
CherryPy Documentation

Release 10.0.1.dev2+ng34d9d70.d20170120

CherryPy Team

January 20, 2017

1	Foreword	1
1.1	Why CherryPy?	1
1.2	Success Stories	2
2	Installation	5
2.1	Requirements	5
2.2	Supported python version	5
2.3	Installing	5
2.4	Run it	6
3	Tutorials	9
3.1	Tutorial 1: A basic web application	9
3.2	Tutorial 2: Different URLs lead to different functions	10
3.3	Tutorial 3: My URLs have parameters	11
3.4	Tutorial 4: Submit this form	12
3.5	Tutorial 5: Track my end-user’s activity	12
3.6	Tutorial 6: What about my javascripts, CSS and images?	13
3.7	Tutorial 7: Give us a REST	15
3.8	Tutorial 8: Make it smoother with Ajax	17
3.9	Tutorial 9: Data is all my life	19
3.10	Tutorial 10: Make it a modern single-page application with React.js	22
3.11	Tutorial 11: Organize my code	25
4	Basics	27
4.1	The one-minute application example	28
4.2	Hosting one or more applications	28
4.3	Logging	29
4.4	Configuring	32
4.5	Cookies	33
4.6	Using sessions	34
4.7	Static content serving	35
4.8	Dealing with JSON	37
4.9	Authentication	37
4.10	Favicon	38
5	Advanced	41
5.1	Set aliases to page handlers	41
5.2	RESTful-style dispatching	42

5.3	Error handling	44
5.4	Streaming the response body	45
5.5	Response timeouts	46
5.6	Deal with signals	47
5.7	Securing your server	47
5.8	Multiple HTTP servers support	48
5.9	WSGI support	48
5.10	WebSocket support	49
5.11	Database support	50
5.12	HTML Templating support	50
5.13	Testing your application	50
6	Configure	53
6.1	Architecture	53
6.2	Declaration	55
6.3	Namespaces	57
7	Extend	61
7.1	Server-wide functions	61
7.2	Per-request functions	67
7.3	Tailored dispatchers	70
7.4	Request body processors	71
8	Deploy	73
8.1	Run as a daemon	73
8.2	Run as a different user	74
8.3	PID files	74
8.4	Systemd socket activation	74
8.5	Control via Supervisor	74
8.6	SSL support	75
8.7	WSGI servers	76
8.8	Virtual Hosting	79
8.9	Reverse-proxying	80
9	Support	83
9.1	I have a question	83
9.2	I have found a bug	83
9.3	I have a feature request	83
9.4	I want to converse	83
10	Contribute	85
10.1	StackOverflow	85
10.2	Filing Bug Reports	85
10.3	Fixing Bugs	85
10.4	Writing Pull Requests	85
11	Testing	87
12	Glossary	89
13	History	91
13.1	v10.0.0	91
13.2	v9.0.0	91
13.3	v8.9.1	91
13.4	v8.9.0	91

13.5	v8.8.0	92
13.6	v8.7.0	92
13.7	v8.6.0	92
13.8	v8.5.0	92
13.9	v8.4.0	92
13.10	v8.3.1	92
13.11	v8.3.0	93
13.12	v8.2.0	93
13.13	v8.1.3	93
13.14	v8.1.2	93
13.15	v8.1.1	93
13.16	v8.1.0	93
13.17	v8.0.1	94
13.18	v8.0.0	94
13.19	7.1.0	94
13.20	7.0.0	94
13.21	6.2.1	94
13.22	6.2.0	95
13.23	6.1.1	95
13.24	6.1.0	95
13.25	6.0.2	95
13.26	6.0.1	95
13.27	6.0.0	95
13.28	5.6.0	95
13.29	5.5.0	96
13.30	5.4.0	96
13.31	5.3.0	96
13.32	5.2.0	96
13.33	5.1.0	96
13.34	5.0.1	97
13.35	5.0.0	97
13.36	4.0.0	97
13.37	3.8.2	97
13.38	3.8.0	97
13.39	3.7.0	97
13.40	3.6.0	98
13.41	3.5.0	98
13.42	3.4.0	98
13.43	3.3.0	98

1.1 Why CherryPy?

CherryPy is among the oldest web framework available for Python, yet many people aren't aware of its existence. One of the reasons for this is that CherryPy is not a complete stack with built-in support for a multi-tier architecture. It doesn't provide frontend utilities nor will it tell you how to speak with your storage. Instead, CherryPy's take is to let the developer make those decisions. This is a contrasting position compared to other well-known frameworks.

CherryPy has a clean interface and does its best to stay out of your way whilst providing a reliable scaffolding for you to build from.

Typical use-cases for CherryPy go from regular web application with user frontends (think blogging, CMS, portals, ecommerce) to web-services only.

Here are some reasons you would want to choose CherryPy:

1. Simplicity

Developing with CherryPy is a simple task. "Hello, world" is only a few lines long, and does not require the developer to learn the entire (albeit very manageable) framework all at once. The framework is very pythonic; that is, it follows Python's conventions very nicely (code is sparse and clean).

Contrast this with J2EE and Python's most popular and visible web frameworks: Django, Zope, Pylons, and Turbogears. In all of them, the learning curve is massive. In these frameworks, "Hello, world" requires the programmer to set up a large scaffold which spans multiple files and to type a lot of boilerplate code. CherryPy succeeds because it does not include the bloat of other frameworks, allowing the programmer to write their web application quickly while still maintaining a high level of organization and scalability.

CherryPy is also very modular. The core is fast and clean, and extension features are easy to write and plug in using code or the elegant config system. The primary components (server, engine, request, response, etc.) are all extendable (even replaceable) and well-managed.

In short, CherryPy empowers the developer to work with the framework, not against or around it.

2. Power

CherryPy leverages all of the power of Python. Python is a dynamic language which allows for rapid development of applications. Python also has an extensive built-in API which simplifies web app development. Even more extensive, however, are the third-party libraries available for Python. These range from object-relational mappers to form libraries, to an automatic Python optimizer, a Windows exe generator, imaging libraries, email support, HTML templating engines, etc. CherryPy applications are just like regular Python applications. CherryPy does not stand in your way if you want to use these brilliant tools.

CherryPy also provides *tools* and *plugins*, which are powerful extension points needed to develop world-class web applications.

3. Maturity

Maturity is extremely important when developing a real-world application. Unlike many other web frameworks, CherryPy has had many final, stable releases. It is fully bugtested, optimized, and proven reliable for real-world use. The API will not suddenly change and break backwards compatibility, so your applications are assured to continue working even through subsequent updates in the current version series.

CherryPy is also a “3.0” project: the first edition of CherryPy set the tone, the second edition made it work, and the third edition makes it beautiful. Each version built on lessons learned from the previous, bringing the developer a superior tool for the job.

4. Community

CherryPy has an devoted community that develops deployed CherryPy applications and are willing and ready to assist you on the CherryPy mailing list or IRC (#cherrypy on OFTC). The developers also frequent the list and often answer questions and implement features requested by the end-users.

5. Deployability

Unlike many other Python web frameworks, there are cost-effective ways to deploy your CherryPy application.

Out of the box, CherryPy includes its own production-ready HTTP server to host your application. CherryPy can also be deployed on any WSGI-compliant gateway (a technology for interfacing numerous types of web servers): `mod_wsgi`, `FastCGI`, `SCGI`, `IIS`, `uwsgi`, `tornado`, etc. Reverse proxying is also a common and easy way to set it up.

In addition, CherryPy is pure-python and is compatible with Python 2.3. This means that CherryPy will run on all major platforms that Python will run on (Windows, MacOSX, Linux, BSD, etc).

webfaction.com, run by the inventor of CherryPy, is a commercial web host that offers CherryPy hosting packages (in addition to several others).

6. It’s free!

All of CherryPy is licensed under the open-source BSD license, which means CherryPy can be used commercially for ZERO cost.

7. Where to go from here?

Check out the [tutorials](#) to start enjoying the fun!

1.2 Success Stories

You are interested in CherryPy but you would like to hear more from people using it, or simply check out products or application running it.

If you would like to have your CherryPy powered website or product listed here, contact us via our [mailing list](#) or IRC (#cherrypy on OFTC).

1.2.1 Websites running atop CherryPy

[Hulu DeeJay and Hulu Sod](#) - Hulu uses CherryPy for some projects. “The service needs to be very high performance. Python, together with CherryPy, [gunicorn](#), and [gevent](#) more than provides for this.”

[Netflix](#) - Netflix uses CherryPy as a building block in their infrastructure: “Restful APIs to large applications with requests, providing web interfaces with CherryPy and [Bottle](#), and crunching data with [scipy](#).”

[Urbanility](#) - French website for local neighbourhood assets in Rennes, France.

MROP Supply - Webshop for industrial equipment, developed using CherryPy 3.2.2 utilizing Python 3.2, with libs: Jinja2-2.6, davispuh-MySQL-for-Python-3-3403794, pyenchant-1.6.5 (for search spelling). “I’m coming over from .net development and found Python and CherryPy to be surprisingly minimalistic. No unnecessary overhead - build everything you need without the extra fluff. I’m a fan!”

CherryMusic - A music streaming server written in python: Stream your own music collection to all your devices! CherryMusic is open source.

YouGov Global - International market research firm, conducts millions of surveys on CherryPy yearly.

Aculab Cloud - Voice and fax applications on the cloud. A simple telephony API for Python, C#, C++, VB, etc... The website and all front-end and back-end web services are built with CherryPy, fronted by nginx (just handling the ssh and reverse-proxy), and running on AWS in two regions.

Learnit Training - Dutch website for an IT, Management and Communication training company. Built on CherryPy 3.2.0 and Python 2.7.3, with `oursql` and `DBUtils` libraries, amongst others.

Linstic - Sticky Notes in your browser (with linking).

Almad’s Homepage - Simple homepage with blog.

Fight.Watch - Twitch.tv web portal for fighting games. Built on CherryPy 3.3.0 and Python 2.7.3 with Jinja 2.7.2 and SQLAlchemy 0.9.4.

1.2.2 Products based on CherryPy

SABnzbd - Open Source Binary Newsreader written in Python.

Headphones - Third-party add-on for SABnzbd.

SickBeard - “Sick Beard is a PVR for newsgroup users (with limited torrent support). It watches for new episodes of your favorite shows and when they are posted it downloads them, sorts and renames them, and optionally generates metadata for them.”

TurboGears - The rapid web development megaframework. Turbogears 1.x used CherryPy. “CherryPy is the underlying application server for TurboGears. It is responsible for taking the requests from the user’s browser, parses them and turns them into calls into the Python code of the web application. Its role is similar to application servers used in other programming languages”.

Indigo - “An intelligent home control server that integrates home control hardware modules to provide control of your home. Indigo’s built-in Web server and client/server architecture give you control and access to your home remotely from other Macs, PCs, internet tablets, PDAs, and mobile phones.”

SlikiWiki - Wiki built on CherryPy and featuring WikiWords, automatic backlinking, site map generation, full text search, locking for concurrent edits, RSS feed embedding, per page access control lists, and page formatting using PyTextile markup.”

read4me - read4me is a Python feed-reading web service.

Firebird QA tools - Firebird QA tools are based on CherryPy.

salt-api - A REST API for Salt, the infrastructure orchestration tool.

1.2.3 Products inspired by CherryPy

OOWeb - “OOWeb is a lightweight, embedded HTTP server for Java applications that maps objects to URL directories, methods to pages and form/querystring arguments as method parameters. OOWeb was originally inspired by CherryPy.”

Installation

CherryPy is a pure Python library. This has various consequences:

- It can run anywhere Python runs
- It does not require a C compiler
- It can run on various implementations of the Python language: CPython, IronPython, Jython and PyPy

Contents

- *Installation*
 - *Requirements*
 - *Supported python version*
 - *Installing*
 - * *Test your installation*
 - *Run it*
 - * *cherryd*
 - *Command-Line Options*

2.1 Requirements

CherryPy does not have any mandatory requirements. However certain features it comes with will require you install certain packages. To simplify installing additional dependencies CherryPy enables you to specify extras in your requirements (e.g. `cherryypy[json, routes_dispatcher, ssl]`): - `doc` – for documentation related stuff - `json` – for custom JSON processing library - `routes_dispatcher` – `routes` for declarative URL mapping dispatcher - `ssl` – for OpenSSL bindings, useful in Python environments not having the builtin `ssl` module - `testing` - `memcached_session` – enables memcached backend session - `xcgi`

2.2 Supported python version

CherryPy supports Python 2.7 through to 3.5.

2.3 Installing

CherryPy can be easily installed via common Python package managers such as `setuptools` or `pip`.

```
$ easy_install cherrypy
```

```
$ pip install cherrypy
```

You may also get the latest CherryPy version by grabbing the source code from Github:

```
$ git clone https://github.com/cherrypy/cherrypy
$ cd cherrypy
$ python setup.py install
```

2.3.1 Test your installation

CherryPy comes with a set of simple tutorials that can be executed once you have deployed the package.

```
$ python -m cherrypy.tutorial.tut01_helloworld
```

Point your browser at <http://127.0.0.1:8080> and enjoy the magic.

Once started the above command shows the following logs:

```
[15/Feb/2014:21:51:22] ENGINE Listening for SIGHUP.
[15/Feb/2014:21:51:22] ENGINE Listening for SIGTERM.
[15/Feb/2014:21:51:22] ENGINE Listening for SIGUSR1.
[15/Feb/2014:21:51:22] ENGINE Bus STARTING
[15/Feb/2014:21:51:22] ENGINE Started monitor thread 'Autoreloader'.
[15/Feb/2014:21:51:22] ENGINE Started monitor thread '_TimeoutMonitor'.
[15/Feb/2014:21:51:22] ENGINE Serving on http://127.0.0.1:8080
[15/Feb/2014:21:51:23] ENGINE Bus STARTED
```

We will explain what all those lines mean later on, but suffice to know that once you see the last two lines, your server is listening and ready to receive requests.

2.4 Run it

During development, the easiest path is to run your application as follow:

```
$ python myapp.py
```

As long as *myapp.py* defines a “`__main__`” section, it will run just fine.

2.4.1 cherryd

Another way to run the application is through the `cherryd` script which is installed along side CherryPy.

Note: This utility command will not concern you if you embed your application with another framework.

Command-Line Options

-c, --config
Specify config file(s)

- d** Run the server as a daemon
- e, --environment** Apply the given config environment (defaults to None)
- f** Start a FastCGI server instead of the default HTTP server
- s** Start a SCGI server instead of the default HTTP server
- i, --import** Specify modules to import
- p, --pidfile** Store the process id in the given file (defaults to None)
- P, --Path** Add the given paths to sys.path

This tutorial will walk you through basic but complete CherryPy applications that will show you common concepts as well as slightly more advanced ones.

Contents

- *Tutorials*
 - *Tutorial 1: A basic web application*
 - *Tutorial 2: Different URLs lead to different functions*
 - *Tutorial 3: My URLs have parameters*
 - *Tutorial 4: Submit this form*
 - *Tutorial 5: Track my end-user's activity*
 - *Tutorial 6: What about my javascripts, CSS and images?*
 - *Tutorial 7: Give us a REST*
 - *Tutorial 8: Make it smoother with Ajax*
 - *Tutorial 9: Data is all my life*
 - *Tutorial 10: Make it a modern single-page application with React.js*
 - *Tutorial 11: Organize my code*
 - * *Dispatchers*
 - * *Tools*
 - * *Plugins*

3.1 Tutorial 1: A basic web application

The following example demonstrates the most basic application you could write with CherryPy. It starts a server and hosts an application that will be served at request reaching <http://127.0.0.1:8080/>

```
1 import cherrypy
2
3
4 class HelloWorld(object):
5     @cherrypy.expose
6     def index(self):
7         return "Hello world!"
8
9
10 if __name__ == '__main__':
11     cherrypy.quickstart(HelloWorld())
```

Store this code snippet into a file named *tut01.py* and execute it as follows:

```
$ python tut01.py
```

This will display something along the following:

```
1 [24/Feb/2014:21:01:46] ENGINE Listening for SIGHUP.
2 [24/Feb/2014:21:01:46] ENGINE Listening for SIGTERM.
3 [24/Feb/2014:21:01:46] ENGINE Listening for SIGUSR1.
4 [24/Feb/2014:21:01:46] ENGINE Bus STARTING
5 CherryPy Checker:
6 The Application mounted at '' has an empty config.
7
8 [24/Feb/2014:21:01:46] ENGINE Started monitor thread 'Autoreloader'.
9 [24/Feb/2014:21:01:46] ENGINE Started monitor thread '_TimeoutMonitor'.
10 [24/Feb/2014:21:01:46] ENGINE Serving on http://127.0.0.1:8080
11 [24/Feb/2014:21:01:46] ENGINE Bus STARTED
```

This tells you several things. The first three lines indicate the server will handle *signal* for you. The next line tells you the current state of the server, as that point it is in *STARTING* stage. Then, you are notified your application has no specific configuration set to it. Next, the server starts a couple of internal utilities that we will explain later. Finally, the server indicates it is now ready to accept incoming communications as it listens on the address *127.0.0.1:8080*. In other words, at that stage your application is ready to be used.

Before moving on, let's discuss the message regarding the lack of configuration. By default, CherryPy has a feature which will review the syntax correctness of settings you could provide to configure the application. When none are provided, a warning message is thus displayed in the logs. That log is harmless and will not prevent CherryPy from working. You can refer to [the documentation above](#) to understand how to set the configuration.

3.2 Tutorial 2: Different URLs lead to different functions

Your applications will obviously handle more than a single URL. Let's imagine you have an application that generates a random string each time it is called:

```
1 import random
2 import string
3
4 import cherrypy
5
6
7 class StringGenerator(object):
8     @cherrypy.expose
9     def index(self):
10         return "Hello world!"
11
12     @cherrypy.expose
13     def generate(self):
14         return ''.join(random.sample(string.hexdigits, 8))
15
16
17 if __name__ == '__main__':
18     cherrypy.quickstart(StringGenerator())
```

Save this into a file named *tut02.py* and run it as follows:

```
$ python tut02.py
```


Go now to <http://localhost:8080/generate> and your browser will display a random string.

