
Baseplate.py

Jun 22, 2020

Contents

1	Table of Contents	3
1.1	Tutorial	3
1.2	API Documentation	13
1.3	CLI Tools	94
1.4	Linters	100
2	Appendix	105
	Python Module Index	107
	Index	109

It's much easier to manage a bunch of services when they all have the same shape: the way they're developed, the way they interact with the infrastructure they run on, and the way they interact with each other. Baseplate is reddit's specification for the common shape of our services. This library, Baseplate.py, is the Python implementation of that specification.

Baseplate.py integrates with existing application frameworks and provides battle-tested libraries to give you everything you need to build a well-behaving production service without having to reinvent the wheel.

Here's a simple Baseplate.py HTTP service built using the [Pyramid web framework](#):

```
from baseplate import Baseplate
from baseplate.clients.sqlalchemy import SQLAlchemySession
from baseplate.frameworks.pyramid import BaseplateConfigurator
from pyramid.config import Configurator
from pyramid.view import view_config

@view_config(route_name="hello_world", renderer="json")
def hello_world(request):
    result = request.db.execute("SELECT date('now');")
    return {"Hello": "World", "Now": result.scalar()}

def make_wsgi_app(app_config):
    baseplate = Baseplate(app_config)
    baseplate.configure_observers()
    baseplate.configure_context({"db": SQLAlchemySession()})

    configurator = Configurator(settings=app_config)
    configurator.include(BaseplateConfigurator(baseplate).include_me)
    configurator.add_route("hello_world", "/", request_method="GET")
    configurator.scan()
    return configurator.make_wsgi_app()
```

Every request to this example service will automatically emit telemetry that allows you to dig into how the service is performing under the hood:

- Timers for how long the whole request took and how long was spent talking to the database.
- Counters for the success/failure of the whole request and each query to the database.
- Distributed tracing spans (including carrying over trace metadata from upstream services and onwards to downstream ones).
- Reporting of stack traces to Sentry on crash.

And you don't have to write any of that.

To get started, [dive into the tutorial](#). Or if you need an API reference, look below.

1.1 Tutorial

1.1.1 A tiny “Hello, World” service

In this tutorial, we’re going to build up a simple service to show off various aspects of Baseplate.py.

Prerequisites

This tutorial expects you to be familiar with Python and the basics of web application development. We will use [Python 3.7](#) and [virtual environments](#). To get set up, see [this guide on installing Python](#).

Make a home for our service

First, let’s create a folder and virtual environment to isolate the code and dependencies for this project.

```
$ mkdir tutorial
$ cd tutorial
$ virtualenv --python=python3.7 venv
Running virtualenv with interpreter /usr/bin/python3.7
Using base prefix '/usr'
New python executable in /home/user/tutorial/venv/bin/python3.7
Also creating executable in /home/user/tutorial/venv/bin/python
Installing setuptools, pkg_resources, pip, wheel...done.
$ source venv/bin/activate
```

Build a simple Pyramid service

Pyramid is a mature web framework for Python that we build HTTP services with. We’ll start our service out by using it without Baseplate.py at all:

```
$ pip install pyramid
```

Now let's write a tiny Pyramid service, open your editor and put the following in `helloworld.py`:

```
from wsgiref.simple_server import make_server

from pyramid.config import Configurator
from pyramid.view import view_config

@view_config(route_name="hello_world", renderer="json")
def hello_world(request):
    return {"Hello": "World"}

def make_wsgi_app():
    configurator = Configurator()
    configurator.add_route("hello_world", "/", request_method="GET")
    configurator.scan()
    return configurator.make_wsgi_app()

if __name__ == "__main__":
    app = make_wsgi_app()
    server = make_server("127.0.0.1", 9090, app)
    server.serve_forever()
```

Then run it:

```
$ python helloworld.py
```

Now that you have got a server running, let's try talking to it. From another terminal:

```
$ curl localhost:9090
{"Hello": "World"}
```

and the server should have logged about that request:

```
127.0.0.1 - - [06/Aug/2019 23:32:40] "GET / HTTP/1.1" 200 18
```

Great! It does not do much, but we have got a very basic service up and running now.

Breaking it down

See also:

You can get way more detail about what's going on in Pyramid in [Pyramid's own tutorial](#).

There are three things going on in this tiny service. Following how the code actually runs, we start out at the end of the file with the creation of the HTTP server:

```
if __name__ == "__main__":
    app = make_wsgi_app()
    server = make_server("127.0.0.1", 9090, app)
    server.serve_forever()
```


This is using the `wsgiref` module from the Python standard library to run a basic development server. WSGI is the Python standard interface between HTTP servers and applications. Pyramid applications are WSGI applications and can be run on any WSGI server.

Note: This server will do fine for this quick start, but we won't want to stick with it as we scale up as it can't handle multiple requests at the same time.

This server code calls our `make_wsgi_app` function to get the actual application. Let's look at that next:

```
def make_wsgi_app():
    configurator = Configurator()
    configurator.add_route("hello_world", "/", request_method="GET")
    configurator.scan()
    return configurator.make_wsgi_app()
```

The real workhorse here is the `Configurator` object from Pyramid. This object helps us configure and build an application.

```
configurator.add_route("hello_world", "/", request_method="GET")
```

First off, we add a route that maps the URL path `/` to the route named `hello_world` when the HTTP verb is `GET`. This means that when a request comes in that matches those criteria, Pyramid will try to find a “view” function that is registered for that route name.

```
configurator.scan()
```

Then we tell Pyramid to scan the current module for `declarative registrations`. Because of the `@view_config` decorator, Pyramid will find the `hello_world` function in our service and recognize that we have registered it to handle the `hello_world` route.

```
return configurator.make_wsgi_app()
```

Finally, we ask the configurator to build a WSGI application based on what we have configured and return that to the server.

At this point, we have done the one-time application startup and handed off our application to the server which is ready to call into it when requests come in. Now it's time to look at the code that actually runs on each request.

```
@view_config(route_name="hello_world", renderer="json")
def hello_world(request):
    return {"Hello": "World"}
```

This function gets called each time a matching request comes in. Pyramid will build a `Request` object and pass it into our function as `request`. This contains all the extra information about the request, like form fields and header values. Whatever gets returned from this function will be rendered by the `renderer` we specified in the `@view_config` and then sent to the client.

Summary

We have built a tiny service on Pyramid and understand how the code all fits together. So far, there's been no `Baseplate.py` at all. Next up, we'll look at what's involved with adding it in.

1.1.2 Servers and Configuration files

In this chapter, we'll lay the foundation for using Baseplate.py in the service.

Install Baseplate.py

First off, let's install Baseplate.py in your virtual environment so we can start using its components.

```
$ pip install 'git+https://github.com/apache/thrift#egg=thrift&subdirectory=lib/py'
$ pip install git+https://github.com/reddit/baseplate.py
```

In the previous chapter, we made our service run its own HTTP/WSGI server. Now we're going to use Baseplate.py's server instead which is run with **baseplate-serve**.

```
$ baseplate-serve
usage: baseplate-serve [-h] [--debug] [--reload] [--app-name NAME]
                      [--server-name NAME] [--bind ENDPOINT]
                      config_file
baseplate-serve: error: the following arguments are required: config_file
```

Uh oh! `config_file`!? I guess we have got some more to do first.

A configuration file

Baseplate services rely on configuration to allow them to behave differently in different environments (development, staging, production, etc.). For Baseplate.py, configuration is stored in a file in standard Python INI file format as understood by `configparser`.

Open a new `helloworld.ini` in the tutorial directory and copy this into it:

```
[app:main]
factory = helloworld:make_wsgi_app

[server:main]
factory = baseplate.server.wsgi
```

Breaking it down, there are two sections to this configuration file, `[app:main]` and `[server:main]`.

```
[app:main]
factory = helloworld:make_wsgi_app
```

The first section defines the entrypoint and settings for the application itself. The `factory` is a function that returns an application object. In this case, it lives in the Python module `helloworld` and the function is called `make_wsgi_app`.

```
[server:main]
factory = baseplate.server.wsgi
```

The second section defines what kind of server we'll run and the settings for that server. Since our application is built for HTTP/WSGI, we use the WSGI server in Baseplate.py.

Note: You might notice that both application and server sections have `:main` in their names. By default, Baseplate.py tools like **baseplate-serve** will look for the sections with `main` in them, but you can override this

with `--app-name=foo` to look up `[app:foo]` or `--server-name` similarly. This allows you to have multiple applications and servers defined in the same configuration file.

OK! Now let's try **baseplate-serve** with our configuration file.

```
$ baseplate-serve helloworld.ini
Traceback (most recent call last):
  File "/home/user/tutorial/venv/bin/baseplate-serve", line 14, in <module>
    load_app_and_run_server()
  File "/home/user/tutorial/venv/lib/python3.7/site-packages/baseplate/server/__init__
↪.py", line 226, in load_app_and_run_server
    app = make_app(config.app)
  File "/home/user/tutorial/venv/lib/python3.7/site-packages/baseplate/server/__init__
↪.py", line 180, in make_app
    return factory(app_config)
TypeError: make_wsgi_app() takes 0 positional arguments but 1 was given
```

It looks like we'll need a little bit more.

Run the service with **baseplate-serve**

In the previous section, we learned that the `[app:main]` section both tells **baseplate-serve** where to find the application *and* holds configuration for that application. The function that we specify in `factory` needs to take a dictionary of the raw configuration values as an argument. Let's add that to our service.

```
from pyramid.config import Configurator
from pyramid.view import view_config

@view_config(route_name="hello_world", renderer="json")
def hello_world(request):
    return {"Hello": "World"}

def make_wsgi_app(app_config):
    configurator = Configurator(settings=app_config)
    configurator.add_route("hello_world", "/", request_method="GET")
    configurator.scan()
    return configurator.make_wsgi_app()
```

All we had to do was add one parameter. We also pass it through to Pyramid's Configurator so any framework-specific settings can be picked up.

Since we're not using the `wsgiref` server anymore, we can drop the whole `if __name__ == "__main__":` section at the end of the file now.

Alright, third time's the charm, right?

```
$ baseplate-serve --debug helloworld.ini
12593:MainThread:baseplate.server.runtime_monitor:INFO:No metrics client configured.
↪Server metrics will not be sent.
12593:MainThread:baseplate.server:INFO:Listening on ('127.0.0.1', 9090)
```

Success! The `--debug` flag will turn on some extra log messages, so we can see a request log when we try hitting the service with **curl** again.

```
$ curl localhost:9090
{"Hello": "World"}
```

And something shows up in the server's logs:

```
12593:DummyThread-1:baseplate.server.wsgi:DEBUG:127.0.0.1 - - [2019-08-07 23:42:32]
↪ "GET / HTTP/1.1" 200 147 0.007743
```

You'll notice the logs look a bit different from before. **baseplate-serve** adds some extra info to help give context to your log entries. That `DummyThread-1` is pretty useless though, so we'll make it useful in the next chapter.

Summary

We have now made a configuration file and made it possible to run our service with **baseplate-serve**.

