E -bb ./Lib/test/regrtest.py -uall -rW
```

Note: Running `Lib/test/regrtest.py` is exactly equivalent to running `-m test`.

10.18.4 Ordering-dependent failures

Sometimes the failure is even subtler, as it relies on the order in which the tests are run. The buildbots *randomize* test order (by using the `-r` option to the test runner) to maximize the probability that potential interferences between library modules are exercised; the downside is that it can make for seemingly sporadic failures.

The `--randseed` option makes it easy to reproduce the exact randomization used in a given build. Again, open the `stdio` link for the failing test run, and check the beginning of the test output proper.

Let's assume, for the sake of example, that the output starts with:

```
./python -Wd -E -bb Lib/test/regrtest.py -uall -rwW
== CPython 3.3a0 (default:22ae2b002865, Mar 30 2011, 13:58:40) [GCC 4.4.5]
== Linux-2.6.36-gentoo-r5-x86_64-AMD_Athlon-tm-_64_X2_Dual_Core_Processor_4400+-
↳with-gentoo-1.12.14 little-endian
== /home/buildbot/buildarea/3.x.ochtman-gentoo-amd64/build/build/test_python_29628
Testing with flags: sys.flags(debug=0, inspect=0, interactive=0, optimize=0, dont_
↳write_bytecode=0, no_user_site=0, no_site=0, ignore_environment=1, verbose=0, bytes_
↳warning=2, quiet=0)
Using random seed 2613169
[ 1/353] test_augassign
[ 2/353] test_functools
```

You can reproduce the exact same order using:

```
./python -Wd -E -bb -m test -uall -rwW --randseed 2613169
```

It will run the following sequence (trimmed for brevity):

```
[ 1/353] test_augassign
[ 2/353] test_functools
[ 3/353] test_bool
[ 4/353] test_contains
[ 5/353] test_compileall
[ 6/353] test_unicode
```

If this is enough to reproduce the failure on your setup, you can then bisect the test sequence to look for the specific interference causing the failure. Copy and paste the test sequence in a text file, then use the `--fromfile` (or `-f`) option of the test runner to run the exact sequence recorded in that text file:

```
./python -Wd -E -bb -m test -uall -rwW --fromfile mytestsequence.txt
```

In the example sequence above, if `test_unicode` had failed, you would first test the following sequence:

```
[ 1/353] test_augassign
[ 2/353] test_functools
[ 3/353] test_bool
[ 6/353] test_unicode
```

And, if it succeeds, the following one instead (which, hopefully, shall fail):

```
[ 4/353] test_contains
[ 5/353] test_compileall
[ 6/353] test_unicode
```

Then, recursively, narrow down the search until you get a single pair of tests which triggers the failure. It is very rare that such an interference involves more than **two** tests. If this is the case, we can only wish you good luck!

Note: You cannot use the `-j` option (for parallel testing) when diagnosing ordering-dependent failures. Using `-j` isolates each test in a pristine subprocess and, therefore, prevents you from reproducing any interference between tests.

10.18.5 Transient failures

While we try to make the test suite as reliable as possible, some tests do not reach a perfect level of reproducibility. Some of them will sometimes display spurious failures, depending on various conditions. Here are common offenders:

- Network-related tests, such as `test_poplib`, `test_urllibnet`, etc. Their failures can stem from adverse network conditions, or imperfect thread synchronization in the test code, which often has to run a server in a separate thread.
- Tests dealing with delicate issues such as inter-thread or inter-process synchronization, or Unix signals: `test_multiprocessing`, `test_threading`, `test_subprocess`, `test_threadsignals`.

When you think a failure might be transient, it is recommended you confirm by waiting for the next build. Still, even if the failure does turn out sporadic and unpredictable, the issue should be reported on the bug tracker; even better if it can be diagnosed and suppressed by fixing the test's implementation, or by making its parameters - such as a timeout - more robust.

10.18.6 Custom builders

When working on a platform-specific issue, you may want to test your changes on the buildbot fleet rather than just on Travis and AppVeyor. To do so, you can make use of the [custom builders](#). These builders track the `buildbot-custom` short-lived branch of the `python/cpython` repository, which is only accessible to core developers.

To start a build on the custom builders, push the commit you want to test to the `buildbot-custom` branch:

```
$ git push upstream <local_branch_name>:buildbot-custom
```

You may run into conflicts if another developer is currently using the custom builders or forgot to delete the branch when they finished. In that case, make sure the other developer is finished and either delete the branch or force-push (add the `-f` option) over it.

When you have gotten the results of your tests, delete the branch:

```
$ git push upstream :buildbot-custom      # or use the GitHub UI
```

If you are interested in the results of a specific test file only, we recommend you change (temporarily, of course) the contents of the `buildbottest` clause in `Makefile.pre.in`; or, for Windows builders, the `Tools/buildbot/test.bat` script.

See also:

[Running a buildbot worker](#)

10.19 Adding to the Stdlib

While the stdlib contains a great amount of useful code, sometimes you want more than is provided. This document is meant to explain how you can get either a new addition to a pre-existing module in the stdlib or add an entirely new module.

Changes to pre-existing code is not covered as that is considered a bugfix and thus is treated as a bug that should be filed on the [issue tracker](#).

10.19.1 Adding to a pre-existing module

If you have found that a function, method, or class is useful and you believe it would be useful to the general Python community, there are some steps to go through in order to see it added to the stdlib.

First is you should gauge the usefulness of the code. Typically this is done by sharing the code publicly. You have a couple of options for this. One is to post it online at the [Python Cookbook](#). Based on feedback or reviews of the recipe you can see if others find the functionality as useful as you do. A search of the issue tracker for previous suggestions related to the proposed addition may turn up a rejected issue that explains why the suggestion will not be accepted. Another is to do a blog post about the code and see what kind of responses you receive. Posting to python-list (see [Following Python's Development](#) for where to find the list and other mailing lists) to discuss your code also works. Finally, asking on a specific SIG (special interest group) from mail.python.org or python-ideas is also acceptable. This is not a required step but it is suggested.

If you have found general acceptance and usefulness for your code from people, you can open an issue on the [issue tracker](#) with the code attached as a *patch*. If possible, also submit a *contributor agreement*.

If a core developer decides that your code would be useful to the general Python community, they will then commit your code. If your code is not picked up by a core developer and committed then please do not take this personally. Through your public sharing of your code in order to gauge community support for it you at least can know that others will come across it who may find it useful.

10.19.2 Adding a new module

It must be stated upfront that getting a new module into the stdlib is very difficult. Adding any significant amount of code to the stdlib increases the burden placed upon core developers. It also means that the module somewhat becomes “sanctioned” by the core developers as a good way to do something, typically leading to the rest of the Python community to using the new module over other available solutions. All of this means that additions to the stdlib are not taken lightly.

Acceptable Types of Modules

Typically two types of modules get added to the stdlib. One type is a module which implements something that is difficult to get right. A good example of this is the `multiprocessing` package. Working out the various OS issues, working through concurrency issues, etc. are all very difficult to get right.

The second type of module is one that implements something that people re-implement constantly. The `itertools` module is a good example of this type as its constituent parts are not necessarily complex, but are used regularly in a wide range of programs and can be a little tricky to get right. Modules that parse widely used data formats also fall under this type of module that the stdlib consists of.

While a new stdlib module does not need to appeal to all users of Python, it should be something that a large portion of the community will find useful. This makes sure that the developer burden placed upon core developers is worth it.

Requirements

In order for a module to even be considered for inclusion into the stdlib, a couple of requirements must be met.

The most basic is that the code must meet *standard patch requirements*. For code that has been developed outside the stdlib typically this means making sure the coding style guides are followed and that the proper tests have been written.

The module needs to have been out in the community for at least a year. Because of Python's conservative nature when it comes to backwards-compatibility, when a module is added to the stdlib its API becomes frozen. This means that a module should only enter the stdlib when it is mature and gone through its "growing pains".

The module needs to be considered best-of-breed. When something is included in the stdlib it tends to be chosen first for products over other third-party solutions. By virtue of having been available to the public for at least a year, a module needs to have established itself as (one of) the top choices by the community for solving the problem the module is intended for.

The development of the module must move into Python's infrastructure (i.e., the module is no longer directly maintained outside of Python). This prevents a divergence between the code that is included in the stdlib and that which is released outside the stdlib (typically done to provide the module to older versions of Python). It also removes the burden of forcing core developers to have to redirect bug reports or patches to an external issue tracker and VCS.

Someone involved with the development of the module must promise to help maintain the module in the stdlib for two years. This not only helps out other core developers by alleviating workload from bug reports that arrive from the first Python release containing the module, but also helps to make sure that the overall design of the module continues to be uniform.

Proposal Process

If the module you want to propose adding to the stdlib meets the proper requirements, you may propose its inclusion. To start, you should email `python-list` or `python-ideas` to make sure the community in general would support the inclusion of the module (see *Following Python's Development*).

If the feedback from the community is positive overall, you will need to write a PEP for the module's inclusion. It should outline what the module's overall goal is, why it should be included in the stdlib, and specify the API of the module. See the [PEP index](#) for PEPs that have been accepted before that proposed a module for inclusion.

Once your PEP is written, send it to `python-ideas` for basic vetting. Be prepared for extensive feedback and lots of discussion (not all of it positive). This will help make the PEP be of good quality and properly formatted.

When you have listened to, responded, and integrated as appropriate the feedback from `python-ideas` into your PEP, you may send it to `python-dev`. You will once again receive a large amount of feedback and discussion. A PEP dictator will be assigned who makes the final call on whether the PEP will be accepted or not. If the PEP dictator agrees to accept your PEP (which typically means that the core developers end up agreeing in general to accepting your PEP) then the module will be added to the stdlib once the creators of the module sign *contributor agreements*.

10.20 Changing the Python Language

On occasion people come up with an idea on how to change or improve Python as a programming language. This document is meant to explain exactly what changes have a reasonable chance of being considered and what the process is to propose changes to the language.

10.20.1 What Qualifies

First and foremost, it must be understood that changes to the Python programming language are difficult to make. When the language changes, **every** Python programmer already in existence and all Python programmers to come will end up eventually learning about the change you want to propose. Books will need updating, code will be changed, and a new way to do things will need to be learned. Changes to the Python programming language are never taken lightly.

Because of the seriousness that language changes carry, any change must be beneficial to a large proportion of Python users. If the change only benefits a small percentage of Python developers then the change will not be made. A good way to see if your idea would work for a large portion of the Python community is to ask on *python-list* or *python-ideas*. You can also go through Python's stdlib and find examples of code which would benefit from your proposed change (which helps communicate the usefulness of your change to others). For further guidance, see *Suggesting new features and language changes*.

Your proposed change also needs to be *Pythonic*. While Guido is the only person who can truly classify something as Pythonic, you can read the *Zen of Python* for guidance.

10.20.2 PEP Process

Once you are certain you have a language change proposal which will appeal to the general Python community, you can begin the process of officially proposing the change. This process is the Python Enhancement Proposal (PEP) process. **PEP 1** describes it in detail.

You will first need a PEP that you will present to python-ideas. You may be a little hazy on the technical details as various core developers can help with that, but do realize that if you do not present your idea to python-ideas or python-list ahead of time you may find out it is technically not possible (e.g., Python's parser will not support the grammar change as it is an LL(1) parser). Expect extensive comments on the PEP, some of which will be negative.

Once your PEP has been modified to be of proper quality and to take into account comments made on python-ideas, it may proceed to python-dev. There it will be assigned a PEP dictator and another general discussion will occur. Once again, you will need to modify your PEP to incorporate the large amount of comments you will receive.

The PEP dictator decides if your PEP is accepted (typically based on whether most core developers support the PEP). If that occurs then your proposed language change will be introduced in the next release of Python. Otherwise your PEP will be recorded as rejected along with an explanation as to why so that others do not propose the same language change in the future.

10.20.3 Suggesting new features and language changes

The *python-ideas* mailing list is specifically intended for discussion of new features and language changes. Please don't be disappointed if your idea isn't met with universal approval: as the long list of Rejected and Withdrawn PEPs in the *PEP Index* attests, and as befits a reasonably mature programming language, getting significant changes into Python isn't a simple task.

If the idea is reasonable, someone will suggest posting it as a feature request on the *issue tracker*, or, for larger changes, writing it up as a *draft PEP*.

Sometimes core developers will differ in opinion, or merely be collectively unconvinced. When there isn't an obvious victor then the *Status Quo Wins a Stalemate* as outlined in the linked post.

For some examples on language changes that were accepted please read *Justifying Python Language Changes*.

10.21 Experts Index

This document has tables that list Python Modules, Tools, Platforms and Interest Areas and names for each item that indicate a maintainer or an expert in the field. This list is intended to be used by issue submitters, issue triage people, and other issue participants to find people to add to the nosy list or to contact directly by email for help and decisions on feature requests and bug fixes. People on this list may be asked to render final judgement on a feature or bug. If no active maintainer is listed for a given module, then questionable changes should go to python-dev, while any other issues can and should be decided by any committer.

Unless a name is followed by a '*', you should never assign an issue to that person, only make them nosy. Names followed by a '*' may be assigned issues involving the module or topic.

The Platform and Interest Area tables list broader fields in which various people have expertise. These people can also be contacted for help, opinions, and decisions when issues involve their areas.

If a listed maintainer does not respond to requests for comment for an extended period (three weeks or more), they should be marked as inactive in this list by placing the word 'inactive' in parenthesis behind their tracker id. They are of course free to remove that inactive mark at any time.

Committers should update these tables as their areas of expertise widen. New topics may be added to the Interest Area table at will.

The existence of this list is not meant to indicate that these people *must* be contacted for decisions; it is, rather, a resource to be used by non-committers to find responsible parties, and by committers who do not feel qualified to make a decision in a particular context.

See also [PEP 291](#) and [PEP 360](#) for information about certain modules with special rules.

10.21.1 Stdlib

Module	Maintainers
<code>__future__</code>	
<code>__main__</code>	gvanrossum, ncoghlan
<code>_dummy_thread</code>	brett.cannon
<code>_thread</code>	
<code>_testbuffer</code>	skrah
<code>abc</code>	
<code>aifc</code>	r.david.murray
<code>argparse</code>	rhettinger*
<code>array</code>	
<code>ast</code>	benjamin.peterson, pablogsal
<code>asynchat</code>	josiahcarlson, giampaolo.rodola*, stutzbach
<code>asyncio</code>	yseivanov, asvetlov
<code>asyncore</code>	josiahcarlson, giampaolo.rodola*, stutzbach
<code>atexit</code>	
<code>audioop</code>	serhiy.storchaka
<code>base64</code>	
<code>bdb</code>	
<code>binascii</code>	
<code>binhex</code>	
<code>bisect</code>	rhettinger*
<code>builtins</code>	
<code>bz2</code>	

Continued on next page

Table 2 – continued from previous page

Module	Maintainers
calendar	rhettinger*
cgi	ethan.furman*
cglib	ethan.furman*
chunk	
cmath	mark.dickinson
cmd	
code	
codecs	lemburg, doerwalter
codeop	
collections	rhettinger*
collections.abc	rhettinger*, stutzbach
colorsys	
compileall	
concurrent.futures	pitrou, bquinlan
configparser	lukasz.langa*
contextlib	ncoghlan, yselivanov
contextvars	
copy	alexandre.vassalotti
copyreg	alexandre.vassalotti
cProfile	
crypt	jafo*
csv	skip.montanaro (inactive)
ctypes	theller (inactive), belopolsky, amaury.forgeotdarc, meador.inge
curses	twouters
dataclasses	eric.smith
datetime	belopolsky, p-ganssle
dbm	
decimal	facundobatista, rhettinger, mark.dickinson, skrah
difflib	tim.peters (inactive)
dis	yselivanov
distutils	eric.araujo, dstufft
doctest	tim.peters (inactive)
dummy_threading	brett.cannon
email	barry, r.david.murray*, maxking
encodings	lemburg
ensurepip	ncoghlan, dstufft, pradyunsg
enum	eli.bendersky*, barry, ethan.furman*
errno	twouters
faulthandler	vstinner
fcntl	twouters
filecmp	
fileinput	
fnmatch	
formatter	
fractions	mark.dickinson, rhettinger
ftplib	giampaolo.rodola*
functools	rhettinger*
gc	pitrou, pablogsal
getopt	

Continued on next page

Table 2 – continued from previous page

Module	Maintainers
getpass	
gettext	
glob	
grp	
gzip	
hashlib	christian.heimes, gregory.p.smith
heapq	rhettinger*, stutzbach
hmac	christian.heimes, gregory.p.smith
html	ezio.melotti
http	
idlelib	kbk (inactive), terry.reedy*, roger.serwy (inactive), taleinat
imaplib	
imghdr	
imp	
importlib	brett.cannon
inspect	yseivanov
io	benjamin.peterson, stutzbach
ipaddress	pmoody
itertools	rhettinger*
json	bob.ippolito (inactive), ezio.melotti, rhettinger
keyword	
lib2to3	benjamin.peterson
libmpdec	skrah
linecache	
locale	lemburg
logging	vinay.sajip
lzma	
mailbox	
mailcap	
marshal	
math	mark.dickinson, rhettinger, stutzbach
mimetypes	
mmap	twouters
modulefinder	theller (inactive), jvr
msilib	
msvcr	
multiprocessing	davin*, pitrou, jnoller (inactive), sbt (inactive)
netrc	
nis	
nntplib	
numbers	
operator	
optparse	aronacher
os	
os.path	serhiy.storchaka
ossaudiodev	
parser	benjamin.peterson, pablogsal
pathlib	
pdb	

Continued on next page

Table 2 – continued from previous page

Module	Maintainers
pickle	alexandre.vassalotti
pickletools	alexandre.vassalotti
pipes	
pkgutil	
platform	lemburg
plistlib	
poplib	
posix	larry
pprint	fdrake
profile	
pstats	
pty	twouters*
pwd	
py_compile	
pyclbr	
pydoc	
queue	rhettinger*
quopri	
random	rhettinger, mark.dickinson
re	effbot (inactive), ezio.melotti, serhiy.storchaka
readline	twouters
reprlib	
resource	twouters
rlcompleter	
runpy	ncoghlan
sched	
secrets	
select	
selectors	neologix, giampaolo.rodola
shelve	
shlex	
shutil	tarek, giampaolo.rodola
signal	
site	
smtpd	giampaolo.rodola
smtplib	
sndhdr	
socket	
socketserver	
spwd	
sqlite3	ghaering
ssl	janssen, christian.heimes, dstufft, alex
stat	christian.heimes
statistics	steven.daprano, rhettinger, taleinat
string	
stringprep	
struct	mark.dickinson, meador.inge
subprocess	astrand (inactive), giampaolo.rodola
sunau	

Continued on next page

Table 2 – continued from previous page

Module	Maintainers
symbol	
syntable	benjamin.peterson
sys	
sysconfig	tarek
syslog	jafo*
tabnanny	tim.peters (inactive)
tarfile	lars.gustaebel
telnetlib	
tempfile	
termios	twouters
test	ezio.melotti
textwrap	
threading	pitrou
time	belopolsky, p-ganssle
timeit	
tkinter	gpolo, serhiy.storchaka
token	
tokenize	meador.inge
trace	belopolsky
traceback	
tracemalloc	vstinner
tty	twouters*
turtle	gregorlingl, willingc
types	yselivanov
typing	gvanrossum, levkivskyi*
unicodedata	lemburg, ezio.melotti
unittest	michael.foord*, ezio.melotti, rbcollins
unittest.mock	michael.foord*
urllib	orsenthil
uu	
uuid	
venv	vinay.sajip
warnings	
wave	
weakref	fdrake
webbrowser	
winreg	stutzbach
winsound	effbot (inactive)
wsgiref	pje
xdrlib	
xml.dom	
xml.dom.minidom	
xml.dom.pulldom	
xml.etree	effbot (inactive), eli.bendersky*, scoder
xml.parsers.expat	
xml.sax	
xml.sax.handler	
xml.sax.saxutils	
xml.sax.xmlreader	

Continued on next page

Table 2 – continued from previous page

Module	Maintainers
xmlrpc	
zipapp	paul.moore
zipfile	alanmcintyre, serhiy.storchaka, twouters
zipimport	twouters*
zlib	twouters

10.21.2 Tools

Tool	Maintainers
Argument Clinic	larry

10.21.3 Platforms

Platform	Maintainers
AIX	David.Edelsohn
Cygwin	jlt63, stutzbach
FreeBSD	
HP-UX	
Linux	
Mac OS X	ronaldoussoren, ned.deily
NetBSD1	
OS2/EMX	aimacintyre
Solaris/OpenIndiana	jcea
Windows	tim.golden, zach.ware, steve.dower, paul.moore
JVM/Java	frank.wierzbicki

10.21.4 Miscellaneous

Interest Area	Maintainers
algorithms	rhettinger*
argument clinic	larry
ast/compiler	benjamin.peterson, brett.cannon, yselivanov, pablogsal
autoconf/makefiles	twouters*
bsd	
bug tracker	ezio.melotti
buildbots	zach.ware, pablogsal
bytecode	benjamin.peterson, yselivanov
context managers	ncoghlan
core workflow	mariatta
coverity scan	christian.heimes, brett.cannon, twouters
cryptography	gregory.p.smith, dstufft
data formats	mark.dickinson
database	lemburg
devguide	eric.araujo, ezio.melotti, willingc, mariatta
documentation	ezio.melotti, eric.araujo, mdk, willingc

Continued on next page

Table 3 – continued from previous page

Interest Area	Maintainers
emoji	mariatta
extension modules	petr.viktorin, ncoghlan
filesystem	giampaolo.rodola
f-strings	eric.smith*
GUI	
i18n	lemburg, eric.araujo
import machinery	brett.cannon, ncoghlan, eric.snow
io	benjamin.peterson, stutzbach
locale	lemburg
mathematics	mark.dickinson, lemburg, stutzbach, rhettinger
memory management	tim.peters, lemburg, twouters
memoryview	skrah
networking	giampaolo.rodola,
object model	benjamin.peterson, twouters
packaging	tarek, lemburg, alexis, eric.araujo, dstufft, paul.moore
performance	brett.cannon, vstinner, serhiy.storchaka, yselivanov, rhettinger
pip	ncoghlan, dstufft, paul.moore, Marcus.Smith, pradyunsg
py3 transition	benjamin.peterson
release management	tarek, lemburg, benjamin.peterson, barry, gvanrossum, anthonybaxter, eric.araujo, ned.deily, georg.brandl, mo
str.format	eric.smith*
testing	michael.foord, ezio.melotti
test coverage	
threads	
time and dates	lemburg, belopolsky, p-ganssle
unicode	lemburg, ezio.melotti, benjamin.peterson,
version control	eric.araujo, ezio.melotti

10.21.5 Documentation Translations

Translation	Coordinator
French	mdk
Japanese	inada.naoki
Korean	flowdas
Bengali India	kushal.das
Hungarian	gbtami
Portuguese	rougeth
Chinese (TW)	adrianliaw
Chinese (CN)	zhsj

10.22 gdb Support

If you experience low-level problems such as crashes or deadlocks (e.g. when tinkering with parts of CPython which are written in C), it can be convenient to use a low-level debugger such as `gdb` in order to diagnose and fix the issue. By default, however, `gdb` (or any of its front-ends) doesn't know about high-level information specific to the CPython interpreter, such as which Python function is currently executing, or what type or value has a given Python object represented by a standard `PyObject *` pointer. We hereafter present two ways to overcome this limitation.

10.22.1 gdb 7 and later

In gdb 7, support for [extending gdb with Python](#) was added. When CPython is built you will notice a `python-gdb.py` file in the root directory of your checkout. Read the module docstring for details on how to use the file to enhance gdb for easier debugging of a CPython process.

To activate support, you must add the directory containing `python-gdb.py` to GDB's "auto-load-safe-path". Put this in your `~/.gdbinit` file:

```
add-auto-load-safe-path /path/to/checkout
```

You can also add multiple paths, separated by `:`.

This is what a backtrace looks like (truncated) when this extension is enabled:

```
#0 0x00000000041a6b1 in PyObject_Malloc (nbytes=Cannot access memory at address_
↳0x7fffff7fefe8
) at Objects/obmalloc.c:748
#1 0x00000000041b7c0 in _PyObject_DebugMallocApi (id=111 'o', nbytes=24) at Objects/
↳obmalloc.c:1445
#2 0x00000000041b717 in _PyObject_DebugMalloc (nbytes=24) at Objects/obmalloc.c:1412
#3 0x00000000044060a in _PyUnicode_New (length=11) at Objects/unicodeobject.c:346
#4 0x0000000004466aa in PyUnicodeUCS2_DecodeUTF8Stateful (s=0x5c2b8d "__lltrace__",_
↳size=11, errors=0x0, consumed=
0x0) at Objects/unicodeobject.c:2531
#5 0x000000000446647 in PyUnicodeUCS2_DecodeUTF8 (s=0x5c2b8d "__lltrace__", size=11,
↳errors=0x0)
at Objects/unicodeobject.c:2495
#6 0x000000000440d1b in PyUnicodeUCS2_FromStringAndSize (u=0x5c2b8d "__lltrace__",_
↳size=11)
at Objects/unicodeobject.c:551
#7 0x000000000440d94 in PyUnicodeUCS2_FromString (u=0x5c2b8d "__lltrace__") at_
↳Objects/unicodeobject.c:569
#8 0x000000000584abd in PyDict_GetItemString (v=
{'Yuck': <type at remote 0xad4730>, '__builtins__': <module at remote_
↳0x7ffff7fd5ee8>, '__file__': 'Lib/test/crashers/nasty_eq_vs_dict.py', '__package__
↳': None, 'y': <Yuck(i=0) at remote 0xaacd80>, 'dict': {0: 0, 1: 1, 2: 2, 3: 3}, '__
↳cached__': None, '__name__': '__main__', 'z': <Yuck(i=0) at remote 0xaace60>, '__
↳doc__': None}, key=
0x5c2b8d "__lltrace__") at Objects/dictobject.c:2171
```

(Notice how the dictionary argument to `PyDict_GetItemString` is displayed as its `repr()`, rather than an opaque `PyObject *` pointer.)

The extension works by supplying a custom printing routine for values of type `PyObject *`. If you need to access lower-level details of an object, then cast the value to a pointer of the appropriate type. For example:

```
(gdb) p globals
$1 = {'__builtins__': <module at remote 0x7ffff7fb1868>, '__name__':
'__main__', 'ctypes': <module at remote 0x7ffff7f14360>, '__doc__': None,
'__package__': None}

(gdb) p *(PyDictObject*)globals
$2 = {ob_refcnt = 3, ob_type = 0x3dbdf85820, ma_fill = 5, ma_used = 5,
ma_mask = 7, ma_table = 0x63d0f8, ma_lookup = 0x3dbdc7ea70
<lookdict_string>, ma_smalltable = {me_hash = 7065186196740147912,
me_key = '__builtins__', me_value = <module at remote 0x7ffff7fb1868>},
me_hash = -368181376027291943, me_key = '__name__',
```

(continues on next page)

(continued from previous page)

```
me_value = '__main__'}, {me_hash = 0, me_key = 0x0, me_value = 0x0},
{me_hash = 0, me_key = 0x0, me_value = 0x0},
{me_hash = -9177857982131165996, me_key = 'ctypes',
me_value = <module at remote 0x7ffff7f14360>},
{me_hash = -8518757509529533123, me_key = '__doc__', me_value = None},
{me_hash = 0, me_key = 0x0, me_value = 0x0}, {
  me_hash = 6614918939584953775, me_key = '__package__', me_value = None}}
```

The pretty-printers try to closely match the `repr()` implementation of the underlying implementation of Python, and thus vary somewhat between Python 2 and Python 3.

An area that can be confusing is that the custom printer for some types look a lot like gdb's built-in printer for standard types. For example, the pretty-printer for a Python 3 `int` gives a `repr()` that is not distinguishable from a printing of a regular machine-level integer:

```
(gdb) p some_machine_integer
$3 = 42

(gdb) p some_python_integer
$4 = 42

(gdb) p *(PyLongObject*)some_python_integer
$5 = {ob_base = {ob_base = {ob_refcnt = 8, ob_type = 0x3dad39f5e0}, ob_size = 1},
ob_digit = {42}}
```

A similar confusion can arise with the `str` type, where the output looks a lot like gdb's built-in printer for `char *`:

```
(gdb) p ptr_to_python_str
$6 = '__builtins__'
```

The pretty-printer for `str` instances defaults to using single-quotes (as does Python's `repr` for strings) whereas the standard printer for `char *` values uses double-quotes and contains a hexadecimal address:

```
(gdb) p ptr_to_char_star
$7 = 0x6d72c0 "hello world"
```

Here's how to see the implementation details of a `str` instance (for Python 3, where a `str` is a `PyUnicodeObject *`):

```
(gdb) p *(PyUnicodeObject*)$6
$8 = {ob_base = {ob_refcnt = 33, ob_type = 0x3dad3a95a0}, length = 12,
str = 0x7ffff2128500, hash = 7065186196740147912, state = 1, defenc = 0x0}
```

As well as adding pretty-printing support for `PyObject *`, the extension adds a number of commands to gdb:

py-list List the Python source code (if any) for the current frame in the selected thread. The current line is marked with a ">":

```
(gdb) py-list
901     if options.profile:
902         options.profile = False
903         profile_me()
904         return
905
>906     u = UI()
907     if not u.quit:
908         try:
```

(continues on next page)

(continued from previous page)

```

909             gtk.main()
910         except KeyboardInterrupt:
911             # properly quit on a keyboard interrupt...

```

Use `py-list START` to list at a different line number within the python source, and `py-list START, END` to list a specific range of lines within the python source.

py-up and py-down The `py-up` and `py-down` commands are analogous to `gdb`'s regular up and down commands, but try to move at the level of CPython frames, rather than C frames.

`gdb` is not always able to read the relevant frame information, depending on the optimization level with which CPython was compiled. Internally, the commands look for C frames that are executing `PyEval_EvalFrameEx` (which implements the core bytecode interpreter loop within CPython) and look up the value of the related `PyFrameObject *`.

They emit the frame number (at the C level) within the thread.

For example:

```

(gdb) py-up
#37 Frame 0x9420b04, for file /usr/lib/python2.6/site-packages/
gnome_sudoku/main.py, line 906, in start_game ()
    u = UI()
(gdb) py-up
#40 Frame 0x948e82c, for file /usr/lib/python2.6/site-packages/
gnome_sudoku/gnome_sudoku.py, line 22, in start_game(main=<module at remote_
↳0xb771b7f4>)
    main.start_game()
(gdb) py-up
Unable to find an older python frame

```

so we're at the top of the python stack. Going back down:

```

(gdb) py-down
#37 Frame 0x9420b04, for file /usr/lib/python2.6/site-packages/gnome_sudoku/main.
↳py, line 906, in start_game ()
    u = UI()
(gdb) py-down
#34 (unable to read python frame information)
(gdb) py-down
#23 (unable to read python frame information)
(gdb) py-down
#19 (unable to read python frame information)
(gdb) py-down
#14 Frame 0x99262ac, for file /usr/lib/python2.6/site-packages/gnome_sudoku/game_
↳selector.py, line 201, in run_swallowed_dialog (self=
↳<NewOrSavedGameSelector(new_game_model=<gtk.ListStore at remote 0x98fab44>,
↳puzzle=None, saved_games=[{'gsd.auto_fills': 0, 'tracking': {}, 'trackers': {}},
↳'notes': [], 'saved_at': 1270084485, 'game': '7 8 0 0 0 0 5 6 0 0 9 0 8 0 1 0
↳0 0 4 6 0 0 0 7 0 6 5 0 0 0 4 7 9 2 0 0 0 9 0 1 0 0 0 3 9 7 6 0 0 0 1 8 0 6 0
↳0 0 0 2 8 0 0 0 5 0 4 0 6 0 0 2 1 0 0 0 0 0 4 5\n7 8 0 0 0 0 0 5 6 0 0 9 0 8 0
↳1 0 0 0 4 6 0 0 0 7 0 6 5 1 8 3 4 7 9 2 0 0 0 9 0 1 0 0 0 3 9 7 6 0 0 0 1 8 0
↳6 0 0 0 2 8 0 0 0 5 0 4 0 6 0 0 2 1 0 0 0 0 0 4 5', 'gsd.impossible_hints': 0,
↳ 'timer.__absolute_start_time__': <float at remote 0x984b474>, 'gsd.hints': 0,
↳'timer.active_time': <float at remote 0x984b494>, 'timer.total_time': <float at
↳remote 0x984b464>}], dialog=<gtk.Dialog at remote 0x98faaa4>, saved_game_model=
↳<gtk.ListStore at remote 0x98fad24>, sudoku_maker=<SudokuMaker(terminated=False,
↳played=[], batch_siz...(truncated)

```

(continues on next page)

(continued from previous page)

```

        swallower.run_dialog(self.dialog)
(gdb) py-down
#11 Frame 0x9aead74, for file /usr/lib/python2.6/site-packages/gnome_sudoku/
↳dialog_swallower.py, line 48, in run_dialog (self=<SwappableArea(running=<gtk.
↳Dialog at remote 0x98faaa4>, main_page=0) at remote 0x98fa6e4>, d=<gtk.Dialog_
↳at remote 0x98faaa4>)
        gtk.main()
(gdb) py-down
#8 (unable to read python frame information)
(gdb) py-down
Unable to find a newer python frame

```

and we're at the bottom of the python stack.

py-bt The `py-bt` command attempts to display a Python-level backtrace of the current thread.

For example:

```

(gdb) py-bt
#8 (unable to read python frame information)
#11 Frame 0x9aead74, for file /usr/lib/python2.6/site-packages/gnome_sudoku/
↳dialog_swallower.py, line 48, in run_dialog (self=<SwappableArea(running=<gtk.
↳Dialog at remote 0x98faaa4>, main_page=0) at remote 0x98fa6e4>, d=<gtk.Dialog_
↳at remote 0x98faaa4>)
        gtk.main()
#14 Frame 0x99262ac, for file /usr/lib/python2.6/site-packages/gnome_sudoku/game_
↳selector.py, line 201, in run_swallowed_dialog (self=
↳<NewOrSavedGameSelector(new_game_model=<gtk.ListStore at remote 0x98fab44>,
↳puzzle=None, saved_games=[{'gsd.auto_fills': 0, 'tracking': {}, 'trackers': {},
↳'notes': [], 'saved_at': 1270084485, 'game': '7 8 0 0 0 0 0 5 6 0 0 9 0 8 0 1 0_
↳0 0 4 6 0 0 0 0 7 0 6 5 0 0 0 4 7 9 2 0 0 0 9 0 1 0 0 0 3 9 7 6 0 0 0 1 8 0 6 0_
↳0 0 0 2 8 0 0 0 5 0 4 0 6 0 0 2 1 0 0 0 0 0 4 5\n7 8 0 0 0 0 0 5 6 0 0 9 0 8 0_
↳1 0 0 0 4 6 0 0 0 0 7 0 6 5 1 8 3 4 7 9 2 0 0 0 9 0 1 0 0 0 3 9 7 6 0 0 0 1 8 0_
↳6 0 0 0 0 2 8 0 0 0 5 0 4 0 6 0 0 2 1 0 0 0 0 0 4 5', 'gsd.impossible_hints': 0,
↳ 'timer.__absolute_start_time__': <float at remote 0x984b474>, 'gsd.hints': 0,
↳ 'timer.active_time': <float at remote 0x984b494>, 'timer.total_time': <float at_
↳remote 0x984b464>}], dialog=<gtk.Dialog at remote 0x98faaa4>, saved_game_model=
↳<gtk.ListStore at remote 0x98fad24>, sudoku_maker=<SudokuMaker(terminated=False,
↳played=[], batch_siz...(truncated)
        swallower.run_dialog(self.dialog)
#19 (unable to read python frame information)
#23 (unable to read python frame information)
#34 (unable to read python frame information)
#37 Frame 0x9420b04, for file /usr/lib/python2.6/site-packages/gnome_sudoku/main.
↳py, line 906, in start_game ()
        u = UI()
#40 Frame 0x948e82c, for file /usr/lib/python2.6/site-packages/gnome_sudoku/gnome_
↳sudoku.py, line 22, in start_game (main=<module at remote 0xb771b7f4>)
        main.start_game()

```

The frame numbers correspond to those displayed by `gdb`'s standard `backtrace` command.

py-print The `py-print` command looks up a Python name and tries to print it. It looks in locals within the current thread, then globals, then finally builtins:

```

(gdb) py-print self
local 'self' = <SwappableArea(running=<gtk.Dialog at remote 0x98faaa4>,
main_page=0) at remote 0x98fa6e4>

```

(continues on next page)

(continued from previous page)

```
(gdb) py-print __name__
global '__name__' = 'gnome_sudoku.dialog_swallower'
(gdb) py-print len
builtin 'len' = <built-in function len>
(gdb) py-print scarlet_pimpernel
'scarlet_pimpernel' not found
```

py-locals The `py-locals` command looks up all Python locals within the current Python frame in the selected thread, and prints their representations:

```
(gdb) py-locals
self = <SwappableArea(running=<gtk.Dialog at remote 0x98faaa4>,
main_page=0) at remote 0x98fa6e4>
d = <gtk.Dialog at remote 0x98faaa4>
```

You can of course use other `gdb` commands. For example, the `frame` command takes you directly to a particular frame within the selected thread. We can use it to go a specific frame shown by `py-bt` like this:

```
(gdb) py-bt
(output snipped)
#68 Frame 0xaa4560, for file Lib/test/regrtest.py, line 1548, in <module> ()
    main()
(gdb) frame 68
#68 0x0000000004cd1e6 in PyEval_EvalFrameEx (f=Frame 0xaa4560, for file Lib/test/
↳ regrtest.py, line 1548, in <module> (), throwflag=0) at Python/ceval.c:2665
2665             x = call_function(&sp, oparg);
(gdb) py-list
1543         # Run the tests in a context manager that temporary changes the CWD to a
1544         # temporary and writable directory. If it's not possible to create or
1545         # change the CWD, the original CWD will be used. The original CWD is
1546         # available from test_support.SAVEDCWD.
1547         with test_support.temp_cwd(TESTCWD, quiet=True):
>1548             main()
```

The `info threads` command will give you a list of the threads within the process, and you can use the `thread` command to select a different one:

```
(gdb) info threads
 105 Thread 0x7ffffefa18710 (LWP 10260)  sem_wait () at ../nptl/sysdeps/unix/sysv/
↳ linux/x86_64/sem_wait.S:86
 104 Thread 0x7ffffdf5fe710 (LWP 10259)  sem_wait () at ../nptl/sysdeps/unix/sysv/
↳ linux/x86_64/sem_wait.S:86
* 1 Thread 0x7ffff7fe2700 (LWP 10145)  0x00000038e46d73e3 in select () at ../sysdeps/
↳ unix/syscall-template.S:82
```

You can use `thread apply all COMMAND` or `(t a a COMMAND` for short) to run a command on all threads. You can use this with `py-bt` to see what every thread is doing at the Python level:

```
(gdb) t a a py-bt

Thread 105 (Thread 0x7ffffefa18710 (LWP 10260)):
#5 Frame 0x7ffffd00019d0, for file /home/david/coding/python-svn/Lib/threading.py,
↳ line 155, in _acquire_restore (self=<_RLock(_Verbose__verbose=False, _RLock__
↳ owner=140737354016512, _RLock__block=<thread.lock at remote 0x858770>, _RLock__
↳ count=1) at remote 0xd7ff40>, count_owner=(1, 140737213728528), count=1,
↳ owner=140737213728528)
```

(continues on next page)

(continued from previous page)

```

        self.__block.acquire()
#8 Frame 0x7fffac001640, for file /home/david/coding/python-svn/Lib/threading.py,
↳line 269, in wait (self=<_Condition(_Condition__lock=<_RLock(_Verbose__
↳verbose=False, _RLock__owner=140737354016512, _RLock__block=<thread.lock at remote_
↳0x858770>, _RLock__count=1) at remote 0xd7ff40>, acquire=<instancemethod at remote_
↳0xd80260>, _is_owned=<instancemethod at remote 0xd80160>, _release_save=
↳<instancemethod at remote 0xd803e0>, release=<instancemethod at remote 0xd802e0>, _
↳acquire_restore=<instancemethod at remote 0xd7ee60>, _Verbose__verbose=False, _
↳Condition__waiters=[]) at remote 0xd7fd10>, timeout=None, waiter=<thread.lock at_
↳remote 0x858a90>, saved_state=(1, 140737213728528))
        self.__acquire_restore(saved_state)
#12 Frame 0x7ffffb8001a10, for file /home/david/coding/python-svn/Lib/test/lock_tests.
↳py, line 348, in f ()
        cond.wait()
#16 Frame 0x7ffffb8001c40, for file /home/david/coding/python-svn/Lib/test/lock_tests.
↳py, line 37, in task (tid=140737213728528)
        f()

Thread 104 (Thread 0x7ffffdf5fe710 (LWP 10259)):
#5 Frame 0x7ffffe4001580, for file /home/david/coding/python-svn/Lib/threading.py,
↳line 155, in _acquire_restore (self=<_RLock(_Verbose__verbose=False, _RLock__
↳owner=140737354016512, _RLock__block=<thread.lock at remote 0x858770>, _RLock__
↳count=1) at remote 0xd7ff40>, count_owner=(1, 140736940992272), count=1,
↳owner=140736940992272)
        self.__block.acquire()
#8 Frame 0x7fffc8002090, for file /home/david/coding/python-svn/Lib/threading.py,
↳line 269, in wait (self=<_Condition(_Condition__lock=<_RLock(_Verbose__
↳verbose=False, _RLock__owner=140737354016512, _RLock__block=<thread.lock at remote_
↳0x858770>, _RLock__count=1) at remote 0xd7ff40>, acquire=<instancemethod at remote_
↳0xd80260>, _is_owned=<instancemethod at remote 0xd80160>, _release_save=
↳<instancemethod at remote 0xd803e0>, release=<instancemethod at remote 0xd802e0>, _
↳acquire_restore=<instancemethod at remote 0xd7ee60>, _Verbose__verbose=False, _
↳Condition__waiters=[]) at remote 0xd7fd10>, timeout=None, waiter=<thread.lock at_
↳remote 0x858860>, saved_state=(1, 140736940992272))
        self.__acquire_restore(saved_state)
#12 Frame 0x7fffac001c90, for file /home/david/coding/python-svn/Lib/test/lock_tests.
↳py, line 348, in f ()
        cond.wait()
#16 Frame 0x7fffac0011c0, for file /home/david/coding/python-svn/Lib/test/lock_tests.
↳py, line 37, in task (tid=140736940992272)
        f()

Thread 1 (Thread 0x7ffff7fe2700 (LWP 10145)):
#5 Frame 0xcb5380, for file /home/david/coding/python-svn/Lib/test/lock_tests.py,
↳line 16, in _wait ()
        time.sleep(0.01)
#8 Frame 0x7fffd00024a0, for file /home/david/coding/python-svn/Lib/test/lock_tests.
↳py, line 378, in _check_notify (self=<ConditionTests(_testMethodName='test_notify',
↳_resultForDoCleanups=<TestResult(_original_stdout=<cStringIO.StringO at remote_
↳0xc191e0>, skipped=[], _mirrorOutput=False, testsRun=39, buffer=False, _original_
↳stderr=<file at remote 0x7ffff7fc6340>, _stdout_buffer=<cStringIO.StringO at remote_
↳0xc9c7f8>, _stderr_buffer=<cStringIO.StringO at remote 0xc9c790>, _
↳moduleSetUpFailed=False, expectedFailures=[], errors=[], _previousTestClass=<type_
↳at remote 0x928310>, unexpectedSuccesses=[], failures=[], shouldStop=False,
↳failfast=False) at remote 0xc185a0>, _threads=(0,), _cleanups=[], _type_equality_
↳funcs={<type at remote 0x7eba00>: <instancemethod at remote 0xd750e0>, <type at_
↳remote 0x7e7820>: <instancemethod at remote 0xd75160>, <type at remote 0x7e30e0>:
↳<instancemethod at remote 0xd75060>, <type at remote 0x7e7d20>: <instancemethod at_
↳remote 0xd751e0>, <type at remote 0x7f19e0... (truncated)

```

(continued from previous page)

```
_wait()
```

Note: This is only available for Python 2.7, 3.2 and higher.

10.22.2 gdb 6 and earlier

The file at `Misc/gdbinit` contains a `gdb` configuration file which provides extra commands when working with a CPython process. To register these commands permanently, either copy the commands to your personal `gdb` configuration file or symlink `~/gdbinit` to `Misc/gdbinit`. To use these commands from a single `gdb` session without registering them, type `source Misc/gdbinit` from your `gdb` session.

10.22.3 Updating auto-load-safe-path to allow test_gdb to run

`test_gdb` attempts to automatically load additional Python specific hooks into `gdb` in order to test them. Unfortunately, the command line options it uses to do this aren't always supported correctly.

If `test_gdb` is being skipped with an "auto-loading has been declined" message, then it is necessary to identify any Python build directories as auto-load safe. One way to achieve this is to add a line like the following to `~/gdbinit` (edit the specific list of paths as appropriate):

```
add-auto-load-safe-path ~/devel/py3k:~/devel/py32:~/devel/py27
```

10.23 Exploring CPython's Internals

This is a quick guide for people who are interested in learning more about CPython's internals. It provides a summary of the source code structure and contains references to resources providing a more in-depth view.

10.23.1 CPython Source Code Layout

This guide gives an overview of CPython's code structure. It serves as a summary of file locations for modules and builtins.

For Python modules, the typical layout is:

- `Lib/<module>.py`
- `Modules/_<module>.c` (if there's also a C accelerator module)
- `Lib/test/test_<module>.py`
- `Doc/library/<module>.rst`

For extension-only modules, the typical layout is:

- `Modules/<module>module.c`
- `Lib/test/test_<module>.py`
- `Doc/library/<module>.rst`

For builtin types, the typical layout is:

- `Objects/<builtin>object.c`
- `Lib/test/test_<builtin>.py`
- `Doc/library/stdtypes.rst`

For builtin functions, the typical layout is:

- `Python/bltinmodule.c`
- `Lib/test/test_builtin.py`
- `Doc/library/functions.rst`

Some exceptions:

- builtin type `int` is at `Objects/longobject.c`
- builtin type `str` is at `Objects/unicodeobject.c`
- builtin module `sys` is at `Python/sysmodule.c`
- builtin module `marshal` is at `Python/marshal.c`
- Windows-only module `winreg` is at `PC/winreg.c`

10.23.2 Additional References

For over 20 years the CPython code base has been changing and evolving. Here's a sample of resources about the architecture of CPython aimed at building your understanding of both the 2.x and 3.x versions of CPython:

Table 4: **Current references**

Title	Brief	Author	Version
A guide from parser to objects, observed using GDB	Code walk from Parser, AST, Sym Table and Objects	Louie Lu	3.7.a0
Green Tree Snakes	The missing Python AST docs	Thomas Kluyver	3.6
Yet another guided tour of CPython	A guide for how CPython REPL works	Guido van Rossum	3.5
Python Asynchronous I/O Walkthrough	How CPython async I/O, generator and coroutine works	Philip Guo	3.5
Coding Patterns for Python Extensions	Reliable patterns of coding Python Extensions in C	Paul Ross	3.4
Your Guide to the CPython Source Code	Your Guide to the CPython Source Code	Anthony Shaw	3.8

Table 5: **Historical references**

Title	Brief	Author	Version
Python's Innards Series	ceval, objects, pystate and miscellaneous topics	Yaniv Akinin	3.1
Eli Bendersky's Python Internals	Objects, Symbol tables and miscellaneous topics	Eli Bendersky	3.x
A guide from parser to objects, observed using Eclipse	Code walk from Parser, AST, Sym Table and Objects	Prashanth Raghu	2.7.12
CPython internals: A ten-hour codewalk through the Python interpreter source code	Code walk from source code to generators	Philip Guo	2.7.8

10.24 Changing CPython's Grammar

10.24.1 Abstract

There's more to changing Python's grammar than editing `Grammar/python.gram`. Here's a checklist.

NOTE: These instructions are for Python 3.9 and beyond. Earlier versions use a different parser technology. You probably shouldn't try to change the grammar of earlier Python versions, but if you really want to, use GitHub to track down the earlier version of this file in the devguide. (Python 3.9 itself actually supports both parsers; the old parser can be invoked by passing `-X oldparser`.)

10.24.2 Checklist

Note: sometimes things mysteriously don't work. Before giving up, try `make clean`.

- `Grammar/python.gram`: The grammar, with actions that build AST nodes. After changing it, run `make regen-pegen`, to regenerate `Parser/parser.c`. (This runs Python's parser generator, `Tools/peg_generator`).
- `Grammar/Tokens` is a place for adding new token types. After changing it, run `make regen-token` to regenerate `Include/token.h`, `Parser/token.c`, `Lib/token.py` and `Doc/library/token-list.inc`. If you change both `python.gram` and `Tokens`, run `make regen-token` before `make regen-pegen`.
- `Parser/Python.asdl` may need changes to match the grammar. Then run `make regen-ast` to regenerate `Include/Python-ast.h` and `Python/Python-ast.c`.
- `Parser/tokenizer.c` contains the tokenization code. This is where you would add a new type of comment or string literal, for example.
- `Python/ast.c` will need changes to validate AST objects involved with the grammar change.
- `Python/ast_unparse.c` will need changes to unparse AST objects involved with the grammar change ("unparsing" is used to turn annotations into strings per [PEP 563](#)).
- The *Design of CPython's Compiler* has its own page.
- `_Unparser` in the `Lib/ast.py` file may need changes to accommodate any modifications in the AST nodes.
- Add some usage of your new syntax to `test_grammar.py`.
- Certain changes may require tweaks to the library module `pyclbr`.
- `Lib/tokenize.py` needs changes to match changes to the tokenizer.
- `Lib/lib2to3/Grammar.txt` may need changes to match the Grammar.
- Documentation must be written! Specifically, one or more of the pages in `Doc/reference/` will need to be updated.

10.25 Design of CPython's Compiler

10.25.1 Abstract

In CPython, the compilation from source code to bytecode involves several steps:

1. Tokenize the source code (`Parser/tokenizer.c`)

2. Parse the stream of tokens into an Abstract Syntax Tree (`Parser/parser.c`)
3. Transform AST into a Control Flow Graph (`Python/compile.c`)
4. Emit bytecode based on the Control Flow Graph (`Python/compile.c`)

The purpose of this document is to outline how these steps of the process work.

This document does not touch on how parsing works beyond what is needed to explain what is needed for compilation. It is also not exhaustive in terms of the how the entire system works. You will most likely need to read some source to have an exact understanding of all details.

10.25.2 Parsing

As of Python 3.9, Python's parser is a PEG parser of a somewhat unusual design (since its input is a stream of tokens rather than a stream of characters as is more common with PEG parsers).

The grammar file for Python can be found in `Grammar/python.gram`. The definitions for literal tokens (such as `:`, numbers, etc.) can be found in `Grammar/Tokens`. Various C files, including `Parser/parser.c` are generated from these (see *Changing CPython's Grammar*).

10.25.3 Abstract Syntax Trees (AST)

Green Tree Snakes

See also [Green Tree Snakes - the missing Python AST docs](#) by Thomas Kluyver.

The abstract syntax tree (AST) is a high-level representation of the program structure without the necessity of containing the source code; it can be thought of as an abstract representation of the source code. The specification of the AST nodes is specified using the Zephyr Abstract Syntax Definition Language (ASDL) [Wang97].

The definition of the AST nodes for Python is found in the file `Parser/Python.asdl`.

Each AST node (representing statements, expressions, and several specialized types, like list comprehensions and exception handlers) is defined by the ASDL. Most definitions in the AST correspond to a particular source construct, such as an 'if' statement or an attribute lookup. The definition is independent of its realization in any particular programming language.

The following fragment of the Python ASDL construct demonstrates the approach and syntax:

```
module Python
{
    stmt = FunctionDef(identifier name, arguments args, stmt* body,
                      expr* decorators)
          | Return(expr? value) | Yield(expr? value)
          attributes (int lineno)
}
```

The preceding example describes two different kinds of statements and an expression: function definitions, return statements, and yield expressions. All three kinds are considered of type `stmt` as shown by `|` separating the various kinds. They all take arguments of various kinds and amounts.

Modifiers on the argument type specify the number of values needed; `?` means it is optional, `*` means 0 or more, while no modifier means only one value for the argument and it is required. `FunctionDef`, for instance, takes an `identifier` for the *name*, `arguments` for *args*, zero or more `stmt` arguments for *body*, and zero or more `expr` arguments for *decorators*.

Do notice that something like 'arguments', which is a node type, is represented as a single AST node and not as a sequence of nodes as with `stmt` as one might expect.

All three kinds also have an 'attributes' argument; this is shown by the fact that 'attributes' lacks a 'l' before it.

The statement definitions above generate the following C structure type:

```
typedef struct _stmt *stmt_ty;

struct _stmt {
    enum { FunctionDef_kind=1, Return_kind=2, Yield_kind=3 } kind;
    union {
        struct {
            identifier name;
            arguments_ty args;
            asdl_seq *body;
        } FunctionDef;

        struct {
            expr_ty value;
        } Return;

        struct {
            expr_ty value;
        } Yield;
    } v;
    int lineno;
}
```

Also generated are a series of constructor functions that allocate (in this case) a `stmt_ty` struct with the appropriate initialization. The `kind` field specifies which component of the union is initialized. The `FunctionDef()` constructor function sets 'kind' to `FunctionDef_kind` and initializes the `name`, `args`, `body`, and `attributes` fields.

10.25.4 Memory Management

Before discussing the actual implementation of the compiler, a discussion of how memory is handled is in order. To make memory management simple, an arena is used. This means that a memory is pooled in a single location for easy allocation and removal. What this gives us is the removal of explicit memory deallocation. Because memory allocation for all needed memory in the compiler registers that memory with the arena, a single call to free the arena is all that is needed to completely free all memory used by the compiler.

In general, unless you are working on the critical core of the compiler, memory management can be completely ignored. But if you are working at either the very beginning of the compiler or the end, you need to care about how the arena works. All code relating to the arena is in either `Include/pyarena.h` or `Python/pyarena.c`.

`PyArena_New()` will create a new arena. The returned `PyArena` structure will store pointers to all memory given to it. This does the bookkeeping of what memory needs to be freed when the compiler is finished with the memory it used. That freeing is done with `PyArena_Free()`. This only needs to be called in strategic areas where the compiler exits.

As stated above, in general you should not have to worry about memory management when working on the compiler. The technical details have been designed to be hidden from you for most cases.

The only exception comes about when managing a `PyObject`. Since the rest of Python uses reference counting, there is extra support added to the arena to cleanup each `PyObject` that was allocated. These cases are very rare. However, if you've allocated a `PyObject`, you must tell the arena about it by calling `PyArena_AddPyObject()`.

10.25.5 Parse Tree to AST

The AST is generated from the parse tree (see `Python/ast.c`) using the function `PyAST_FromNode()`.

The function begins a tree walk of the parse tree, creating various AST nodes as it goes along. It does this by allocating all new nodes it needs, calling the proper AST node creation functions for any required supporting functions, and connecting them as needed.

Do realize that there is no automated nor symbolic connection between the grammar specification and the nodes in the parse tree. No help is directly provided by the parse tree as in yacc.

For instance, one must keep track of which node in the parse tree one is working with (e.g., if you are working with an 'if' statement you need to watch out for the ':' token to find the end of the conditional).

The functions called to generate AST nodes from the parse tree all have the name `ast_for_xx` where `xx` is the grammar rule that the function handles (`alias_for_import_name` is the exception to this). These in turn call the constructor functions as defined by the ASDL grammar and contained in `Python/Python-ast.c` (which was generated by `Parser/asdl_c.py`) to create the nodes of the AST. This all leads to a sequence of AST nodes stored in `asdl_seq` structs.

Function and macros for creating and using `asdl_seq * types` as found in `Python/asdl.c` and `Include/asdl.h` are as follows:

`Py_asdl_seq_new(Py_ssize_t, PyArena *)` Allocate memory for an `asdl_seq` for the specified length

`asdl_seq_GET(asdl_seq *, int)` Get item held at a specific position in an `asdl_seq`

`asdl_seq_SET(asdl_seq *, int, stmt_ty)` Set a specific index in an `asdl_seq` to the specified value

`asdl_seq_LEN(asdl_seq *)` Return the length of an `asdl_seq`

If you are working with statements, you must also worry about keeping track of what line number generated the statement. Currently the line number is passed as the last parameter to each `stmt_ty` function.

10.25.6 Control Flow Graphs

A control flow graph (often referenced by its acronym, CFG) is a directed graph that models the flow of a program using basic blocks that contain the intermediate representation (abbreviated "IR", and in this case is Python bytecode) within the blocks. Basic blocks themselves are a block of IR that has a single entry point but possibly multiple exit points. The single entry point is the key to basic blocks; it all has to do with jumps. An entry point is the target of something that changes control flow (such as a function call or a jump) while exit points are instructions that would change the flow of the program (such as jumps and 'return' statements). What this means is that a basic block is a chunk of code that starts at the entry point and runs to an exit point or the end of the block.

As an example, consider an 'if' statement with an 'else' block. The guard on the 'if' is a basic block which is pointed to by the basic block containing the code leading to the 'if' statement. The 'if' statement block contains jumps (which are exit points) to the true body of the 'if' and the 'else' body (which may be `NULL`), each of which are their own basic blocks. Both of those blocks in turn point to the basic block representing the code following the entire 'if' statement.

CFGs are usually one step away from final code output. Code is directly generated from the basic blocks (with jump targets adjusted based on the output order) by doing a post-order depth-first search on the CFG following the edges.

10.25.7 AST to CFG to Bytecode

With the AST created, the next step is to create the CFG. The first step is to convert the AST to Python bytecode without having jump targets resolved to specific offsets (this is calculated when the CFG goes to final bytecode). Essentially, this transforms the AST into Python bytecode with control flow represented by the edges of the CFG.

Conversion is done in two passes. The first creates the namespace (variables can be classified as local, free/cell for closures, or global). With that done, the second pass essentially flattens the CFG into a list and calculates jump offsets for final output of bytecode.

The conversion process is initiated by a call to the function `PyAST_CompileObject()` in `Python/compile.c`. This function does both the conversion of the AST to a CFG and outputting final bytecode from the CFG. The AST to CFG step is handled mostly by two functions called by `PyAST_CompileObject()`; `PySymtable_BuildObject()` and `compiler_mod()`. The former is in `Python/symtable.c` while the latter is in `Python/compile.c`.

`PySymtable_BuildObject()` begins by entering the starting code block for the AST (passed-in) and then calling the proper `symtable_visit_xx` function (with `xx` being the AST node type). Next, the AST tree is walked with the various code blocks that delineate the reach of a local variable as blocks are entered and exited using `symtable_enter_block()` and `symtable_exit_block()`, respectively.

Once the symbol table is created, it is time for CFG creation, whose code is in `Python/compile.c`. This is handled by several functions that break the task down by various AST node types. The functions are all named `compiler_visit_xx` where `xx` is the name of the node type (such as `stmt`, `expr`, etc.). Each function receives a `struct compiler *` and `xx_ty` where `xx` is the AST node type. Typically these functions consist of a large 'switch' statement, branching based on the kind of node type passed to it. Simple things are handled inline in the 'switch' statement with more complex transformations farmed out to other functions named `compiler_xx` with `xx` being a descriptive name of what is being handled.

When transforming an arbitrary AST node, use the `VISIT()` macro. The appropriate `compiler_visit_xx` function is called, based on the value passed in for <node type> (so `VISIT(c, expr, node)` calls `compiler_visit_expr(c, node)`). The `VISIT_SEQ()` macro is very similar, but is called on AST node sequences (those values that were created as arguments to a node that used the '*' modifier). There is also `VISIT_SLICE()` just for handling slices.

Emission of bytecode is handled by the following macros:

ADDOP(struct compiler *, int) add a specified opcode

ADDOP_I(struct compiler *, int, Py_ssize_t) add an opcode that takes an integer argument

ADDOP_O(struct compiler *, int, PyObject *, TYPE) add an opcode with the proper argument based on the position of the specified PyObject in PyObject sequence object, but with no handling of mangled names; used for when you need to do named lookups of objects such as globals, consts, or parameters where name mangling is not possible and the scope of the name is known; *TYPE* is the name of PyObject sequence (*names* or *varnames*)

ADDOP_N(struct compiler *, int, PyObject *, TYPE) just like `ADDOP_O`, but steals a reference to PyObject

ADDOP_NAME(struct compiler *, int, PyObject *, TYPE) just like `ADDOP_O`, but name mangling is also handled; used for attribute loading or importing based on name

ADDOP_LOAD_CONST(struct compiler *, PyObject *) add the `LOAD_CONST` opcode with the proper argument based on the position of the specified PyObject in the consts table.

ADDOP_LOAD_CONST_NEW(struct compiler *, PyObject *) just like `ADDOP_LOAD_CONST_NEW`, but steals a reference to PyObject

ADDOP_JABS(struct compiler *, int, basicblock *) create an absolute jump to a basic block

ADDOP_JREL(struct compiler *, int, basicblock *) create a relative jump to a basic block

Several helper functions that will emit bytecode and are named `compiler_xx()` where `xx` is what the function helps with (`list`, `boolop`, etc.). A rather useful one is `compiler_nameop()`. This function looks up the scope of a variable and, based on the expression context, emits the proper opcode to load, store, or delete the variable.

As for handling the line number on which a statement is defined, this is handled by `compiler_visit_stmt()` and thus is not a worry.

In addition to emitting bytecode based on the AST node, handling the creation of basic blocks must be done. Below are the macros and functions used for managing basic blocks:

NEXT_BLOCK(struct compiler *) create an implicit jump from the current block to the new block

compiler_new_block(struct compiler *) create a block but don't use it (used for generating jumps)

compiler_use_next_block(struct compiler *, basicblock *block) set a previously created block as a current block

Once the CFG is created, it must be flattened and then final emission of bytecode occurs. Flattening is handled using a post-order depth-first search. Once flattened, jump offsets are backpatched based on the flattening and then a `PyCodeObject` is created. All of this is handled by calling `assemble()`.

10.25.8 Introducing New Bytecode

Sometimes a new feature requires a new opcode. But adding new bytecode is not as simple as just suddenly introducing new bytecode in the AST->bytecode step of the compiler. Several pieces of code throughout Python depend on having correct information about what bytecode exists.

First, you must choose a name and a unique identifier number. The official list of bytecode can be found in `Lib/opcode.py`. If the opcode is to take an argument, it must be given a unique number greater than that assigned to `HAVE_ARGUMENT` (as found in `Lib/opcode.py`).

Once the name/number pair has been chosen and entered in `Lib/opcode.py`, you must also enter it into `Doc/library/dis.rst`, and regenerate `Include/opcode.h` and `Python/opcode_targets.h` by running `make regen-opcode regen-opcode-targets`.

With a new bytecode you must also change what is called the magic number for `.pyc` files. The variable `MAGIC_NUMBER` in `Lib/importlib/_bootstrap_external.py` contains the number. Changing this number will lead to all `.pyc` files with the old `MAGIC_NUMBER` to be recompiled by the interpreter on import. Whenever `MAGIC_NUMBER` is changed, the ranges in the `magic_values` array in `PC/launcher.c` must also be updated. Changes to `Lib/importlib/_bootstrap_external.py` will take effect only after running `make regen-importlib`. Running this command before adding the new bytecode target to `Python/ceval.c` will result in an error. You should only run `make regen-importlib` after the new bytecode target has been added.

Finally, you need to introduce the use of the new bytecode. Altering `Python/compile.c` and `Python/ceval.c` will be the primary places to change. You must add the case for a new opcode into the 'switch' statement in the `stack_effect()` function in `Python/compile.c`. If the new opcode has a jump target, you will need to update macros and 'switch' statements in `Python/peephole.c`. If it affects a control flow or the block stack, you may have to update the `frame_setlineno()` function in `Objects/frameobject.c`. `Lib/dis.py` may need an update if the new opcode interprets its argument in a special way (like `FORMAT_VALUE` or `MAKE_FUNCTION`).

If you make a change here that can affect the output of bytecode that is already in existence and you do not change the magic number constantly, make sure to delete your old `.py(clo)` files! Even though you will end up changing the magic number if you change the bytecode, while you are debugging your work you will be changing the bytecode output without constantly bumping up the magic number. This means you end up with stale `.pyc` files that will not be recreated. Running `find . -name '*.py[co]' -exec rm -f '{}' +` should delete all `.pyc` files you have, forcing new ones to be created and thus allow you test out your new bytecode properly. Run `make regen-importlib` for updating the bytecode of frozen `importlib` files. You have to run `make` again after this for recompiling generated C files.

10.25.9 Code Objects

The result of `PyAST_CompileObject()` is a `PyCodeObject` which is defined in `Include/code.h`. And with that you now have executable Python bytecode!

The code objects (byte code) are executed in `Python/ceval.c`. This file will also need a new case statement for the new opcode in the big switch statement in `_PyEval_EvalFrameDefault()`.

10.25.10 Important Files

- Parser/

Python.asdl ASDL syntax file

asdl.py Parser for ASDL definition files. Reads in an ASDL description and parses it into an AST that describes it.

asdl_c.py “Generate C code from an ASDL description.” Generates `Python/Python-ast.c` and `Include/Python-ast.h`.

- Python/

Python-ast.c Creates C structs corresponding to the ASDL types. Also contains code for marshalling AST nodes (core ASDL types have marshalling code in `asdl.c`). “File automatically generated by `Parser/asdl_c.py`”. This file must be committed separately after every grammar change is committed since the `__version__` value is set to the latest grammar change revision number.

asdl.c Contains code to handle the ASDL sequence type. Also has code to handle marshalling the core ASDL types, such as number and identifier. Used by `Python-ast.c` for marshalling AST nodes.

ast.c Converts Python's parse tree into the abstract syntax tree.

ast_opt.c Optimizes the AST.

ast_unparse.c Converts the AST expression node back into a string (for string annotations).

ceval.c Executes byte code (aka, eval loop).

compile.c Emits bytecode based on the AST.

symtable.c Generates a symbol table from AST.

peephole.c Optimizes the bytecode.

pyarena.c Implementation of the arena memory manager.

wordcode_helpers.h Helpers for generating bytecode.

opcode_targets.h One of the files that must be modified if `Lib/opcode.py` is.

- Include/

Python-ast.h Contains the actual definitions of the C structs as generated by `Python/Python-ast.c`. “Automatically generated by `Parser/asdl_c.py`”.

asdl.h Header for the corresponding `Python/ast.c`.

ast.h Declares `PyAST_FromNode()` external (from `Python/ast.c`).

code.h Header file for `Objects/codeobject.c`; contains definition of `PyCodeObject`.

symtable.h Header for Python/symtable.c. struct symtable and PySTEntryObject are defined here.

pyarena.h Header file for the corresponding Python/pyarena.c.

opcode.h One of the files that must be modified if Lib/opcode.py is.

- Objects/

codeobject.c Contains PyCodeObject-related code (originally in Python/compile.c).

frameobject.c Contains the frame_setlineno() function which should determine whether it is allowed to make a jump between two points in a bytecode.

- Lib/

opcode.py Master list of bytecode; if this file is modified you must modify several other files accordingly (see “*Introducing New Bytecode*”)

importlib/_bootstrap_external.py Home of the magic number (named MAGIC_NUMBER) for bytecode versioning.

10.25.11 Known Compiler-related Experiments

This section lists known experiments involving the compiler (including bytecode).

Skip Montanaro presented a paper at a Python workshop on a peephole optimizer¹.

Michael Hudson has a non-active SourceForge project named Bytecodehacks² that provides functionality for playing with bytecode directly.

An opcode to combine the functionality of LOAD_ATTR/CALL_FUNCTION was created named CALL_ATTR³. Currently only works for classic classes and for new-style classes rough benchmarking showed an actual slowdown thanks to having to support both classic and new-style classes.

10.25.12 References

10.26 Design of CPython's Garbage Collector

Author Pablo Galindo Salgado

10.26.1 Abstract

The main garbage collection algorithm used by CPython is reference counting. The basic idea is that CPython counts how many different places there are that have a reference to an object. Such a place could be another object, or a global (or static) C variable, or a local variable in some C function. When an object's reference count becomes zero, the object is deallocated. If it contains references to other objects, their reference counts are decremented. Those other objects may be deallocated in turn, if this decrement makes their reference count become zero, and so on. The reference count field can be examined using the `sys.getrefcount` function (notice that the value returned by this function is always 1 more as the function also has a reference to the object when called):

¹ Skip Montanaro's Peephole Optimizer Paper (<https://drive.google.com/open?id=0B2InO7qBBGRXQXIDM3FVdWZxQWc>)

² Bytecodehacks Project (<http://bytecodehacks.sourceforge.net/bch-docs/bch/index.html>)

³ CALL_ATTR opcode (<https://bugs.python.org/issue709744>)

```

>>> x = object()
>>> sys.getrefcount(x)
2
>>> y = x
>>> sys.getrefcount(x)
3
>>> del y
>>> sys.getrefcount(x)
2

```

The main problem with the reference counting scheme is that it does not handle reference cycles. For instance, consider this code:

```

>>> container = []
>>> container.append(container)
>>> sys.getrefcount(container)
3
>>> del container

```

In this example, `container` holds a reference to itself, so even when we remove our reference to it (the variable “`container`”) the reference count never falls to 0 because it still has its own internal reference. Therefore it would never be cleaned just by simple reference counting. For this reason some additional machinery is needed to clean these reference cycles between objects once they become unreachable. This is the cyclic garbage collector, usually called just Garbage Collector (GC), even though reference counting is also a form of garbage collection.

10.26.2 Memory layout and object structure

Normally the C structure supporting a regular Python object looks as follows:

```

object ----> +-----+-----+-----+-----+-----+-----+ \
              |                                     | |
              |               ob_refcnt             | |
              +-----+-----+-----+-----+-----+ | PyObject_HEAD
              |                                     | |
              |               *ob_type              | |
              +-----+-----+-----+-----+-----+ /
              |                                     |
              |               ...                    |

```

In order to support the garbage collector, the memory layout of objects is altered to accommodate extra information **before** the normal layout:

```

              +-----+-----+-----+-----+-----+ \
              |               *_gc_next              | |
              +-----+-----+-----+-----+-----+ | PyGC_Head
              |               *_gc_prev              | |
object ----> +-----+-----+-----+-----+-----+ /
              |                                     | \
              |               ob_refcnt             | |
              +-----+-----+-----+-----+-----+ | PyObject_HEAD
              |                                     | |
              |               *ob_type              | |
              +-----+-----+-----+-----+-----+ /
              |                                     |
              |               ...                    |

```

In this way the object can be treated as a normal python object and when the extra information associated to the GC is needed the previous fields can be accessed by a simple type cast from the original object: `((PyGC_Head *) (the_object)-1)`.

As is explained later in the *Optimization: reusing fields to save memory* section, these two extra fields are normally used to keep doubly linked lists of all the objects tracked by the garbage collector (these lists are the GC generations,

more on that in the *Optimization: generations* section), but they are also reused to fulfill other purposes when the full doubly linked list structure is not needed as a memory optimization.

Doubly linked lists are used because they efficiently support most frequently required operations. In general, the collection of all objects tracked by GC are partitioned into disjoint sets, each in its own doubly linked list. Between collections, objects are partitioned into “generations”, reflecting how often they’ve survived collection attempts. During collections, the generation(s) being collected are further partitioned into, e.g., sets of reachable and unreachable objects. Doubly linked lists support moving an object from one partition to another, adding a new object, removing an object entirely (objects tracked by GC are most often reclaimed by the refcounting system when GC isn’t running at all!), and merging partitions, all with a small constant number of pointer updates. With care, they also support iterating over a partition while objects are being added to - and removed from - it, which is frequently required while GC is running.

Specific APIs are offered to allocate, deallocate, initialize, track, and untrack objects with GC support. These APIs can be found in the [Garbage Collector C API documentation](#).

Apart from this object structure, the type object for objects supporting garbage collection must include the `Py_TPFLAGS_HAVE_GC` in its `tp_flags` slot and provide an implementation of the `tp_traverse` handler. Unless it can be proven that the objects cannot form reference cycles with only objects of its type or unless the type is immutable, a `tp_clear` implementation must also be provided.

10.26.3 Identifying reference cycles

The algorithm that CPython uses to detect those reference cycles is implemented in the `gc` module. The garbage collector **only focuses** on cleaning container objects (i.e. objects that can contain a reference to one or more objects). These can be arrays, dictionaries, lists, custom class instances, classes in extension modules, etc. One could think that cycles are uncommon but the truth is that many internal references needed by the interpreter create cycles everywhere. Some notable examples:

- Exceptions contain traceback objects that contain a list of frames that contain the exception itself.
- Module-level functions reference the module’s dict (which is needed to resolve globals), which in turn contains entries for the module-level functions.
- Instances have references to their class which itself references its module, and the module contains references to everything that is inside (and maybe other modules) and this can lead back to the original instance.
- When representing data structures like graphs, it is very typical for them to have internal links to themselves.

To correctly dispose of these objects once they become unreachable, they need to be identified first. Inside the function that identifies cycles, two double-linked lists are maintained: one list contains all objects to be scanned, and the other will contain all objects “tentatively” unreachable.

To understand how the algorithm works, Let’s take the case of a circular linked list which has one link referenced by a variable A, and one self-referencing object which is completely unreachable:

```
>>> import gc

>>> class Link:
...     def __init__(self, next_link=None):
...         self.next_link = next_link

>>> link_3 = Link()
>>> link_2 = Link(link_3)
>>> link_1 = Link(link_2)
>>> link_3.next_link = link_1
>>> A = link_1
>>> del link_1, link_2, link_3
```

(continues on next page)

(continued from previous page)

```

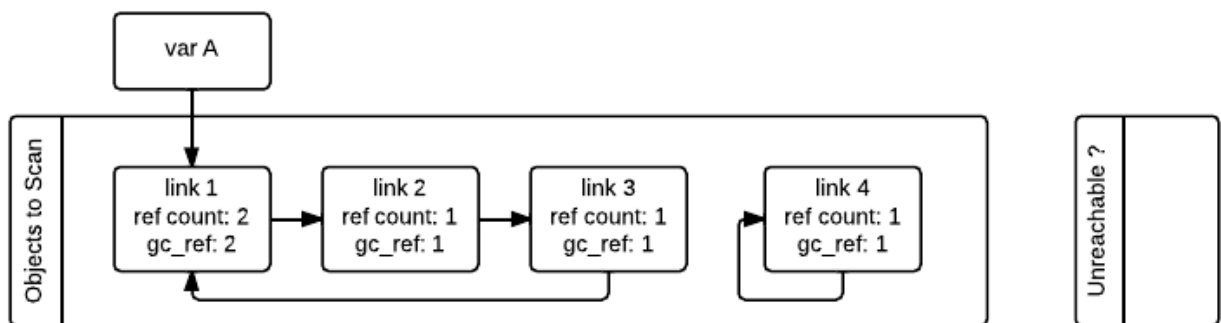
>>> link_4 = Link()
>>> link_4.next_link = link_4

>>> gc.collect()
2

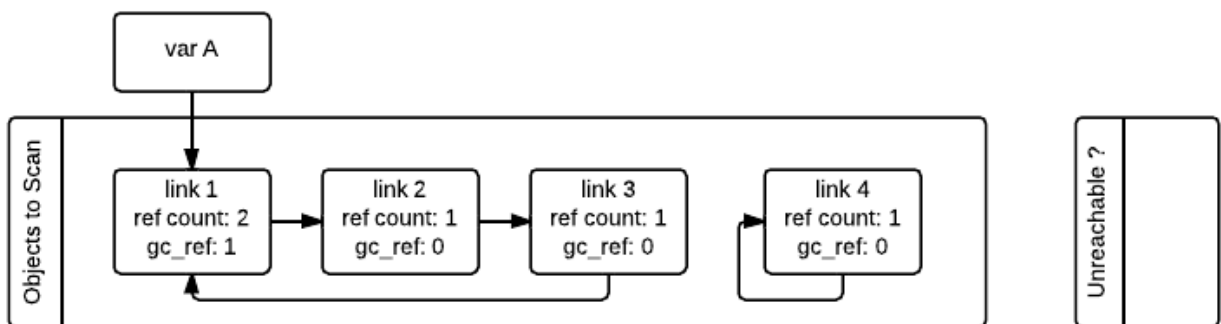
```

When the GC starts, it has all the container objects it wants to scan on the first linked list. The objective is to move all the unreachable objects. Since most objects turn out to be reachable, it is much more efficient to move the unreachable as this involves fewer pointer updates.

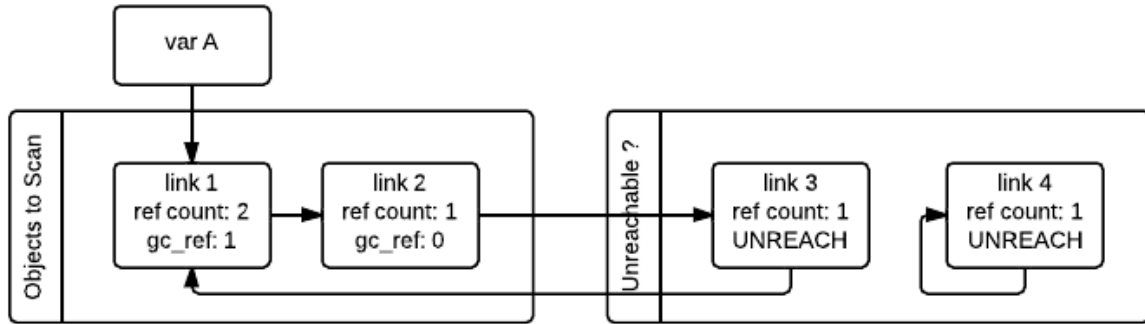
Every object that supports garbage collection will have an extra reference count field initialized to the reference count (`gc_ref` in the figures) of that object when the algorithm starts. This is because the algorithm needs to modify the reference count to do the computations and in this way the interpreter will not modify the real reference count field.



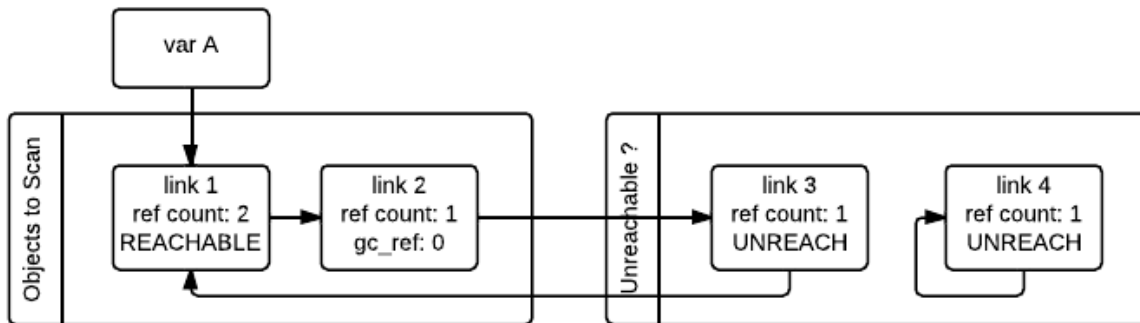
The GC then iterates over all containers in the first list and decrements by one the `gc_ref` field of any other object that container is referencing. Doing this makes use of the `tp_traverse` slot in the container class (implemented using the C API or inherited by a superclass) to know what objects are referenced by each container. After all the objects have been scanned, only the objects that have references from outside the “objects to scan” list will have `gc_refs > 0`.



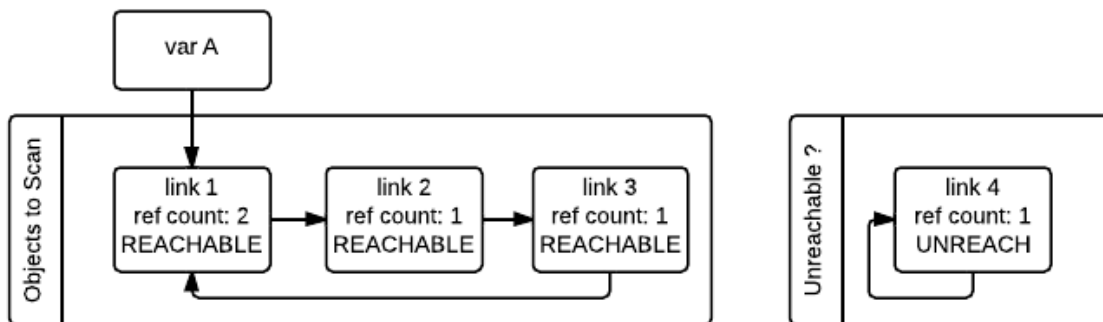
Notice that having `gc_refs == 0` does not imply that the object is unreachable. This is because another object that is reachable from the outside (`gc_refs > 0`) can still have references to it. For instance, the `link_2` object in our example ended having `gc_refs == 0` but is referenced still by the `link_1` object that is reachable from the outside. To obtain the set of objects that are really unreachable, the garbage collector re-scans the container objects using the `tp_traverse` slot; this time with a different traverse function that marks objects with `gc_refs == 0` as “tentatively unreachable” and then moves them to the tentatively unreachable list. The following image depicts the state of the lists in a moment when the GC processed the `link_3` and `link_4` objects but has not processed `link_1` and `link_2` yet.



Then the GC scans the next `link_1` object. Because it has `gc_refs == 1`, the gc does not do anything special because it knows it has to be reachable (and is already in what will become the reachable list):



When the GC encounters an object which is reachable (`gc_refs > 0`), it traverses its references using the `tp_traverse` slot to find all the objects that are reachable from it, moving them to the end of the list of reachable objects (where they started originally) and setting its `gc_refs` field to 1. This is what happens to `link_2` and `link_3` below as they are reachable from `link_1`. From the state in the previous image and after examining the objects referred to by `link_1` the GC knows that `link_3` is reachable after all, so it is moved back to the original list and its `gc_refs` field is set to one so if the GC visits it again, it does know that is reachable. To avoid visiting a object twice, the GC marks all objects that are already visited once (by unsetting the `PREV_MASK_COLLECTING` flag) so if an object that has already been processed is referred by some other object, the GC does not process it twice.



Notice that once an object that was marked as “tentatively unreachable” and later is moved back to the reachable list, it will be visited again by the garbage collector as now all the references that that objects has need to be processed as well. This process is really a breadth first search over the object graph. Once all the objects are scanned, the GC knows that all container objects in the tentatively unreachable list are really unreachable and can thus be garbage collected.

Pragmatically, it's important to note that no recursion is required by any of this, and neither does it in any other way require additional memory proportional to the number of objects, number of pointers, or the lengths of pointer chains. Apart from $O(1)$ storage for internal C needs, the objects themselves contain all the storage the GC algorithms require.

Why moving unreachable objects is better

It sounds logical to move the unreachable objects under the premise that most objects are usually reachable, until you think about it: the reason it pays isn't actually obvious.

Suppose we create objects A, B, C in that order. They appear in the young generation in the same order. If B points to A, and C to B, and C is reachable from outside, then the adjusted refcounts after the first step of the algorithm runs will be 0, 0, and 1 respectively because the only reachable object from the outside is C.

When the next step of the algorithm finds A, A is moved to the unreachable list. The same for B when it's first encountered. Then C is traversed, B is moved *back* to the reachable list. B is eventually traversed, and then A is moved back to the reachable list.

So instead of not moving at all, the reachable objects B and A are each moved twice. Why is this a win? A straightforward algorithm to move the reachable objects instead would move A, B, and C once each. The key is that this dance leaves the objects in order C, B, A - it's reversed from the original order. On all *subsequent* scans, none of them will move. Since most objects aren't in cycles, this can save an unbounded number of moves across an unbounded number of later collections. The only time the cost can be higher is the first time the chain is scanned.

10.26.4 Destroying unreachable objects

Once the GC knows the list of unreachable objects, a very delicate process starts with the objective of completely destroying these objects. Roughly, the process follows these steps in order:

1. Handle and clean weak references (if any). If an object that is in the unreachable set is going to be destroyed and has weak references with callbacks, these callbacks need to be honored. This process is **very** delicate as any error can cause objects that will be in an inconsistent state to be resurrected or reached by some python functions invoked from the callbacks. In addition, weak references that also are part of the unreachable set (the object and its weak reference are in cycles that are unreachable) need to be cleaned immediately, without executing the callback. Otherwise it will be triggered later, when the `tp_clear` slot is called, causing havoc. Ignoring the weak reference's callback is fine because both the object and the weakref are going away, so it's legitimate to say the weak reference is going away first.
2. If an object has legacy finalizers (`tp_del` slot) move them to the `gc.garbage` list.
3. Call the finalizers (`tp_finalize` slot) and mark the objects as already finalized to avoid calling them twice if they resurrect or if other finalizers have removed the object first.
4. Deal with resurrected objects. If some objects have been resurrected, the GC finds the new subset of objects that are still unreachable by running the cycle detection algorithm again and continues with them.
5. Call the `tp_clear` slot of every object so all internal links are broken and the reference counts fall to 0, triggering the destruction of all unreachable objects.

10.26.5 Optimization: generations

In order to limit the time each garbage collection takes, the GC uses a popular optimization: generations. The main idea behind this concept is the assumption that most objects have a very short lifespan and can thus be collected shortly after their creation. This has proven to be very close to the reality of many Python programs as many temporarily objects are created and destroyed very fast. The older an object is the less likely it is to become unreachable.

To take advantage of this fact, all container objects are segregated across three spaces/generations. Every new object starts in the first generation (generation 0). The previous algorithm is executed only over the objects of a particular generation and if an object survives a collection of its generation it will be moved to the next one (generation 1), where it will be surveyed for collection less often. If the same object survives another GC round in this new generation (generation 1) it will be moved to the last generation (generation 2) where it will be surveyed the least often.

Generations are collected when the number of objects that they contain reach some predefined threshold, which is unique for each generation and is lower the older the generations are. These thresholds can be examined using the `gc.get_threshold` function:

```
>>> import gc
>>> gc.get_threshold()
(700, 10, 10)
```

The content of these generations can be examined using the `gc.get_objects (generation=NUM)` function and collections can be triggered specifically in a generation by calling `gc.collect (generation=NUM)`.

```
>>> import gc
>>> class MyObj:
...     pass
...

# Move everything to the last generation so its easier to inspect
# the younger generations.

>>> gc.collect()
0

# Create a reference cycle

>>> x = MyObj()
>>> x.self = x

# Initially the object is in the youngest generation.

>>> gc.get_objects(generation=0)
[... , <__main__.MyObj object at 0x7fbcc12a3400>, ...]

# After a collection of the youngest generation the object
# moves to the next generation.

>>> gc.collect(generation=0)
0
>>> gc.get_objects(generation=0)
[]
>>> gc.get_objects(generation=1)
[... , <__main__.MyObj object at 0x7fbcc12a3400>, ...]
```

Collecting the oldest generation

In addition to the various configurable thresholds, the GC only triggers a full collection of the oldest generation if the ratio `long_lived_pending / long_lived_total` is above a given value (hardwired to 25%). The reason is that, while “non-full” collections (i.e., collections of the young and middle generations) will always examine roughly the same number of objects (determined by the aforementioned thresholds) the cost of a full collection is proportional to the total number of long-lived objects, which is virtually unbounded. Indeed, it has been remarked that doing a full collection every <constant number> of object creations entails a dramatic performance degradation

in workloads which consist of creating and storing lots of long-lived objects (e.g. building a large list of GC-tracked objects would show quadratic performance, instead of linear as expected). Using the above ratio, instead, yields amortized linear performance in the total number of objects (the effect of which can be summarized thusly: “each full garbage collection is more and more costly as the number of objects grows, but we do fewer and fewer of them”).

10.26.6 Optimization: reusing fields to save memory

In order to save memory, the two linked list pointers in every object with GC support are reused for several purposes. This is a common optimization known as “fat pointers” or “tagged pointers”: pointers that carry additional data, “folded” into the pointer, meaning stored inline in the data representing the address, taking advantage of certain properties of memory addressing. This is possible as most architectures align certain types of data to the size of the data, often a word or multiple thereof. This discrepancy leaves a few of the least significant bits of the pointer unused, which can be used for tags or to keep other information – most often as a bit field (each bit a separate tag) – as long as code that uses the pointer masks out these bits before accessing memory. E.g., on a 32-bit architecture (for both addresses and word size), a word is 32 bits = 4 bytes, so word-aligned addresses are always a multiple of 4, hence end in 00, leaving the last 2 bits available; while on a 64-bit architecture, a word is 64 bits word = 8 bytes, so word-aligned addresses end in 000, leaving the last 3 bits available.

The CPython GC makes use of two fat pointers that corresponds to the extra fields of `PyGC_Head` discussed in the *Memory layout and object structure* section:

Warning: Because the presence of extra information, “tagged” or “fat” pointers cannot be dereferenced directly and the extra information must be stripped off before obtaining the real memory address. Special care needs to be taken with functions that directly manipulate the linked lists, as these functions normally assume the pointers inside the lists are in a consistent state.

- The `_gc_prev`` field is normally used as the “previous” pointer to maintain the doubly linked list but its lowest two bits are used to keep the flags `PREV_MASK_COLLECTING` and `_PyGC_PREV_MASK_FINALIZED`. Between collections, the only flag that can be present is `_PyGC_PREV_MASK_FINALIZED` that indicates if an object has been already finalized. During collections `_gc_prev` is temporarily used for storing a copy of the reference count (`gc_refs`), in addition to two flags, and the GC linked list becomes a singly linked list until `_gc_prev` is restored.
- The `_gc_next` field is used as the “next” pointer to maintain the doubly linked list but during collection its lowest bit is used to keep the `NEXT_MASK_UNREACHABLE` flag that indicates if an object is tentatively unreachable during the cycle detection algorithm. This is a drawback to using only doubly linked lists to implement partitions: while most needed operations are constant-time, there is no efficient way to determine which partition an object is currently in. Instead, when that’s needed, ad hoc tricks (like the `NEXT_MASK_UNREACHABLE` flag) are employed.

10.26.7 Optimization: delay tracking containers

Certain types of containers cannot participate in a reference cycle, and so do not need to be tracked by the garbage collector. Untracking these objects reduces the cost of garbage collections. However, determining which objects may be untracked is not free, and the costs must be weighed against the benefits for garbage collection. There are two possible strategies for when to untrack a container:

1. When the container is created.
2. When the container is examined by the garbage collector.

As a general rule, instances of atomic types aren’t tracked and instances of non-atomic types (containers, user-defined objects...) are. However, some type-specific optimizations can be present in order to suppress the garbage collector footprint of simple instances. Some examples of native types that benefit from delayed tracking:

- Tuples containing only immutable objects (integers, strings etc, and recursively, tuples of immutable objects) do not need to be tracked. The interpreter creates a large number of tuples, many of which will not survive until garbage collection. It is therefore not worthwhile to untrack eligible tuples at creation time. Instead, all tuples except the empty tuple are tracked when created. During garbage collection it is determined whether any surviving tuples can be untracked. A tuple can be untracked if all of its contents are already not tracked. Tuples are examined for untracking in all garbage collection cycles. It may take more than one cycle to untrack a tuple.
- Dictionaries containing only immutable objects also do not need to be tracked. Dictionaries are untracked when created. If a tracked item is inserted into a dictionary (either as a key or value), the dictionary becomes tracked. During a full garbage collection (all generations), the collector will untrack any dictionaries whose contents are not tracked.

The garbage collector module provides the python function `is_tracked(obj)`, which returns the current tracking status of the object. Subsequent garbage collections may change the tracking status of the object.