Let's take a minute to decompose what's happening here. This is the URL that you have typed into your browser: <http://localhost:8080/generate>

This URL contains various parts:

- `http://` which roughly indicates it's a URL using the HTTP protocol (see [RFC 2616](#)).
- `localhost:8080` is the server's address. It's made of a hostname and a port.
- `/generate` which is the path segment of the URL. This is what CherryPy uses to locate an *exposed* function or method to respond.

Here CherryPy uses the `index()` method to handle `/` and the `generate()` method to handle `/generate`

3.3 Tutorial 3: My URLs have parameters

In the previous tutorial, we have seen how to create an application that could generate a random string. Let's now assume you wish to indicate the length of that string dynamically.

```

1 import random
2 import string
3
4 import cherrypy
5
6
7 class StringGenerator(object):
8     @cherrypy.expose
9     def index(self):
10         return "Hello world!"
11
12     @cherrypy.expose
13     def generate(self, length=8):
14         return ''.join(random.sample(string.hexdigits, int(length)))
15
16
17 if __name__ == '__main__':
18     cherrypy.quickstart(StringGenerator())

```

Save this into a file named `tut03.py` and run it as follows:

```
$ python tut03.py
```

Go now to <http://localhost:8080/generate?length=16> and your browser will display a generated string of length 16. Notice how we benefit from Python's default arguments' values to support URLs such as <http://localhost:8080/generate> still.

In a URL such as this one, the section after `?` is called a query-string. Traditionally, the query-string is used to contextualize the URL by passing a set of (key, value) pairs. The format for those pairs is `key=value`. Each pair being separated by a `&` character.

Notice how we have to convert the given `length` value to an integer. Indeed, values are sent out from the client to our server as strings.

Much like CherryPy maps URL path segments to exposed functions, query-string keys are mapped to those exposed function parameters.

3.4 Tutorial 4: Submit this form

CherryPy is a web framework upon which you build web applications. The most traditional shape taken by applications is through an HTML user-interface speaking to your CherryPy server.

Let's see how to handle HTML forms via the following example.

```
1 import random
2 import string
3
4 import cherrypy
5
6
7 class StringGenerator(object):
8     @cherrypy.expose
9     def index(self):
10        return """<html>
11            <head></head>
12            <body>
13                <form method="get" action="generate">
14                    <input type="text" value="8" name="length" />
15                    <button type="submit">Give it now!</button>
16                </form>
17            </body>
18        </html>"""
19
20    @cherrypy.expose
21    def generate(self, length=8):
22        return ''.join(random.sample(string.hexdigits, int(length)))
23
24
25 if __name__ == '__main__':
26    cherrypy.quickstart(StringGenerator())
```

Save this into a file named *tut04.py* and run it as follows:

```
$ python tut04.py
```

Go now to <http://localhost:8080/> and your browser and this will display a simple input field to indicate the length of the string you want to generate.

Notice that in this example, the form uses the *GET* method and when you pressed the *Give it now!* button, the form is sent using the same URL as in the *previous* tutorial. HTML forms also support the *POST* method, in that case the query-string is not appended to the URL but it sent as the body of the client's request to the server. However, this would not change your application's exposed method because CherryPy handles both the same way and uses the exposed's handler parameters to deal with the query-string (key, value) pairs.

3.5 Tutorial 5: Track my end-user's activity

It's not uncommon that an application needs to follow the user's activity for a while. The usual mechanism is to use a *session identifier* that is carried during the conversation between the user and your application.

```
1 import random
2 import string
3
4 import cherrypy
```

```

5
6
7 class StringGenerator(object):
8     @cherrypy.expose
9     def index(self):
10        return """<html>
11            <head></head>
12            <body>
13                <form method="get" action="generate">
14                    <input type="text" value="8" name="length" />
15                    <button type="submit">Give it now!</button>
16                </form>
17            </body>
18        </html>"""
19
20    @cherrypy.expose
21    def generate(self, length=8):
22        some_string = ''.join(random.sample(string.hexdigits, int(length)))
23        cherrypy.session['mystring'] = some_string
24        return some_string
25
26    @cherrypy.expose
27    def display(self):
28        return cherrypy.session['mystring']
29
30
31 if __name__ == '__main__':
32     conf = {
33         '/': {
34             'tools.sessions.on': True
35         }
36     }
37     cherrypy.quickstart(StringGenerator(), '/', conf)

```

Save this into a file named *tut05.py* and run it as follows:

```
$ python tut05.py
```

In this example, we generate the string as in the *previous* tutorial but also store it in the current session. If you go to <http://localhost:8080/>, generate a random string, then go to <http://localhost:8080/display>, you will see the string you just generated.

The lines 30-34 show you how to enable the session support in your CherryPy application. By default, CherryPy will save sessions in the process's memory. It supports more persistent *backends* as well.

3.6 Tutorial 6: What about my javascripts, CSS and images?

Web applications are usually also made of static content such as javascript, CSS files or images. CherryPy provides support to serve static content to end-users.

Let's assume, you want to associate a stylesheet with your application to display a blue background color (why not?).

First, save the following stylesheet into a file named *style.css* and stored into a local directory *public/css*.

```

1 body {
2     background-color: blue;
3 }

```

Now let's update the HTML code so that we link to the stylesheet using the <http://localhost:8080/static/css/style.css> URL.

```
1 import os, os.path
2 import random
3 import string
4
5 import cherrypy
6
7
8 class StringGenerator(object):
9     @cherrypy.expose
10    def index(self):
11        return """<html>
12            <head>
13                <link href="/static/css/style.css" rel="stylesheet">
14            </head>
15            <body>
16                <form method="get" action="generate">
17                    <input type="text" value="8" name="length" />
18                    <button type="submit">Give it now!</button>
19                </form>
20            </body>
21        </html>"""
22
23    @cherrypy.expose
24    def generate(self, length=8):
25        some_string = ''.join(random.sample(string.hexdigits, int(length)))
26        cherrypy.session['mystring'] = some_string
27        return some_string
28
29    @cherrypy.expose
30    def display(self):
31        return cherrypy.session['mystring']
32
33
34 if __name__ == '__main__':
35     conf = {
36         '/': {
37             'tools.sessions.on': True,
38             'tools.staticdir.root': os.path.abspath(os.getcwd())
39         },
40         '/static': {
41             'tools.staticdir.on': True,
42             'tools.staticdir.dir': './public'
43         }
44     }
45     cherrypy.quickstart(StringGenerator(), '/', conf)
```

Save this into a file named *tut06.py* and run it as follows:

```
$ python tut06.py
```

Going to <http://localhost:8080/>, you should be greeted by a flashy blue color.

CherryPy provides support to serve a single file or a complete directory structure. Most of the time, this is what you'll end up doing so this is what the code above demonstrates. First, we indicate the *root* directory of all of our static content. This must be an absolute path for security reason. CherryPy will complain if you provide only relative paths when looking for a match to your URLs.

Then we indicate that all URLs which path segment starts with */static* will be served as static content. We map that URL to the *public* directory, a direct child of the *root* directory. The entire sub-tree of the *public* directory will be served as static content. CherryPy will map URLs to path within that directory. This is why */static/css/style.css* is found in *public/css/style.css*.

3.7 Tutorial 7: Give us a REST

It's not unusual nowadays that web applications expose some sort of datamodel or computation functions. Without going into its details, one strategy is to follow the [REST principles edicted by Roy T. Fielding](#).

Roughly speaking, it assumes that you can identify a resource and that you can address that resource through that identifier.

“What for?” you may ask. Well, mostly, these principles are there to ensure that you decouple, as best as you can, the entities your application expose from the way they are manipulated or consumed. To embrace this point of view, developers will usually design a web API that expose pairs of (*URL, HTTP method, data, constraints*).

Note: You will often hear REST and web API together. The former is one strategy to provide the latter. This tutorial will not go deeper in that whole web API concept as it's a much more engaging subject, but you ought to read more about it online.

Lets go through a small example of a very basic web API mildly following REST principles.

```

1 import random
2 import string
3
4 import cherrypy
5
6
7 @cherrypy.expose
8 class StringGeneratorWebService(object):
9
10     @cherrypy.tools.accept(media='text/plain')
11     def GET(self):
12         return cherrypy.session['mystring']
13
14     def POST(self, length=8):
15         some_string = ''.join(random.sample(string.hexdigits, int(length)))
16         cherrypy.session['mystring'] = some_string
17         return some_string
18
19     def PUT(self, another_string):
20         cherrypy.session['mystring'] = another_string
21
22     def DELETE(self):
23         cherrypy.session.pop('mystring', None)
24
25
26 if __name__ == '__main__':
27     conf = {
28         '/': {
29             'request.dispatch': cherrypy.dispatch.MethodDispatcher(),
30             'tools.sessions.on': True,
31             'tools.response_headers.on': True,
32             'tools.response_headers.headers': [('Content-Type', 'text/plain')],

```

```
33     }
34 }
35 cherrypy.quickstart(StringGeneratorWebService(), '/', conf)
```

Save this into a file named *tut07.py* and run it as follows:

```
$ python tut07.py
```

Before we see it in action, let's explain a few things. Until now, CherryPy was creating a tree of exposed methods that were used to match URLs. In the case of our web API, we want to stress the role played by the actual requests' HTTP methods. So we created methods that are named after them and they are all exposed at once by decorating the class itself with *cherrypy.expose*.

However, we must then switch from the default mechanism of matching URLs to method for one that is aware of the whole HTTP method shenanigan. This is what goes on line 27 where we create a *MethodDispatcher* instance.

Then we force the responses *content-type* to be *text/plain* and we finally ensure that *GET* requests will only be responded to clients that accept that *content-type* by having a *Accept: text/plain* header set in their request. However, we do this only for that HTTP method as it wouldn't have much meaning on the other methods.

For the purpose of this tutorial, we will be using a Python client rather than your browser as we wouldn't be able to actually try our web API otherwise.

Please install [requests](#) through the following command:

```
$ pip install requests
```

Then fire up a Python terminal and try the following commands:

```
1 >>> import requests
2 >>> s = requests.Session()
3 >>> r = s.get('http://127.0.0.1:8080/')
4 >>> r.status_code
5 500
6 >>> r = s.post('http://127.0.0.1:8080/')
7 >>> r.status_code, r.text
8 (200, u'04A92138')
9 >>> r = s.get('http://127.0.0.1:8080/')
10 >>> r.status_code, r.text
11 (200, u'04A92138')
12 >>> r = s.get('http://127.0.0.1:8080/', headers={'Accept': 'application/json'})
13 >>> r.status_code
14 406
15 >>> r = s.put('http://127.0.0.1:8080/', params={'another_string': 'hello'})
16 >>> r = s.get('http://127.0.0.1:8080/')
17 >>> r.status_code, r.text
18 (200, u'hello')
19 >>> r = s.delete('http://127.0.0.1:8080/')
20 >>> r = s.get('http://127.0.0.1:8080/')
21 >>> r.status_code
22 500
```

The first and last 500 responses stem from the fact that, in the first case, we haven't yet generated a string through *POST* and, on the latter case, that it doesn't exist after we've deleted it.

Lines 12-14 show you how the application reacted when our client requested the generated string as a JSON format. Since we configured the web API to only support plain text, it returns the appropriate [HTTP error code](#).

Note: We use the [Session](#) interface of *requests* so that it takes care of carrying the session id stored in the request

cookie in each subsequent request. That is handy.

Important: It's all about RESTful URLs these days, isn't it?

It is likely your URL will be made of dynamic parts that you will not be able to match to page handlers. For example, `/library/12/book/15` cannot be directly handled by the default CherryPy dispatcher since the segments `12` and `15` will not be matched to any Python callable.

This can be easily workarounded with two handy CherryPy features explained in the [advanced section](#).

3.8 Tutorial 8: Make it smoother with Ajax

In the recent years, web applications have moved away from the simple pattern of “HTML forms + refresh the whole page”. This traditional scheme still works very well but users have become used to web applications that don't refresh the entire page. Broadly speaking, web applications carry code performed client-side that can speak with the backend without having to refresh the whole page.

This tutorial will involve a little more code this time around. First, let's see our CSS stylesheet located in `public/css/style.css`.

```
1 body {
2   background-color: blue;
3 }
4
5 #the-string {
6   display: none;
7 }
```

We're adding a simple rule about the element that will display the generated string. By default, let's not show it up. Save the following HTML code into a file named `index.html`.

```
1 <!DOCTYPE html>
2 <html>
3   <head>
4     <link href="/static/css/style.css" rel="stylesheet">
5     <script src="http://code.jquery.com/jquery-2.0.3.min.js"></script>
6     <script type="text/javascript">
7       $(document).ready(function() {
8
9         $("#generate-string").click(function(e) {
10          $.post("/generator", {"length": $("#input[name='length']").val()})
11            .done(function(string) {
12              $("#the-string").show();
13              $("#the-string input").val(string);
14            });
15          e.preventDefault();
16        });
17
18        $("#replace-string").click(function(e) {
19          $.ajax({
20            type: "PUT",
21            url: "/generator",
22            data: {"another_string": $("#the-string input").val()}
23          })
24          .done(function() {
```

```

25     alert("Replaced!");
26     });
27     e.preventDefault();
28     });
29
30     $("#delete-string").click(function(e) {
31     $.ajax({
32         type: "DELETE",
33         url: "/generator"
34     })
35     .done(function() {
36         $("#the-string").hide();
37     });
38     e.preventDefault();
39     });
40
41     });
42 </script>
43 </head>
44 <body>
45     <input type="text" value="8" name="length"/>
46     <button id="generate-string">Give it now!</button>
47     <div id="the-string">
48         <input type="text" />
49         <button id="replace-string">Replace</button>
50         <button id="delete-string">Delete it</button>
51     </div>
52 </body>
53 </html>

```

We'll be using the jQuery framework out of simplicity but feel free to replace it with your favourite tool. The page is composed of simple HTML elements to get user input and display the generated string. It also contains client-side code to talk to the backend API that actually performs the hard work.

Finally, here's the application's code that serves the HTML page above and responds to requests to generate strings. Both are hosted by the same application server.

```

1  import os, os.path
2  import random
3  import string
4
5  import cherrypy
6
7
8  class StringGenerator(object):
9      @cherrypy.expose
10     def index(self):
11         return open('index.html')
12
13
14 @cherrypy.expose
15 class StringGeneratorWebService(object):
16
17     @cherrypy.tools.accept(media='text/plain')
18     def GET(self):
19         return cherrypy.session['mystring']
20
21     def POST(self, length=8):
22         some_string = ''.join(random.sample(string.hexdigits, int(length)))

```



```

23     cherrypy.session['mystring'] = some_string
24     return some_string
25
26     def PUT(self, another_string):
27         cherrypy.session['mystring'] = another_string
28
29     def DELETE(self):
30         cherrypy.session.pop('mystring', None)
31
32
33 if __name__ == '__main__':
34     conf = {
35         '/': {
36             'tools.sessions.on': True,
37             'tools.staticdir.root': os.path.abspath(os.getcwd())
38         },
39         '/generator': {
40             'request.dispatch': cherrypy.dispatch.MethodDispatcher(),
41             'tools.response_headers.on': True,
42             'tools.response_headers.headers': [('Content-Type', 'text/plain')],
43         },
44         '/static': {
45             'tools.staticdir.on': True,
46             'tools.staticdir.dir': './public'
47         }
48     }
49     webapp = StringGenerator()
50     webapp.generator = StringGeneratorWebService()
51     cherrypy.quickstart(webapp, '/', conf)

```

Save this into a file named *tut08.py* and run it as follows:

```
$ python tut08.py
```

Go to <http://127.0.0.1:8080/> and play with the input and buttons to generate, replace or delete the strings. Notice how the page isn't refreshed, simply part of its content.

Notice as well how your frontend converses with the backend using a straightforward, yet clean, web service API. That same API could easily be used by non-HTML clients.

3.9 Tutorial 9: Data is all my life

Until now, all the generated strings were saved in the session, which by default is stored in the process memory. Though, you can persist sessions on disk or in a distributed memory store, this is not the right way of keeping your data on the long run. Sessions are there to identify your user and carry as little amount of data as necessary for the operation carried by the user.

To store, persist and query data you need a proper database server. There exist many to choose from with various paradigm support:

- relational: PostgreSQL, SQLite, MariaDB, Firebird
- column-oriented: HBase, Cassandra
- key-store: redis, memcached
- document oriented: Couchdb, MongoDB
- graph-oriented: neo4j

Let's focus on the relational ones since they are the most common and probably what you will want to learn first.

For the sake of reducing the number of dependencies for these tutorials, we will go for the `sqlite` database which is directly supported by Python.

Our application will replace the storage of the generated string from the session to a SQLite database. The application will have the same HTML code as *tutorial 08*. So let's simply focus on the application code itself:

```
1 import os, os.path
2 import random
3 import sqlite3
4 import string
5 import time
6
7 import cherrypy
8
9 DB_STRING = "my.db"
10
11
12 class StringGenerator(object):
13     @cherrypy.expose
14     def index(self):
15         return open('index.html')
16
17
18 @cherrypy.expose
19 class StringGeneratorWebService(object):
20
21     @cherrypy.tools.accept(media='text/plain')
22     def GET(self):
23         with sqlite3.connect(DB_STRING) as c:
24             cherrypy.session['ts'] = time.time()
25             r = c.execute("SELECT value FROM user_string WHERE session_id=?",
26                           [cherrypy.session.id])
27             return r.fetchone()
28
29     def POST(self, length=8):
30         some_string = ''.join(random.sample(string.hexdigits, int(length)))
31         with sqlite3.connect(DB_STRING) as c:
32             cherrypy.session['ts'] = time.time()
33             c.execute("INSERT INTO user_string VALUES (?, ?)",
34                       [cherrypy.session.id, some_string])
35         return some_string
36
37     def PUT(self, another_string):
38         with sqlite3.connect(DB_STRING) as c:
39             cherrypy.session['ts'] = time.time()
40             c.execute("UPDATE user_string SET value=? WHERE session_id=?",
41                       [another_string, cherrypy.session.id])
42
43     def DELETE(self):
44         cherrypy.session.pop('ts', None)
45         with sqlite3.connect(DB_STRING) as c:
46             c.execute("DELETE FROM user_string WHERE session_id=?",
47                       [cherrypy.session.id])
48
49
50 def setup_database():
51     """
```

```

52  Create the `user_string` table in the database
53  on server startup
54  """
55  with sqlite3.connect(DB_STRING) as con:
56      con.execute("CREATE TABLE user_string (session_id, value)")
57
58
59  def cleanup_database():
60      """
61      Destroy the `user_string` table from the database
62      on server shutdown.
63      """
64      with sqlite3.connect(DB_STRING) as con:
65          con.execute("DROP TABLE user_string")
66
67
68  if __name__ == '__main__':
69      conf = {
70          '/': {
71              'tools.sessions.on': True,
72              'tools.staticdir.root': os.path.abspath(os.getcwd())
73          },
74          '/generator': {
75              'request.dispatch': cherrypy.dispatch.MethodDispatcher(),
76              'tools.response_headers.on': True,
77              'tools.response_headers.headers': [('Content-Type', 'text/plain')],
78          },
79          '/static': {
80              'tools.staticdir.on': True,
81              'tools.staticdir.dir': './public'
82          }
83      }
84
85      cherrypy.engine.subscribe('start', setup_database)
86      cherrypy.engine.subscribe('stop', cleanup_database)
87
88      webapp = StringGenerator()
89      webapp.generator = StringGeneratorWebService()
90      cherrypy.quickstart(webapp, '/', conf)

```

Save this into a file named *tut09.py* and run it as follows:

```
$ python tut09.py
```

Let's first see how we create two functions that create and destroy the table within our database. These functions are registered to the CherryPy's server on lines 85-86, so that they are called when the server starts and stops.