So what did any of this do for us? **baseplate-serve** is how we let production infrastructure run our application and interact with it. It knows how to process multiple requests simultaneously and will handle things like the infrastructure asking it to gracefully shut down.

But the real fun of `Baseplate.py` comes when we start using its framework integration to get some visibility into the guts of the application. Let's see what that looks like in the next chapter.

1.1.3 Observers — Looking under the hood

In the previous chapter we took our basic Pyramid service and made it compatible with **baseplate-serve**. Now we'll start using `Baseplate.py` inside the service itself so we can see what's going on while handling requests.

Wire up the Baseplate object

The heart of `Baseplate.py`'s framework is the `Baseplate` object. No matter what kind of service you're writing—Pyramid, Thrift, etc.—this class works the same. The magic happens when we use one of `Baseplate.py`'s framework integrations to connect the two things together. Let's do that in our service now.

```
from baseplate import Baseplate
from baseplate.frameworks.pyramid import BaseplateConfigurator
from pyramid.config import Configurator
from pyramid.view import view_config

@view_config(route_name="hello_world", renderer="json")
def hello_world(request):
    return {"Hello": "World"}

def make_wsgi_app(app_config):
    baseplate = Baseplate(app_config)
    baseplate.configure_observers()

    configurator = Configurator(settings=app_config)
    configurator.include(BaseplateConfigurator(baseplate).includeme)
    configurator.add_route("hello_world", "/", request_method="GET")
    configurator.scan()
    return configurator.make_wsgi_app()
```

This is all we need to do to get basic observability in our service. Line-by-line:

```
baseplate = Baseplate(app_config)
```

We create a *Baseplate* object during application startup.

```
baseplate.configure_observers()
```

Then we call *configure_observers()* and pass in the application configuration. We'll talk more about this in a moment.

```
configurator.include(BaseplateConfigurator(baseplate).includeme)
```

Finally we connect up with Pyramid's framework to integrate it all together.

We can now run our server again and make some requests to it to see what's different.

```
$ baseplate-serve --debug helloworld.ini
1072:MainThread:baseplate:DEBUG:The following observers are unconfigured and won't_
↳run: metrics, tracing, error_reporter
1072:MainThread:baseplate.server.runtime_monitor:INFO:No metrics client configured.
↳Server metrics will not be sent.
1072:MainThread:baseplate.server:INFO:Listening on ('127.0.0.1', 9090)
1072:3776872808671626432:baseplate.server.wsgi:DEBUG:127.0.0.1 - - [2019-08-08_
↳04:46:58] "GET / HTTP/1.1" 200 147 0.008789
```

It still works and things don't look too different. The first thing you'll see is the *observers are unconfigured* line. This is there because we called *configure_observers()*. We did not add anything to our configuration file so of course they're all unconfigured!

There is one other change even with those unconfigured observers. The log line for the test request we sent no longer says *DummyThread-1* but instead has a really long number in its place. That's the Trace ID of the request. You'll see in the next section that when a single request causes multiple log lines, they'll all have the same Trace ID which helps correlate them.

Note: In fact, the Trace ID will be the same across all services involved in handling a single end-user request. We'll talk more about this in a later chapter.

That's sort of useful but we can do better. Next, let's configure another observer to get more visibility.

Configure the metrics observer

After the previous section, our application is now wired up to use Baseplate.py with Pyramid. Now we'll turn on an observer to see it in action.

One of the available observers sends metrics to *StatsD*. We'll turn that on, but since we don't actually have a StatsD running we'll leave it in debug mode that just prints the metrics it would send out to the logs instead.

```
[app:main]
factory = helloworld:make_wsgi_app

metrics.namespace = helloworld

[server:main]
factory = baseplate.server.wsgi
```

To turn it on, we just add one line to our configuration file. This tells the observer what base name it should use for the metrics it sends. Once we have done that, we can start the server up again.

```
$ baseplate-serve --debug helloworld.ini
1104:MainThread:baseplate:DEBUG:The following observers are unconfigured and won't_
↳run: tracing, error_reporter
1104:MainThread:baseplate.server:INFO:Listening on ('127.0.0.1', 9090)
1104:2115808718993382189:baseplate.lib.metrics:DEBUG:Would send metric b'helloworld.
↳server.hello_world:3.53074|ms'
1104:2115808718993382189:baseplate.lib.metrics:DEBUG:Would send metric b'helloworld.
↳server.hello_world.success:1|c'
1104:2115808718993382189:baseplate.server.wsgi:DEBUG:127.0.0.1 - - [2019-08-08_
↳04:47:32] "GET / HTTP/1.1" 200 147 0.009720
```

If it worked right, `metrics` won't be listed as an unconfigured observer anymore. Now when you make requests to your service you'll see a few extra log lines that say `Would send metric . . .`. These are the metrics the observer would be sending if we had a StatsD server set up. Also note that the trace ID is the same on all these log lines.

Since our service is super simple, we only get two metrics on each request. The first is a timer that tracks how long the endpoint took to respond to the request. The second metric increments a success or failure counter every time the endpoint responds or crashes.

If you leave your server running long enough, you'll also see some extra metrics appear periodically:

```
1104:Server Monitoring:baseplate.lib.metrics:DEBUG:Would send metric b'helloworld.
↳runtime.reddit.PID1104.active_requests:0|g'
1104:Server Monitoring:baseplate.lib.metrics:DEBUG:Would send metric b'helloworld.
↳runtime.reddit.PID1104.gc.gen0.collections:154|g'
1104:Server Monitoring:baseplate.lib.metrics:DEBUG:Would send metric b'helloworld.
↳runtime.reddit.PID1104.gc.gen0.collected:8244|g'
. . .
```

These metrics come out of the server itself and track information that's not specific to an individual request but rather about the overall health of the service. This includes things like statistics from Python's garbage collector, the state of any connection pools, and how many concurrent requests your application is handling.

Summary

We have integrated `Baseplate.py`'s tools into our service and started seeing some of the benefit of its observers. Our service is pretty simple still though, it's about time it actually talks to something else. In the next chapter, we'll add a database and see what that looks like.

1.1.4 Clients — Talking to the outside world

In the previous chapter, we integrated `Baseplate.py` into our Pyramid application and got some observers working to see how things were performing. Most applications are not so self-contained though and have to talk to other services to do their job. In this chapter, we'll add a dependency on a database to see what that looks like.

Adding a database

We're going to use a popular Python ORM called [SQLAlchemy](#) to talk to our database. Let's install that to get started:

```
$ pip install sqlalchemy
```

Now that's installed, we can use `Baseplate.py`'s helpers to add `SQLAlchemy` to our service.

```

from baseplate import Baseplate
from baseplate.clients.sqlalchemy import SQLAlchemySession
from baseplate.frameworks.pyramid import BaseplateConfigurator
from pyramid.config import Configurator
from pyramid.view import view_config

@view_config(route_name="hello_world", renderer="json")
def hello_world(request):
    result = request.db.execute("SELECT date('now');")
    return {"Hello": "World", "Now": result.scalar()}

def make_wsgi_app(app_config):
    baseplate = Baseplate(app_config)
    baseplate.configure_observers()
    baseplate.configure_context({"db": SQLAlchemySession()})

    configurator = Configurator(settings=app_config)
    configurator.include(BaseplateConfigurator(baseplate).includeme)
    configurator.add_route("hello_world", "/", request_method="GET")
    configurator.scan()
    return configurator.make_wsgi_app()

```

Pretty simple, but there's something subtle going on here. Let's dig into it.

```
baseplate.configure_context({"db": SQLAlchemySession()})
```

This call to `configure_context()`, during application startup, tells Baseplate.py that we want to add a SQLAlchemy `Session` to the “context” with the name `db`.

What exactly the “context” is depends on what framework you're using, but for Pyramid applications it's the `request` object that Pyramid gives to every request handler.

Note: Why do we pass in the context configuration as a dictionary? It's possible to set up multiple clients at the same time this way. You can even do more complicated things like nesting dictionaries to organize the stuff you add to the context. See `configure_context()` for more info.

```
result = request.db.execute("SELECT date('now');")
```

Since we have connected the `Baseplate` object with Pyramid and told it to configure the context like this, we'll now see a `db` attribute on the `request` that has that SQLAlchemy session we wanted.

OK. Now we have got that wired up, let's try running it.

```

$ baseplate-serve --debug helloworld.ini
...
baseplate.lib.config.ConfigurationError: db.url: no value specified

```

Ah! It looks like we have got some configuring to do.

Configure the new client

Telling Baseplate.py that we wanted to add the SQLAlchemy session to our context did not actually give it any hint about how that session should be configured. SQLAlchemy can transparently handle different SQL databases for us and the location at which to find them will be different depending on if we're running in production, staging, or development. So it's time for the configuration file again.

```
[app:main]
factory = helloworld:make_wsgi_app

metrics.namespace = helloworld

db.url = sqlite:///

[server:main]
factory = baseplate.server.wsgi
```

To wire up the database, all we need is to add a SQLAlchemy [URL](#) to the configuration file. Because we configured the session to use the name `db` the relevant configuration line is prefixed with that name.

We're just going to use an in-memory [SQLite](#) database here because it's built into Python and we don't have to install anything else.

Now when we fire up the service, it launches and returns requests with our new data.

```
$ curl localhost:9090
{"Hello": "World", "Now": "2019-08-09"}
```

Great! If you look at the server logs when you make a request, you'll notice there are new metrics:

```
$ baseplate-serve --debug helloworld.ini
...
17905:7296338476964580186:baseplate.lib.metrics:DEBUG:Would send metric b'helloworld.
↪clients.db.execute:0.0824928|ms'
17905:7296338476964580186:baseplate.lib.metrics:DEBUG:Would send metric b'helloworld.
↪server.hello_world:4.16493|ms'
17905:7296338476964580186:baseplate.lib.metrics:DEBUG:Would send metric b'helloworld.
↪clients.db.execute.success:1|c'
17905:7296338476964580186:baseplate.lib.metrics:DEBUG:Would send metric b'helloworld.
↪server.hello_world.success:1|c'
...
```

The Baseplate.py SQLAlchemy integration automatically tracked usage of our database and reports timers, counters, and other goodies to our monitoring systems.

Summary

We have now hooked up our service to a simple database and when we run queries Baseplate.py automatically tracks them and emits telemetry.

1.2 API Documentation

1.2.1 Observability Framework

The observability framework is the part of Baseplate.py that integrates with other Python application frameworks to bring automatic telemetry to your application.

baseplate

The heart of the Baseplate framework is its telemetry system. Here's an incomplete example of an application built with the framework:

```
def do_something(request):
    result = request.db.execute("SELECT date('now');")

def make_app(app_config):
    ... snip ...

    baseplate = Baseplate(app_config)
    baseplate.configure_observers()
    baseplate.configure_context({"db": SQLAlchemySession()})

    ... snip ...
```

When a request is made that routes to the `do_something` handler, a *ServerSpan* is automatically created to track the time spent processing the request in our application. If the incoming request has trace headers, the constructed server span will have the same IDs as the upstream service's child span or else it will start a new trace with randomized values.