```
>>> gc.is_tracked(0)
False
>>> gc.is_tracked("a")
False
>>> gc.is_tracked([])
True
>>> gc.is_tracked({})
False
>>> gc.is_tracked({"a": 1})
False
>>> gc.is_tracked({"a": []})
True
```

10.27 Updating standard library extension modules

In this section, we could explain how to write a CPython extension with the C language, but the topic can take a complete book.

For this reason, we prefer to give you some links where you can read a very good documentation.

Read the following references:

- <https://docs.python.org/dev/c-api/>
- <https://docs.python.org/dev/extending/>
- <https://www.python.org/dev/peps/pep-0399/>
- <https://pythonextensionpatterns.readthedocs.io/en/latest/>

10.28 Coverity Scan

Coverity Scan is a free service for static code analysis of Open Source projects. It is based on Coverity's commercial product and is able to analyze C, C++ and Java code.

Coverity's static code analysis doesn't run the code. Instead of that it uses abstract interpretation to gain information about the code's control flow and data flow. It's able to follow all possible code paths that a program may take. For example the analyzer understands that `malloc()` returns a memory that must be freed with `free()` later. It follows all branches and function calls to see if all possible combinations free the memory. The analyzer is able to detect all sorts of issues like resource leaks (memory, file descriptors), NULL dereferencing, use after free, unchecked return values, dead code, buffer overflows, integer overflows, uninitialized variables, and many more.

10.28.1 Access to analysis reports

The results are available on the [Coverity Scan](#) website. In order to access the results you have to create an account yourself. Then go to *Projects using Scan* and add yourself to the Python project. New members must be approved by an admin (see *Contact*).

Access is restricted to Python core developers only. Other individuals may be given access at our own discretion, too. Every now and then Coverity detects a critical issue in Python's code – new analyzers may even find new bugs in mature code. We don't want to disclose issues prematurely.

10.28.2 Building and uploading analysis

The process is automated. A script checks out the code, runs `cov-build` and uploads the latest analysis to Coverity. Since Coverity has limited the maximum number of builds per week Python is analyzed every second day. The build runs on a dedicated virtual machine on PSF's infrastructure at OSU Open Source Labs. The process is maintained by Christian Heimes (see *Contact*). At present only the tip is analyzed with the 64bit Linux tools.

10.28.3 Known limitations

Some aspects of Python's C code are not yet understood by Coverity.

False positives

Py_BuildValue("N", PyObject*) Coverity doesn't understand that N format char passes the object along without touching its reference count. On this ground the analyzer detects a resource leak. CID 719685

PyLong_FromLong() for negative values Coverity claims that `PyLong_FromLong()` and other `PyLong_From*`() functions cannot handle a negative value because the value might be used as an array index in `get_small_int()`. CID 486783

PyLong_FromLong() for n in [-5 ... +255] For integers in the range of Python's small int cache the `PyLong_From*`() function can never fail and never returns NULL. CID 1058291

PyArg_ParseTupleAndKeywords(args, kwargs, "s#", &data, &length) Some functions use the format char combination such as `s#`, `u#` or `z#` to get data and length of a character array. Coverity doesn't recognize the relation between data and length. Sometimes it detects a buffer overflow if data is written to a fixed size buffer although `length <= sizeof(buffer)`. CID 486613

path_converter() dereferencing after null check The `path_converter()` function in `posixmodule.c` makes sure that either `path_t.narrow` or `path_t.wide` is filled unless `path_t.nullable` is explicitly enabled. CID 719648

Intentionally

Py_VA_COPY() Python is written in C89 (ANSI C), therefore it can't use C99 features such as `va_copy()`. Python's own variant `Py_VA_COPY()` uses `memcpy()` to make a copy of a `va_list` variable. Coverity detects two issues in this approach: "Passing argument "lva" of type "va_list" and sizeof(va_list) to function memcpy() is suspicious." CID 486405 and "Uninitialized pointer read" CID 486630.

10.28.4 Modeling

Modeling is explained in the *Coverity Help Center* which is available in the help menu of *Coverity Connect*. `coverity_model.c` contains a copy of Python's modeling file for Coverity. Please keep the copy in sync with the model file in *Analysis Settings* of *Coverity Scan*.

10.28.5 Workflow

False positive and intentional issues

If the problem is listed under *Known limitations* then please set the classification to either "False positive" or "Intentional", the action to "Ignore", owner to your own account and add a comment why the issue is considered false positive or intentional.

If you think it's a new false positive or intentional then please contact an admin. The first step should be an updated to Python's *Modeling* file.

Positive issues

You should always create an issue unless it's really a trivial case. Please add the full url to the ticket under *Ext. Reference* and add the CID (Coverity ID) to both the ticket and the checkin message. It makes it much easier to understand the relation between tickets, fixes and Coverity issues.

10.28.6 Contact

Please include both Brett and Christian in any mail regarding Coverity. Mails to Coverity should go through Brett or Christian, too.

Christian Heimes <[christian \(at\) python \(dot\) org](mailto:christian@python.org)> admin, maintainer of build machine, intermediary between Python and Coverity

Brett Cannon <[brett \(at\) python \(dot\) org](mailto:brett@python.org)> co-admin

Dakshesh Vyas <scan-admin@coverity.com> Technical Manager - Coverity Scan

See also:

[Coverity Scan FAQ](#)

10.29 Dynamic Analysis with Clang

This document describes how to use Clang to perform analysis on Python and its libraries. In addition to performing the analysis, the document will cover downloading, building and installing the latest Clang/LLVM combination (which is currently 3.4).

This document does not cover interpreting the findings. For a discussion of interpreting results, see Marshall Clow's [Testing libc++ with -fsanitize=undefined](#). The blog posting is a detailed examinations of issues uncovered by Clang in `libc++`.

10.29.1 What is Clang?

Clang is the C, C++ and Objective C front-end for the LLVM compiler. The front-end provides access to LLVM's optimizer and code generator. The sanitizers - or checkers - are hooks into the code generation phase to instrument compiled code so suspicious behavior is flagged.

10.29.2 What are Sanitizers?

Clang sanitizers are runtime checkers used to identify suspicious and undefined behavior. The checking occurs at runtime with actual runtime parameters so false positives are kept to a minimum.

There are a number of sanitizers available, but two that should be used on a regular basis are the Address Sanitizer (or ASan) and the Undefined Behavior Sanitizer (or UBSan). ASan is invoked with the compiler option `-fsanitize=address`, and UBSan is invoked with `-fsanitize=undefined`. The flags are passed through `CFLAGS` and `CXXFLAGS`, and sometimes through `CC` and `CXX` (in addition to the compiler).

A complete list of sanitizers can be found at [Controlling Code Generation](#).

Note: Because sanitizers operate at runtime on real program parameters, its important to provide a complete set of positive and negative self tests.

Clang and its sanitizers have strengths (and weaknesses). Its just one tool in the war chest to uncovering bugs and improving code quality. Clang should be used to compliment other methods, including Code Reviews, Valgrind, Coverity, etc.

10.29.3 Clang/LLVM Setup

This portion of the document covers downloading, building and installing Clang and LLVM. There are three components to download and build. They are the LLVM compiler, the compiler front end and the compiler runtime library.

In preparation you should create a scratch directory. Also ensure you are using Python 2 and not Python 3. Python 3 will cause the build to fail.

Download, Build and Install

Perform the following to download, build and install the Clang/LLVM 3.4.

```
# Download
wget https://llvm.org/releases/3.4/llvm-3.4.src.tar.gz
wget https://llvm.org/releases/3.4/clang-3.4.src.tar.gz
wget https://llvm.org/releases/3.4/compiler-rt-3.4.src.tar.gz

# LLVM
tar xvf llvm-3.4.src.tar.gz
cd llvm-3.4/tools

# Clang Front End
tar xvf ../../clang-3.4.src.tar.gz
mv clang-3.4 clang

# Compiler RT
cd ../projects
tar xvf ../../compiler-rt-3.4.src.tar.gz
```

(continues on next page)

(continued from previous page)

```
mv compiler-rt-3.4/ compiler-rt

# Build
cd ..
./configure --enable-optimized --prefix=/usr/local
make -j4
sudo make install
```

Note: If you receive an error 'LibraryDependencies.inc' file not found, then ensure you are utilizing Python 2 and not Python 3. If you encounter the error after switching to Python 2, then delete everything and start over.

After `make install` executes, the compilers will be installed in `/usr/local/bin` and the various libraries will be installed in `/usr/local/lib/clang/3.4/lib/linux/`:

```
$ ls /usr/local/lib/clang/3.4/lib/linux/
libclang_rt.asan-x86_64.a   libclang_rt.profile-x86_64.a
libclang_rt.dfsan-x86_64.a libclang_rt.san-x86_64.a
libclang_rt.full-x86_64.a  libclang_rt.tsan-x86_64.a
libclang_rt.lsan-x86_64.a  libclang_rt.ubsan_cxx-x86_64.a
libclang_rt.msan-x86_64.a  libclang_rt.ubsan-x86_64.a
```

On Mac OS X, the libraries are installed in `/usr/local/lib/clang/3.3/lib/darwin/`:

```
$ ls /usr/local/lib/clang/3.3/lib/darwin/
libclang_rt.10.4.a           libclang_rt.ios.a
libclang_rt.asan_osx.a       libclang_rt.osx.a
libclang_rt.asan_osx_dynamic.dylib libclang_rt.profile_ios.a
libclang_rt.cc_kext.a        libclang_rt.profile_osx.a
libclang_rt.cc_kext_ios5.a   libclang_rt.ubsan_osx.a
libclang_rt.eprintf.a
```

Note: You should never have to add the libraries to a project. Clang will handle it for you. If you find you cannot pass the `-fsanitize=XXX` flag through `make`'s implicit variables (`CFLAGS`, `CXXFLAGS`, `CC`, `CXXFLAGS`, `LDFLAGS`) during `configure`, then you should modify the `makefile` after configuring to ensure the flag is passed through the compiler.

The installer does not install all the components needed on occasion. For example, you might want to run a `scan-build` or examine the results with `scan-view`. You can copy the components by hand with:

```
sudo mkdir /usr/local/bin/scan-build
sudo cp -r llvm-3.4/tools/clang/tools/scan-build /usr/local/bin
sudo mkdir /usr/local/bin/scan-view
sudo cp -r llvm-3.4/tools/clang/tools/scan-view /usr/local/bin
```

Note: Because the installer does not install all the components needed on occasion, you should not delete the scratch directory until you are sure things work as expected. If a library is missing, then you should search for it in the Clang/LLVM build directory.

10.29.4 Python Build Setup

This portion of the document covers invoking Clang and LLVM with the options required so the sanitizers analyze Python with under its test suite. Two checkers are used - ASan and UBSan.

Because the sanitizers are runtime checkers, its best to have as many positive and negative self tests as possible. You can never have enough self tests.

The general idea is to compile and link with the sanitizer flags. At link time, Clang will include the needed runtime libraries. However, you can't use CFLAGS and CXXFLAGS to pass the options through the compiler to the linker because the makefile rules for BUILDPYTHON, _testembed and _freeze_importlib don't use the implicit variables.

As a workaround to the absence of flags to the linker, you can pass the sanitizer options by way of the compilers - CC and CXX. Passing the flags though the compiler is used below, but passing them through LDFLAGS is also reported to work.

Building Python

To begin, export the variables of interest with the desired sanitizers. Its OK to specify both sanitizers:

```
# ASan
export CC="/usr/local/bin/clang -fsanitize=address"
export CXX="/usr/local/bin/clang++ -fsanitize=address -fno-sanitize=vptr"
```

Or:

```
# UBSan
export CC="/usr/local/bin/clang -fsanitize=undefined"
export CXX="/usr/local/bin/clang++ -fsanitize=undefined -fno-sanitize=vptr"
```

The `-fno-sanitize=vptr` removes vtable checks that are part of UBSan from C++ projects due to noise. Its not needed with Python, but you will likely need it for other C++ projects.

After exporting CC and CXX, configure as normal:

```
$ ./configure
checking build system type... x86_64-unknown-linux-gnu
checking host system type... x86_64-unknown-linux-gnu
checking for --enable-universalsdk... no
checking for --with-universal-archs... 32-bit
checking MACHDEP... linux
checking for --without-gcc... no
checking for gcc... /usr/local/bin/clang -fsanitize=undefined
checking whether the C compiler works... yes
...
```

Next is a standard make (formatting added for clarity):

```
$ make
/usr/local/bin/clang -fsanitize=undefined -c -Wno-unused-result
-DNDEBUG -g -fwrapv -O3 -Wall -Wstrict-prototypes -I.
-IInclude -I./Include -DPy_BUILD_CORE -o Modules/python.o
./Modules/python.c
/usr/local/bin/clang -fsanitize=undefined -c -Wno-unused-result
-DNDEBUG -g -fwrapv -O3 -Wall -Wstrict-prototypes -I.
-IInclude -I./Include -DPy_BUILD_CORE -o Parser/acceler.o
```

(continues on next page)

(continued from previous page)

```
Parser/acceler.c
...
```

Finally is make test (formatting added for clarity):

```
Objects/longobject.c:39:42: runtime error: index -1 out of bounds
    for type 'PyLongObject [262]'
Objects/tupleobject.c:188:13: runtime error: member access within
    misaligned address 0x2b76be018078 for type 'PyGC_Head' (aka
    'union_gc_head'), which requires 16 byte alignment
    0x2b76be018078: note: pointer points here
    00 00 00 00  40 53 5a b6 76 2b 00 00  60 52 5a b6 ...
                   ^
...
```

If you are using the address sanitizer, its important to pipe the output through `asan_symbolize.py` to get a good trace. For example, from Issue 20953 during compile (formatting added for clarity):

```
$ make test 2>&1 | asan_symbolize.py
...