Next, notice how we replaced all the session code with calls to the database. We use the session id to identify the user's string within our database. Since the session will go away after a while, it's probably not the right approach. A better idea would be to associate the user's login or more resilient unique identifier. For the sake of our demo, this should do.

Important: In this example, we must still set the session to a dummy value so that the session is not `discarded` on each request by CherryPy. Since we now use the database to store the generated string, we simply store a dummy timestamp inside the session.

Note: Unfortunately, `sqlite` in Python forbids us to share a connection between threads. Since CherryPy is a multi-threaded server, this would be an issue. This is the reason why we open and close a connection to the database on each call. This is clearly not really production friendly, and it is probably advisable to either use a more capable database engine or a higher level library, such as [SQLAlchemy](#), to better support your application's needs.

3.10 Tutorial 10: Make it a modern single-page application with React.js

In the recent years, client-side single-page applications (SPA) have gradually eaten server-side generated content web applications's lunch.

This tutorial demonstrates how to integrate with [React.js](#), a Javascript library for SPA released by Facebook in 2013. Please refer to [React.js](#) documentation to learn more about it.

To demonstrate it, let's use the code from [tutorial 09](#). However, we will be replacing the HTML and Javascript code.

First, let's see how our HTML code has changed:

```
1 <!DOCTYPE html>
2 <html>
3   <head>
4     <link href="/static/css/style.css" rel="stylesheet">
5     <script src="https://cdnjs.cloudflare.com/ajax/libs/react/0.13.3/react.js"></script>
6     <script src="http://code.jquery.com/jquery-2.1.1.min.js"></script>
7     <script src="https://cdnjs.cloudflare.com/ajax/libs/babel-core/5.8.23/browser.min.js"></script>
8   </head>
9   <body>
10    <div id="generator"></div>
11    <script type="text/babel" src="static/js/gen.js"></script>
12  </body>
13 </html>
```

Basically, we have removed the entire Javascript code that was using jQuery. Instead, we load the `React.js` library as well as a new, local, Javascript module, named `gen.js` and located in the `public/js` directory:

```
1 var StringGeneratorBox = React.createClass({
2   handleGenerate: function() {
3     var length = this.state.length;
4     this.setState(function() {
5       $.ajax({
6         url: this.props.url,
7         dataType: 'text',
8         type: 'POST',
9         data: {
10          "length": length
11        },
12        success: function(data) {
13          this.setState({
14            length: length,
15            string: data,
16            mode: "edit"
17          });
18        }.bind(this),
19        error: function(xhr, status, err) {
20          console.error(this.props.url,
```

```
21         status, err.toString()
22     );
23     }.bind(this)
24 });
25 });
26 },
27 handleEdit: function() {
28     var new_string = this.state.string;
29     this.setState(function() {
30         $.ajax({
31             url: this.props.url,
32             type: 'PUT',
33             data: {
34                 "another_string": new_string
35             },
36             success: function() {
37                 this.setState({
38                     length: new_string.length,
39                     string: new_string,
40                     mode: "edit"
41                 });
42             }.bind(this),
43             error: function(xhr, status, err) {
44                 console.error(this.props.url,
45                     status, err.toString()
46                 );
47             }.bind(this)
48         });
49     });
50 },
51 handleDelete: function() {
52     this.setState(function() {
53         $.ajax({
54             url: this.props.url,
55             type: 'DELETE',
56             success: function() {
57                 this.setState({
58                     length: "8",
59                     string: "",
60                     mode: "create"
61                 });
62             }.bind(this),
63             error: function(xhr, status, err) {
64                 console.error(this.props.url,
65                     status, err.toString()
66                 );
67             }.bind(this)
68         });
69     });
70 },
71 handleLengthChange: function(length) {
72     this.setState({
73         length: length,
74         string: "",
75         mode: "create"
76     });
77 },
78 handleStringChange: function(new_string) {
```

```

79     this.setState({
80         length: new_string.length,
81         string: new_string,
82         mode: "edit"
83     });
84 },
85 getInitialState: function() {
86     return {
87         length: "8",
88         string: "",
89         mode: "create"
90     };
91 },
92 render: function() {
93     return (
94         <div className="stringGenBox">
95             <StringGeneratorForm onCreateString={this.handleGenerate}
96                                 onReplaceString={this.handleEdit}
97                                 onDeleteString={this.handleDelete}
98                                 onLengthChange={this.handleLengthChange}
99                                 onStringChange={this.handleStringChange}
100                                mode={this.state.mode}
101                                length={this.state.length}
102                                string={this.state.string}/>
103         </div>
104     );
105 }
106 });
107
108 var StringGeneratorForm = React.createClass({
109     handleCreate: function(e) {
110         e.preventDefault();
111         this.props.onCreateString();
112     },
113     handleReplace: function(e) {
114         e.preventDefault();
115         this.props.onReplaceString();
116     },
117     handleDelete: function(e) {
118         e.preventDefault();
119         this.props.onDeleteString();
120     },
121     handleLengthChange: function(e) {
122         e.preventDefault();
123         var length = React.findDOMNode(this.refs.length).value.trim();
124         this.props.onLengthChange(length);
125     },
126     handleStringChange: function(e) {
127         e.preventDefault();
128         var string = React.findDOMNode(this.refs.string).value.trim();
129         this.props.onStringChange(string);
130     },
131     render: function() {
132         if (this.props.mode == "create") {
133             return (
134                 <div>
135                     <input type="text" ref="length" defaultValue="8" value={this.props.length} onChange={this
136                     <button onClick={this.handleCreate}>Give it now!</button>

```

```

137     </div>
138     );
139   } else if (this.props.mode == "edit") {
140     return (
141       <div>
142         <input type="text" ref="string" value={this.props.string} onChange={this.handleStringChange} />
143         <button onClick={this.handleReplace}>Replace</button>
144         <button onClick={this.handleDelete}>Delete it</button>
145       </div>
146     );
147   }
148
149   return null;
150 }
151 });
152
153 React.render(
154   <StringGeneratorBox url="/generator" />,
155   document.getElementById('generator')
156 );

```

Wow! What a lot of code for something so simple, isn't it? The entry point is the last few lines where we indicate that we want to render the HTML code of the `StringGeneratorBox` React.js class inside the `generator` div.

When the page is rendered, so is that component. Notice how it is also made of another component that renders the form itself.

This might be a little over the top for such a simple example but hopefully will get you started with React.js in the process.

There is not much to say and, hopefully, the meaning of that code is rather clear. The component has an internal `state` in which we store the current string as generated/modified by the user.

When the user [changes the content of the input boxes](#), the state is updated on the client side. Then, when a button is clicked, that state is sent out to the backend server using the API endpoint and the appropriate action takes places. Then, the state is updated and so is the view.

3.11 Tutorial 11: Organize my code

CherryPy comes with a powerful architecture that helps you organizing your code in a way that should make it easier to maintain and more flexible.

Several mechanisms are at your disposal, this tutorial will focus on the three main ones:

- *dispatchers*
- *tools*
- *plugins*

In order to understand them, let's imagine you are at a superstore:

- You have several tills and people queuing for each of them (those are your requests)
- You have various sections with food and other stuff (these are your data)
- Finally you have the superstore people and their daily tasks to make sure sections are always in order (this is your backend)

In spite of being really simplistic, this is not far from how your application behaves. CherryPy helps you structure your application in a way that mirrors these high-level ideas.

3.11.1 Dispatchers

Coming back to the superstore example, it is likely that you will want to perform operations based on the till:

- Have a till for baskets with less than ten items
- Have a till for disabled people
- Have a till for pregnant women
- Have a till where you can only using the store card

To support these use-cases, CherryPy provides a mechanism called a *dispatcher*. A dispatcher is executed early during the request processing in order to determine which piece of code of your application will handle the incoming request. Or, to continue on the store analogy, a dispatcher will decide which till to lead a customer to.

3.11.2 Tools

Let's assume your store has decided to operate a discount spree but, only for a specific category of customers. CherryPy will deal with such use case via a mechanism called a *tool*.

A tool is a piece of code that runs on a per-request basis in order to perform additional work. Usually a tool is a simple Python function that is executed at a given point during the process of the request by CherryPy.

3.11.3 Plugins

As we have seen, the store has a crew of people dedicated to manage the stock and deal with any customers' expectation.

In the CherryPy world, this translates into having functions that run outside of any request life-cycle. These functions should take care of background tasks, long lived connections (such as those to a database for instance), etc.

Plugins are called that way because they work along with the CherryPy *engine* and extend it with your operations.

Basics

The following sections will drive you through the basics of a CherryPy application, introducing some essential concepts.

Contents

- *Basics*
 - *The one-minute application example*
 - *Hosting one or more applications*
 - * *Single application*
 - * *Multiple applications*
 - *Logging*
 - * *Disable logging*
 - * *Play along with your other loggers*
 - *Configuring*
 - * *Global server configuration*
 - * *Per-application configuration*
 - * *Additional application settings*
 - *Cookies*
 - *Using sessions*
 - * *Filesystem backend*
 - * *Memcached backend*
 - * *Other backends*
 - *Static content serving*
 - * *Serving a single file*
 - * *Serving a whole directory*
 - * *Specifying an index file*
 - * *Allow files downloading*
 - *Dealing with JSON*
 - * *Decoding request*
 - * *Encoding response*
 - *Authentication*
 - * *Basic*
 - * *Digest*
 - *Favicon*

4.1 The one-minute application example

The most basic application you can write with CherryPy involves almost all its core concepts.

```
1 import cherrypy
2
3 class Root(object):
4     @cherrypy.expose
5     def index(self):
6         return "Hello World!"
7
8 if __name__ == '__main__':
9     cherrypy.quickstart(Root(), '/')
```

First and foremost, for most tasks, you will never need more than a single import statement as demonstrated in line 1.

Before discussing the meat, let's jump to line 9 which shows, how to host your application with the CherryPy application server and serve it with its builtin HTTP server at the `/` path. All in one single line. Not bad.

Let's now step back to the actual application. Even though CherryPy does not mandate it, most of the time your applications will be written as Python classes. Methods of those classes will be called by CherryPy to respond to client requests. However, CherryPy needs to be aware that a method can be used that way, we say the method needs to be *exposed*. This is precisely what the `cherrypy.expose()` decorator does in line 4.

Save the snippet in a file named *myapp.py* and run your first CherryPy application:

```
$ python myapp.py
```

Then point your browser at <http://127.0.0.1:8080>. Tada!

Note: CherryPy is a small framework that focuses on one single task: take a HTTP request and locate the most appropriate Python function or method that match the request's URL. Unlike other well-known frameworks, CherryPy does not provide a built-in support for database access, HTML templating or any other middleware nifty features.

In a nutshell, once CherryPy has found and called an *exposed* method, it is up to you, as a developer, to provide the tools to implement your application's logic.

CherryPy takes the opinion that you, the developer, know best.

Warning: The previous example demonstrated the simplicity of the CherryPy interface but, your application will likely contain a few other bits and pieces: static service, more complex structure, database access, etc. This will be developed in the tutorial section.

CherryPy is a minimal framework but not a bare one, it comes with a few basic tools to cover common usages that you would expect.

4.2 Hosting one or more applications

A web application needs an HTTP server to be accessed to. CherryPy provides its own, production ready, HTTP server. There are two ways to host an application with it. The simple one and the almost-as-simple one.

4.2.1 Single application

The most straightforward way is to use `cherrypy.quickstart()` function. It takes at least one argument, the instance of the application to host. Two other settings are optionals. First, the base path at which the application will be accessible from. Second, a config dictionary or file to configure your application.

```
cherrypy.quickstart(Blog())
cherrypy.quickstart(Blog(), '/blog')
cherrypy.quickstart(Blog(), '/blog', {'/': {'tools.gzip.on': True}})
```

The first one means that your application will be available at `http://hostname:port/` whereas the other two will make your blog application available at `http://hostname:port/blog`. In addition, the last one provides specific settings for the application.

Note: Notice in the third case how the settings are still relative to the application, not where it is made available at, hence the `{'/': ... }` rather than a `{'/blog': ... }`

4.2.2 Multiple applications

The `cherrypy.quickstart()` approach is fine for a single application, but lacks the capacity to host several applications with the server. To achieve this, one must use the `cherrypy.tree.mount` function as follows:

```
cherrypy.tree.mount(Blog(), '/blog', blog_conf)
cherrypy.tree.mount(Forum(), '/forum', forum_conf)

cherrypy.engine.start()
cherrypy.engine.block()
```

Essentially, `cherrypy.tree.mount` takes the same parameters as `cherrypy.quickstart()`: an *application*, a hosting path segment and a configuration. The last two lines are simply starting application server.

Important: `cherrypy.quickstart()` and `cherrypy.tree.mount` are not exclusive. For instance, the previous lines can be written as:

```
cherrypy.tree.mount(Blog(), '/blog', blog_conf)
cherrypy.quickstart(Forum(), '/forum', forum_conf)
```

Note: You can also *host foreign WSGI application*.

4.3 Logging

Logging is an important task in any application. CherryPy will log all incoming requests as well as protocol errors.

To do so, CherryPy manages two loggers:

- an access one that logs every incoming requests
- an application/error log that traces errors or other application-level messages

Your application may leverage that second logger by calling `cherrypy.log()`.

```
cherrypy.log("hello there")
```

You can also log an exception:

```
try:
    ...
except:
    cherrypy.log("kaboom!", traceback=True)
```

Both logs are writing to files identified by the following keys in your configuration:

- `log.access_file` for incoming requests using the [common log format](#)
- `log.error_file` for the other log

See also:

Refer to the `cherrypy._cplogging` module for more details about CherryPy's logging architecture.

4.3.1 Disable logging

You may be interested in disabling either logs.

To disable file logging, simply set an empty string to the `log.access_file` or `log.error_file` keys in your *global configuration*.

To disable, console logging, set `log.screen` to `False`.

```
cherrypy.config.update({'log.screen': False,
                       'log.access_file': '',
                       'log.error_file': ''})
```

4.3.2 Play along with your other loggers

Your application may obviously already use the `logging` module to trace application level messages. Below is a simple example on setting it up.

```
import logging
import logging.config

import cherrypy

logger = logging.getLogger()
db_logger = logging.getLogger('db')

LOG_CONF = {
    'version': 1,

    'formatters': {
        'void': {
            'format': ''
        },
        'standard': {
            'format': '%(asctime)s [%(levelname)s] %(name)s: %(message)s'
        },
    },
    'handlers': {
        'default': {
```

```

        'level': 'INFO',
        'class': 'logging.StreamHandler',
        'formatter': 'standard',
        'stream': 'ext://sys.stdout'
    },
    'cherrypy_console': {
        'level': 'INFO',
        'class': 'logging.StreamHandler',
        'formatter': 'void',
        'stream': 'ext://sys.stdout'
    },
    'cherrypy_access': {
        'level': 'INFO',
        'class': 'logging.handlers.RotatingFileHandler',
        'formatter': 'void',
        'filename': 'access.log',
        'maxBytes': 10485760,
        'backupCount': 20,
        'encoding': 'utf8'
    },
    'cherrypy_error': {
        'level': 'INFO',
        'class': 'logging.handlers.RotatingFileHandler',
        'formatter': 'void',
        'filename': 'errors.log',
        'maxBytes': 10485760,
        'backupCount': 20,
        'encoding': 'utf8'
    },
},
'loggers': {
    '': {
        'handlers': ['default'],
        'level': 'INFO'
    },
    'db': {
        'handlers': ['default'],
        'level': 'INFO',
        'propagate': False
    },
    'cherrypy.access': {
        'handlers': ['cherrypy_access'],
        'level': 'INFO',
        'propagate': False
    },
    'cherrypy.error': {
        'handlers': ['cherrypy_console', 'cherrypy_error'],
        'level': 'INFO',
        'propagate': False
    },
},
}

class Root(object):
    @cherrypy.expose
    def index(self):

        logger.info("boom")

```

```
db_logger.info("bam")
cherry.py.log("bang")

    return "hello world"

if __name__ == '__main__':
    cherry.py.config.update({'log.screen': False,
                           'log.access_file': '',
                           'log.error_file': ''})
cherry.py.engine.unsubscribe('graceful', cherry.py.log.reopen_files)
logging.config.dictConfig(LOG_CONF)
cherry.py.quickstart(Root())
```

In this snippet, we create a `configuration dictionary` that we pass on to the logging module to configure our loggers:

- the default root logger is associated to a single stream handler
- a logger for the db backend with also a single stream handler

In addition, we re-configure the CherryPy loggers:

- the top-level `cherry.py.access` logger to log requests into a file
- the `cherry.py.error` logger to log everything else into a file and to the console

We also prevent CherryPy from trying to open its log files when the autoreloader kicks in. This is not strictly required since we do not even let CherryPy open them in the first place. But, this avoids wasting time on something useless.

4.4 Configuring

CherryPy comes with a fine-grained configuration mechanism and settings can be set at various levels.