When we call `request.db.execute(...)` in the handler, Baseplate creates a child *Span* object to represent the time taken running the query.

The creation of the server and child spans trigger callbacks on all the *ServerSpanObserver* and *SpanObserver* objects registered. Because we called `configure_observers()` in our setup, this means we have observers that send telemetry about how our service is functioning.

Framework Configuration

At the root of each application is a single instance of *Baseplate*. This object can be integrated with various other frameworks (e.g. Thrift, Pyramid, etc.) using one of *the integrations*.

class `baseplate.Baseplate` (*app_config=None*)

The core of the Baseplate framework.

This class coordinates monitoring and tracing of service calls made to and from this service. See *baseplate.frameworks* for how to integrate it with the application framework you are using.

__init__ (*app_config=None*)

Initialize the core observability framework.

Parameters `app_config` (*Optional[Dict[str, str]]*) – The raw configuration dictionary for your application as supplied by **baseplate-serve** or **baseplate-script**. In addition to allowing framework-level configuration (described next), if this is supplied you do not need to pass the configuration again when calling `configure_observers()` or `configure_context()`.

Baseplate services can identify themselves to downstream services in requests. The name a service identifies as defaults to the Python module the `Baseplate` object is instantiated in. To override the default, make sure you are passing in an `app_config` and configure the name in your INI file:

```
[app:main]
baseplate.service_name = foo_service
...
```

Return type `None`

configure_observers (*app_config=None, module_name=None*)

Configure diagnostics observers based on application configuration.

This installs all the currently supported observers that have settings in the configuration file.

See `baseplate.observers` for the configuration settings available for each observer.

Parameters

- **app_config** (`Optional[Dict[str, str]]`) – The application configuration which should have settings for the error reporter. If not specified, the config must be passed to the `Baseplate()` constructor.
- **module_name** (`Optional[str]`) – Name of the root package of the application. If not specified, will be guessed from the package calling this function.

Return type `None`

configure_context (*args, **kwargs)

Add a number of objects to each request's context object.

Configure and attach multiple clients to the `RequestContext` in one place. This takes a full configuration spec like `baseplate.lib.config.parse_config()` and will attach the specified structure onto the context object each request.

For example, a configuration like:

```
baseplate = Baseplate(app_config)
baseplate.configure_context({
    "cfg": {
        "doggo_is_good": config.Boolean,
    },
    "cache": MemcachedClient(),
    "cassandra": {
        "foo": CassandraClient("foo_keyspace"),
        "bar": CassandraClient("bar_keyspace"),
    },
})
```

would build a context object that could be used like:

```
assert context.cfg.doggo_is_good == True
context.cache.get("example")
context.cassandra.foo.execute()
```

Parameters

- **app_config** – The raw stringy configuration dictionary.
- **context_spec** – A specification of what the configuration should look like.

Return type `None`

add_to_context (*name*, *context_factory*)

Add an attribute to each request's context object.

On each request, the factory will be asked to create an appropriate object to attach to the `RequestContext`.

Parameters

- **name** (`str`) – The attribute on the context object to attach the created object to. This may also be used for metric/tracing purposes so it should be descriptive.
- **context_factory** (`baseplate.clients.ContextFactory`) – A factory.

Return type `None`

Per-request Context

Each time a new request comes in to be served, the time taken to handle the request is represented by a new `ServerSpan` instance. During the course of handling that request, our application might make calls to remote services or do expensive calculations, the time spent can be represented by child `Span` instances.

Spans have names and IDs and track their parent relationships. When calls are made to remote services, the information that identifies the local child span representing that service call is passed along to the remote service and becomes the server span in the remote service. This allows requests to be traced across the infrastructure.

Small bits of data, called annotations, can be attached to spans as well. This could be the URL fetched, or how many items were sent in a batch, or whatever else might be helpful.

class `baseplate.RequestContext` (*context_config*, *prefix=None*, *span=None*, *wrapped=None*)

The request context object.

The context object is passed into each request handler by the framework you're using. In some cases (e.g. Pyramid) the request object will also inherit from another base class and carry extra framework-specific information.

Clients and configuration added to the context via `configure_context()` or `add_to_context()` will be available as an attribute on this object. To take advantage of Baseplate's automatic monitoring, any interactions with external services should be done through these clients.

`Baseplate.make_context_object()`

Make a context object for the request.

Return type `RequestContext`

`Baseplate.make_server_span` (*context*, *name*, *trace_info=None*)

Return a server span representing the request we are handling.

In a server, a server span represents the time spent on a single incoming request. Any calls made to downstream services will be new child spans of the server span, and the server span will in turn be the child span of whatever upstream request it is part of, if any.

Parameters

- **context** (`RequestContext`) – The `RequestContext` for this request.
- **name** (`str`) – A name to identify the type of this request, e.g. a route or RPC method name.
- **trace_info** (`Optional[TraceInfo]`) – The trace context of this request as passed in from upstream. If `None`, a new trace context will be generated.

Return type `ServerSpan`

`Baseplate.server_context(name)`

Create a server span and return a context manager for its lifecycle.

This is a convenience wrapper around a common pattern seen outside of servers handling requests. For example, simple cron jobs or one-off scripts might want to create a temporary span and access the context object. Instead of calling `make_context_object()` followed by `make_server_span()` manually, this method bundles it all up for you:

```
with baseplate.server_context("foo") as context:
    context.redis.ping()
```

Note: This should not be used within an existing span context (such as during request processing) as it creates a new span unrelated to any other ones.

Return type `Iterator[RequestContext]`

`class baseplate.ServerSpan(trace_id, parent_id, span_id, sampled, flags, name, context, baseplate=None)`

A server span represents a request this server is handling.

The server span is available on the `RequestContext` during requests as the `trace` attribute.

`finish(exc_info=None)`

Record the end of the span.

Parameters `exc_info` (`Optional[Tuple[Optional[Type[BaseException]], Optional[BaseException], Optional[traceback]]]`) – If the span ended because of an exception, this is the exception information. The default is `None` which indicates normal exit.

Return type `None`

`incr_tag(key, delta=1)`

Increment a tag value on the span.

This is useful to count instances of an event in your application. In addition to showing up as a tag on the span, the value may also be aggregated separately as an independent counter.

Parameters

- **key** (`str`) – The name of the tag.
- **value** – The amount to increment the value. Defaults to 1.

Return type `None`

`log(name, payload=None)`

Add a log entry to the span.

Log entries are timestamped events recording notable moments in the lifetime of a span.

Parameters

- **name** (`str`) – The name of the log entry. This should be a stable identifier that can apply to multiple span instances.
- **payload** (`Optional[Any]`) – Optional log entry payload. This can be arbitrary data.

Return type `None`

make_child (*name*, *local=False*, *component_name=None*)

Return a child Span whose parent is this Span.

The child span can either be a local span representing an in-request operation or a span representing an outbound service call.

In a server, a local span represents the time spent within a local component performing an operation or set of operations. The local component is some grouping of business logic, which is then split up into operations which could each be wrapped in local spans.

Parameters

- **name** (*str*) – Name to identify the operation this span is recording.
- **local** (*bool*) – Make this span a LocalSpan if True, otherwise make this span a base Span.
- **component_name** (*Optional[str]*) – Name to identify local component this span is recording in if it is a local span.

Return type *Span*

register (*observer*)

Register an observer to receive events from this span.

Return type *None*

set_tag (*key*, *value*)

Set a tag on the span.

Tags are arbitrary key/value pairs that add context and meaning to the span, such as a hostname or query string. Observers may interpret or ignore tags as they desire.

Parameters

- **key** (*str*) – The name of the tag.
- **value** (*Any*) – The value of the tag.

Return type *None*

start ()

Record the start of the span.

This notifies any observers that the span has started, which indicates that timers etc. should start ticking.

Spans also support the [context manager protocol](#), for use with Python's `with` statement. When the context is entered, the span calls `start()` and when the context is exited it automatically calls `finish()`.

Return type *None*

class `baseplate.Span` (*trace_id*, *parent_id*, *span_id*, *sampled*, *flags*, *name*, *context*, *baseplate=None*)

A span represents a single RPC within a system.

register (*observer*)

Register an observer to receive events from this span.

Return type *None*

start ()

Record the start of the span.

This notifies any observers that the span has started, which indicates that timers etc. should start ticking.

Spans also support the [context manager protocol](#), for use with Python's `with` statement. When the context is entered, the span calls `start()` and when the context is exited it automatically calls `finish()`.

Return type `None`

set_tag (*key*, *value*)

Set a tag on the span.

Tags are arbitrary key/value pairs that add context and meaning to the span, such as a hostname or query string. Observers may interpret or ignore tags as they desire.

Parameters

- **key** (`str`) – The name of the tag.
- **value** (`Any`) – The value of the tag.

Return type `None`

incr_tag (*key*, *delta=1*)

Increment a tag value on the span.

This is useful to count instances of an event in your application. In addition to showing up as a tag on the span, the value may also be aggregated separately as an independent counter.

Parameters

- **key** (`str`) – The name of the tag.
- **value** – The amount to increment the value. Defaults to 1.

Return type `None`

log (*name*, *payload=None*)

Add a log entry to the span.

Log entries are timestamped events recording notable moments in the lifetime of a span.

Parameters

- **name** (`str`) – The name of the log entry. This should be a stable identifier that can apply to multiple span instances.
- **payload** (`Optional[Any]`) – Optional log entry payload. This can be arbitrary data.

Return type `None`

finish (*exc_info=None*)

Record the end of the span.

Parameters **exc_info** (`Optional[Tuple[Optional[Type[BaseException]], Optional[BaseException], Optional[traceback]]]`) – If the span ended because of an exception, this is the exception information. The default is `None` which indicates normal exit.

Return type `None`

make_child (*name*, *local=False*, *component_name=None*)

Return a child Span whose parent is this Span.

Return type `Span`

class `baseplate.TraceInfo`

Trace context for a span.

If this request was made at the behest of an upstream service, the upstream service should have passed along trace information. This class is used for collecting the trace context and passing it along to the server span.

trace_id

The ID of the whole trace. This will be the same for all downstream requests.

parent_id

The ID of the parent span, or None if this is the root span.

span_id

The ID of the current span. Should be unique within a trace.

sampled

True if this trace was selected for sampling. Will be propagated to child spans.

flags

A bit field of extra flags about this trace.

classmethod new()

Generate IDs for a new initial server span.

This span has no parent and has a random ID. It cannot be correlated with any upstream requests.

Return type *TraceInfo*

classmethod from_upstream(trace_id, parent_id, span_id, sampled, flags)

Build a *TraceInfo* from individual headers.

Parameters

- **trace_id** (*Optional[int]*) – The ID of the trace.
- **parent_id** (*Optional[int]*) – The ID of the parent span.
- **span_id** (*Optional[int]*) – The ID of this span within the tree.
- **sampled** (*Optional[bool]*) – Boolean flag to determine request sampling.
- **flags** (*Optional[int]*) – Bit flags for communicating feature flags downstream

Raises *ValueError* if any of the values are inappropriate.