/usr/local/bin/clang -fsanitize=address -Xlinker -export-dynamic
-o python Modules/python.o libpython3.3m.a -ldl -lutil
/usr/local/ssl/lib/libssl.a /usr/local/ssl/lib/libcrypto.a -lm
./python -E -S -m sysconfig --generate-posix-vars
=====
==24064==ERROR: AddressSanitizer: heap-buffer-overflow on address
0x619000004020 at pc 0x4ed4b2 bp 0x7fff80fff010 sp 0x7fff80fff008
READ of size 4 at 0x619000004020 thread T0
#0 0x4ed4b1 in PyObject_Free Python-3.3.5/./Objects/obmalloc.c:987
#1 0x7a2141 in code_dealloc Python-3.3.5/./Objects/codeobject.c:359
#2 0x620c00 in PyImport_ImportFrozenModuleObject
    Python-3.3.5/./Python/import.c:1098
#3 0x620d5c in PyImport_ImportFrozenModule
    Python-3.3.5/./Python/import.c:1114
#4 0x63fd07 in import_init Python-3.3.5/./Python/pythonrun.c:206
#5 0x63f636 in _Py_InitializeEx_Private
    Python-3.3.5/./Python/pythonrun.c:369
#6 0x681d77 in Py_Main Python-3.3.5/./Modules/main.c:648
#7 0x4e6894 in main Python-3.3.5/./Modules/python.c:62
#8 0x2abf9a525eac in __libc_start_main
    /home/aurel32/eglibc/eglibc-2.13/csu/libc-start.c:244
#9 0x4e664c in _start (Python-3.3.5/./python+0x4e664c)

AddressSanitizer can not describe address in more detail (wild
memory access suspected).
SUMMARY: AddressSanitizer: heap-buffer-overflow
    Python-3.3.5/./Objects/obmalloc.c:987 PyObject_Free
Shadow bytes around the buggy address:
 0x0c327fff87b0: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
 0x0c327fff87c0: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
 0x0c327fff87d0: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
 0x0c327fff87e0: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
 0x0c327fff87f0: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
=>0x0c327fff8800: fa fa fa fa[fa]fa fa fa fa fa fa fa fa fa fa fa
 0x0c327fff8810: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
```

(continues on next page)

(continued from previous page)

```

0x0c327fff8820: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
0x0c327fff8830: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
0x0c327fff8840: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
0x0c327fff8850: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
Shadow byte legend (one shadow byte represents 8 application bytes):
Addressable:          00
Partially addressable: 01 02 03 04 05 06 07
Heap left redzone:    fa
Heap right redzone:   fb
Freed heap region:    fd
Stack left redzone:   f1
Stack mid redzone:    f2
Stack right redzone:  f3
Stack partial redzone: f4
Stack after return:   f5
Stack use after scope: f8
Global redzone:       f9
Global init order:    f6
Poisoned by user:     f7
ASan internal:        fe
==24064==ABORTING
make: *** [pybuilddir.txt] Error 1

```

Note: `asan_symbolize.py` is supposed to be installed during `make install`. If its not installed, then look in the Clang/LLVM build directory for it and copy it to `/usr/local/bin`.

Blacklisting (Ignoring) Findings

Clang allows you to alter the behavior of sanitizer tools for certain source-level by providing a special blacklist file at compile-time. The blacklist is needed because it reports every instance of an issue, even if the issue is reported 10's of thousands of time in un-managed library code.

You specify the blacklist with `-fsanitize-blacklist=XXX`. For example:

```
-fsanitize-blacklist=my_blacklist.txt
```

`my_blacklist.txt` would then contain entries such as the following. The entry will ignore a bug in `libc++`'s `ios` formatting functions:

```
fun: _Ios_Fmtflags
```

As an example with Python 3.4.0, `audioop.c` will produce a number of findings:

```

./Modules/audioop.c:422:11: runtime error: left shift of negative value -1
./Modules/audioop.c:446:19: runtime error: left shift of negative value -1
./Modules/audioop.c:476:19: runtime error: left shift of negative value -1
./Modules/audioop.c:504:16: runtime error: left shift of negative value -1
./Modules/audioop.c:533:22: runtime error: left shift of negative value -128
./Modules/audioop.c:775:19: runtime error: left shift of negative value -70
./Modules/audioop.c:831:19: runtime error: left shift of negative value -70
./Modules/audioop.c:881:19: runtime error: left shift of negative value -1
./Modules/audioop.c:920:22: runtime error: left shift of negative value -70
./Modules/audioop.c:967:23: runtime error: left shift of negative value -70

```

(continues on next page)

(continued from previous page)

```
./Modules/audioop.c:968:23: runtime error: left shift of negative value -70
...
```

One of the function of interest is `audioop_getsample_impl` (flagged at line 422), and the blacklist entry would include:

```
fun:audioop_getsample_imp
```

Or, you could ignore the entire file with:

```
src:Modules/audioop.c
```

Unfortunately, you won't know what to blacklist until you run the sanitizer.

The documentation is available at [Sanitizer special case list](#).

10.30 Running a buildbot worker

Python's *Continuous Integration* system was discussed earlier. We sometimes refer to the collection of *build workers* as our "buildbot fleet". The machines that comprise the fleet are voluntarily contributed resources. Many are run by individual volunteers out of their own pockets and time, while others are supported by corporations. Even the corporate sponsored buildbots, however, tend to exist because some individual championed them, made them a reality, and is committed to maintaining them.

Anyone can contribute a buildbot to the fleet. This chapter describes how to go about setting up a buildbot worker, getting it added, and some hints about buildbot maintenance.

Anyone running a buildbot that is part of the fleet should subscribe to the [python-buildbots](#) mailing list. This mailing list is also the place to contact if you want to contribute a buildbot but have questions.

As for what kind of buildbot to run... take a look at our [current fleet](#). Pretty much anything that isn't on that list would be interesting: different Linux/UNIX distributions, different versions of the various OSes, other OSes if you or someone are prepared to make the test suite actually pass on that new OS. Even if you only want to run an OS that's already on our list there may be utility in setting it up; we also need to build and test python under various alternate build configurations. Post to the mailing list and talk about what you'd like to contribute.

10.30.1 Preparing for buildbot worker setup

Since the goal is to build Python from source, the system will need to have everything required to do normal python development: a compiler, a linker, and (except on windows) the "development" headers for any of the optional modules (zlib, OpenSSL, etc) supported by the platform. Follow the steps outlined in *Getting Started* for the target platform, all the way through to having a working compiled python.

In order to set up the buildbot software, you will need to obtain an identifier and password for your worker so it can join the fleet. Email python-buildbots@python.org to discuss adding your worker and to obtain the needed workername and password. You can do some of the steps that follow before having the credentials, but it is easiest to have them before the "buildbot worker" step below.

10.30.2 Setting up the buildbot worker

Conventional always-on machines

You need a recent version of the `buildbot` software, and you will probably want a separate ‘buildbot’ user to run the buildbot software. You may also want to set the buildbot up using a virtual environment, depending on how you manage your system. We won’t cover how to that here; it doesn’t differ from setting up a virtual environment for any other software, but you’ll need to modify the sequence of steps below as appropriate if you choose that path.

For Linux:

- If your package manager provides the buildbot worker software, that is probably the best way to install it; it may create the buildbot user for you, in which case you can skip that step. Otherwise, do `pip install buildbot-worker`.
- Create a buildbot user (using, eg: `useradd`) if necessary.
- Log in as the buildbot user.

For Mac:

- Create a buildbot user using the OS/X control panel user admin. It should be a “standard” user.
- Log in as the buildbot user.
- Install the buildbot worker¹ by running `pip install buildbot-worker`.

For Windows:

- Create a buildbot user as a “standard” user.
- Install the latest version of Python from python.org.
- Open a Command Prompt.
- Execute `python -m pip install pypiwin32 buildbot-worker` (note that `python.exe` is not added to `PATH` by default, making the `python` command accessible is left as an exercise for the user).

In a terminal window for the buildbot user, issue the following commands (you can put the `buildarea` wherever you want to):

```
mkdir buildarea
buildbot-worker create-worker buildarea buildbot.python.org:9020 workername_
↪workerpasswd
```

(Note that on Windows, the `buildbot-worker` command will be in the `Scripts` directory of your Python installation.)

Once this initial worker setup completes, you should edit the files `buildarea/info/admin` and `buildarea/info/host` to provide your contact info and information on the host configuration, respectively. This information will be presented in the buildbot web pages that display information about the builders running on your worker.

You will also want to make sure that the worker is started when the machine reboots:

For Linux:

- Add the following line to `/etc/crontab`:

```
@reboot buildbot-worker restart /path/to/buildarea
```

Note that we use `restart` rather than `start` in case a crash has left a `twistd.pid` file behind.

For OSX:

¹ If the buildbot is going to do Framework builds, it is better to use the Apple-shipped Python so as to avoid any chance of the buildbot picking up components from the installed `python.org` python.

- Create a bin directory for your buildbot user:

```
mkdir bin
```

- Place the following script, named `run_worker.sh`, into that directory:

```
#!/bin/bash
export PATH=/usr/local/bin:/Library/Frameworks/Python.framework/Versions/
↳2.7/bin:$PATH
export LC_CTYPE=en_US.utf-8
cd /Users/buildbot/buildarea
twistd --nodaemon --python=buildbot.tac --logfile=buildbot.log --
↳prefix=worker
```

If you use pip with Apple's system python, add `'/System'` to the front of the path to the Python bin directory.

- Place a file with the following contents into `/Library/LaunchDaemons`:

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE plist PUBLIC "-//Apple Computer//DTD PLIST 1.0//EN"
  "http://www.apple.com/DTDs/PropertyList-1.0.dtd">
<plist version="1.0">
<dict>
  <key>Label</key>
  <string>net.buildbot.worker</string>
  <key>UserName</key>
  <string>buildbot</string>
  <key>WorkingDirectory</key>
  <string>/Users/buildbot/buildarea</string>
  <key>ProgramArguments</key>
  <array>
    <string>/Users/buildbot/bin/run_worker.sh</string>
  </array>
  <key>StandardOutPath</key>
  <string>twistd.log</string>
  <key>StandardErrorPath</key>
  <string>twistd.log</string>
  <key>KeepAlive</key>
  <true/>
  <key>SessionCreate</key>
  <true/>
</dict>
</plist>
```

The recommended name for the file is `net.buildbot.worker`.

For Windows:

- Add a Scheduled Task to run `buildbot-worker start buildarea` as the buildbot user “when the computer starts up”. It is best to provide absolute paths to the `buildbot-worker` command and the `buildarea` directory. It is also recommended to set the task to run in the directory that contains the `buildarea` directory.
- Alternatively (note: don't do both!), set up the worker service as described in the [buildbot documentation](#).

To start the worker running for your initial testing, you can do:

```
buildbot-worker start buildarea
```

Then you can either wait for someone to make a commit, or you can pick a builder associated with your worker from the [list of builders](#) and force a build.

In any case you should initially monitor builds on your builders to make sure the tests are passing and to resolve any platform issues that may be revealed by tests that fail. Unfortunately we do not currently have a way to notify you only of failures on your builders, so doing periodic spot checks is also a good idea.

Latent workers

We also support running [latent workers](#) on the AWS EC2 service. To set up such a worker:

- Start an instance of your chosen base AMI and set it up as a conventional worker.
- After the instance is fully set up as a conventional worker (including worker name and password, and admin and host information), create an AMI from the instance and stop the instance.
- Contact the buildmaster administrator who gave you your worker name and password and give them the following information:
 - Instance size (such as `m4.large`)
 - Full region specification (such as `us-west-2`)
 - AMI ID (such as `ami-1234beef`)
 - An Access Key ID and Access Key. It is recommended to set up a separate IAM user with full access to EC2 and provide the access key information for that user rather than for your main account.

The buildmaster cannot guarantee that it will always shut down your instance(s), so it is recommended to periodically check and make sure there are no “zombie” instances running on your account, created by the buildbot master. Also, if you notice that your worker seems to have been down for an unexpectedly long time, please ping the [python-buildbots](#) list to request that the master be restarted.

Latent workers should also be updated periodically to include operating system or other software updates, but when to do such maintenance is largely up to you as the worker owner. There are a couple different options for doing such updates:

- Start an instance from your existing AMI, do updates on that instance, and save a new AMI from the updated instance. Note that (especially for Windows workers) you should do at least one restart of the instance after doing updates to be sure that any post-reboot update work is done before creating the new AMI.
- Create an entirely new setup from a newer base AMI using your existing worker name and password.

Whichever way you choose to update your AMI, you'll need to provide the buildmaster administrators with the new AMI ID.

10.30.3 Buildbot worker operation

Most of the time, running a worker is a “set and forget” operation, depending on the level of involvement you want to have in resolving bugs revealed by your builders. There are, however, times when it is helpful or even necessary for you to get involved. As noted above, you should be subscribed to python-buildbots@python.org so that you will be made aware of any fleet-wide issues.

Necessary tasks include, obviously, keeping the buildbot running. Currently the system for notifying buildbot owners when their workers go offline is not working; this is something we hope to resolve. So currently it is helpful if you periodically check the status of your worker. We will also contact you via your contact address in `buildarea/info/admin` when we notice there is a problem that has not been resolved for some period of time and you have not responded to a posting on the [python-buildbots](#) list about it.

We currently do not have a minimum version requirement for the worker software. However, this is something we will probably establish as we tune the fleet, so another task will be to occasionally upgrade the buildbot worker software. Coordination for this will be done via `python-buildbots@python.org`.

The most interesting extra involvement is when your worker reveals a unique or almost-unique problem: a test that is failing on your system but not on other systems. In this case you should be prepared to offer debugging help to the people working on the bug: running tests by hand on the worker machine or, if possible, providing ssh access to a committer to run experiments to try to resolve the issue.

10.30.4 Required Ports

The worker operates as a *client* to the *buildmaster*. This means that all network connections are *outbound*. This is true also for the network tests in the test suite. Most consumer firewalls will allow any outbound traffic, so normally you do not need to worry about what ports the buildbot uses. However, corporate firewalls are sometimes more restrictive, so here is a table listing all of the outbound ports used by the buildbot and the python test suite (this list may not be complete as new tests may have been added since this table was last vetted):

Port	Host	Description
20, 21	ftp.debian.org	test_urllib2net
53	your DNS server	test_socket, and others implicitly
80	python.org example.com	(several tests)
119	news.gmane.org	test_nntplib
443	(various)	test_ssl
465	smtp.gmail.com	test_smtpnet
587	smtp.gmail.com	test_smtpnet
9020	python.org	connection to buildmaster

Many tests will also create local TCP sockets and connect to them, usually using either `localhost` or `127.0.0.1`.

10.30.5 Required Resources

Based on the last time we did a [survey](#) on buildbot requirements, the recommended resource allocations for a python buildbot are at least:

- 2 CPUs
- 512 MB RAM
- 30 GB free disk space

The bigmem tests won't run in this configuration, since they require substantially more memory, but these resources should be sufficient to ensure that Python compiles correctly on the platform and can run the rest of the test suite.

10.30.6 Security Considerations

We only allow builds to be triggered against commits to the [CPython repository on GitHub](#). This means that the code your buildbot will run will have been vetted by a committer. However, mistakes and bugs happen, as could a compromise, so keep this in mind when siting your buildbot on your network and establishing the security around it. Treat the buildbot like you would any resource that is public facing and might get hacked (use a VM and/or jail/chroot/solaris zone, put it in a DMZ, etc). While the buildbot does not have any ports open for inbound traffic (and is not public facing in that sense), committer mistakes do happen, and security flaws are discovered in both released and unreleased code, so treating the buildbot as if it were fully public facing is a good policy.

Code runs differently as privileged and unprivileged users. We would love to have builders running as privileged accounts, but security considerations do make that difficult, as access to root can provide access to surprising resources (such as spoofed IP packets, changes in MAC addresses, etc) even on a VM setup. But if you are confident in your setup, we'd love to have a buildbot that runs python as root.

Note that the above is a summary of a [discussion](#) on python-dev about buildbot security that includes examples of the tests for which privilege matters. There was no final consensus, but the information is useful as a point of reference.

10.31 Core Developer Motivations and Affiliations

CPython core developers participate in the core development process for a variety of reasons. Being accepted as a core developer indicates that an individual is interested in acquiring those responsibilities, has the ability to collaborate effectively with existing core developers, and has had the time available to demonstrate both that interest and that ability.

This page allows core developers that choose to do so to provide more information to the rest of the Python community regarding their personal situation (such as their general location and professional affiliations), as well as any personal motivations that they consider particularly relevant.

Core developers that wish to provide this additional information add a new entry to the *Published entries* section below. Guidelines relating to content and layout are included as comments in the source code for this page.

Core developers that are available for training, consulting, contract, or full-time work, or are seeking crowdfunding support for their community contributions, may also choose to provide that information here (including linking out to commercial sites with the relevant details).

For more information on the origins and purpose of this page, see *Goals of this page*.

10.31.1 Published entries

The following core developers have chosen to provide additional details regarding their professional affiliations and (optionally) other reasons for participating in the CPython core development process:

Brett Cannon (Canada)

- Personal site: [snarky.ca](#)
- [Extended bio](#)
- Microsoft (Software Developer)
- Python Software Foundation (Fellow)

Nick Coghlan (Australia)

- Personal site: [Curious Efficiency](#)
- [Extended bio](#)
- [Tritium](#) (Software Developer)
- Python Software Foundation (Fellow, Packaging Working Group)

Nick began using Python as a testing and prototyping language while working for Boeing Defence Australia, and continues to use it for that purpose today.

As a core developer, he is primarily interested in helping to ensure Python's continued suitability for educational, testing and data analysis use cases, as well as in encouraging good architectural practices when assembling Python applications and test harnesses from open source components.

Steve Dower (United States/Australia)

- Microsoft (Software Developer)
- Personal site: stevedower.id.au
- Speaking: stevedower.id.au/speaking
- Work blog: aka.ms/pythonblog
- Email address: steve.dower@python.org

Steve started with Python while automating a test harness for medical devices, and now works for Microsoft on anything that makes Python more accessible to developers on any platform.

As a core developer, his focus is on maintaining the already excellent Windows support and improving Python's ability to be embedded in other applications.

Christian Heimes (Germany)

- Red Hat (Software Developer, Security Engineering / Identity Management)
- Python Software Foundation (Fellow)

Mariatta (Canada)

- Personal site: mariatta.ca
- Works as a [Software Engineer](#) in Vancouver, helps organize [Vancouver PyLadies](#) meetup on the side, and sometimes [speaks](#) at conferences.
- Email address: mariatta@python.org
- [Sponsor Mariatta on GitHub](#)
- [Patreon](#)

Support Mariatta by [becoming a sponsor](#), sending her a [happiness packet](#), or [paypal](#).

R. David Murray (United States)

- Personal site: bitdance.com
- Available for [Python and Internet Services Consulting](#) and [Python contract programming](#)

David has been involved in the Internet since the days when the old IBM BITNET and the ARPANet got cross connected, and in Python programming since he first discovered it around the days of Python 1.4. After transitioning from being Director of Operations for dialup Internet providers (when that business started declining) to being a full

time independent consultant, David started contributing directly to CPython development. He became a committer in 2009. He subsequently took over primary maintenance of the email package from Barry Warsaw, and contributed the unicode oriented API. David is also active in mentoring new contributors and, when time is available, working on the infrastructure that supports CPython development, specifically the Roundup-based bug tracker and the buildbot system.

David currently does both proprietary and open source development work, primarily in Python, through the company in which he is a partner, [Murray & Walker, Inc.](#) He has done contract work focused specifically on CPython development both through the PSF (the kickstart of the email unicode API development) and directly funded by interested corporations (additional development work on email funded by QNX, and work on CPython ICC support funded by Intel). He would like to spend more of his (and his company's) time on open source work, and so is actively seeking additional such contract opportunities.

Antoine Pitrou (France)

- LinkedIn: <https://www.linkedin.com/in/pitrou/> (Senior Software Engineer)
- Independent (currently Ursa Labs / R Studio)
- Python Software Foundation (Fellow)
- Available for open source contract work
- Email address: antoine@python.org

Antoine started working with Python in 2005 in order to implement a decentralized virtual world protocol. He started contributing to CPython in 2007 and became a core developer in 2008. His motivations have been driven both by the abstract desire to make Python better for the whole world, and by the concrete roadblocks he was hitting in professional settings. Topics of choice have included interpreter optimizations, garbage collection, network programming, system programming and concurrent programming (such as maintaining `multiprocessing`).