See also:

Once you have reviewed the basics, please refer to the *in-depth discussion* around configuration.

4.4.1 Global server configuration

To configure the HTTP and application servers, use the `cherry.py.config.update()` method.

```
cherry.py.config.update({'server.socket_port': 9090})
```

The `cherry.py.config` object is a dictionary and the update method merges the passed dictionary into it.

You can also pass a file instead (assuming a `server.conf` file):

```
[global]
server.socket_port: 9090
```

```
cherry.py.config.update("server.conf")
```

Warning: `cherry.py.config.update()` is not meant to be used to configure the application. It is a common mistake. It is used to configure the server and engine.

4.4.2 Per-application configuration

To configure your application, pass in a dictionary or a file when you associate your application to the server.

```
cherry.py.quickstart(myapp, '/', {'/': {'tools.gzip.on': True}})
```

or via a file (called *app.conf* for instance):

```
[/]  
tools.gzip.on: True
```

```
cherry.py.quickstart(myapp, '/', "app.conf")
```

Although, you can define most of your configuration in a global fashion, it is sometimes convenient to define them where they are applied in the code.

```
class Root(object):  
    @cherry.py.expose  
    @cherry.py.tools.gzip()  
    def index(self):  
        return "hello world!"
```

A variant notation to the above:

```
class Root(object):  
    @cherry.py.expose  
    def index(self):  
        return "hello world!"  
    index._cp_config = {'tools.gzip.on': True}
```

Both methods have the same effect so pick the one that suits your style best.

4.4.3 Additional application settings

You can add settings that are not specific to a request URL and retrieve them from your page handler as follows:

```
[/]  
tools.gzip.on: True  
  
[googleapi]  
key = "..."  
appid = "..."
```

```
class Root(object):  
    @cherry.py.expose  
    def index(self):  
        google_appid = cherry.py.request.app.config['googleapi']['appid']  
        return "hello world!"  
  
cherry.py.quickstart(Root(), '/', "app.conf")
```

4.5 Cookies

CherryPy uses the `Cookie` module from python and in particular the `Cookie.SimpleCookie` object type to handle cookies.

- To send a cookie to a browser, set `cherry.py.response.cookie[key] = value`.
- To retrieve a cookie sent by a browser, use `cherry.py.request.cookie[key]`.
- To delete a cookie (on the client side), you must *send* the cookie with its expiration time set to 0:

```
cherry.py.response.cookie[key] = value
cherry.py.response.cookie[key]['expires'] = 0
```

It's important to understand that the request cookies are **not** automatically copied to the response cookies. Clients will send the same cookies on every request, and therefore `cherry.py.request.cookie` should be populated each time. But the server doesn't need to send the same cookies with every response; therefore, `cherry.py.response.cookie` will usually be empty. When you wish to "delete" (expire) a cookie, therefore, you must set `cherry.py.response.cookie[key] = value` first, and then set its `expires` attribute to 0.

Extended example:

```
import cherry.py

class MyCookieApp(object):
    @cherry.py.expose
    def set(self):
        cookie = cherry.py.response.cookie
        cookie['cookieName'] = 'cookieValue'
        cookie['cookieName']['path'] = '/'
        cookie['cookieName']['max-age'] = 3600
        cookie['cookieName']['version'] = 1
        return "<html><body>Hello, I just sent you a cookie</body></html>"

    @cherry.py.expose
    def read(self):
        cookie = cherry.py.request.cookie
        res = ""<html><body>Hi, you sent me %s cookies.<br />
            Here is a list of cookie names/values:<br />"" % len(cookie)
        for name in cookie.keys():
            res += "name: %s, value: %s<br>" % (name, cookie[name].value)
        return res + "</body></html>"

if __name__ == '__main__':
    cherry.py.quickstart(MyCookieApp(), '/cookie')
```

4.6 Using sessions

Sessions are one of the most common mechanism used by developers to identify users and synchronize their activity. By default, CherryPy does not activate sessions because it is not a mandatory feature to have, to enable it simply add the following settings in your configuration:

```
[/]
tools.sessions.on: True
```

```
cherry.py.quickstart(myapp, '/', "app.conf")
```

Sessions are, by default, stored in RAM so, if you restart your server all of your current sessions will be lost. You can store them in memcached or on the filesystem instead.

Using sessions in your applications is done as follows:


```
import cherrypy

@cherrypy.expose
def index(self):
    if 'count' not in cherrypy.session:
        cherrypy.session['count'] = 0
    cherrypy.session['count'] += 1
```

In this snippet, everytime the the index page handler is called, the current user's session has its 'count' key incremented by 1.

CherryPy knows which session to use by inspecting the cookie sent alongside the request. This cookie contains the session identifier used by CherryPy to load the user's session from the storage.

See also:

Refer to the `cherrypy.lib.sessions` module for more details about the session interface and implementation. Notably you will learn about sessions expiration.

4.6.1 Filesystem backend

Using a filesystem is a simple to not lose your sessions between reboots. Each session is saved in its own file within the given directory.

```
[/]
tools.sessions.on: True
tools.sessions.storage_class = cherrypy.lib.sessions.FileSession
tools.sessions.storage_path = "/some/directory"
```

4.6.2 Memcached backend

[Memcached](#) is a popular key-store on top of your RAM, it is distributed and a good choice if you want to share sessions outside of the process running CherryPy.

Requires that the Python [memcached](#) package is installed, which may be indicated by installing `cherrypy[memcached_session]`.

```
[/]
tools.sessions.on: True
tools.sessions.storage_class = cherrypy.lib.sessions.MemcachedSession
```

4.6.3 Other backends

Any other library may implement a session backend. Simply subclass `cherrypy.lib.sessions.Session` and indicate that subclass as `tools.sessions.storage_class`.

4.7 Static content serving

CherryPy can serve your static content such as images, javascript and CSS resources, etc.

Note: CherryPy uses the `mimetypes` module to determine the best content-type to serve a particular resource. If the choice is not valid, you can simply set more media-types as follows:

```
import mimetypes
mimetypes.types_map['.csv'] = 'text/csv'
```

4.7.1 Serving a single file

You can serve a single file as follows:

```
[/style.css]
tools.staticfile.on = True
tools.staticfile.filename = "/home/site/style.css"
```

CherryPy will automatically respond to URLs such as *http://hostname/style.css*.

4.7.2 Serving a whole directory

Serving a whole directory is similar to a single file:

```
[/static]
tools.staticdir.on = True
tools.staticdir.dir = "/home/site/static"
```

Assuming you have a file at *static/js/my.js*, CherryPy will automatically respond to URLs such as *http://hostname/static/js/my.js*.

Note: CherryPy always requires the absolute path to the files or directories it will serve. If you have several static sections to configure but located in the same root directory, you can use the following shortcut:

```
[/]
tools.staticdir.root = "/home/site"

[/static]
tools.staticdir.on = True
tools.staticdir.dir = "static"
```

4.7.3 Specifying an index file

By default, CherryPy will respond to the root of a static directory with an 404 error indicating the path */* was not found. To specify an index file, you can use the following:

```
[/static]
tools.staticdir.on = True
tools.staticdir.dir = "/home/site/static"
tools.staticdir.index = "index.html"
```

Assuming you have a file at *static/index.html*, CherryPy will automatically respond to URLs such as *http://hostname/static/* by returning its contents.

4.7.4 Allow files downloading

Using "application/x-download" response content-type, you can tell a browser that a resource should be downloaded onto the user's machine rather than displayed.

You could for instance write a page handler as follows:

```
from cherrypy.lib.static import serve_file

@cherrypy.expose
def download(self, filepath):
    return serve_file(filepath, "application/x-download", "attachment")
```

Assuming the filepath is a valid path on your machine, the response would be considered as a downloadable content by the browser.

Warning: The above page handler is a security risk on its own since any file of the server could be accessed (if the user running the server had permissions on them).

4.8 Dealing with JSON

CherryPy has built-in support for JSON encoding and decoding of the request and/or response.

4.8.1 Decoding request

To automatically decode the content of a request using JSON:

```
class Root(object):
    @cherrypy.expose
    @cherrypy.tools.json_in()
    def index(self):
        data = cherrypy.request.json
```

The *json* attribute attached to the request contains the decoded content.

4.8.2 Encoding response

To automatically encode the content of a response using JSON:

```
class Root(object):
    @cherrypy.expose
    @cherrypy.tools.json_out()
    def index(self):
        return {'key': 'value'}
```

CherryPy will encode any content returned by your page handler using JSON. Not all type of objects may natively be encoded.

4.9 Authentication

CherryPy provides support for two very simple authentication mechanisms, both described in [RFC 2617](#): Basic and Digest. They are most commonly known to trigger a browser's popup asking users their name and password.

4.9.1 Basic

Basic authentication is the simplest form of authentication however it is not a secure one as the user's credentials are embedded into the request. We advise against using it unless you are running on SSL or within a closed network.

```
from cherrypy.lib import auth_basic

USERS = {'jon': 'secret'}

def validate_password(realm, username, password):
    if username in USERS and USERS[username] == password:
        return True
    return False

conf = {
    '/protected/area': {
        'tools.auth_basic.on': True,
        'tools.auth_basic.realm': 'localhost',
        'tools.auth_basic.checkpassword': validate_password
    }
}

cherrypy.quickstart(myapp, '/', conf)
```

Simply put, you have to provide a function that will be called by CherryPy passing the username and password decoded from the request.

The function can read its data from any source it has to: a file, a database, memory, etc.

4.9.2 Digest

Digest authentication differs by the fact the credentials are not carried on by the request so it's a little more secure than basic.

CherryPy's digest support has a similar interface to the basic one explained above.

```
from cherrypy.lib import auth_digest

USERS = {'jon': 'secret'}

conf = {
    '/protected/area': {
        'tools.auth_digest.on': True,
        'tools.auth_digest.realm': 'localhost',
        'tools.auth_digest.get_ha1': auth_digest.get_ha1_dict_plain(USERS),
        'tools.auth_digest.key': 'a565c27146791cfb'
    }
}

cherrypy.quickstart(myapp, '/', conf)
```

4.10 Favicon

CherryPy serves its own sweet red cherrypy as the default [favicon](#) using the static file tool. You can serve your own favicon as follows:

```
import cherrypy

class HelloWorld(object):
    @cherrypy.expose
    def index(self):
        return "Hello World!"

if __name__ == '__main__':
    cherrypy.quickstart(HelloWorld(), '/',
        {
            '/favicon.ico':
                {
                    'tools.staticfile.on': True,
                    'tools.staticfile.filename': '/path/to/myfavicon.ico'
                }
        }
    )
```

Please refer to the *static serving* section for more details.

You can also use a file to configure it:

```
[/favicon.ico]
tools.staticfile.on: True
tools.staticfile.filename: "/path/to/myfavicon.ico"
```

```
import cherrypy

class HelloWorld(object):
    @cherrypy.expose
    def index(self):
        return "Hello World!"

if __name__ == '__main__':
    cherrypy.quickstart(HelloWorld(), '/', app.conf)
```


CherryPy has support for more advanced features that these sections will describe.

Contents

- *Advanced*
 - *Set aliases to page handlers*
 - *RESTful-style dispatching*
 - * *The special `_cp_dispatch` method*
 - * *The `popargs` decorator*
 - *Error handling*
 - *Streaming the response body*
 - * *The “normal” CherryPy response process*
 - * *How “streaming output” works with CherryPy*
 - *Response timeouts*
 - * *Timeout Monitor*
 - *Deal with signals*
 - * *Windows Console Events*
 - *Securing your server*
 - *Multiple HTTP servers support*
 - *WSGI support*
 - * *Make your CherryPy application a WSGI application*
 - * *Host a foreign WSGI application in CherryPy*
 - * *No need for the WSGI interface?*
 - *WebSocket support*
 - *Database support*
 - *HTML Templating support*
 - *Testing your application*

5.1 Set aliases to page handlers

A fairly unknown, yet useful, feature provided by the `cherrypy.expose()` decorator is to support aliases.

Let's use the template provided by *tutorial 03*:

```
import random
import string

import cherrypy
```

```
class StringGenerator(object):
    @cherry.py.expose(['generer', 'generar'])
    def generate(self, length=8):
        return ''.join(random.sample(string.hexdigits, int(length)))

if __name__ == '__main__':
    cherry.py.quickstart(StringGenerator())
```

In this example, we create localized aliases for the page handler. This means the page handler will be accessible via:

- /generate
- /generer (French)
- /generar (Spanish)

Obviously, your aliases may be whatever suits your needs.

Note: The alias may be a single string or a list of them.

5.2 RESTful-style dispatching

The term *RESTful URL* is sometimes used to talk about friendly URLs that nicely map to the entities an application exposes.

Important: We will not enter the debate around what is restful or not but we will showcase two mechanisms to implement the usual idea in your CherryPy application.

Let's assume you wish to create an application that exposes music bands and their records. Your application will probably have the following URLs:

- [http://hostname/<artist>/](http://hostname/<artist>)
- [http://hostname/<artist>/albums/<album_title>/](http://hostname/<artist>/albums/<album_title>)

It's quite clear you would not create a page handler named after every possible band in the world. This means you will need a page handler that acts as a proxy for all of them.

The default dispatcher cannot deal with that scenario on its own because it expects page handlers to be explicitly declared in your source code. Luckily, CherryPy provides ways to support those use cases.

See also:

This section extends from this [stackoverflow response](#).

5.2.1 The special `_cp_dispatch` method

`_cp_dispatch` is a special method you declare in any of your *controller* to massage the remaining segments before CherryPy gets to process them. This offers you the capacity to remove, add or otherwise handle any segment you wish and, even, entirely change the remaining parts.

```
import cherry.py

class Band(object):
```



```

def __init__(self):
    self.albums = Album()

def _cp_dispatch(self, vpath):
    if len(vpath) == 1:
        cherrypy.request.params['name'] = vpath.pop()
        return self

    if len(vpath) == 3:
        cherrypy.request.params['artist'] = vpath.pop(0) # /band name/
        vpath.pop(0) # /albums/
        cherrypy.request.params['title'] = vpath.pop(0) # /album title/
        return self.albums

    return vpath

@cherrypy.expose
def index(self, name):
    return 'About %s...' % name

class Album(object):
    @cherrypy.expose
    def index(self, artist, title):
        return 'About %s by %s...' % (title, artist)

if __name__ == '__main__':
    cherrypy.quickstart(Band())

```

Notice how the controller defines `_cp_dispatch`, it takes a single argument, the URL path info broken into its segments. The method can inspect and manipulate the list of segments, removing any or adding new segments at any position. The new list of segments is then sent to the dispatcher which will use it to locate the appropriate resource.

In the above example, you should be able to go to the following URLs:

- <http://localhost:8080/nirvana/>
- <http://localhost:8080/nirvana/albums/nevermind/>

The `/nirvana/` segment is associated to the band and the `/nevermind/` segment relates to the album.

To achieve this, our `_cp_dispatch` method works on the idea that the default dispatcher matches URLs against page handler signatures and their position in the tree of handlers.

In this case, we take the dynamic segments in the URL (band and record names), we inject them into the request parameters and we remove them from the segment lists as if they had never been there in the first place.

In other words, `_cp_dispatch` makes it as if we were working on the following URLs:

- <http://localhost:8080/?artist=nirvana>
- <http://localhost:8080/albums/?artist=nirvana&title=nevermind>

5.2.2 The `popargs` decorator

`cherrypy.popargs()` is more straightforward as it gives a name to any segment that CherryPy wouldn't be able to interpret otherwise. This makes the matching of segments with page handler signatures easier and helps CherryPy understand the structure of your URL.

```
import cherrypy

@cherrypy.popargs('band_name')
class Band(object):
    def __init__(self):
        self.albums = Album()

    @cherrypy.expose
    def index(self, band_name):
        return 'About %s...' % band_name

@cherrypy.popargs('album_title')
class Album(object):
    @cherrypy.expose
    def index(self, band_name, album_title):
        return 'About %s by %s...' % (album_title, band_name)

if __name__ == '__main__':
    cherrypy.quickstart(Band())
```

This works similarly to `_cp_dispatch` but, as said above, is more explicit and localized. It says:

- take the first segment and store it into a parameter named `band_name`
- take again the first segment (since we removed the previous first) and store it into a parameter named `album_title`

Note that the decorator accepts more than a single binding. For instance:

```
@cherrypy.popargs('album_title')
class Album(object):
    def __init__(self):
        self.tracks = Track()

@cherrypy.popargs('track_num', 'track_title')
class Track(object):
    @cherrypy.expose
    def index(self, band_name, album_title, track_num, track_title):
        ...
```

This would handle the following URL:

- <http://localhost:8080/nirvana/albums/nevermind/tracks/06/polly>

Notice finally how the whole stack of segments is passed to each page handler so that you have the full context.

5.3 Error handling

CherryPy's `HTTPError` class supports raising immediate responses in the case of errors.

```
class Root:
    @cherrypy.expose
    def thing(self, path):
        if not authorized():
            raise cherrypy.HTTPError(401, 'Unauthorized')
        try:
            file = open(path)
        except FileNotFoundError:
            raise cherrypy.HTTPError(404)
```

`HTTPError.handle` is a context manager which supports translating exceptions raised in the app into an appropriate HTTP response, as in the second example.

```
class Root:
    @cherry.py.expose
    def thing(self, path):
        with cherry.py.HTTPError.handle(FileNotFoundError, 404):
            file = open(path)
```

5.4 Streaming the response body

CherryPy handles HTTP requests, packing and unpacking the low-level details, then passing control to your application's *page handler*, which produce the body of the response. CherryPy allows you to return body content in a variety of types: a string, a list of strings, a file. CherryPy also allows you to *yield* content, rather than *return* content. When you use “yield”, you also have the option of streaming the output.

In general, it is safer and easier to not stream output. Therefore, streaming output is off by default. Streaming output and also using sessions requires a good understanding of how `session locks` work.

5.4.1 The “normal” CherryPy response process

When you provide content from your page handler, CherryPy manages the conversation between the HTTP server and your code like this:

Notice that the HTTP server gathers all output first and then writes everything to the client at once: status, headers, and body. This works well for static or simple pages, since the entire response can be changed at any time, either in your application code, or by the CherryPy framework.