Return type *TraceInfo*

Observers

To actually do something with all these spans, Baseplate provides observer interfaces which receive notification of events happening in the application via calls to various methods.

The base type of observer is *BaseplateObserver* which can be registered with the root *Baseplate* instance using the *register()* method. Whenever a new server span is created in your application (i.e. a new request comes in to be served) the observer has its *on_server_span_created()* method called with the relevant details. This method can register *ServerSpanObserver* instances with the new server span to receive events as they happen.

Spans can be notified of five common events:

- *on_start()*, the span started.
- *on_set_tag()*, a tag was set on the span.
- *on_log()*, a log was entered on the span.
- *on_finish()*, the span finished.
- *on_child_span_created()*, a new child span was created.

New child spans are created in the application automatically by various client library wrappers e.g. for a call to a remote service or database, and can also be created explicitly for local actions like expensive computations. The handler can register new *SpanObserver* instances with the new child span to receive events as they happen.

It's up to the observers to attach meaning to these events. For example, the metrics observer would start a timer `on_start()` and record the elapsed time to StatsD `on_finish()`.

`Baseplate.register(observer)`

Register an observer.

Parameters `observer` (*BaseplateObserver*) – An observer.

Return type `None`

class `baseplate.BaseplateObserver`

Interface for an observer that watches Baseplate.

`on_server_span_created(context, server_span)`

Do something when a server span is created.

Baseplate calls this when a new request begins.

Parameters

- **context** (*RequestContext*) – The *RequestContext* for this request.
- **server_span** (*ServerSpan*) – The span representing this request.

Return type `None`

class `baseplate.ServerSpanObserver`

Interface for an observer that watches the server span.

`on_child_span_created(span)`

Do something when a child span is created.

SpanObserver objects call this when a new child span is created.

Parameters `span` (*Span*) – The new child span.

Return type `None`

`on_finish(exc_info)`

Do something when the observed span is finished.

Parameters `exc_info` (`Optional[Tuple[Optional[Type[BaseException]], Optional[BaseException], Optional[traceback]]]`) – If the span ended because of an exception, the exception info. Otherwise, `None`.

Return type `None`

`on_incr_tag(key, delta)`

Do something when a tag value is incremented on the observed span.

Return type `None`

`on_log(name, payload)`

Do something when a log entry is added to the span.

Return type `None`

`on_set_tag(key, value)`

Do something when a tag is set on the observed span.

Return type `None`

`on_start()`

Do something when the observed span is started.

Return type `None`

```

class baseplate.SpanObserver
    Interface for an observer that watches a span.

    on_start ()
        Do something when the observed span is started.

        Return type None

    on_set_tag (key, value)
        Do something when a tag is set on the observed span.

        Return type None

    on_incr_tag (key, delta)
        Do something when a tag value is incremented on the observed span.

        Return type None

    on_log (name, payload)
        Do something when a log entry is added to the span.

        Return type None

    on_finish (exc_info)
        Do something when the observed span is finished.

        Parameters exc_info (Optional[Tuple[Optional[Type[BaseException]],
            Optional[BaseException], Optional[traceback]]]) – If the span ended because
            of an exception, the exception info. Otherwise, None.

        Return type None

    on_child_span_created (span)
        Do something when a child span is created.

        SpanObserver objects call this when a new child span is created.

        Parameters span (Span) – The new child span.

        Return type None

```

Legacy Methods

```

Baseplate.configure_logging ()
    Add request context to the logging system.

    Deprecated since version 1.0: Use configure_observers () instead.

    Return type None

Baseplate.configure_metrics (metrics_client)
    Send timing metrics to the given client.

    This also adds a baseplate.lib.metrics.Batch object to the metrics attribute on the
    RequestContext where you can add your own application-specific metrics. The batch is automatically
    flushed at the end of the request.

    Deprecated since version 1.0: Use configure_observers () instead.

    Parameters metrics_client (Client) – Metrics client to send request metrics to.

    Return type None

```

`Baseplate.configure_tracing(tracing_client)`

Collect and send span information for request tracing.

When configured, this will send tracing information automatically collected by Baseplate to the configured distributed tracing service.

Deprecated since version 1.0: Use `configure_observers()` instead.

Parameters `tracing_client` (`baseplate.observers.tracing.TracingClient`) – Tracing client to send request traces to.

Return type `None`

`Baseplate.configure_error_reporting(client)`

Send reports for unexpected exceptions to the given client.

This also adds a `raven.Client` object to the `sentry` attribute on the `RequestContext` where you can send your own application-specific events.

Deprecated since version 1.0: Use `configure_observers()` instead.

Parameters `client` (`raven.Client`) – A configured raven client.

Return type `None`

`baseplate.clients`

Helpers that integrate common client libraries with Baseplate.py.

This package contains modules which integrate various client libraries with Baseplate.py's instrumentation. When using these client library integrations, trace information is passed on and metrics are collected automatically.

Instrumented Client Libraries

`baseplate.clients.cassandra`

`Cassandra` is a database designed for high-availability, high write throughput, and eventual consistency.

Baseplate.py supports both the base `Python Cassandra driver` and the `Cassandra ORM`, `CQLMapper`.

Example

To integrate the `Cassandra driver` with your application, add the appropriate client declaration to your context configuration:

```
baseplate.configure_context(
    app_config,
    {
        ...
        "foo": CassandraClient("mykeyspace"),
        ...
    }
)
```

configure it in your application's configuration file:

```
[app:main]

...

# required: a comma-delimited list of hosts to contact to find the ring
foo.contact_points = cassandra-01.local, cassandra-02.local

# optional: the port to connect to on each cassandra server
# (default: 9042)
foo.port = 9999

# optional: the name of a CredentialSecret holding credentials for
# authenticating to cassandra
foo.credential_secret = secret/my_service/cassandra-foo

...
```

and then use the attached `Session`-like object in request:

```
def my_method(request):
    request.foo.execute("SELECT 1;")
```

Configuration

class `baseplate.clients.cassandra.CassandraClient` (*keyspace*, ***kwargs*)
Configure a Cassandra client.

This is meant to be used with `baseplate.Baseplate.configure_context()`.

See `cluster_from_config()` for available configuration settings.

Parameters **keyspace** (*str*) – Which keyspace to set as the default for operations.

class `baseplate.clients.cassandra.CQLMapperClient` (*keyspace*, ***kwargs*)
Configure a CQLMapper client.

This is meant to be used with `baseplate.Baseplate.configure_context()`.

See `cluster_from_config()` for available configuration settings.

Parameters **keyspace** (*str*) – Which keyspace to set as the default for operations.

`baseplate.clients.cassandra.cluster_from_config` (*app_config*, *secrets=None*,
prefix='cassandra.', *execution_profiles=None*, ***kwargs*)

Make a Cluster from a configuration dictionary.

The keys useful to `cluster_from_config()` should be prefixed, e.g. `cassandra.contact_points` etc. The `prefix` argument specifies the prefix used to filter keys. Each key is mapped to a corresponding keyword argument on the `Cluster` constructor. Any keyword arguments given to this function will be passed through to the `Cluster` constructor. Keyword arguments take precedence over the configuration file.

Supported keys:

- `contact_points` (required): comma delimited list of contact points to try connecting for cluster discovery
- `port`: The server-side port to open connections to.

- **credentials_secret (optional):** the key used to retrieve the database credentials from secrets as a *CredentialSecret*.

Parameters `execution_profiles` (`Optional[Dict[str, ExecutionProfile]]`) – Configured execution profiles to provide to the rest of the application.

Return type `Cluster`

Classes

class `baseplate.clients.cassandra.CassandraContextFactory` (`session`)
Cassandra session context factory.

This factory will attach a proxy object which acts like a `cassandra.cluster.Session` to an attribute on the *RequestContext*. The `execute()`, `execute_async()` and `prepare()` methods will automatically record diagnostic information.

Parameters `session` (`cassandra.cluster.Session`) – A configured session object.

class `baseplate.clients.cassandra.CQLMapperContextFactory` (`session`)
CQLMapper ORM connection context factory.

This factory will attach a new CQLMapper `cqlmapper.connection.Connection` to an attribute on the *RequestContext*. This `Connection` object will use the same proxy object that `CassandraContextFactory` attaches to a context to run queries so the `execute` command will automatically record diagnostic information.

Parameters `session` (`cassandra.cluster.Session`) – A configured session object.

`baseplate.clients.hvac`

Vault is a high-security store for secret tokens, credentials, and other sensitive information. *HVAC* is a Python client library for Vault.

Note: The *SecretsStore* handles the most common use case of Vault in a Baseplate application: secure retrieval of secret tokens. This client is only necessary when taking advantage of more advanced features of Vault such as the *Transit backend* or *Cubbyholes*. If these don't sound familiar, check out the secrets store before digging in here.

Example

To integrate HVAC with your application, add the appropriate client declaration to your context configuration:

```
baseplate.configure_context(  
    app_config,  
    {  
        ...  
        "foo": HvacClient(),  
        ...  
    }  
)
```

configure it in your application's configuration file:

```
[app:main]

...

# optional: how long to wait for calls to vault
foo.timeout = 300 milliseconds

...
```

and then use it in request:

```
def my_method(request):
    request.foo.is_initialized()
```

See [HVAC's README](#) for documentation on the methods available from its client.

Configuration

class `baseplate.clients.hvac.HvacClient` (*secrets*)

Configure an HVAC client.

This is meant to be used with `baseplate.Baseplate.configure_context()`.

See `hvac_factory_from_config()` for available configuration settings.

Parameters `secrets` (*SecretsStore*) – The configured secrets store for this application.

`baseplate.clients.hvac.hvac_factory_from_config` (*app_config*, *secrets_store*, *prefix*=*'vault.'*)

Make an HVAC client factory from a configuration dictionary.

The keys useful to `hvac_factory_from_config()` should be prefixed, e.g. `vault.timeout`. The `prefix` argument specifies the prefix used to filter keys.

Supported keys:

- **timeout:** How long to wait for calls to Vault. (*Timespan()*)

Parameters

- **app_config** (*Dict[str, str]*) – The raw application configuration.
- **secrets_store** (*SecretsStore*) – A configured secrets store from which we can get a Vault authentication token.
- **prefix** (*str*) – The prefix for configuration keys.

Return type *HvacContextFactory*

Classes

class `baseplate.clients.hvac.HvacContextFactory` (*secrets_store*, *timeout*)

HVAC client context factory.

This factory will attach a proxy object which acts like an `hvac.Client` to an attribute on the *RequestContext*. All methods that talk to Vault will be automatically instrumented for tracing and diagnostic metrics.

Parameters

- **secrets_store** (`baseplate.lib.secrets.SecretsStore`) – Configured secrets store from which we can get a Vault authentication token.
- **timeout** (`datetime.timedelta`) – How long to wait for calls to Vault.

make_object_for_context (*name*, *span*)

Return an object that can be added to the context object.