As a professional, Antoine has been first specializing in network programming, and more lately in open source data science infrastructure such as Dask, Numba, Apache Arrow.

Victor Stinner (France)

- [Personal website](#)
- Red Hat (Senior Software Engineer)

Victor is paid by Red Hat to maintain Python upstream and downstream (RHEL, CentOS, Fedora & Software collections). See [Victor's contributions to Python](#).

Kushal Das (India)

- [Personal website](#)
- [Freedom of the Press Foundation](#) (Staff)
- Python Software Foundation (Fellow)

Barry Warsaw (United States)

- [LinkedIn](#): (Senior Staff Software Engineer - Python Foundation team)
- Personal site: barry.warsaw.us
- Blog: [We Fear Change](#)
- Email address: barry@python.org
- Python Software Foundation (Fellow)

Barry has been working in, with, and on Python since 1994. He attended the first Python workshop at NBS (now NIST) in Gaithersburg, MD in 1994, where he met Guido and several other early Python adopters. Barry subsequently worked with Guido for 8 years while at CNRI. From 2007 until 2017, Barry worked for Canonical, corporate sponsor of Ubuntu Linux, primarily on the Python ecosystem, and is both an Ubuntu and a Debian uploading developer. Barry has served as Python's postmaster, webmaster, release manager, Language Summit co-chair, Jython project leader, GNU Mailman project leader, and probably lots of other things he shouldn't admit to.

Eric Snow (United States)

- Microsoft (Software Developer)
- Python Software Foundation (Fellow)

Dino Viehland (United States)

- Microsoft: <https://notebooks.azure.com/> (Software Engineer)
- Email address: dinov@microsoft.com

Dino started working with Python in 2005 by working on IronPython, an implementation of Python running on .NET. He was one of the primary developers on the project for 6 years. After that he started the Python Tools for Visual Studio project focusing on providing advanced code completion and debugging features for Python. Today he works on Azure Notebooks bringing the Python based Jupyter notebook as a hosted on-line service.

Carol Willing (United States)

- Personal site: [Willing Consulting](#)
- [Extended bio](#)
- Project Jupyter, Cal Poly SLO (Research Software Engineer)
- Python Software Foundation (Fellow)

Carol is focused on Python's usage in education and scientific research. She is interested in organizational development, operational workflows, and sustainability of open source projects.

10.31.2 Goals of this page

The [issue metrics](#) automatically collected by the CPython issue tracker strongly suggest that the current core development process is bottlenecked on core developer time - this is most clearly indicated in the first metrics graph, which

shows both the number of open issues and the number of patches awaiting review growing steadily over time, despite CPython being one of the most active open source projects in the world. This bottleneck then impacts not only resolving open issues and applying submitted patches, but also the process of identifying, nominating and mentoring new core developers.

The core commit statistics monitored by sites like [OpenHub](#) provide a good record as to *who* is currently handling the bulk of the review and maintenance work, but don't provide any indication as to the factors currently influencing people's ability to spend time on reviewing proposed changes, or mentoring new contributors.

This page aims to provide at least some of that missing data by encouraging core developers to highlight professional affiliations in the following two cases (even if not currently paid for time spent participating in the core development process):

- developers working for vendors that distribute a commercially supported Python runtime
- developers working for Sponsor Members of the Python Software Foundation

These are cases where documenting our affiliations helps to improve the overall transparency of the core development process, as well as making it easier for staff at these organisations to locate colleagues that can help them to participate in and contribute effectively to supporting the core development process.

Core developers working for organisations with a vested interest in the sustainability of the CPython core development process are also encouraged to seek opportunities to spend work time on mentoring potential new core developers, whether through the general [core mentorship program](#), through mentoring colleagues, or through more targeted efforts like Outreachy's paid [internships](#) and Google's [Summer of Code](#).

Core developers that are available for consulting or contract work on behalf of the Python Software Foundation or other organisations are also encouraged to provide that information here, as this will help the PSF to better facilitate funding of core development work by organisations that don't directly employ any core developers themselves.

Finally, some core developers seeking to increase the time they have available to contribute to CPython may wish to pursue crowdfunding efforts that allow their contributions to be funded directly by the community, rather than relying on institutional sponsors allowing them to spend some or all of their work time contributing to CPython development.

10.31.3 Limitations on scope

- Specific technical areas of interest for core developers should be captured in the [Experts Index](#).
- This specific listing is limited to CPython core developers (since it's focused on the specific constraint that is core developer time), but it would be possible to create a more expansive listing on the Python wiki that also covers issue triagers, and folks seeking to become core developers.
- Changes to the software and documentation maintained by core developers, together with related design discussions, all take place in public venues, and hence are inherently subject to full public review. Accordingly, core developers are NOT required to publish their motivations and affiliations if they do not choose to do so. This helps to ensure that core contribution processes remain open to anyone that is in a position to sign the [Contributor Licensing Agreement](#), the details of which are filed privately with the Python Software Foundation, rather than publicly.

10.32 Git Bootcamp and Cheat Sheet

Note: This section provides instructions on common tasks in CPython's workflow. It's designed to assist new contributors who have some familiarity with git and GitHub.

If you are new to git and GitHub, please become comfortable with these instructions before submitting a pull request. As there are several ways to accomplish these tasks using git and GitHub, this section reflects one method suitable for new contributors. Experienced contributors may desire a different approach.

In this section, we will go over some commonly used Git commands that are relevant to CPython's workflow.

Note: Setting up git aliases for common tasks can be useful to you. You can get more information about that in [git documentation](#)

Contents

- *Git Bootcamp and Cheat Sheet*
 - *Forking CPython GitHub Repository*
 - *Cloning a Forked CPython Repository*
 - *Listing the Remote Repositories*
 - *Setting Up Your Name and Email Address*
 - *Enabling autocrlf on Windows*
 - *Creating and Switching Branches*
 - *Deleting Branches*
 - *Staging and Committing Files*
 - *Reverting Changes*
 - *Stashing Changes*
 - *Committing Changes*
 - *Pushing Changes*
 - *Creating a Pull Request*
 - *Syncing with Upstream*
 - *Applying a Patch from Mercurial to Git*
 - *Downloading Other's Patches*
 - *Accepting and Merging a Pull Request*
 - *Backporting Merged Changes*
 - *Editing a Pull Request Prior to Merging*

10.32.1 Forking CPython GitHub Repository

You will only need to do this once.

1. Go to <https://github.com/python/cpython>.
2. Press **Fork** on the top right.
3. When asked where to fork the repository, choose to fork it to your username.

4. Your forked CPython repository will be created at <https://github.com/<username>/cpython>.

10.32.2 Cloning a Forked CPython Repository

You will only need to do this once. From your command line:

```
git clone git@github.com:<username>/cpython.git
```

It is also recommended to configure an `upstream` remote repository:

```
cd cpython
git remote add upstream git@github.com:python/cpython.git
```

You can also use SSH-based or HTTPS-based URLs.

10.32.3 Listing the Remote Repositories

To list the remote repositories that are configured, along with their URLs:

```
git remote -v
```

You should have two remote repositories: `origin` pointing to your forked CPython repository, and `upstream` pointing to the official CPython repository:

```
origin  git@github.com:<username>/cpython.git (fetch)
origin  git@github.com:<username>/cpython.git (push)
upstream      git@github.com:python/cpython.git (fetch)
upstream      git@github.com:python/cpython.git (push)
```

10.32.4 Setting Up Your Name and Email Address

```
git config --global user.name "Your Name"
git config --global user.email your.email@example.com
```

The `--global` flag sets these parameters globally while the `--local` flag sets them only for the current project.

10.32.5 Enabling `autocrlf` on Windows

The `autocrlf` option will fix automatically any Windows-specific line endings. This should be enabled on Windows, since the public repository has a hook which will reject all changesets having the wrong line endings:

```
git config --global core.autocrlf input
```

10.32.6 Creating and Switching Branches

Important: Never commit directly to the `master` branch.

Create a new branch and switch to it:

```
# creates a new branch off master and switch to it
git checkout -b <branch-name> master
```

This is equivalent to:

```
# create a new branch from master, without checking it out
git branch <branch-name> master
# check out the branch
git checkout <branch-name>
```

To find the branch you are currently on:

```
git branch
```

The current branch will have an asterisk next to the branch name. Note, this will only list all of your local branches.

To list all the branches, including the remote branches:

```
git branch -a
```

To switch to a different branch:

```
git checkout <another-branch-name>
```

Other releases are just branches in the repository. For example, to work on the 2.7 release from the `upstream` remote:

```
git checkout -b 2.7 upstream/2.7
```

10.32.7 Deleting Branches

To delete a **local** branch that you no longer need:

```
git checkout master
git branch -D <branch-name>
```

To delete a **remote** branch:

```
git push origin -d <branch-name>
```

You may specify more than one branch for deletion.

10.32.8 Staging and Committing Files

1. To show the current changes:

```
git status
```

2. To stage the files to be included in your commit:

```
git add <filename1> <filename2>
```

3. To commit the files that have been staged (done in step 2):

```
git commit -m "bpo-XXXX: This is the commit message."
```


10.32.9 Reverting Changes

To revert changes to a file that has not been committed yet:

```
git checkout <filename>
```

If the change has been committed, and now you want to reset it to whatever the origin is at:

```
git reset --hard HEAD
```

10.32.10 Stashing Changes

To stash away changes that are not ready to be committed yet:

```
git stash
```

To re-apply the last stashed change:

```
git stash pop
```

10.32.11 Committing Changes

Add the files you want to commit:

```
git add <filename>
```

Commit the files:

```
git commit -m "<message>"
```

10.32.12 Pushing Changes

Once your changes are ready for a review or a pull request, you will need to push them to the remote repository.

```
git checkout <branch-name>  
git push origin <branch-name>
```

10.32.13 Creating a Pull Request

1. Go to <https://github.com/python/cpython>.
2. Press the `New pull request` button.
3. Click the `compare across forks` link.
4. Select the base repository: `python/cpython` and base branch: `master`.
5. Select the head repository: `<username>/cpython` and head branch: the branch containing your changes.
6. Press the `Create pull request` button.

10.32.14 Syncing with Upstream

Scenario:

- You forked the CPython repository some time ago.
- Time passes.
- There have been new commits made in the upstream CPython repository.
- Your forked CPython repository is no longer up to date.
- You now want to update your forked CPython repository to be the same as the upstream CPython repository.

Solution:

```
git checkout master
git pull upstream master
git push origin master
```

Another scenario:

- You created `some-branch` some time ago.
- Time passes.
- You made some commits to `some-branch`.
- Meanwhile, there are recent changes from the upstream CPython repository.
- You want to incorporate the recent changes from the upstream CPython repository into `some-branch`.

Solution:

```
git checkout some-branch
git fetch upstream
git merge upstream/master
git push origin some-branch
```

You may see error messages like “CONFLICT” and “Automatic merge failed;” when you run `git merge upstream/master`.

When it happens, you need to resolve conflict. See these articles about resolving conflicts:

- [About merge conflicts](#)
- [Resolving a merge conflict using the command line](#)

10.32.15 Applying a Patch from Mercurial to Git

Scenario:

- A Mercurial patch exists but there is no pull request for it.

Solution:

1. Download the patch locally.
2. Apply the patch:

```
git apply /path/to/issueNNNN-git.patch
```

If there are errors, update to a revision from when the patch was created and then try the `git apply` again:

```
git checkout $(git rev-list -n 1 --before="yyyy-mm-dd hh:mm:ss" master)
git apply /path/to/issueNNNN-git.patch
```

If the patch still won't apply, then a patch tool will not be able to apply the patch and it will need to be re-implemented manually.

3. If the apply was successful, create a new branch and switch to it.
4. Stage and commit the changes.
5. If the patch was applied to an old revision, it needs to be updated and merge conflicts need to be resolved:

```
git rebase master
git mergetool
```

6. Push the changes and open a pull request.

10.32.16 Downloading Other's Patches

Scenario:

- A contributor made a pull request to CPython.
- Before merging it, you want to be able to test their changes locally.

On Unix and MacOS, set up the following git alias:

```
$ git config --global alias.pr '!sh -c "git fetch upstream pull/${1}/head:pr_${1} && \
↳git checkout pr_${1}" -'
```

On Windows, reverse the single (') and double (") quotes:

```
git config --global alias.pr "!sh -c 'git fetch upstream pull/${1}/head:pr_${1} && \
↳git checkout pr_${1}' -"
```

The alias only needs to be done once. After the alias is set up, you can get a local copy of a pull request as follows:

```
git pr <pr_number>
```

Note: `hub` command line utility makes this workflow very easy. You can check out the branch by `hub pr checkout <pr_number> [<branch_name>]`. This command configures remote URL for the branch too. So you can `git push` if the pull request author checked "Allow edits from maintainers" when creating the pull request.

10.32.17 Accepting and Merging a Pull Request

Pull requests can be accepted and merged by a Python Core Developer.

1. At the bottom of the pull request page, click the `Squash and merge` button.
2. Replace the reference to GitHub pull request `#NNNN` with `GH-NNNN`. If the title is too long, the pull request number can be added to the message body.
3. Adjust and clean up the commit message.

Example of good commit message:

```
bpo-12345: Improve the spam module (GH-777)

* Add method A to the spam module
* Update the documentation of the spam module
```

Example of bad commit message:

```
bpo-12345: Improve the spam module (#777)

* Improve the spam module
* merge from master
* adjust code based on review comment
* rebased
```

Note: [How to Write a Git Commit Message](#) is a nice article describing how to write a good commit message.

4. Press the `Confirm squash and merge` button.

10.32.18 Backporting Merged Changes

A pull request may need to be backported into one of the maintenance branches after it has been accepted and merged into `master`. It is usually indicated by the label `needs backport to X.Y` on the pull request itself.

Use the utility script `cherry_picker.py` from the `core-workflow` repository to backport the commit.

The commit hash for backporting is the squashed commit that was merged to the `master` branch. On the merged pull request, scroll to the bottom of the page. Find the event that says something like:

```
<core_developer> merged commit <commit_sha1> into python:master <sometime> ago.
```

By following the link to `<commit_sha1>`, you will get the full commit hash.

Alternatively, the commit hash can also be obtained by the following git commands:

```
git fetch upstream
git rev-parse ":/bpo-12345"
```

The above commands will print out the hash of the commit containing `"bpo-12345"` as part of the commit message.

When formatting the commit message for a backport commit: leave the original one as is and delete the number of the backport pull request.

Example of good backport commit message:

```
bpo-12345: Improve the spam module (GH-777)

* Add method A to the spam module
* Update the documentation of the spam module

(cherry picked from commit 62adc55)
```

Example of bad backport commit message:

```
bpo-12345: Improve the spam module (GH-777) (#888)
```

(continues on next page)

(continued from previous page)

- ```
* Add method A to the spam module
* Update the documentation of the spam module
```

### 10.32.19 Editing a Pull Request Prior to Merging

When a pull request submitter has enabled the [Allow edits from maintainers](#) option, Python Core Developers may decide to make any remaining edits needed prior to merging themselves, rather than asking the submitter to do them. This can be particularly appropriate when the remaining changes are bookkeeping items like updating `Misc/ACKS`.

To edit an open pull request that targets `master`:

1. In the pull request page, under the description, there is some information about the contributor's forked CPython repository and branch name that will be useful later:

```
<contributor> wants to merge 1 commit into python:master from <contributor>:
↔<branch_name>
```

2. Fetch the pull request, using the `git pr` alias:

```
git pr <pr_number>
```

This will checkout the contributor's branch at `<pr_number>`.

3. Make and commit your changes on the branch. For example, merge in changes made to `master` since the PR was submitted (any merge commits will be removed by the later `Squash and Merge` when accepting the change):

```
git fetch upstream
git merge upstream/master
git add <filename>
git commit -m "<message>"
```

4. Push the changes back to the contributor's PR branch:

```
git push git@github.com:<contributor>/cpython <pr_number>:<branch_name>
```

5. Optionally, *delete the PR branch*.

## 10.33 Appendix: Topics

### 10.33.1 Basics for contributors

---

**Note: Recommended reading**

- *Getting Started*
- *Lifecycle of a Pull Request*

- 
- *Where to Get Help*
  - *Following Python's Development*

### 10.33.2 Core developers

---

**Note: Recommended reading**

- *Getting Started*
  - *Lifecycle of a Pull Request*
  - *Accepting Pull Requests*
- 

- *How to Become a Core Developer*
- *Developer Log*
- *Core Developer Motivations and Affiliations*
- *Experts Index*

### 10.33.3 Development workflow for contributors

- *Development Cycle*
- *Lifecycle of a Pull Request*
- *Fixing “easy” Issues (and Beyond)*

### 10.33.4 Documenting Python and style guide

- *Helping with Documentation*
- *Documenting Python*

### 10.33.5 Issue tracking and triaging

- *Issue Tracking*
- *Triaging an Issue*

### 10.33.6 Language development in depth

- *Exploring CPython's Internals*
- *Changing CPython's Grammar*
- *Design of CPython's Compiler*
- *Design of CPython's Garbage Collector*
- *Adding to the Stdlib*
- *Changing the Python Language*
- *Porting Python to a new platform*

### 10.33.7 Testing and continuous integration

- *Running & Writing Tests*
- *Increase Test Coverage*
- *Silence Warnings From the Test Suite*
- *Continuous Integration*
- *Running a buildbot worker*
- *Coverity Scan*





---

## Bibliography

---

- [Wang97] Daniel C. Wang, Andrew W. Appel, Jeff L. Korn, and Chris S. Serra. *The Zephyr Abstract Syntax Description Language*. In *Proceedings of the Conference on Domain-Specific Languages*, pp. 213–227, 1997.



## P

PEP process, 101

Python Enhancement Proposals

PEP 361, 93

PEP 373, 93

PEP 375, 93

PEP 392, 93

PEP 398, 93

PEP 429, 93

PEP 478, 7

PEP 494, 7

PEP 537, 7

PEP 569, 7

PEP 596, 7

PEP 619, 7

Python Enhancement Proposals

PEP 1, 101

PEP 13, 81

PEP 291, 102

PEP 360, 102

PEP 387, 32

PEP 545, 64, 65

PEP 563, 117

PEP 7, 15, 32

PEP 8, 15, 32