5.4.2 How “streaming output” works with CherryPy

When you set the config entry “response.stream” to True (and use “yield”), CherryPy manages the conversation between the HTTP server and your code like this:

When you stream, your application doesn't immediately pass raw body content back to CherryPy or to the HTTP server. Instead, it passes back a generator. At that point, CherryPy finalizes the status and headers, **before** the generator has been consumed, or has produced any output. This is necessary to allow the HTTP server to send the headers and pieces of the body as they become available.

Once CherryPy has set the status and headers, it sends them to the HTTP server, which then writes them out to the client. From that point on, the CherryPy framework mostly steps out of the way, and the HTTP server essentially requests content directly from your application code (your page handler method).

Therefore, when streaming, if an error occurs within your page handler, CherryPy will not catch it—the HTTP server will catch it. Because the headers (and potentially some of the body) have already been written to the client, the server *cannot* know a safe means of handling the error, and will therefore simply close the connection (the current, builtin servers actually write out a short error message in the body, but this may be changed, and is not guaranteed behavior for all HTTP servers you might use with CherryPy).

In addition, you cannot manually modify the status or headers within your page handler if that handler method is a streaming generator, because the method will not be iterated over until after the headers have been written to the client. **This includes raising exceptions like `HTTPError`, `NotFound`, `InternalRedirect` and `HTTPRedirect`.** To use a streaming generator while modifying headers, you would have to return a generator that is separate from (or embedded in) your page handler. For example:

```
class Root:
    @cherrypy.expose
    def thing(self):
        cherrypy.response.headers['Content-Type'] = 'text/plain'
        if not authorized():
            raise cherrypy.NotFound()
        def content():
            yield "Hello, "
            yield "world"
        return content()
    thing._cp_config = {'response.stream': True}
```

Streaming generators are sexy, but they play havoc with HTTP. CherryPy allows you to stream output for specific situations: pages which take many minutes to produce, or pages which need a portion of their content immediately output to the client. Because of the issues outlined above, **it is usually better to flatten (buffer) content rather than stream content**. Do otherwise only when the benefits of streaming outweigh the risks.

5.5 Response timeouts

CherryPy responses include 3 attributes related to time:

- `response.time`: the `time.time()` at which the response began
- `response.timeout`: the number of seconds to allow responses to run
- `response.timed_out`: a boolean indicating whether the response has timed out (default False).

The request processing logic inspects the value of `response.timed_out` at various stages; if it is ever True, then `TimeoutError` is raised. You are free to do the same within your own code.

Rather than calculate the difference by hand, you can call `response.check_timeout` to set `timed_out` for you.

Note: The default response timeout is 300 seconds.

5.5.1 Timeout Monitor

In addition, CherryPy includes a `cherrypy.engine.timeout_monitor` which monitors all active requests in a separate thread; periodically, it calls `check_timeout` on them all. It is subscribed by default. To turn it off:

```
[global]
engine.timeout_monitor.on: False
```

or:

```
cherrypy.engine.timeout_monitor.unsubscribe()
```

You can also change the interval (in seconds) at which the timeout monitor runs:

```
[global]
engine.timeout_monitor.frequency: 60 * 60
```

The default is once per minute. The above example changes that to once per hour.

5.6 Deal with signals

This *engine plugin* is instantiated automatically as `cherry.py.engine.signal_handler`. However, it is only *subscribed* automatically by `cherry.py.quickstart()`. So if you want signal handling and you're calling:

```
tree.mount()
engine.start()
engine.block()
```

on your own, be sure to add before you start the engine:

```
engine.signals.subscribe()
```

5.6.1 Windows Console Events

Microsoft Windows uses console events to communicate some signals, like Ctrl-C. Deploying CherryPy on Windows platforms requires [Python for Windows Extensions](#), which are installed automatically, being provided an extra dependency with environment marker. With that installed, CherryPy will handle Ctrl-C and other console events (CTRL_C_EVENT, CTRL_LOGOFF_EVENT, CTRL_BREAK_EVENT, CTRL_SHUTDOWN_EVENT, and CTRL_CLOSE_EVENT) automatically, shutting down the bus in preparation for process exit.

5.7 Securing your server

Note: This section is not meant as a complete guide to securing a web application or ecosystem. Please review the various guides provided at [OWASP](#).

There are several settings that can be enabled to make CherryPy pages more secure. These include:

Transmitting data:

1. Use Secure Cookies

Rendering pages:

1. Set HttpOnly cookies
2. Set XFrame options
3. Enable XSS Protection
4. Set the Content Security Policy

An easy way to accomplish this is to set headers with a tool and wrap your entire CherryPy application with it:

```
import cherryypy

# set the priority according to your needs if you are hooking something
# else on the 'before_finalize' hook point.
@cherryypy.tools.register('before_finalize', priority=60)
def secureheaders():
    headers = cherryypy.response.headers
    headers['X-Frame-Options'] = 'DENY'
    headers['X-XSS-Protection'] = '1; mode=block'
    headers['Content-Security-Policy'] = "default-src='self'"
```

Note: Read more about those headers.

Then, in the *configuration file* (or any other place that you want to enable the tool):

```
[/]  
tools.secureheaders.on = True
```

If you use *sessions* you can also enable these settings:

```
[/]  
tools.sessions.on = True  
# increase security on sessions  
tools.sessions.secure = True  
tools.sessions.httponly = True
```

If you use SSL you can also enable Strict Transport Security:

```
# add this to secureheaders():  
# only add Strict-Transport headers if we're actually using SSL; see the ietf spec  
# "An HSTS Host MUST NOT include the STS header field in HTTP responses  
# conveyed over non-secure transport"  
# http://tools.ietf.org/html/draft-ietf-websec-strict-transport-sec-14#section-7.2  
if (cherrypy.server.ssl_certificate != None and cherrypy.server.ssl_private_key != None):  
headers['Strict-Transport-Security'] = 'max-age=31536000' # one year
```

Next, you should probably use *SSL*.

5.8 Multiple HTTP servers support

CherryPy starts its own HTTP server whenever you start the engine. In some cases, you may wish to host your application on more than a single port. This is easily achieved:

```
from cherrypy._cpserver import Server  
server = Server()  
server.socket_port = 8090  
server.subscribe()
```

You can create as many `server` instances as you need, once *subscribed*, they will follow the CherryPy engine's life-cycle.

5.9 WSGI support

CherryPy supports the WSGI interface defined in [PEP 333](#) as well as its updates in [PEP 3333](#). It means the following:

- You can host a foreign WSGI application with the CherryPy server
- A CherryPy application can be hosted by another WSGI server

5.9.1 Make your CherryPy application a WSGI application

A WSGI application can be obtained from your application as follows:

```
import cherrypy
wsgiapp = cherrypy.Application(StringGenerator(), '/', config=myconf)
```

Simply use the `wsgiapp` instance in any WSGI-aware server.

5.9.2 Host a foreign WSGI application in CherryPy

Assuming you have a WSGI-aware application, you can host it in your CherryPy server using the `cherrypy.tree.graft` facility.

```
def raw_wsgi_app(environ, start_response):
    status = '200 OK'
    response_headers = [('Content-type', 'text/plain')]
    start_response(status, response_headers)
    return ['Hello world!']

cherrypy.tree.graft(raw_wsgi_app, '/')
```

Important: You cannot use tools with a foreign WSGI application. However, you can still benefit from the *CherryPy bus*.

5.9.3 No need for the WSGI interface?

The default CherryPy HTTP server supports the WSGI interfaces defined in [PEP 333](#) and [PEP 3333](#). However, if your application is a pure CherryPy application, you can switch to a HTTP server that by-passes the WSGI layer altogether. It will provide a slight performance increase.

```
import cherrypy

class Root(object):
    @cherrypy.expose
    def index(self):
        return "Hello World!"

if __name__ == '__main__':
    from cherrypy._cpnative_server import CPHTTPServer
    cherrypy.server.httpserver = CPHTTPServer(cherrypy.server)

    cherrypy.quickstart(Root(), '/')
```

Important: Using the native server, you will not be able to graft a WSGI application as shown in the previous section. Doing so will result in a server error at runtime.

5.10 WebSocket support

[WebSocket](#) is a recent application protocol that came to life from the HTML5 working-group in response to the needs for bi-directional communication. Various hacks had been proposed such as Comet, polling, etc.

WebSocket is a socket that starts its life from a HTTP upgrade request. Once the upgrade is performed, the underlying socket is kept opened but not used in a HTTP context any longer. Instead, both connected endpoints may use the socket to push data to the other end.

CherryPy itself does not support WebSocket, but the feature is provided by an external library called `ws4py`.

5.11 Database support

CherryPy does not bundle any database access but its architecture makes it easy to integrate common database interfaces such as the DB-API specified in [PEP 249](#). Alternatively, you can also use an ORM such as `SQLAlchemy` or `SQLObject`.

You will find [here](#) a recipe on how integrating SQLAlchemy using a mix of *plugins* and *tools*.

5.12 HTML Templating support

CherryPy does not provide any HTML template but its architecture makes it easy to integrate one. Popular ones are `Mako` or `Jinja2`.

You will find [here](#) a recipe on how to integrate them using a mix *plugins* and *tools*.

5.13 Testing your application

Web applications, like any other kind of code, must be tested. CherryPy provides a `helper` class to ease writing functional tests.

Here is a simple example for a basic echo application:

```
import cherrypy
from cherrypy.test import helper

class SimpleCPTest(helper.CPWebCase):
    def setup_server():
        class Root(object):
            @cherrypy.expose
            def echo(self, message):
                return message

        cherrypy.tree.mount(Root())
        setup_server = staticmethod(setup_server)

    def test_message_should_be_returned_as_is(self):
        self.getPage("/echo?message=Hello%20world")
        self.assertStatus('200 OK')
        self.assertHeader('Content-Type', 'text/html; charset=utf-8')
        self.assertBody('Hello world')

    def test_non_utf8_message_will_fail(self):
        """
        CherryPy defaults to decode the query-string
        using UTF-8, trying to send a query-string with
        a different encoding will raise a 404 since
        it considers it's a different URL.
        """
```



```
"""
self.getPage("/echo?message=A+bient%F4t",
             headers=[
                 ('Accept-Charset', 'ISO-8859-1,utf-8'),
                 ('Content-Type', 'text/html;charset=ISO-8859-1')
             ])
self.assertStatus('404 Not Found')
```

As you can see the, test inherits from that helper class. You should setup your application and mount it as per-usual. Then, define your various tests and call the helper `getPage()` method to perform a request. Simply use the various specialized `assert*` methods to validate your workflow and data.

You can then run the test using `py.test` as follows:

```
$ py.test -s test_echo_app.py
```

The `-s` is necessary because the CherryPy class also wraps `stdin` and `stdout`.

Note: Although they are written using the typical pattern the `unittest` module supports, they are not bare unit tests. Indeed, a whole CherryPy stack is started for you and runs your application. If you want to really unit test your CherryPy application, meaning without having to start a server, you may want to have a look at this [recipe](#).

Configure

Configuration in CherryPy is implemented via dictionaries. Keys are strings which name the mapped value; values may be of any type.

In CherryPy 3, you use configuration (files or dicts) to set attributes directly on the engine, server, request, response, and log objects. So the best way to know the full range of what's available in the config file is to simply import those objects and see what `help(obj)` tells you.

Note: If you are new to CherryPy, please refer first to the simpler *basic config* section first.

Contents

- *Configure*
 - *Architecture*
 - * *Global config*
 - * *Application config*
 - * *Request config*
 - *Declaration*
 - * *Configuration files*
 - * *_cp_config: attaching config to handlers*
 - *Namespaces*
 - * *Builtin namespaces*
 - * *Custom config namespaces*
 - * *Environments*

6.1 Architecture

The first thing you need to know about CherryPy 3's configuration is that it separates *global* config from *application* config. If you're deploying multiple *applications* at the same *site* (and more and more people are, as Python web apps are tending to decentralize), you need to be careful to separate the configurations, as well. There's only ever one "global config", but there is a separate "app config" for each app you deploy.

CherryPy *Requests* are part of an *Application*, which runs in a *global* context, and configuration data may apply to any of those three scopes. Let's look at each of those scopes in turn.

6.1.1 Global config

Global config entries apply everywhere, and are stored in `cherrypy.config`. This flat dict only holds global config data; that is, “site-wide” config entries which affect all mounted applications.

Global config is stored in the `cherrypy.config` dict, and you therefore update it by calling `cherrypy.config.update(conf)`. The `conf` argument can be either a filename, an open file, or a dict of config entries. Here’s an example of passing a dict argument:

```
cherrypy.config.update({'server.socket_host': '64.72.221.48',
                       'server.socket_port': 80,
                       })
```

The `server.socket_host` option in this example determines on which network interface CherryPy will listen. The `server.socket_port` option declares the TCP port on which to listen.

6.1.2 Application config

Application entries apply to a single mounted application, and are stored on each Application object itself as `app.config`. This is a two-level dict where each top-level key is a path, or “relative URL” (for example, `"/` or `/my/page`), and each value is a dict of config entries. The URL’s are relative to the script name (mount point) of the Application. Usually, all this data is provided in the call to `tree.mount(root(), script_name='/path/to', config=conf)`, although you may also use `app.merge(conf)`. The `conf` argument can be either a filename, an open file, or a dict of config entries.

Configuration file example:

```
[/]
tools.trailing_slash.on = False
request.dispatch: cherrypy.dispatch.MethodDispatcher()
```

or, in python code:

```
config = {'/':
          {
            'request.dispatch': cherrypy.dispatch.MethodDispatcher(),
            'tools.trailing_slash.on': False,
          }
        }
cherrypy.tree.mount(Root(), config=config)
```

CherryPy only uses sections that start with `"/` (except `[global]`, see below). That means you can place your own configuration entries in a CherryPy config file by giving them a section name which does not start with `"/`. For example, you might include database entries like this:

```
[global]
server.socket_host: "0.0.0.0"

[Databases]
driver: "postgres"
host: "localhost"
port: 5432

[/path]
response.timeout: 6000
```

Then, in your application code you can read these values during request time via `cherrypy.request.app.config['Databases']`. For code that is outside the request process, you’ll

have to pass a reference to your Application around.

6.1.3 Request config

Each Request object possesses a single `request.config` dict. Early in the request process, this dict is populated by merging Global config, Application config, and any config acquired while looking up the page handler (see next). This dict contains only those config entries which apply to the given request.

Note: when you do an `InternalRedirect`, this config attribute is recalculated for the new path.

6.2 Declaration

Configuration data may be supplied as a Python dictionary, as a filename, or as an open file object.

6.2.1 Configuration files

When you supply a filename or file, CherryPy uses Python's builtin `ConfigParser`; you declare Application config by writing each path as a section header, and each entry as a "key: value" (or "key = value") pair:

```
[/path/to/my/page]
response.stream: True
tools.trailing_slash.extra = False
```

Combined Configuration Files

If you are only deploying a single application, you can make a single config file that contains both global and app entries. Just stick the global entries into a config section named `[global]`, and pass the same file to both `config.update` and `tree.mount` `<cherry.py._cptree.Tree.mount()`. If you're calling `cherry.py.quickstart(app root, script name, config)`, it will pass the config to both places for you. But as soon as you decide to add another application to the same site, you need to separate the two config files/dicts.

Separate Configuration Files

If you're deploying more than one application in the same process, you need (1) file for global config, plus (1) file for *each* Application. The global config is applied by calling `cherry.py.config.update`, and application config is usually passed in a call to `cherry.py.tree.mount`.

In general, you should set global config first, and then mount each application with its own config. Among other benefits, this allows you to set up global logging so that, if something goes wrong while trying to mount an application, you'll see the tracebacks. In other words, use this order:

```
# global config
cherry.py.config.update({'environment': 'production',
                        'log.error_file': 'site.log',
                        # ...
                        })

# Mount each app and pass it its own config
```

```
cherry.py.tree.mount(root1, "", appconf1)
cherry.py.tree.mount(root2, "/forum", appconf2)
cherry.py.tree.mount(root3, "/blog", appconf3)

if hasattr(cherry.py.engine, 'block'):
    # 3.1 syntax
    cherry.py.engine.start()
    cherry.py.engine.block()
else:
    # 3.0 syntax
    cherry.py.server.quickstart()
    cherry.py.engine.start()
```

Values in config files use Python syntax

Config entries are always a key/value pair, like `server.socket_port = 8080`. The key is always a name, and the value is always a Python object. That is, if the value you are setting is an `int` (or other number), it needs to look like a Python `int`; for example, `8080`. If the value is a string, it needs to be quoted, just like a Python string. Arbitrary objects can also be created, just like in Python code (assuming they can be found/imported). Here's an extended example, showing you some of the different types:

```
[global]
log.error_file: "/home/fumanchu/myapp.log"
environment = 'production'
server.max_request_body_size: 1200

[/myapp]
tools.trailing_slash.on = False
request.dispatch: cherry.py.dispatch.MethodDispatcher()
```

6.2.2 `_cp_config`: attaching config to handlers

Config files have a severe limitation: values are always keyed by URL. For example:

```
[/path/to/page]
methods_with_bodies = ("POST", "PUT", "PROPPATCH")
```

It's obvious that the extra method is the norm for that path; in fact, the code could be considered broken without it. In CherryPy, you can attach that bit of config directly on the page handler:

```
@cherry.py.expose
def page(self):
    return "Hello, world!"
page._cp_config = {"request.methods_with_bodies": ("POST", "PUT", "PROPPATCH")}
```

`_cp_config` is a reserved attribute which the dispatcher looks for at each node in the object tree. The `_cp_config` attribute must be a CherryPy config dictionary. If the dispatcher finds a `_cp_config` attribute, it merges that dictionary into the rest of the config. The entire merged config dictionary is placed in `cherry.py.request.config`.

This can be done at any point in the tree of objects; for example, we could have attached that config to a class which contains the page method:

```
class SetOPages:
    _cp_config = {"request.methods_with_bodies": ("POST", "PUT", "PROPPATCH")}
```

```
@cherry.py.expose
def page(self):
    return "Hullo, World!"
```

Note: This behavior is only guaranteed for the default dispatcher. Other dispatchers may have different restrictions on where you can attach `_cp_config` attributes.

This technique allows you to:

- Put config near where it’s used for improved readability and maintainability.
- Attach config to objects instead of URL’s. This allows multiple URL’s to point to the same object, yet you only need to define the config once.
- Provide defaults which are still overridable in a config file.