Parameters

- **name** (*str*) – The name assigned to this object on the context.
- **span** (*Span*) – The current span this object is being made for.

Return type `InstrumentedHvacClient`

`baseplate.clients.kombu`

This integration adds support for sending messages to queue brokers (like [RabbitMQ](#)) via [Kombu](#). If you are looking to consume messages, check out the `baseplate.frameworks.queue_consumer` framework integration instead.

Example

To integrate it with your application, add the appropriate client declaration to your context configuration:

```
baseplate.configure_context(  
    app_config,  
    {  
        ...  
        "foo": KombuProducer(),  
        ...  
    }  
)
```

configure it in your application's configuration file:

```
[app:main]  
  
...  
  
# required: where to find the queue broker  
foo.hostname = rabbit.local  
  
# optional: the rabbitmq virtual host to use  
foo.virtual_host = /  
  
# required: which type of exchange to use  
foo.exchange_type = topic  
  
# optional: the name of the exchange to use (default is no name)  
foo.exchange_name = bar  
  
...
```

and then use the attached `Producer`-like object in request:

```
def my_method(request):
    request.foo.publish("boo!", routing_key="route_me")
```

Serialization

This integration also supports adding custom serializers to `Kombu` via the `baseplate.clients.kombu.KombuSerializer` interface and the `baseplate.clients.kombu.register_serializer` function. This serializer can be passed to the `baseplate.clients.kombu.KombuProducerContextFactory` for use by the `baseplate.clients.kombu.KombuProducer` to allow for automatic serialization when publishing.

In order to use a custom serializer, you must first register it with Kombu using the provided `baseplate.clients.kombu.register_serializer` function.

In-addition to the base interface, we also provide a serializer for Thrift objects: `baseplate.clients.kombu.KombuThriftSerializer`.

Example

```
serializer = KombuThriftSerializer[ThriftStruct](ThriftStruct)
register_serializer(serializer)
```

Interface

class `baseplate.clients.kombu.KombuSerializer`

Interface for wrapping non-built-in serializers for Kombu.

`baseplate.clients.kombu.register_serializer(serializer)`

Register *serializer* with the Kombu serialization registry.

The serializer will be registered using *serializer.name* and will be sent to the message broker with the header “application/x-*{serializer.name}*”. You need to call *register_serializer* before you can use a serializer for automatic serialization when publishing and deserializing when consuming.

Return type `None`

Serializers

class `baseplate.clients.kombu.KombuThriftSerializer` (*thrift_class*, *protocol_factory*=`<thrift.protocol.TBinaryProtocol.TBinaryProtocol object>`)

Thrift object serializer for Kombu.

Configuration

class `baseplate.clients.kombu.KombuProducer` (*max_connections*=`None`, *serializer*=`None`)

Configure a Kombu producer.

This is meant to be used with `baseplate.Baseplate.configure_context()`.

See `connection_from_config()` and `exchange_from_config()` for available configuration settings.

Parameters `max_connections` (`Optional[int]`) – The maximum number of connections.

```
baseplate.clients.kombu.connection_from_config(app_config,    prefix,    secrets=None,
                                                **kwargs)
```

Make a Connection from a configuration dictionary.

The keys useful to `connection_from_config()` should be prefixed, e.g. `amqp.hostname` etc. The `prefix` argument specifies the prefix used to filter keys. Each key is mapped to a corresponding keyword argument on the `Connection` constructor. Any keyword arguments given to this function will be passed through to the `Connection` constructor. Keyword arguments take precedence over the configuration file.

Supported keys:

- credentials_secret
- hostname
- virtual_host

Return type `Connection`

```
baseplate.clients.kombu.exchange_from_config(app_config, prefix, **kwargs)
```

Make an Exchange from a configuration dictionary.

The keys useful to `exchange_from_config()` should be prefixed, e.g. `amqp.exchange_name` etc. The `prefix` argument specifies the prefix used to filter keys. Each key is mapped to a corresponding keyword argument on the `Exchange` constructor. Any keyword arguments given to this function will be passed through to the `Exchange` constructor. Keyword arguments take precedence over the configuration file.

Supported keys:

- exchange_name
- exchange_type

Return type Exchange

Classes

```
class baseplate.clients.kombu.KombuProducerContextFactory(connection, exchange,  
max_connections=None,  
serializer=None)
```

KombuProducer context factory.

This factory will attach a proxy object which acts like a `kombu.Producer` to an attribute on the `RequestContext`. The `publish()` method will automatically record diagnostic information.

Parameters

- **connection** (`Connection`) – A configured connection object.
- **exchange** (`Exchange`) – A configured exchange object
- **max_connections** (`Optional[int]`) – The maximum number of connections.

baseplate.clients.memcache

Memcached is a high-performance in-memory key value store frequently used for caching. **Pymemcache** is a Python client library for it.

Example

To integrate pymemcache with your application, add the appropriate client declaration to your context configuration:

```
baseplate.configure_context(
    app_config,
    {
        ...
        "foo": MemcacheClient(),
        ...
    }
)
```

configure it in your application's configuration file:

```
[app:main]

...

# required: the host and port to connect to
foo.endpoint = localhost:11211

# optional: the maximum size of the connection pool (default 2147483648)
foo.max_pool_size = 99

# optional: how long to wait for connections to establish
foo.connect_timeout = .5 seconds

# optional: how long to wait for a memcached command
foo.timeout = 100 milliseconds

...
```

and then use the attached `PooledClient`-like object in request:

```
def my_method(request):
    request.foo.incr("bar")
```

Configuration

class baseplate.clients.memcache.**MemcacheClient** (*serializer=None, deserializer=None*)

Configure a Memcached client.

This is meant to be used with `baseplate.Baseplate.configure_context()`.

See `pool_from_config()` for available configuration settings.

Parameters

- **serializer** (Optional[Callable[[str, Any], Tuple[bytes, int]]]) – function to serialize values to strings suitable for being stored in memcached. An example is `make_dump_and_compress_fn()`.
- **deserializer** (Optional[Callable[[str, bytes, int], Any]]) – function to convert strings returned from memcached to arbitrary objects, must be compatible with `serializer`. An example is `decompress_and_load()`.

`baseplate.clients.memcache.pool_from_config(app_config, prefix='memcache.', serializer=None, deserializer=None)`

Make a PooledClient from a configuration dictionary.

The keys useful to `pool_from_config()` should be prefixed, e.g. `memcache.endpoint`, `memcache.max_pool_size`, etc. The `prefix` argument specifies the prefix used to filter keys. Each key is mapped to a corresponding keyword argument on the `PooledClient` constructor.

Supported keys:

- **endpoint (required):** a string representing a host and port to connect to memcached service, e.g. `localhost:11211` or `127.0.0.1:11211`.
- **max_pool_size:** an integer for the maximum pool size to use, by default this is 2147483648.
- **connect_timeout:** how long (as `Timespan()`) to wait for a connection to memcached server. Defaults to the underlying socket default timeout.
- **timeout:** how long (as `Timespan()`) to wait for calls on the socket connected to memcache. Defaults to the underlying socket default timeout.

Parameters

- **app_config** (Dict[str, str]) – the raw application configuration
- **prefix** (str) – prefix for configuration keys
- **serializer** (Optional[Callable[[str, Any], Tuple[bytes, int]]]) – function to serialize values to strings suitable for being stored in memcached. An example is `make_dump_and_compress_fn()`.
- **deserializer** (Optional[Callable[[str, bytes, int], Any]]) – function to convert strings returned from memcached to arbitrary objects, must be compatible with `serializer`. An example is `decompress_and_load()`.

Return type `PooledClient`

Returns `pymemcache.client.base.PooledClient`

Classes

class `baseplate.clients.memcache.MemcacheContextFactory` (*pooled_client*)

Memcache client context factory.

This factory will attach a `MonitoredMemcacheConnection` to an attribute on the `RequestContext`. When memcache commands are executed via this connection object, they will use connections from the provided `PooledClient` and automatically record diagnostic information.

Parameters `pooled_client` (`PooledClient`) – A pooled client.

make_object_for_context (*name, span*)

Return an object that can be added to the context object.

Parameters

- **name** (*str*) – The name assigned to this object on the context.
- **span** (*Span*) – The current span this object is being made for.

Return type *MonitoredMemcacheConnection*

```
class baseplate.clients.memcache.MonitoredMemcacheConnection (context_name,  
                                                             server_span,  
                                                             pooled_client)
```

Memcache connection that collects diagnostic information.

This connection acts like a *PooledClient* except that operations are wrapped with diagnostic collection. Some methods may not yet be wrapped with monitoring. Please request assistance if any needed methods are not being monitored.

Serialization/deserialization helpers

Memcache serialization/deserialization (and compression) helper methods.

Memcached can only store strings, so to store arbitrary objects we need to serialize them to strings and be able to deserialize them back to their original form.

New services should use `dump_and_compress()` and `decompress_and_load()`.

Services that need to read and write to the same memcache instances as r2 should use `pickle_and_compress()` and `decompress_and_unpickle()`.

```
baseplate.clients.memcache.lib.decompress_and_load (key, serialized, flags)
```

Deserialize data.

This should be paired with `make_dump_and_compress_fn()`.

Parameters

- **key** (*str*) – the memcached key.
- **serialized** (*bytes*) – the serialized object returned from memcached.
- **flags** (*int*) – value stored and returned from memcached for the client to use to indicate how the value was serialized.

Return type *Any*

Returns The deserialized value.

```
baseplate.clients.memcache.lib.make_dump_and_compress_fn (min_compress_length=0,  
                                                         compress_level=1)
```

Make a serializer.

This should be paired with `decompress_and_load()`.

The resulting method is a chain of `json.loads()` and `zlib` compression. Values that are not JSON serializable will result in a `TypeError`.

Parameters

- **min_compress_length** (*int*) – the minimum serialized string length to enable `zlib` compression. 0 disables compression.
- **compress_level** (*int*) – `zlib` compression level. 0 disables compression and 9 is the maximum value.

Return type *Callable[[str, Any], Tuple[bytes, int]]*

Returns The serializer.

`baseplate.clients.memcache.lib.decompress_and_unpickle` (*key*, *serialized*, *flags*)
Deserialize data stored by pylibmc.

Warning: This should only be used when sharing caches with applications using pylibmc (like r2). New applications should use the safer and future proofed `decompress_and_load()`.

Parameters

- **key** (*str*) – the memcached key.
- **serialized** (*bytes*) – the serialized object returned from memcached.
- **flags** (*int*) – value stored and returned from memcached for the client to use to indicate how the value was serialized.

Return type *Any*

Returns the deserialized value.

`baseplate.clients.memcache.lib.make_pickle_and_compress_fn` (*min_compress_length=0*,
compress_level=1)

Make a serializer compatible with pylibmc readers.

The resulting method is a chain of `pickle.dumps()` and `zlib` compression. This should be paired with `decompress_and_unpickle()`.

Warning: This should only be used when sharing caches with applications using pylibmc (like r2). New applications should use the safer and future proofed `make_dump_and_compress_fn()`.

Parameters

- **min_compress_length** (*int*) – the minimum serialized string length to enable `zlib` compression. 0 disables compression.
- **compress_level** (*int*) – `zlib` compression level. 0 disables compression and 9 is the maximum value.

Return type `Callable[[str, Any], Tuple[bytes, int]]`

Returns the serializer method.

baseplate.clients.redis

`Redis` is an in-memory data structure store used where speed is necessary but complexity is beyond simple key-value operations. (If you're just doing caching, prefer `memcached`). `Redis-py` is a Python client library for `Redis`.

Example

To integrate `redis-py` with your application, add the appropriate client declaration to your context configuration:

```
baseplate.configure_context(  
    app_config,  
    {  
        ...
```

(continues on next page)

(continued from previous page)

```

        "foo": RedisClient(),
        ...
    }
)

```

configure it in your application's configuration file:

```

[app:main]

...

# required: what redis instance to connect to
foo.url = redis://localhost:6379/0

# optional: the maximum size of the connection pool
foo.max_connections = 99

# optional: how long to wait for a connection to establish
foo.socket_connect_timeout = 3 seconds

# optional: how long to wait for a command to execute
foo.socket_timeout = 200 milliseconds

...

```

and then use the attached `Redis`-like object in request:

```

def my_method(request):
    request.foo.ping()

```

Configuration

class `baseplate.clients.redis.RedisClient` (***kwargs*)
 Configure a Redis client.

This is meant to be used with `baseplate.Baseplate.configure_context()`.

See `pool_from_config()` for available configuration settings.

`baseplate.clients.redis.pool_from_config` (*app_config*, *prefix='redis.'*, ***kwargs*)
 Make a `ConnectionPool` from a configuration dictionary.

The keys useful to `pool_from_config()` should be prefixed, e.g. `redis.url`, `redis.max_connections`, etc. The `prefix` argument specifies the prefix used to filter keys. Each key is mapped to a corresponding keyword argument on the `redis.ConnectionPool` constructor.

Supported keys:

- `url` (required): a URL like `redis://localhost/0`.
- `max_connections`: an integer maximum number of connections in the pool
- **`socket_connect_timeout`: how long to wait for sockets to connect. e.g. 200 milliseconds** (`Timespan()`)
- **`socket_timeout`: how long to wait for socket operations, e.g. 200 milliseconds** (`Timespan()`)

Return type `ConnectionPool`

Classes

class `baseplate.clients.redis.RedisContextFactory` (*connection_pool*)

Redis client context factory.

This factory will attach a `MonitoredRedisConnection` to an attribute on the `RequestContext`. When Redis commands are executed via this connection object, they will use connections from the provided `redis.ConnectionPool` and automatically record diagnostic information.

Parameters `connection_pool` (`ConnectionPool`) – A connection pool.

make_object_for_context (*name*, *span*)

Return an object that can be added to the context object.

Parameters

- **name** (`str`) – The name assigned to this object on the context.
- **span** (`Span`) – The current span this object is being made for.

Return type `MonitoredRedisConnection`

class `baseplate.clients.redis.MonitoredRedisConnection` (*context_name*, *server_span*,
connection_pool)

Redis connection that collects diagnostic information.

This connection acts like `redis.StrictRedis` except that all operations are automatically wrapped with diagnostic collection.

The interface is the same as that class except for the `pipeline()` method.

Note: Locks and PubSub are currently unsupported.

execute_command (**args*, ***kwargs*)

Execute a command and return a parsed response

Return type `Any`

pipeline (*name*, *transaction=True*, *shard_hint=None*)

Create a pipeline.

This returns an object on which you can call the standard Redis commands. Execution will be deferred until `execute` is called. This is useful for saving round trips.

Parameters

- **name** (`str`) – The name to attach to diagnostics for this pipeline.
- **transaction** (`bool`) – Whether or not the commands in the pipeline are wrapped with a transaction and executed atomically.

Return type `MonitoredRedisPipeline`

transaction (**args*, ***kwargs*)

Not currently implemented.

Return type `Any`

lock (**args*, ***kwargs*)

Not currently implemented.

Return type `Any`

pubsub (**args, **kwargs*)
Not currently implemented.

Return type `Any`

class `baseplate.clients.redis.MessageQueue` (*name, client*)

A Redis-backed variant of `MessageQueue`.

Parameters

- **name** (`str`) – can be any string.
- **client** (`ConnectionPool`) – should be a `redis.ConnectionPool` or `redis.BlockingConnectionPool` from which a client connection can be created from (preferably generated from the `pool_from_config()` helper).

get (*timeout=None*)
Read a message from the queue.

Parameters **timeout** (`Optional[float]`) – If the queue is empty, the call will block up to `timeout` seconds or forever if `None`, if a float is given, it will be rounded up to be an integer

Raises `TimedOutError` The queue was empty for the allowed duration of the call.

Return type `bytes`

put (*message, timeout=None*)
Add a message to the queue.

Parameters **message** (`bytes`) – will be typecast to a string upon storage and will come out of the queue as a string regardless of what type they are when passed into this method.

Return type `None`

unlink ()
Not implemented for Redis variant.

Return type `None`

close ()
Close queue when finished.

Will delete the queue from the Redis server (Note, can still enqueue and dequeue as the actions will recreate the queue)

Return type `None`

`baseplate.clients.sqlalchemy`

`SQLAlchemy` is an ORM and general-purpose SQL engine for Python. It can work with many different SQL database backends. Reddit generally uses it to talk to `PostgreSQL`.

Example

To integrate `SQLAlchemy` with your application, add the appropriate client declaration to your context configuration:

```
baseplate.configure_context (
    app_config,
    {
        ...
        "foo": SQLAlchemySession(),
        ...
    }
)
```

configure it in your application's configuration file:

```
[app:main]

...

# required: sqlalchemy URL describing a database to connect to
foo.url = postgresql://postgres.local:6543/bar

# optional: the name of a CredentialSecret holding credentials for
# authenticating to the database
foo.credentials_secret = secret/my_service/db-foo

...
```

and then use the attached `Session` object in request:

```
def my_method(request):
    request.foo.query(MyModel).filter_by(...).all()
```

Configuration

class `baseplate.clients.sqlalchemy.SQLAlchemySession` (*secrets=None, **kwargs*)
Configure a SQLAlchemy Session.

This is meant to be used with `baseplate.Baseplate.configure_context()`.

See `engine_from_config()` for available configuration settings.

Parameters `secrets` (Optional[`SecretsStore`]) – Required if configured to use credentials to talk to the database.

`baseplate.clients.sqlalchemy.engine_from_config` (*app_config, secrets=None, prefix='database.', **kwargs*)

Make an `Engine` from a configuration dictionary.

The keys useful to `engine_from_config()` should be prefixed, e.g. `database.url`, etc. The `prefix` argument specifies the prefix used to filter keys.

Supported keys:

- **url**: the connection URL to the database, passed to `make_url()` to create the URL used to connect to the database.
- **credentials_secret** (optional): the key used to retrieve the database credentials from `secrets` as a `CredentialSecret`. If this is supplied, any credentials given in `url` will be replaced by these.
- **pool_recycle** (optional): this setting causes the pool to recycle connections after the given number of seconds has passed. It defaults to -1, or no timeout.

- **pool_pre_ping (optional):** when set to true, this setting causes sqlalchemy to perform a liveness-check query each time a connection is checked out of the pool. If the liveness-check fails, the connection is gracefully recycled. This ensures severed connections are handled more gracefully, at the cost of doing a *SELECT 1* at the start of each checkout. When used, this obviates most of the reasons you might use pool_recycle, and as such they shouldn't normally be used simultaneously. Requires SQLAlchemy 1.3.

Return type Engine

Classes

class baseplate.clients.sqlalchemy.**SQLAlchemyEngineContextFactory** (engine)
SQLAlchemy core engine context factory.

This factory will attach a SQLAlchemy `sqlalchemy.engine.Engine` to an attribute on the `RequestContext`. All cursor (query) execution will automatically record diagnostic information.

Additionally, the trace and span ID will be added as a comment to the text of the SQL statement. This is to aid correlation of queries with requests.

See also:

The engine is the low-level SQLAlchemy API. If you want to use the ORM, consider using `SQLAlchemySessionContextFactory` instead.

Parameters engine (Engine) – A configured SQLAlchemy engine.

class baseplate.clients.sqlalchemy.**SQLAlchemySessionContextFactory** (engine)
SQLAlchemy ORM session context factory.

This factory will attach a new SQLAlchemy `sqlalchemy.orm.session.Session` to an attribute on the `RequestContext`. All cursor (query) execution will automatically record diagnostic information.

The session will be automatically closed, but not committed or rolled back, at the end of each request.

See also:

The session is part of the high-level SQLAlchemy ORM API. If you want to do raw queries, consider using `SQLAlchemyEngineContextFactory` instead.

Parameters engine (Engine) – A configured SQLAlchemy engine.

Runtime Metrics

In addition to request-level metrics reported through spans, this wrapper reports connection pool statistics periodically via the *Process-level metrics* system. All metrics are prefixed as follows:

```
{namespace}.runtime.{hostname}.PID{pid}.clients.{name}
```

where `namespace` is the application's namespace, `hostname` and `pid` come from the operating system, and `name` is the name given to `add_to_context()` when registering this context factory.

The following metrics are reported:

pool.size The size limit for the connection pool.

pool.open_and_available How many connections have been established but are sitting available for use in the connection pool.

pool.in_use How many connections have been established and are currently checked out and being used.

pool.overflow How many connections beyond the pool size are currently being used. See `sqlalchemy.pool.QueuePool` for more information.

`baseplate.clients.thrift`

Thrift is a cross-language framework for cross-service communication. Developers write a language-independent definition of a service's API (the "IDL") and Thrift's code generator makes server and client libraries for that API.

Example

To add a Thrift client to your application, ensure that your service has Baseplate.py's Thrift build step integrated in its root `setup.py`:

```
from baseplate.frameworks.thrift.command import ThriftBuildPyCommand

...

setup(
    ...

    cmdclass={
        "build_py": ThriftBuildPyCommand,
    },
)
```

Then add the downstream service's IDL to your application directory:

```
$ cp ~/src/other/other/other.thrift myservice/
```

Baseplate.py will automatically run the Thrift compiler on this when building your application and put the output into a Python package within your application like `yourservice.other_thrift`. Import the client class and add it to your context object:

```
from .other_thrift import OtherService

...

def make_wsgi_app(app_config):
    ...

    baseplate.configure_context(
        app_config,
        {
            ...
            "foo": ThriftClient(OtherService.Client),
            ...
        }
    )
```

then configure it in your application's configuration file:

```
[app:main]
```

(continues on next page)

(continued from previous page)

```

...

# required: the host:port to find the service at
foo.endpoint = localhost:9999

# optional: the size of the connection pool (default 10)
foo.size = 10

# optional: how long a connection can be alive before we
# recycle it (default 1 minute)
foo.max_age = 1 minute

# optional: how long before we time out when connecting
# or doing an RPC (default 1 second)
foo.timeout = 1 second

# optional: how many times we'll retry connecting (default 3)
foo.max_retries = 3

...