6.3 Namespaces

Because config entries usually just set attributes on objects, they’re almost all of the form: `object.attribute`. A few are of the form: `object.subobject.attribute`. They look like normal Python attribute chains, because they work like them. We call the first name in the chain the “*config namespace*”. When you provide a config entry, it is bound as early as possible to the actual object referenced by the namespace; for example, the entry `response.stream` actually sets the `stream` attribute of `cherry.py.response`! In this way, you can easily determine the default value by firing up a python interpreter and typing:

```
>>> import cherry.py
>>> cherry.py.response.stream
False
```

Each config namespace has its own handler; for example, the “request” namespace has a handler which takes your config entry and sets that value on the appropriate “request” attribute. There are a few namespaces, however, which don’t work like normal attributes behind the scenes; however, they still use dotted keys and are considered to “have a namespace”.

6.3.1 Builtin namespaces

Entries from each namespace may be allowed in the global, application root (“/”) or per-path config, or a combination:

Scope	Global	Application Root	App Path
engine	X		
hooks	X	X	X
log	X	X	
request	X	X	X
response	X	X	X
server	X		
tools	X	X	X

engine

Entries in this namespace controls the ‘application engine’. These can only be declared in the global config. Any attribute of `cherry.py.engine` may be set in config; however, there are a few extra entries available in config:

- **Plugin attributes.** Many of the Engine Plugins are themselves attributes of `cherrypy.engine`. You can set any attribute of an attached plugin by simply naming it. For example, there is an instance of the `Autoreloader` class at `engine.autoreload`; you can set its “frequency” attribute via the config entry `engine.autoreload.frequency = 60`. In addition, you can turn such plugins on and off by setting `engine.autoreload.on = True` or `False`.
- `engine.SIGHUP/SIGTERM`: These entries can be used to set the list of listeners for the given channel. Mostly, this is used to turn off the signal handling one gets automatically via `cherrypy.quickstart()`.

hooks

Declares additional request-processing functions. Use this to append your own Hook functions to the request. For example, to add `my_hook_func` to the `before_handler` hookpoint:

```
[/]  
hooks.before_handler = myapp.my_hook_func
```

log

Configures logging. These can only be declared in the global config (for global logging) or `[/]` config (for each application). See `LogManager` for the list of configurable attributes. Typically, the “`access_file`”, “`error_file`”, and “`screen`” attributes are the most commonly configured.

request

Sets attributes on each Request. See the `Request` class for a complete list.

response

Sets attributes on each Response. See the `Response` class for a complete list.

server

Controls the default HTTP server via `cherrypy.server` (see that class for a complete list of configurable attributes). These can only be declared in the global config.

tools

Enables and configures additional request-processing packages. See the `/tutorial/tools` overview for more information.

wsgi

Adds WSGI middleware to an Application’s “pipeline”. These can only be declared in the app’s root config (“/”).

- `wsgi.pipeline`: Appends to the WSGI pipeline. The value must be a list of (name, app factory) pairs. Each app factory must be a WSGI callable class (or callable that returns a WSGI callable); it must take an initial ‘nextapp’ argument, plus any optional keyword arguments. The optional arguments may be configured via `wsgi.<name>.<arg>`.
- `wsgi.response_class`: Overrides the default `Response` class.

checker

Controls the “checker”, which looks for common errors in app state (including config) when the engine starts. You can turn off individual checks by setting them to `False` in config. See `cherry.py._cpchecker.Checker` for a complete list. Global config only.

6.3.2 Custom config namespaces

You can define your own namespaces if you like, and they can do far more than simply set attributes. The `test/test_config` module, for example, shows an example of a custom namespace that coerces incoming params and outgoing body content. The `cherry.py._cpwsgi` module includes an additional, builtin namespace for invoking WSGI middleware.

In essence, a config namespace handler is just a function, that gets passed any config entries in its namespace. You add it to a namespaces registry (a dict), where keys are namespace names and values are handler functions. When a config entry for your namespace is encountered, the corresponding handler function will be called, passing the config key and value; that is, `namespaces[namespace](k, v)`. For example, if you write:

```
def db_namespace(k, v):
    if k == 'connstring':
        orm.connect(v)
cherry.py.config.namespaces['db'] = db_namespace
```

then `cherry.py.config.update({"db.connstring": "Oracle:host=1.10.100.200;sid=TEST"})` will call `db_namespace('connstring', 'Oracle:host=1.10.100.200;sid=TEST')`.

The point at which your namespace handler is called depends on where you add it:

Scope	Namespace dict	Handler is called in
Global	<code>cherry.py.config.namespaces</code>	<code>cherry.py.config.update</code>
Application	<code>app.namespaces</code>	<code>Application.merge</code> (which is called by <code>cherry.py.tree.mount</code>)
Request	<code>app.request_class.namespaces</code>	<code>Request.configure</code> (called for each request, after the handler is looked up)

The name can be any string, and the handler must be either a callable or a (Python 2.5 style) context manager.

If you need additional code to run when all your namespace keys are collected, you can supply a callable context manager in place of a normal function for the handler. Context managers are defined in [PEP 343](#).

6.3.3 Environments

The only key that does not exist in a namespace is the “*environment*” entry. It only applies to the global config, and only when you use `cherry.py.config.update`. This special entry *imports* other config entries from the following template stored in `cherry.py._cpconfig.environments[environment]`.

If you find the set of existing environments (production, staging, etc) too limiting or just plain wrong, feel free to extend them or add new environments:

```
cherry.py._cpconfig.environments['staging']['log.screen'] = False

cherry.py._cpconfig.environments['Greek'] = {
    'tools.encode.encoding': 'ISO-8859-7',
    'tools.decode.encoding': 'ISO-8859-7',
}
```

Extend

CherryPy is truly an open framework, you can extend and plug new functions at will either server-side or on a per-requests basis. Either way, CherryPy is made to help you build your application and support your architecture via simple patterns.

Contents

- *Extend*
 - *Server-wide functions*
 - * *Publish/Subscribe pattern*
 - *Typical pattern*
 - *Implementation details*
 - *Engine as a pubsub bus*
 - *Built-in channels*
 - *Bus API*
 - * *Plugins*
 - *Create a plugin*
 - *Enable a plugin*
 - *Disable a plugin*
 - *Per-request functions*
 - * *Hook point*
 - * *Tools*
 - *Stateful tools*
 - *Tools ordering*
 - *Toolboxes*
 - * *Request parameters manipulation*
 - *Tailored dispatchers*
 - * *Tool or dispatcher?*
 - *Request body processors*

7.1 Server-wide functions

CherryPy can be considered both as a HTTP library as much as a web application framework. In that latter case, its architecture provides mechanisms to support operations accross the whole server instance. This offers a powerful canvas to perform persistent operations as server-wide functions live outside the request processing itself. They are available to the whole process as long as the bus lives.

Typical use cases:

- Keeping a pool of connection to an external server so that you need not to re-open them on each request (database connections for instance).
- Background processing (say you need work to be done without blocking the whole request itself).

7.1.1 Publish/Subscribe pattern

CherryPy's backbone consists of a bus system implementing a simple publish/subscribe messaging pattern. Simply put, in CherryPy everything is controlled via that bus. One can easily picture the bus as a sushi restaurant's belt as in the picture below.



You can subscribe and publish to channels on a bus. A channel is bit like a unique identifier within the bus. When a message is published to a channel, the bus will dispatch the message to all subscribers for that channel.

One interesting aspect of a pubsub pattern is that it promotes decoupling between a caller and the callee. A published message will eventually generate a response but the publisher does not know where that response came from.

Thanks to that decoupling, a CherryPy application can easily access functionalities without having to hold a reference to the entity providing that functionality. Instead, the application simply publishes onto the bus and will receive the appropriate response, which is all that matter.

Typical pattern

Let's take the following dummy application:

```
import cherrypy

class ECommerce(object):
    def __init__(self, db):
        self.mydb = db

    @cherrypy.expose
    def save_kart(self, cart_data):
        cart = Cart(cart_data)
        self.mydb.save(cart)

if __name__ == '__main__':
    cherrypy.quickstart(ECommerce(), '/')
```

The application has a reference to the database but this creates a fairly strong coupling between the database provider and the application.

Another approach to work around the coupling is by using a pubsub workflow:

```
import cherrypy

class ECommerce(object):
    @cherrypy.expose
    def save_kart(self, cart_data):
        cart = Cart(cart_data)
        cherrypy.engine.publish('db-save', cart)

if __name__ == '__main__':
    cherrypy.quickstart(ECommerce(), '/')
```

In this example, we publish a *cart* instance to *db-save* channel. One or many subscribers can then react to that message and the application doesn't have to know about them.

Note: This approach is not mandatory and it's up to you to decide how to design your entities interaction.

Implementation details

CherryPy's bus implementation is simplistic as it registers functions to channels. Whenever a message is published to a channel, each registered function is applied with that message passed as a parameter.

The whole behaviour happens synchronously and, in that sense, if a subscriber takes too long to process a message, the remaining subscribers will be delayed.


CherryPy's bus is not an advanced pubsub messaging broker system such as provided by [zeromq](#) or [RabbitMQ](#). Use it with the understanding that it may have a cost.

Engine as a pubsub bus

As said earlier, CherryPy is built around a pubsub bus. All entities that the framework manages at runtime are working on top of a single bus instance, which is named the *engine*.

The bus implementation therefore provides a set of common channels which describe the application's lifecycle:

```
o
|
v
```



The states' transitions trigger channels to be published to so that subscribers can react to them.

One good example is the HTTP server which will transition from a “*STOPPED*” state to a “*STARTED*” state whenever a message is published to the *start* channel.

Built-in channels

In order to support its life-cycle, CherryPy defines a set of common channels that will be published to at various states:

- “**start**”: When the bus is in the “*STARTING*” state
- “**main**”: Periodically from the CherryPy's mainloop
- “**stop**”: When the bus is in the “*STOPPING*” state
- “**graceful**”: When the bus requests a reload of subscribers
- “**exit**”: When the bus is in the “*EXITING*” state

This channel will be published to by the *engine* automatically. Register therefore any subscribers that would need to react to the transition changes of the *engine*.

In addition, a few other channels are also published to during the request processing.

- “**before_request**”: right before the request is processed by CherryPy
- “**after_request**”: right after it has been processed

Also, from the `cherry.py.process.plugins.ThreadManager` plugin:

- “**acquire_thread**”
- “**start_thread**”
- “**stop_thread**”
- “**release_thread**”

Bus API

In order to work with the bus, the implementation provides the following simple API:

- `cherry.py.engine.publish(channel, *args):`
- The *channel* parameter is a string identifying the channel to which the message should be sent to
- **args* is the message and may contain any valid Python values or objects.
- `cherry.py.engine.subscribe(channel, callable):`
- The *channel* parameter is a string identifying the channel the *callable* will be registered to.
- *callable* is a Python function or method which signature must match what will be published.
- `cherry.py.engine.unsubscribe(channel, callable):`
- The *channel* parameter is a string identifying the channel the *callable* was registered to.

- *callable* is the Python function or method which was registered.

7.1.2 Plugins

Plugins, simply put, are entities that play with the bus, either by publishing or subscribing to channels, usually both at the same time.

Important: Plugins are extremely useful whenever you have functionalities:

- Available across the whole application server
 - Associated to the application's life-cycle
 - You want to avoid being strongly coupled to the application
-

Create a plugin

A typical plugin looks like this:

```
import cherrypy
from cherrypy.process import wspbus, plugins

class DatabasePlugin(plugins.SimplePlugin):
    def __init__(self, bus, db_class):
        plugins.SimplePlugin.__init__(self, bus)
        self.db = db_class()

    def start(self):
        self.bus.log('Starting up DB access')
        self.bus.subscribe("db-save", self.save_it)

    def stop(self):
        self.bus.log('Stopping down DB access')
        self.bus.unsubscribe("db-save", self.save_it)

    def save_it(self, entity):
        self.db.save(entity)
```

The `cherrypy.process.plugins.SimplePlugin` is a helper class provided by CherryPy that will automatically subscribe your *start* and *stop* methods to the related channels.

When the *start* and *stop* channels are published on, those methods are called accordingly.

Notice then how our plugin subscribes to the *db-save* channel so that the bus can dispatch messages to the plugin.

Enable a plugin

To enable the plugin, it has to be registered to the the bus as follows:

```
DatabasePlugin(cherrypy.engine, SQLiteDB).subscribe()
```

The *SQLiteDB* here is a fake class that is used as our database provider.

Disable a plugin

You can also unregister a plugin as follows:

```
someplugin.unsubscribe()
```

This is often used when you want to prevent the default HTTP server from being started by CherryPy, for instance if you run on top of a different HTTP server (WSGI capable):

```
cherry.py.server.unsubscribe()
```

Let's see an example using this default application:

```
import cherry.py

class Root(object):
    @cherry.py.expose
    def index(self):
        return "hello world"

if __name__ == '__main__':
    cherry.py.quickstart(Root())
```

For instance, this is what you would see when running this application:

```
[27/Apr/2014:13:04:07] ENGINE Listening for SIGHUP.
[27/Apr/2014:13:04:07] ENGINE Listening for SIGTERM.
[27/Apr/2014:13:04:07] ENGINE Listening for SIGUSR1.
[27/Apr/2014:13:04:07] ENGINE Bus STARTING
[27/Apr/2014:13:04:07] ENGINE Started monitor thread 'Autoreloader'.
[27/Apr/2014:13:04:07] ENGINE Started monitor thread '_TimeoutMonitor'.
[27/Apr/2014:13:04:08] ENGINE Serving on http://127.0.0.1:8080
[27/Apr/2014:13:04:08] ENGINE Bus STARTED
```

Now let's unsubscribe the HTTP server:

```
import cherry.py

class Root(object):
    @cherry.py.expose
    def index(self):
        return "hello world"

if __name__ == '__main__':
    cherry.py.server.unsubscribe()
    cherry.py.quickstart(Root())
```

This is what we get:

```
[27/Apr/2014:13:08:06] ENGINE Listening for SIGHUP.
[27/Apr/2014:13:08:06] ENGINE Listening for SIGTERM.
[27/Apr/2014:13:08:06] ENGINE Listening for SIGUSR1.
[27/Apr/2014:13:08:06] ENGINE Bus STARTING
[27/Apr/2014:13:08:06] ENGINE Started monitor thread 'Autoreloader'.
[27/Apr/2014:13:08:06] ENGINE Started monitor thread '_TimeoutMonitor'.
[27/Apr/2014:13:08:06] ENGINE Bus STARTED
```

As you can see, the server is not started. The missing:


```
[27/Apr/2014:13:04:08] ENGINE Serving on http://127.0.0.1:8080
```

7.2 Per-request functions

One of the most common task in a web application development is to tailor the request's processing to the runtime context.

Within CherryPy, this is performed via what are called *tools*. If you are familiar with Django or WSGI middlewares, CherryPy tools are similar in spirit. They add functions that are applied during the request/response processing.

7.2.1 Hook point

A hook point is a point during the request/response processing.

Here is a quick rundown of the “hook points” that you can hang your tools on:

- **“on_start_resource”** - The earliest hook; the Request-Line and request headers have been processed and a dispatcher has set `request.handler` and `request.config`.
- **“before_request_body”** - Tools that are hooked up here run right before the request body would be processed.
- **“before_handler”** - Right before the `request.handler` (the *exposed* callable that was found by the dispatcher) is called.
- **“before_finalize”** - This hook is called right after the page handler has been processed and before CherryPy formats the final response object. It helps you for example to check for what could have been returned by your page handler and change some headers if needed.
- **“on_end_resource”** - Processing is complete - the response is ready to be returned. This doesn't always mean that the `request.handler` (the exposed page handler) has executed! It may be a generator. If your tool absolutely needs to run after the page handler has produced the response body, you need to either use `on_end_request` instead, or wrap the `response.body` in a generator which applies your tool as the response body is being generated.
- **“before_error_response”** - Called right before an error response (status code, body) is set.
- **“after_error_response”** - Called right after the error response (status code, body) is set and just before the error response is finalized.
- **“on_end_request”** - The request/response conversation is over, all data has been written to the client, nothing more to see here, move along.

7.2.2 Tools

A tool is a simple callable object (function, method, object implementing a `__call__` method) that is attached to a *hook point*.

Below is a simple tool that is attached to the `before_finalize` hook point, hence after the page handler was called:

```
@cherrypy.tools.register('before_finalize')
def logit():
    print(cherrypy.request.remote.ip)
```

Tools can also be created and assigned manually. The decorator registration is equivalent to:

```
cherrypy.tools.logit = cherrypy.Tool('before_finalize', logit)
```

Using that tool is as simple as follows:

```
class Root(object):
    @cherry.py.expose
    @cherry.py.tools.logit()
    def index(self):
        return "hello world"
```

Obviously the tool may be declared the *other usual ways*.

Note: The name of the tool, technically the attribute set to *cherry.py.tools*, does not have to match the name of the callable. However, it is that name that will be used in the configuration to refer to that tool.

Stateful tools

The tools mechanism is really flexible and enables rich per-request functionalities.

Straight tools as shown in the previous section are usually good enough. However, if your workflow requires some sort of state during the request processing, you will probably want a class-based approach:

```
import time

import cherry.py

class TimingTool(cherry.py.Tool):
    def __init__(self):
        cherry.py.Tool.__init__(self, 'before_handler',
                                self.start_timer,
                                priority=95)

    def _setup(self):
        cherry.py.Tool._setup(self)
        cherry.py.request.hooks.attach('before_finalize',
                                       self.end_timer,
                                       priority=5)

    def start_timer(self):
        cherry.py.request._time = time.time()

    def end_timer(self):
        duration = time.time() - cherry.py.request._time
        cherry.py.log("Page handler took %.4f" % duration)

cherry.py.tools.timeit = TimingTool()
```

This tool computes the time taken by the page handler for a given request. It stores the time at which the handler is about to get called and logs the time difference right after the handler returned its result.

The import bits is that the *cherry.py.Tool* constructor allows you to register to a hook point but, to attach the same tool to a different hook point, you must use the *cherry.py.request.hooks.attach* method. The *cherry.py.Tool._setup* method is automatically called by CherryPy when the tool is applied to the request.