```

and finally use the attached client in request:

```

def my_method(request):
    request.foo.is_healthy()

```

Classes

class `baseplate.clients.thrift.ThriftClient` (*client_cls*, ***kwargs*)

Configure a Thrift client.

This is meant to be used with `baseplate.Baseplate.configure_context()`.

See `baseplate.lib.thrift_pool.thrift_pool_from_config()` for available configuration settings.

Parameters *client_cls* (*Any*) – The class object of a Thrift-generated client class, e.g. `YourService.Client`.

class `baseplate.clients.thrift.ThriftContextFactory` (*pool*, *client_cls*)

Thrift client pool context factory.

This factory will attach a proxy object with the same interface as your thrift client to an attribute on the `RequestContext`. When a thrift method is called on this proxy object, it will check out a connection from the connection pool and execute the RPC, automatically recording diagnostic information.

Parameters

- **pool** (`ThriftConnectionPool`) – The connection pool.
- **client_cls** (*Any*) – The class object of a Thrift-generated client class, e.g. `YourService.Client`.

The proxy object has a `retrying` method which takes the same parameters as `RetryPolicy.new` and acts as a context manager. The context manager returns another proxy object where Thrift service method calls will be automatically retried with the specified retry policy when transient errors occur:

```
with context.my_service.retrying(attempts=3) as svc:
    svc.some_method()
```

Runtime Metrics

In addition to request-level metrics reported through spans, this wrapper reports connection pool statistics periodically via the *Process-level metrics* system. All metrics are prefixed as follows:

```
{namespace}.runtime.{hostname}.PID{pid}.clients.{name}
```

where `namespace` is the application's namespace, `hostname` and `pid` come from the operating system, and `name` is the name given to `add_to_context()` when registering this context factory.

The following metrics are reported:

pool.size The size limit for the connection pool.

pool.in_use How many connections have been established and are currently checked out and being used.

DIY: The Factory

If a library you want is not supported here, it can be added to your own application by implementing *ContextFactory*.

class `baseplate.clients.ContextFactory`

An interface for adding stuff to the context object.

Objects implementing this interface can be passed to `add_to_context()`. The return value of `make_object_for_context()` will be added to the *RequestContext* with the name specified in `add_to_context`.

report_runtime_metrics (*batch*)

Report runtime metrics to the stats system.

Parameters *batch* (`baseplate.lib.metrics.Client`) – A metrics client to report statistics to.

Return type `None`

make_object_for_context (*name*, *span*)

Return an object that can be added to the context object.

Parameters

- **name** (`str`) – The name assigned to this object on the context.
- **span** (`baseplate.Span`) – The current span this object is being made for.

Return type `Any`

To integrate with `configure_context()` for maximum convenience, make a parser that implements `baseplate.lib.config.Parser` and returns your *ContextFactory*.

```
class MyClient(config.Parser):
    def parse(
        self, key_path: str, raw_config: config.RawConfig
    ) -> "MyContextFactory":
        parser = config.SpecParser(
```

(continues on next page)

(continued from previous page)

```

    {
        "foo": config.Integer(),
        "bar": config.Boolean(),
    }
)
result = parser.parse(key_path, raw_config)
return MyContextFactory(foo=result.foo, bar=result.bar)

```

baseplate.frameworks

Helpers for integration with various application frameworks.

This package contains modules which integrate Baseplate with common application frameworks.

Baseplate.py provides integrations with common Python application frameworks. These integrations automatically manage the *ServerSpan* lifecycle for each unit of work the framework processes (requests or messages).

baseplate.frameworks.thrift

Thrift is a cross-language framework for cross-service communication. Developers write a language-independent definition of a service's API (the "IDL") and Thrift's code generator makes server and client libraries for that API.

This module provides a wrapper for a `TProcessor` which integrates Baseplate's facilities into the Thrift request lifecycle.

An abbreviated example of it in use:

```

logger = logging.getLogger(__name__)

def make_processor(app_config):
    baseplate = Baseplate(app_config)

    handler = MyHandler()
    processor = my_thrift.MyService.Processor(handler)
    return baseplateify_processor(processor, logger, baseplate)

```

`baseplate.frameworks.thrift.baseplateify_processor` (*processor*, *logger*, *baseplate*, *edge_context_factory=None*)

Wrap a Thrift Processor with Baseplate's span lifecycle.

Parameters

- **processor** (`TProcessor`) – The service's processor to wrap.
- **logger** (`Logger`) – The logger to use for error and debug logging.
- **baseplate** (`Baseplate`) – The baseplate instance for your application.
- **edge_context_factory** (`Optional[EdgeRequestContextFactory]`) – A configured factory for handling edge request context.

Return type `TProcessor`

baseplate.frameworks.pyramid

Pyramid is a mature web framework for Python that we build HTTP services with.

This module provides a configuration extension for Pyramid which integrates Baseplate's facilities into the Pyramid WSGI request lifecycle.

An abbreviated example of it in use:

```
def make_app(app_config):
    configurator = Configurator()

    baseplate = Baseplate(app_config)
    baseplate_config = BaseplateConfigurator(
        baseplate,
        trust_trace_headers=True,
    )
    configurator.include(baseplate_config.includeme)

    return configurator.make_wsgi_app()
```

Warning: Because of how Baseplate instruments Pyramid, you should not make an `exception view` prevent Baseplate from seeing the unhandled error and reporting it appropriately.

```
class baseplate.frameworks.pyramid.BaseplateConfigurator(baseplate,
                                                           trust_trace_headers=None,
                                                           edge_context_factory=None,
                                                           header_trust_handler=None)
```

Configuration extension to integrate Baseplate into Pyramid.

Parameters

- **baseplate** (*Baseplate*) – The Baseplate instance for your application.
- **edge_context_factory** (*Optional[EdgeRequestContextFactory]*) – A configured factory for handling edge request context.
- **header_trust_handler** (*Optional[HeaderTrustHandler]*) – An object which will be used to verify whether baseplate should parse the request context headers, for example trace ids. See *StaticTrustHandler* for the default implementation.

```
class baseplate.frameworks.pyramid.HeaderTrustHandler
```

Abstract class used by *BaseplateConfigurator* to validate headers.

See *StaticTrustHandler* for the default implementation.

```
should_trust_trace_headers(request)
```

Return whether baseplate should parse the trace headers from the inbound request.

Parameters **request** (*Request*) – The request

Return type *bool*

Returns Whether baseplate should parse the trace headers from the inbound request.

```
should_trust_edge_context_payload(request)
```

Return whether baseplate should trust the edge context headers from the inbound request.

Parameters **request** (*Request*) – The request

Return type *bool*

Returns Whether baseplate should trust the inbound edge context headers

class `baseplate.frameworks.pyramid.StaticTrustHandler` (*trust_headers=False*)
 Default implementation for handling headers.

This class is created automatically by BaseplateConfigurator unless you supply your own HeaderTrustHandler

Parameters `trust_headers` (`bool`) – Whether or not to trust trace and edge context headers from inbound requests. This value will be returned by `should_trust_trace_headers` and `should_trust_edge_context_payload`.

Warning: Do not set `trust_headers` to `True` unless you are sure your application is only accessible by trusted sources (usually backend-only services).

should_trust_trace_headers (*request*)

Return whether baseplate should parse the trace headers from the inbound request.

Parameters `request` (`Request`) – The request

Return type `bool`

Returns Whether baseplate should parse the trace headers from the inbound request.

should_trust_edge_context_payload (*request*)

Return whether baseplate should trust the edge context headers from the inbound request.

Parameters `request` (`Request`) – The request

Return type `bool`

Returns Whether baseplate should trust the inbound edge context headers

Events

Within its Pyramid integration, Baseplate will emit events at various stages of the request lifecycle that services can hook into.

class `baseplate.frameworks.pyramid.ServerSpanInitialized` (*request*)

Event that Baseplate fires after creating the `ServerSpan` for a Request.

This event will be emitted before the Request is passed along to its handler. Baseplate initializes the `ServerSpan` in response to a `pyramid.events.ContextFound` event emitted by Pyramid so while we can guarantee what Baseplate has done when this event is emitted, we cannot guarantee that any other subscribers to `pyramid.events.ContextFound` have been called or not.

`baseplate.frameworks.queue_consumer`

`baseplate.frameworks.queue_consumer.kafka`

This module provides a `QueueConsumerFactory` that allows you to run a `QueueConsumerServer` that integrates Baseplate's facilities with Kafka.

An abbreviated example of it in use:

```
import confluent_kafka
from baseplate import RequestContext
from typing import Any
from typing import Dict

def process_links(
    context: RequestContext,
    data: Dict[str, Any],
    message: confluent_kafka.Message,
):
    print(f"processing {data}")

def make_consumer_factory(app_config):
    baseplate = Baseplate(app_config)
    return InOrderConsumerFactory.new(
        name="kafka_consumer.link_consumer_v0",
        baseplate=baseplate,
        bootstrap_servers="127.0.0.1:9092",
        group_id="service.link_consumer",
        topics=["new_links", "edited_links"],
        handler_fn=process_links,
    )
```

This will create a Kafka consumer group named 'service.link_consumer' that consumes from the topics 'new_links' and 'edited_links'. Messages read from those topics will be fed to process_links.

Factory

```
class baseplate.frameworks.queue_consumer.kafka.InOrderConsumerFactory(name,
                                                                    base-
                                                                    plate,
                                                                    con-
                                                                    sumer,
                                                                    han-
                                                                    dler_fn,
                                                                    kafka_consume_batch_size=
                                                                    mes-
                                                                    sage_unpack_fn=<function
                                                                    loads>,
                                                                    health_check_fn=None)
```

Factory for running a QueueConsumerServer using Kafka.

The *InOrderConsumerFactory* attempts to achieve in order, exactly once message processing.

This will run a single *KafkaConsumerWorker* that reads messages from Kafka and puts them into an internal work queue. Then it will run a single *KafkaMessageHandler* that reads messages from the internal work queue, processes them with the *handler_fn*, and then commits each message's offset to Kafka.

This one-at-a-time, in-order processing ensures that when a failure happens during processing we don't commit its offset (or the offset of any later messages) and that when the server restarts it will receive the failed message and attempt to process it again. Additionally, because each message's offset is committed immediately after processing we should never process a message more than once.

For most cases where you just need a basic consumer with sensible defaults you can use *InOrderConsumerFactory.new*.

If you need more control, you can create the Consumer yourself and use the constructor directly.


```
classmethod new(name, baseplate, bootstrap_servers, group_id, topics, handler_fn,
                 kafka_consume_batch_size=1, message_unpack_fn=<function loads>,
                 health_check_fn=None)
```

Return a new `_BaseKafkaQueueConsumerFactory`.