Next, let's see how to use our tool:

```
class Root(object):
    @cherry.py.expose
    @cherry.py.tools.timeit()
```

```
def index(self):
    return "hello world"
```

Tools ordering

Since you can register many tools at the same hookpoint, you may wonder in which order they will be applied.

CherryPy offers a deterministic, yet so simple, mechanism to do so. Simply set the **priority** attribute to a value from 1 to 100, lower values providing greater priority.

If you set the same priority for several tools, they will be called in the order you declare them in your configuration.

Toolboxes

All of the builtin CherryPy tools are collected into a Toolbox called `cherrypy.tools`. It responds to config entries in the `"tools"` namespace. You can add your own Tools to this Toolbox as described above.

You can also make your own Toolboxes if you need more modularity. For example, you might create multiple Tools for working with JSON, or you might publish a set of Tools covering authentication and authorization from which everyone could benefit (hint, hint). Creating a new Toolbox is as simple as:

```
import cherrypy

# Create a new Toolbox.
newauthtools = cherrypy._cptools.Toolbox("newauth")

# Add a Tool to our new Toolbox.
@newauthtools.register('before_request_body')
def check_access(default=False):
    if not getattr(cherrypy.request, "userid", default):
        raise cherrypy.HTTPError(401)
```

Then, in your application, use it just like you would use `cherrypy.tools`, with the additional step of registering your toolbox with your app. Note that doing so automatically registers the `"newauth"` config namespace; you can see the config entries in action below:

```
import cherrypy

class Root(object):
    @cherrypy.expose
    def default(self):
        return "Hello"

conf = {
    '/demo': {
        'newauth.check_access.on': True,
        'newauth.check_access.default': True,
    }
}

app = cherrypy.tree.mount(Root(), config=conf)
```

7.2.3 Request parameters manipulation

HTTP uses strings to carry data between two endpoints. However your application may make better use of richer object types. As it wouldn't be really readable, nor a good idea regarding maintenance, to let each page handler

deserialize data, it's a common pattern to delegate this functions to tools.

For instance, let's assume you have a user id in the query-string and some user data stored into a database. You could retrieve the data, create an object and pass it on to the page handler instead of the user id.

```
import cherrypy

class UserManager(cherrypy.Tool):
    def __init__(self):
        cherrypy.Tool.__init__(self, 'before_handler',
                               self.load, priority=10)

    def load(self):
        req = cherrypy.request

        # let's assume we have a db session
        # attached to the request somehow
        db = req.db

        # retrieve the user id and remove it
        # from the request parameters
        user_id = req.params.pop('user_id')
        req.params['user'] = db.get(int(user_id))

cherrypy.tools.user = UserManager()

class Root(object):
    @cherrypy.expose
    @cherrypy.tools.user()
    def index(self, user):
        return "hello %s" % user.name
```

In other words, CherryPy give you the power to:

- inject data, that wasn't part of the initial request, into the page handler
- remove data as well
- convert data into a different, more useful, object to remove that burden from the page handler itself

7.3 Tailored dispatchers

Dispatching is the art of locating the appropriate page handler for a given request. Usually, dispatching is based on the request's URL, the query-string and, sometimes, the request's method (GET, POST, etc.).

Based on this, CherryPy comes with various dispatchers already.

In some cases however, you will need a little more. Here is an example of dispatcher that will always ensure the incoming URL leads to a lower-case page handler.

```
import random
import string

import cherrypy
from cherrypy._cpdispatch import Dispatcher

class StringGenerator(object):
    @cherrypy.expose
```

```
def generate(self, length=8):
    return ''.join(random.sample(string.hexdigits, int(length)))

class ForceLowerDispatcher(Dispatcher):
    def __call__(self, path_info):
        return Dispatcher.__call__(self, path_info.lower())

if __name__ == '__main__':
    conf = {
        '/': {
            'request.dispatch': ForceLowerDispatcher(),
        }
    }
    cherrypy.quickstart(StringGenerator(), '/', conf)
```

Once you run this snippet, go to:

- <http://localhost:8080/generate?length=8>
- <http://localhost:8080/GENerAte?length=8>

In both cases, you will be led to the *generate* page handler. Without our home-made dispatcher, the second one would fail and return a 404 error ([RFC 2616#sec10.4.5](#)).

7.3.1 Tool or dispatcher?

In the previous example, why not simply use a tool? Well, the sooner a tool can be called is always after the page handler has been found. In our example, it would be already too late as the default dispatcher would have not even found a match for */GENerAte*.

A dispatcher exists mostly to determine the best page handler to serve the requested resource.

On the other hand, tools are there to adapt the request's processing to the runtime context of the application and the request's content.

Usually, you will have to write a dispatcher only if you have a very specific use case to locate the most adequate page handler. Otherwise, the default ones will likely suffice.

7.4 Request body processors

Since its 3.2 release, CherryPy provides a really elegant and powerful mechanism to deal with a request's body based on its mimetype. Refer to the `cherrypy._cpreqbody` module to understand how to implement your own processors.

Deploy

CherryPy stands on its own, but as an application server, it is often located in shared or complex environments. For this reason, it is not uncommon to run CherryPy behind a reverse proxy or use other servers to host the application.

Note: CherryPy's server has proven reliable and fast enough for years now. If the volume of traffic you receive is average, it will do well enough on its own. Nonetheless, it is common to delegate the serving of static content to more capable servers such as [nginx](#) or CDN.

Contents

- *Deploy*
 - *Run as a daemon*
 - *Run as a different user*
 - *PID files*
 - *Systemd socket activation*
 - *Control via Supervisord*
 - *SSL support*
 - *WSGI servers*
 - * *Embedding into another WSGI framework*
 - * *Tornado*
 - * *Twisted*
 - * *uwsgi*
 - *Virtual Hosting*
 - *Reverse-proxying*
 - * *Apache*
 - * *Nginx*

8.1 Run as a daemon

CherryPy allows you to easily decouple the current process from the parent environment, using the traditional double-fork:

```
from cherrypy.process.plugins import Daemonizer
d = Daemonizer(cherrypy.engine)
d.subscribe()
```

Note: This *engine plugin* is only available on Unix and similar systems which provide *fork()*.

If a startup error occurs in the forked children, the return code from the parent process will still be 0. Errors in the initial daemonizing process still return proper exit codes, but errors after the fork won't. Therefore, if you use this plugin to daemonize, don't use the return code as an accurate indicator of whether the process fully started. In fact, that return code only indicates if the process successfully finished the first fork.

The plugin takes optional arguments to redirect standard streams: `stdin`, `stdout`, and `stderr`. By default, these are all redirected to `/dev/null`, but you're free to send them to log files or elsewhere.

Warning: You should be careful to not start any threads before this plugin runs. The plugin will warn if you do so, because "...the effects of calling functions that require certain resources between the call to `fork()` and the call to an `exec` function are undefined". ([ref](#)). It is for this reason that the Server plugin runs at priority 75 (it starts worker threads), which is later than the default priority of 65 for the Daemonizer.

8.2 Run as a different user

Use this *engine plugin* to start your CherryPy site as root (for example, to listen on a privileged port like 80) and then reduce privileges to something more restricted.

This priority of this plugin's "start" listener is slightly higher than the priority for `server.start` in order to facilitate the most common use: starting on a low port (which requires root) and then dropping to another user.

```
DropPrivileges(cherrypy.engine, uid=1000, gid=1000).subscribe()
```

8.3 PID files

The `PIDFile` *engine plugin* is pretty straightforward: it writes the process id to a file on start, and deletes the file on exit. You must provide a `pidfile` argument, preferably an absolute path:

```
PIDFile(cherrypy.engine, '/var/run/myapp.pid').subscribe()
```

8.4 Systemd socket activation

Socket Activation is a systemd feature that allows to setup a system so that the systemd will sit on a port and start services 'on demand' (a little bit like `inetd` and `xinetd` used to do).

CherryPy has built-in socket activation support, if run from a systemd service file it will detect the `LISTEN_PID` environment variable to know that it should consider fd 3 to be the passed socket.

To read more about socket activation: <http://0pointer.de/blog/projects/socket-activation.html>

8.5 Control via Supervisor

`Supervisord` is a powerful process control and management tool that can perform a lot of tasks around process monitoring.

Below is a simple supervisor configuration for your CherryPy application.


```
[unix_http_server]
file=/tmp/supervisor.sock

[supervisord]
logfile=/tmp/supervisord.log ; (main log file;default $CWD/supervisord.log)
logfile_maxbytes=50MB       ; (max main logfile bytes b4 rotation;default 50MB)
logfile_backups=10          ; (num of main logfile rotation backups;default 10)
loglevel=info               ; (log level;default info; others: debug,warn,trace)
pidfile=/tmp/supervisord.pid ; (supervisord pidfile;default supervisord.pid)
nodaemon=false              ; (start in foreground if true;default false)
minfds=1024                 ; (min. avail startup file descriptors;default 1024)
minprocs=200                ; (min. avail process descriptors;default 200)

[rpcinterface:supervisor]
supervisor.rpcinterface_factory = supervisor.rpcinterface:make_main_rpcinterface

[supervisorctl]
serverurl=unix:///tmp/supervisor.sock

[program:myapp]
command=python server.py
environment=PYTHONPATH=.
directory=.
```

This could control your server via the `server.py` module as the application entry point.

```
import cherrypy

class Root(object):
    @cherrypy.expose
    def index(self):
        return "Hello World!"

cherrypy.config.update({'server.socket_port': 8090,
                       'engine.autoreload.on': False,
                       'log.access_file': './access.log',
                       'log.error_file': './error.log'})
cherrypy.quickstart(Root())
```

To take the configuration (assuming it was saved in a file called `supervisor.conf`) into account:

```
$ supervisord -c supervisord.conf
$ supervisorctl update
```

Now, you can point your browser at <http://localhost:8090/> and it will display *Hello World!*.

To stop supervisor, type:

```
$ supervisorctl shutdown
```

This will obviously shutdown your application.

8.6 SSL support

Note: You may want to test your server for SSL using the services from [Qualys, Inc.](#)

CherryPy can encrypt connections using SSL to create an https connection. This keeps your web traffic secure. Here's how.

1. Generate a private key. We'll use openssl and follow the [OpenSSL Keys HOWTO](#):

```
$ openssl genrsa -out privkey.pem 2048
```

You can create either a key that requires a password to use, or one without a password. Protecting your private key with a password is much more secure, but requires that you enter the password every time you use the key. For example, you may have to enter the password when you start or restart your CherryPy server. This may or may not be feasible, depending on your setup.

If you want to require a password, add one of the `-aes128`, `-aes192` or `-aes256` switches to the command above. You should not use any of the DES, 3DES, or SEED algorithms to protect your password, as they are insecure.

SSL Labs recommends using 2048-bit RSA keys for security (see references section at the end).

2. Generate a certificate. We'll use openssl and follow the [OpenSSL Certificates HOWTO](#). Let's start off with a self-signed certificate for testing:

```
$ openssl req -new -x509 -days 365 -key privkey.pem -out cert.pem
```

openssl will then ask you a series of questions. You can enter whatever values are applicable, or leave most fields blank. The one field you *must* fill in is the 'Common Name': enter the hostname you will use to access your site. If you are just creating a certificate to test on your own machine and you access the server by typing 'localhost' into your browser, enter the Common Name 'localhost'.

3. Decide whether you want to use python's built-in SSL library, or the pyOpenSSL library. CherryPy supports either.

(a) *Built-in*. To use python's built-in SSL, add the following line to your CherryPy config:

```
cherrypy.server.ssl_module = 'builtin'
```

(a) *pyOpenSSL*. Because python did not have a built-in SSL library when CherryPy was first created, the default setting is to use pyOpenSSL. To use it you'll need to install it (we could recommend you install [cython](#) first):

```
$ pip install cython, pyOpenSSL
```

2. Add the following lines in your CherryPy config to point to your certificate files:

```
cherrypy.server.ssl_certificate = "cert.pem"
cherrypy.server.ssl_private_key = "privkey.pem"
```

5. If you have a certificate chain at hand, you can also specify it:

```
cherrypy.server.ssl_certificate_chain = "certchain.pem"
```

6. Start your CherryPy server normally. Note that if you are debugging locally and/or using a self-signed certificate, your browser may show you security warnings.

8.7 WSGI servers

8.7.1 Embedding into another WSGI framework

Though CherryPy comes with a very reliable and fast enough HTTP server, you may wish to integrate your CherryPy application within a different framework. To do so, we will benefit from the WSGI interface defined in [PEP 333](#) and [PEP 3333](#).

Note that you should follow some basic rules when embedding CherryPy in a third-party WSGI server:

- If you rely on the “*main*” channel to be published on, as it would happen within the CherryPy’s mainloop, you should find a way to publish to it within the other framework’s mainloop.
- Start the CherryPy’s engine. This will publish to the “*start*” channel of the bus.

```
cherry.py.engine.start()
```

- Stop the CherryPy’s engine when you terminate. This will publish to the “*stop*” channel of the bus.

```
cherry.py.engine.stop()
```

- Do not call `cherry.py.engine.block()`.
- Disable the built-in HTTP server since it will not be used.

```
cherry.py.server.unsubscribe()
```

- Disable autoreload. Usually other frameworks won’t react well to it, or sometimes, provide the same feature.

```
cherry.py.config.update({'engine.autoreload.on': False})
```

- Disable CherryPy signals handling. This may not be needed, it depends on how the other framework handles them.

```
cherry.py.engine.signals.subscribe()
```

- Use the “*embedded*” environment configuration scheme.

```
cherry.py.config.update({'environment': 'embedded'})
```

Essentially this will disable the following:

- Stdout logging
- Autoreloader
- Configuration checker
- Headers logging on error
- Tracebacks in error
- Mismatched params error during dispatching
- Signals (SIGHUP, SIGTERM)

8.7.2 Tornado

You can use tornado HTTP server as follow:

```
import cherry.py

class Root(object):
    @cherry.py.expose
    def index(self):
        return "Hello World!"

if __name__ == '__main__':
    import tornado
    import tornado.httpserver
    import tornado.wsgi
```

```
# our WSGI application
wsgiapp = cherrypy.tree.mount(Root())

# Disable the autoreload which won't play well
cherrypy.config.update({'engine.autoreload.on': False})

# let's not start the CherryPy HTTP server
cherrypy.server.unsubscribe()

# use CherryPy's signal handling
cherrypy.engine.signals.subscribe()

# Prevent CherryPy logs to be propagated
# to the Tornado logger
cherrypy.log.error_log.propagate = False

# Run the engine but don't block on it
cherrypy.engine.start()

# Run thr tornado stack
container = tornado.wsgi.WSGIContainer(wsgiapp)
http_server = tornado.httpserver.HTTPServer(container)
http_server.listen(8080)
# Publish to the CherryPy engine as if
# we were using its mainloop
tornado.ioloop.PeriodicCallback(lambda: cherrypy.engine.publish('main'), 100).start()
tornado.ioloop.IOLoop.instance().start()
```

8.7.3 Twisted

You can use Twisted HTTP server as follow:

```
import cherrypy

from twisted.web.wsgi import WSGIResource
from twisted.internet import reactor
from twisted.internet import task

# Our CherryPy application
class Root(object):
    @cherrypy.expose
    def index(self):
        return "hello world"

# Create our WSGI app from the CherryPy application
wsgiapp = cherrypy.tree.mount(Root())

# Configure the CherryPy's app server
# Disable the autoreload which won't play well
cherrypy.config.update({'engine.autoreload.on': False})

# We will be using Twisted HTTP server so let's
# disable the CherryPy's HTTP server entirely
cherrypy.server.unsubscribe()

# If you'd rather use CherryPy's signal handler
```

```
# Uncomment the next line. I don't know how well this
# will play with Twisted however
#cherry.py.engine.signals.subscribe()

# Publish periodically onto the 'main' channel as the bus mainloop would do
task.LoopingCall(lambda: cherry.py.engine.publish('main')).start(0.1)

# Tie our app to Twisted
reactor.addSystemEventTrigger('after', 'startup', cherry.py.engine.start)
reactor.addSystemEventTrigger('before', 'shutdown', cherry.py.engine.exit)
resource = WSGIResource(reactor, reactor.getThreadPool(), wsgiapp)
```

Notice how we attach the bus methods to the Twisted's own lifecycle.

Save that code into a module named *cptw.py* and run it as follows:

```
$ twistd -n web --port 8080 --wsgi cptw.wsgiapp
```

8.7.4 uwsgi

You can use uwsgi HTTP server as follow:

```
import cherry.py

# Our CherryPy application
class Root(object):
    @cherry.py.expose
    def index(self):
        return "hello world"

cherry.py.config.update({'engine.autoreload.on': False})
cherry.py.server.unsubscribe()
cherry.py.engine.start()

wsgiapp = cherry.py.tree.mount(Root())
```

Save this into a Python module called *mymod.py* and run it as follows:

```
$ uwsgi --socket 127.0.0.1:8080 --protocol=http --wsgi-file mymod.py --callable wsgiapp
```

8.8 Virtual Hosting

CherryPy has support for virtual-hosting. It does so through a dispatchers that locate the appropriate resource based on the requested domain.

Below is a simple example for it:

```
import cherry.py

class Root(object):
    def __init__(self):
        self.app1 = App1()
        self.app2 = App2()

class App1(object):
    @cherry.py.expose
```

```
def index(self):
    return "Hello world from app1"

class App2(object):
    @cherrypy.expose
    def index(self):
        return "Hello world from app2"

if __name__ == '__main__':
    hostmap = {
        'company.com:8080': '/app1',
        'home.net:8080': '/app2',
    }

    config = {
        'request.dispatch': cherrypy.dispatch.VirtualHost(**hostmap)
    }

    cherrypy.quickstart(Root(), '/', {'/': config})
```

In this example, we declare two domains and their ports:

- company.com:8080
- home.net:8080

Thanks to the `cherrypy.dispatch.VirtualHost` dispatcher, we tell CherryPy which application to dispatch to when a request arrives. The dispatcher looks up the requested domain and call the according application.

Note: To test this example, simply add the following rules to your *hosts* file:

```
127.0.0.1    company.com
127.0.0.1    home.net
```

8.9 Reverse-proxying

8.9.1 Apache

8.9.2 Nginx

nginx is a fast and modern HTTP server with a small footprint. It is a popular choice as a reverse proxy to application servers such as CherryPy.

This section will not cover the whole range of features nginx provides. Instead, it will simply provide you with a basic configuration that can be a good starting point.

```
1 upstream apps {
2     server 127.0.0.1:8080;
3     server 127.0.0.1:8081;
4 }
5
6 gzip_http_version 1.0;
7 gzip_proxied      any;
8 gzip_min_length  500;
9 gzip_disable     "MSIE [1-6]\.";
```

```

10 gzip_types      text/plain text/xml text/css
11                text/javascript
12                application/javascript;
13
14 server {
15     listen 80;
16     server_name  www.example.com;
17
18     access_log   /app/logs/www.example.com.log combined;
19     error_log    /app/logs/www.example.com.log;
20
21     location ^~ /static/ {
22         root /app/static;
23     }
24
25     location / {
26         proxy_pass      http://apps;
27         proxy_redirect  off;
28         proxy_set_header Host $host;
29         proxy_set_header X-Real-IP $remote_addr;
30         proxy_set_header X-Forwarded-For $proxy_add_x_forwarded_for;
31         proxy_set_header X-Forwarded-Host $server_name;
32     }
33 }

```

Edit this configuration to match your own paths. Then, save this configuration into a file under `/etc/nginx/conf.d/` (assuming Ubuntu). The filename is irrelevant. Then run the following commands:

```

$ sudo service nginx stop
$ sudo service nginx start

```

Hopefully, this will be enough to forward requests hitting the nginx frontend to your CherryPy application. The `upstream` block defines the addresses of your CherryPy instances.

It shows that you can load-balance between two application servers. Refer to the nginx documentation to understand how this achieved.

```

upstream apps {
    server 127.0.0.1:8080;
    server 127.0.0.1:8081;
}

```

Later on, this block is used to define the reverse proxy section.

Now, let's see our application:

```

import cherrypy

class Root(object):
    @cherrypy.expose
    def index(self):
        return "hello world"

if __name__ == '__main__':
    cherrypy.config.update({
        'server.socket_port': 8080,
        'tools.proxy.on': True,
        'tools.proxy.base': 'http://www.example.com'
    })
    cherrypy.quickstart(Root())

```

If you run two instances of this code, one on each port defined in the nginx section, you will be able to reach both of them via the load-balancing done by nginx.

Notice how we define the proxy tool. It is not mandatory and used only so that the CherryPy request knows about the true client's address. Otherwise, it would know only about the nginx's own address. This is most visible in the logs.

The `base` attribute should match the `server_name` section of the nginx configuration.

Support

You've read the documentation and you've brushed up on the basics of Python and web development, but you still could use some help. Users have several options.

9.1 I have a question

If you have a question and cannot find an answer for it in issues or the the [documentation](#), [please create an issue](#).

Questions and their answers have great value for the community, and a tip is to really put the effort in and write a good explanation, you will get better and quicker answers. Examples are strongly encouraged.

9.2 I have found a bug

If no one have already, [create an issue](#). Be sure to provide ample information, remember that any help won't be better than your explanation.

Unless something is very obviously wrong, you are likely to be asked to provide a working example, displaying the erroneous behaviour.

Note: While this might feel troublesome, a tip is to always make a separate example that have the same dependencies as your project. It is great for troubleshooting those annoying problems where you don't know if the problem is at your end or the components. Also, you can then easily fork and provide as an example. You will get answers and resolutions way quicker. Also, many other open source projects require it.

9.3 I have a feature request

Good stuff! [Please create an issue!](#) Note: Features are more likely to be added the more users they seem to benefit.

9.4 I want to converse

The [gitter page](#) is good for when you want to discuss in real time or get pointed in the right direction.

Contribute

CherryPy is a community-maintained, open-source project hosted at Github. The project active encourages aspiring and experienced users to dive in and add their best contribution to the project.

How can you contribute? Well, first search the [docs](#) and the [project page](#) to see if someone has already reported your issue.

10.1 StackOverflow

On [StackOverflow](#), there are questions tagged with ‘cherrypy’. Answer unanswered questions, add an improved answer, clarify an answer with a comment, or ask more meaningful questions there. Earn reputation and share experience.

CherryPy also maintains a [StackOverflow Wiki](#) where anyone can publish tricks and techniques and refine others.

10.2 Filing Bug Reports

If you find a bug, an issue where the product doesn’t behave as you expect, you may file a bug report at [the project page](#). Be sure to include what your expectation was, what happened instead, details about your system that might be relevant, and steps that someone else could take to replicate your finding. The more detailed and exact your description, the better one of the volunteers on the project may be able to help resolve your issue.

10.3 Fixing Bugs

CherryPy has a number of open, reported [issues](#). Some of them are complicated and difficult, but others are more straightforward and shovel-ready. Feel free to find one that you think you can solve or introduce yourself and ask for guidance in [our gitter channel](#).

As you work through the issue and commit changes to your clone of the repository, be sure to add issue references to your changes (like “Fixes #999” or “Ref #999”) so your changes link to the issue and vice-versa.

10.4 Writing Pull Requests

To contribute, first read [How to write the perfect pull request](#) and file your contribution with the [CherryPy Project page](#).

Testing

- To run the regression tests, first install tox:

```
pip install 'tox>=2.5'
```

then run it

```
tox
```

- To run individual tests type:

```
tox -- -k test_foo
```

Glossary

application A CherryPy application is simply a class instance containing at least one page handler.

controller Loose name commonly given to a class owning at least one exposed method

exposed A Python function or method which has an attribute called *exposed* set to *True*. This attribute can be set directly or via the `cherrypy.expose()` decorator.

```
@cherrypy.expose
def method(...):
    ...
```

is equivalent to:

```
def method(...):
    ...
method.exposed = True
```

page handler Name commonly given to an exposed method

13.1 v10.0.0

20 Jan 2017

- #1332: CherryPy now uses `portend` for checking and waiting on ports for startup and teardown checks. The following names are no longer present:

- `cherrypy._cpserver.client_host`
- `cherrypy._cpserver.check_port`
- `cherrypy._cpserver.wait_for_free_port`
- `cherrypy._cpserver.wait_for_occupied_port`
- `cherrypy.process.servers.check_port`
- `cherrypy.process.servers.wait_for_free_port`
- `cherrypy.process.servers.wait_for_occupied_port`

Use this functionality from the `portend` package directly.

13.2 v9.0.0

19 Jan 2017

- #1481: Move functionality from `cherrypy.wsgiserver` to the `cheroot 5.0` project.

13.3 v8.9.1

16 Jan 2017

- #1537: Restore dependency on `pywin32` for Python 3.6.

13.4 v8.9.0

13 Jan 2017

- [#1547](#): Replaced `cherryd` distutils script with a `setuptools` console entry point.

When running CherryPy in daemon mode, the forked process no longer changes directory to `/`. If that behavior is something on which your application relied and should rely, please file a ticket with the project.

13.5 v8.8.0

09 Jan 2017

- [#1528](#): Allow a timeout of 0 to server.

13.6 v8.7.0

31 Dec 2016

- [#645](#): Setting a bind port of 0 will bind to an ephemeral port.

13.7 v8.6.0

27 Dec 2016

- [#1538](#) and [#1090](#): Removed cruft from the setup script and instead rely on `include_package_data` to ensure the relevant files are included in the package. Note, this change does cause `LICENSE.md` no longer to be included in the installed package.

13.8 v8.5.0

26 Dec 2016

- The `pyOpenSSL` support is now included on Python 3 builds, removing the last disparity between Python 2 and Python 3 in the CherryPy package. This change is one small step in consideration of [#1399](#). This change also fixes RPM builds, as reported in [#1149](#).

13.9 v8.4.0

26 Dec 2016

- [#1532](#): Also release wheels for Python 2, enabling offline installation.

13.10 v8.3.1

25 Dec 2016

- [#1537](#): Disable dependency on `pywin32` on Python 3.6 until a viable build of `pywin32` can be made on that Python version.

13.11 v8.3.0

24 Dec 2016

- Consolidated some documentation and include the more concise readme in the package long description, as found on PyPI.

13.12 v8.2.0

23 Dec 2016

- #1463: CherryPy tests are now run under pytest and invoked using tox.

13.13 v8.1.3

16 Dec 2016

- #1530: Fix the issue with TypeError being swallowed by decorated handlers.

13.14 v8.1.2

28 Sep 2016

- #1508

13.15 v8.1.1

27 Sep 2016

- #1497: Handle errors thrown by `ssl_module: 'builtin'` when client opens connection to HTTPS port using HTTP.
- #1350: Fix regression introduced in v6.1.0 where environment construction for `WSGIGateway_u0` was passing one parameter and not two.
- Other miscellaneous fixes.

13.16 v8.1.0

04 Sep 2016

- #1473: `HTTPError` now also works as a context manager.
- #1487: The sessions tool now accepts a `storage_class` parameter, which supersedes the new deprecated `storage_type` parameter. The `storage_class` should be the actual `Session` subclass to be used.
- Releases now use `setuptools_scm` to track the release versions. Therefore, releases can be cut by simply tagging a commit in the repo. Versions numbers are now stored in exactly one place.

13.17 v8.0.1

03 Sep 2016

- #1489 via #1493: Additionally reject anything else that's not bytes.
- #1492: systemd socket activation.

13.18 v8.0.0

02 Sep 2016

- #1483: Remove Deprecated constructs:
 - `cherry.py.lib.http` module.
 - `unrepr`, `modules`, and `attributes` in `cherry.py.lib`.
- #1476: Drop support for `python-memcached<1.58`
- #1401: Handle `NoSSLErrors`.
- #1489: In `wsgiserver.WSGIGateway.respond`, the application must now yield bytes and not text, as the spec requires. If text is received, it will now raise a `ValueError` instead of silently encoding using ISO-8859-1.
- Removed unicode filename from the package, working around pip #3894 and `setuptools` #704.

13.19 7.1.0

1458: Implement systemd's socket activation mechanism for CherryPy servers, based on work sponsored by Endless Computers.

Socket Activation allows one to setup a system so that systemd will sit on a port and start services 'on demand' (a little bit like `inetd` and `xinetd` used to do).

13.20 7.0.0

Removed the long-deprecated backward compatibility for legacy config keys in the engine. Use the config for the namespaced-plugins instead:

- `autoreload_on` -> `autoreload.on`
- `autoreload_frequency` -> `autoreload.frequency`
- `autoreload_match` -> `autoreload.match`
- `reload_files` -> `autoreload.files`
- `deadlock_poll_frequency` -> `timeout_monitor.frequency`

13.21 6.2.1

1460: Fix KeyError in Bus.publish when signal handlers set in config.

13.22 6.2.0

- [#1441](#): Added tool to automatically convert request params based on type annotations (primarily in Python 3). For example:

```
@cherry.py.tools.params() def resource(self, limit: int):
    assert isinstance(limit, int)
```

13.23 6.1.1

- [Issue #1411](#): Fix issue where autoreload fails when the host interpreter for CherryPy was launched using `python -m`.

13.24 6.1.0

- Combined `wsgiserver2` and `wsgiserver3` modules into a single module, `cherry.py.wsgiserver`.

13.25 6.0.2

- [Issue #1445](#): Correct additional typos.

13.26 6.0.1

- [Issue #1444](#): Correct typos in `@cherry.py.expose` decorators.

13.27 6.0.0

- `Setuptools` is now required to build CherryPy. Pure distutils installs are no longer supported. This change allows CherryPy to depend on other packages and re-use code from them. It's still possible to install pre-built CherryPy packages (wheels) using `pip` without `Setuptools`.
- `six` is now a requirement and subsequent requirements will be declared in the project metadata.
- [#1440](#): Back out changes from [#1432](#) attempting to fix redirects with Unicode URLs, as it also had the unintended consequence of causing the `'Location'` to be `bytes` on Python 3.
- `cherry.py.expose` now works on classes.
- `cherry.py.config` decorator is now used throughout the code internally.

13.28 5.6.0

- `@cherry.py.expose` now will also set the `exposed` attribute on a class.
- Rewrote all tutorials and internal usage to prefer the decorator usage of `expose` rather than setting the attribute explicitly.

- Removed test-specific code from tutorials.

13.29 5.5.0

- #1397: Fix for filenames with semicolons and quote characters in filenames found in headers.
- #1311: Added decorator for registering tools.
- #1194: Use simpler encoding rules for `SCRIPT_NAME` and `PATH_INFO` environment variables in CherryPy Tree allowing non-latin characters to pass even when `wsgi.version` is not `u.0`.
- #1352: Ensure that multipart fields are decoded even when cached in a file.

13.30 5.4.0

- `cherry.py.test.webtest.WebCase` now honors a `'WEBTEST_INTERACTIVE'` environment variable to disable interactive tests (still enabled by default). Set to `'0'` or `'false'` or `'False'` to disable interactive tests.
- #1408: Fix `AttributeError` when `listiterator` was accessed using the `next` attribute.
- #748: Removed `cherry.py.lib.sessions.PostgresqlSession`.
- #1432: Fix errors with redirects to Unicode URLs.

13.31 5.3.0

- #1202: Add support for specifying a certificate authority when serving SSL using the built-in SSL support.
- Use `ssl.create_default_context` when available.
- #1392: Catch platform-specific socket errors on OS X.
- #1386: Fix parsing of URIs containing `://` in the path part.

13.32 5.2.0

- #1410: Moved hosting to Github ([cherrypy/cherrypy](https://github.com/cherrypy/cherrypy)).

13.33 5.1.0

10 Mar 2016

- Bugfix issue #1315 for `test_HTTP11_pipelining` test in Python 3.5
- Bugfix issue #1382 regarding the keyword arguments support for Python 3 on the config file.
- Bugfix issue #1406 for `test_2_KeyboardInterrupt` test in Python 3.5. by monkey patching the `HTTPRequest` given a bug on CPython that is affecting the testsuite (<https://bugs.python.org/issue23377>).
- Add additional parameter `raise_subcls` to the tests helpers `openURL` and `CPWebCase.getPage` to have finer control on which exceptions can be raised.
- Add support for direct keywords on the calls (e.g. `foo=bar`) on the config file under Python 3.

- Add additional validation to determine if the process is running as a daemon on `cherrypy.process.plugins.SignalHandler` to allow the execution of the testsuite under CI tools.

13.34 5.0.1

05 Feb 2016

- Bugfix for `NameError` following #94.

13.35 5.0.0

04 Feb 2016

- Removed deprecated support for `ssl_certificate` and `ssl_private_key` attributes and implicit construction of SSL adapter on Python 2 WSGI servers.
- Default SSL Adapter on Python 2 is the builtin SSL adapter, matching Python 3 behavior.
- Pull request #94: In proxy tool, defer to Host header for resolving the base if no base is supplied.

13.36 4.0.0

19 Dec 2015

- Drop support for Python 2.5 and earlier.
- No longer build Windows installers by default.

13.37 3.8.2

17 Dec 2015

- Pull Request #116: Correct `InternalServerError` when null bytes in static file path. Now responds with 404 instead.

13.38 3.8.0

26 Jun 2015

- Pull Request #96: Pass `exc_info` to logger as keyword rather than formatting the error and injecting into the message.

13.39 3.7.0

24 Apr 2015

- CherryPy daemon may now be invoked with `python -m cherrypy` in addition to the `cherryd` script.

- [Issue #1298](#): Fix SSL handling on CPython 2.7 with builtin SSL module and pyOpenSSL 0.14. This change will break PyPy for now.
- Several documentation fixes.

13.40 3.6.0

14 Sep 2014

- Fixed HTTP range headers for negative length larger than content size.
- Disabled universal wheel generation as wsgiserver has Python duality.
- [Pull Request #42](#): Correct TypeError in `check_auth` when encrypt is used.
- [Pull Request #59](#): Correct signature of `HandlerWrapperTool`.
- [Pull Request #60](#): Fix error in `SessionAuth` where `login_screen` was incorrectly used.
- [Issue #1077](#): Support keyword-only arguments in dispatchers (Python 3).
- [Issue #1019](#): Allow logging host name in the access log.
- [Pull Request #50](#): Fixed race condition in session cleanup.

13.41 3.5.0

27 Jun 2014

- [Issue #1301](#): When the incoming queue is full, now reject additional connections. This functionality was added to CherryPy 3.0, but unintentionally lost in 3.1.

13.42 3.4.0

27 Jun 2014

- Miscellaneous quality improvements.

13.43 3.3.0

16 Apr 2014

CherryPy adopts semver.

CherryPy is a pythonic, object-oriented web framework.

CherryPy allows developers to build web applications in much the same way they would build any other object-oriented Python program. This results in smaller source code developed in less time.

CherryPy is now more than ten years old and it has proven to be fast and reliable. It is being used in production by many sites, from the simplest to the most demanding.

A CherryPy application typically looks like this:


```
import cherrypy

class HelloWorld(object):
    @cherrypy.expose
    def index(self):
        return "Hello World!"

cherrypy.quickstart(HelloWorld())
```

In order to make the most of CherryPy, you should start with the *tutorials* that will lead you through the most common aspects of the framework. Once done, you will probably want to browse through the *basics* and *advanced* sections that will demonstrate how to implement certain operations. Finally, you will want to carefully read the configuration and *extend* sections that go in-depth regarding the powerful features provided by the framework.

Above all, have fun with your application!

Symbols

- P, -Path
 - cherryd command line option, 7
- c, -config
 - cherryd command line option, 6
- d
 - cherryd command line option, 6
- e, -environment
 - cherryd command line option, 7
- f
 - cherryd command line option, 7
- i, -import
 - cherryd command line option, 7
- p, -pidfile
 - cherryd command line option, 7
- s
 - cherryd command line option, 7

A

application, **89**

C

cherryd command line option

- P, -Path, 7
- c, -config, 6
- d, 6
- e, -environment, 7
- f, 7
- i, -import, 7
- p, -pidfile, 7
- s, 7

controller, **89**

Ctrl-C, 47

E

exposed, **89**

F

FastCGI, 7

P

page handler, **89**

PID file, 7

Python Enhancement Proposals

- PEP 249, 50
- PEP 333, 48, 49, 76
- PEP 3333, 48, 49, 76
- PEP 343, 59

R

RFC

- RFC 2616, 11
- RFC 2616#sec10.4.5, 71
- RFC 2617, 37

S

SCGI, 7

shutdown, 47

W

Windows, 47