This method will create the `Consumer` for you and is appropriate to use in most cases where you just need a basic consumer with sensible defaults.

This method will also enforce naming standards for the Kafka consumer group and the baseplate server span.

Parameters

- **name** (`str`) – A name for your consumer process. Must look like “kafka_consumer.{group_name}”
- **baseplate** (`Baseplate`) – The Baseplate set up for your consumer.
- **bootstrap_servers** (`str`) – A comma delimited string of kafka brokers.
- **group_id** (`str`) – The kafka consumer group id. Must look like “{service_name}.{group_name}” to help prevent collisions between services.
- **topics** (`Sequence[str]`) – An iterable of kafka topics to consume from.
- **handler_fn** (`Callable[[RequestContext, Any, Message], None]`) – A `baseplate.frameworks.queue_consumer.kafka.Handler` function that will process an individual message.
- **kafka_consume_batch_size** (`int`) – The number of messages the `KafkaConsumerWorker` reads from Kafka in each batch. Defaults to 1.
- **message_unpack_fn** (`Callable[[bytes], Any]`) – A function that takes one argument, the `bytes` message body and returns the message in the format the handler expects. Defaults to `json.loads`.
- **health_check_fn** (`Optional[Callable[[Dict[str, Any]], bool]]`) – A `baseplate.server.queue_consumer.HealthcheckCallback` function that can be used to customize your health check.

Return type `_BaseKafkaQueueConsumerFactory`

```
__init__(name, baseplate, consumer, handler_fn, kafka_consume_batch_size=1,
         message_unpack_fn=<function loads>, health_check_fn=None)
    _BaseKafkaQueueConsumerFactory constructor.
```

Parameters

- **name** (`str`) – A name for your consumer process. Must look like “kafka_consumer.{group_name}”
- **baseplate** (`Baseplate`) – The Baseplate set up for your consumer.
- **consumer** (`Consumer`) – An instance of `Consumer`.
- **handler_fn** (`Callable[[RequestContext, Any, Message], None]`) – A `baseplate.frameworks.queue_consumer.kafka.Handler` function that will process an individual message.
- **kafka_consume_batch_size** (`int`) – The number of messages the `KafkaConsumerWorker` reads from Kafka in each batch. Defaults to 1.
- **message_unpack_fn** (`Callable[[bytes], Any]`) – A function that takes one argument, the `bytes` message body and returns the message in the format the handler expects. Defaults to `json.loads`.

- **health_check_fn** (*Optional*[*Callable*[[*Dict*[*str*, *Any*]], *bool*]]) – A *baseplate.server.queue_consumer.HealthcheckCallback* function that can be used to customize your health check.

```
class baseplate.frameworks.queue_consumer.kafka.FastConsumerFactory (name,  
                                                                    base-  
                                                                    plate,  
                                                                    con-  
                                                                    sumer,  
                                                                    han-  
                                                                    dler_fn,  
                                                                    kafka_consume_batch_size=1,  
                                                                    mes-  
                                                                    sage_unpack_fn=<function  
                                                                    loads>,  
                                                                    health_check_fn=None)
```

Factory for running a *QueueConsumerServer* using Kafka.

The *FastConsumerFactory* prioritizes high throughput over exactly once message processing.

This will run a single *KafkaConsumerWorker* that reads messages from Kafka and puts them into an internal work queue. Then it will run multiple *KafkaMessageHandler*'s that read messages from the internal work queue, processes them with the *handler_fn*. The number of *KafkaMessageHandler* processes is controlled by the *max_concurrency* parameter in the *~baseplate.server.queue_consumer.QueueConsumerServer* configuration. Kafka partition offsets are automatically committed by the *confluent_kafka.Consumer* every 5 seconds, so any message that has been read by the *KafkaConsumerWorker* could be committed, regardless of whether it has been processed.

This server should be able to achieve very high message processing throughput due to the multiple *KafkaMessageHandler* processes and less frequent, background partition offset commits. This does come at a price though: messages may be processed out of order, not at all, or multiple times. This is appropriate when processing throughput is important and it's acceptable to skip messages or process messages more than once (maybe there is ratelimiting in the handler or somewhere downstream).

Messages processed out of order: Messages are added to the internal work queue in order, but one worker may finish processing a "later" message before another worker finishes processing an "earlier" message.

Messages never processed: If the server crashes it may not have processed some messages that have already had their offsets automatically committed. When the server restarts it won't read those messages.

Messages processed more than once: If the server crashes it may have processed some messages but not yet committed their offsets. When the server restarts it will reprocess those messages.

For most cases where you just need a basic consumer with sensible defaults you can use *FastConsumerFactory.new*.

If you need more control, you can create the *Consumer* yourself and use the constructor directly.

```
classmethod new (name, baseplate, bootstrap_servers, group_id, topics, handler_fn,  
                 kafka_consume_batch_size=1, message_unpack_fn=<function loads>,  
                 health_check_fn=None)
```

Return a new *_BaseKafkaQueueConsumerFactory*.

This method will create the *Consumer* for you and is appropriate to use in most cases where you just need a basic consumer with sensible defaults.

This method will also enforce naming standards for the Kafka consumer group and the baseplate server span.

Parameters

- **name** (*str*) – A name for your consumer process. Must look like “kafka_consumer.{group_name}”
- **baseplate** (*Baseplate*) – The Baseplate set up for your consumer.
- **bootstrap_servers** (*str*) – A comma delimited string of kafka brokers.
- **group_id** (*str*) – The kafka consumer group id. Must look like “{service_name}.{group_name}” to help prevent collisions between services.
- **topics** (*Sequence[str]*) – An iterable of kafka topics to consume from.
- **handler_fn** (*Callable[[RequestContext, Any, Message], None]*) – A *baseplate.frameworks.queue_consumer.kafka.Handler* function that will process an individual message.
- **kafka_consume_batch_size** (*int*) – The number of messages the *KafkaConsumerWorker* reads from Kafka in each batch. Defaults to 1.
- **message_unpack_fn** (*Callable[[bytes], Any]*) – A function that takes one argument, the *bytes* message body and returns the message in the format the handler expects. Defaults to *json.loads*.
- **health_check_fn** (*Optional[Callable[[Dict[str, Any]], bool]]*) – A *baseplate.server.queue_consumer.HealthcheckCallback* function that can be used to customize your health check.

Return type *_BaseKafkaQueueConsumerFactory*

__init__ (*name, baseplate, consumer, handler_fn, kafka_consume_batch_size=1, message_unpack_fn=<function loads>, health_check_fn=None*)
_BaseKafkaQueueConsumerFactory constructor.

Parameters

- **name** (*str*) – A name for your consumer process. Must look like “kafka_consumer.{group_name}”
- **baseplate** (*Baseplate*) – The Baseplate set up for your consumer.
- **consumer** (*Consumer*) – An instance of *Consumer*.
- **handler_fn** (*Callable[[RequestContext, Any, Message], None]*) – A *baseplate.frameworks.queue_consumer.kafka.Handler* function that will process an individual message.
- **kafka_consume_batch_size** (*int*) – The number of messages the *KafkaConsumerWorker* reads from Kafka in each batch. Defaults to 1.
- **message_unpack_fn** (*Callable[[bytes], Any]*) – A function that takes one argument, the *bytes* message body and returns the message in the format the handler expects. Defaults to *json.loads*.
- **health_check_fn** (*Optional[Callable[[Dict[str, Any]], bool]]*) – A *baseplate.server.queue_consumer.HealthcheckCallback* function that can be used to customize your health check.

`baseplate.frameworks.queue_consumer.kombu`

Kombu is a library for interacting with queue brokers.

This module provides a *QueueConsumerFactory* that allows you to run a *QueueConsumerServer* that integrates Baseplate’s facilities with *Kombu*.

An abbreviated example of it in use:

```
import kombu
from baseplate import RequestContext
from typing import Any

def process_links(
    context: RequestContext,
    body: Any,
    message: kombu.Message,
):
    print(f"processing {body}")

def make_consumer_factory(app_config):
    baseplate = Baseplate(app_config)
    exchange = Exchange("reddit_exchange", "direct")
    connection = Connection(
        hostname="amqp://guest:guest@reddit.local:5672",
        virtual_host="/",
    )
    queue_name = "process_links_q"
    routing_keys = ["link_created"]
    return KombuQueueConsumerFactory.new(
        baseplate=baseplate,
        exchange=exchange,
        connection=connection,
        queue_name=queue_name,
        routing_keys=routing_keys,
        handler_fn=process_links,
    )
```

This will create a queue named 'process_links_q' and bind the routing key 'link_created'. It will then register a consumer for 'process_links_q' to read messages and feed them to process_links.

Factory

```
class baseplate.frameworks.queue_consumer.kombu.KombuQueueConsumerFactory (baseplate,
                                                                              name,
                                                                              con-
                                                                              nec-
                                                                              tion,
                                                                              queues,
                                                                              han-
                                                                              dler_fn,
                                                                              er-
                                                                              ror_handler_fn=None,
                                                                              health_check_fn=None,
                                                                              se-
                                                                              ri-
                                                                              al-
                                                                              izer=None)
```

Factory for running a QueueConsumerServer using Kombu.

For simple cases where you just need a basic queue with all the default parameters for your message broker, you can use *KombuQueueConsumerFactory.new*.

If you need more control, you can create the *Queue*s yourself and use the constructor directly.

classmethod new (*baseplate*, *exchange*, *connection*, *queue_name*, *routing_keys*, *handler_fn*, *error_handler_fn=None*, *health_check_fn=None*, *serializer=None*)
 Return a new *KombuQueueConsumerFactory*.

This method will create the *Queue* s for you and is appropriate to use in simple cases where you just need a basic queue with all the default parameters for your message broker.

Parameters

- **baseplate** (*Baseplate*) – The Baseplate set up for your consumer.
- **exchange** (*Exchange*) – The *kombu.Exchange* that you will bind your *Queue* s to.
- **exchange** – The *kombu.Connection* to your message broker.
- **queue_name** (*str*) – Name for your queue.
- **routing_keys** (*Sequence[str]*) – List of routing keys that you will create *Queue* s to consume from.
- **handler_fn** (*Callable[[RequestContext, Any, Message], None]*) – A function that will process an individual message from a queue.
- **error_handler_fn** (*Optional[Callable[[RequestContext, Any, Message, Exception], None]]*) – A function that will be called when an error is thrown while executing the *handler_fn*. This function will be responsible for calling *message.ack* or *message.requeue* as it will not be automatically called by *KombuMessageHandler*’s *handle* function.
- **health_check_fn** (*Optional[Callable[[Dict[str, Any], bool]]*) – A *baseplate.server.queue_consumer.HealthcheckCallback* function that can be used to customize your health check.
- **serializer** (*Optional[KombuSerializer[~T]]*) – A *baseplate.clients.kombu.KombuSerializer* that should be used to decode the messages you are consuming.

Return type *KombuQueueConsumerFactory*

__init__ (*baseplate*, *name*, *connection*, *queues*, *handler_fn*, *error_handler_fn=None*, *health_check_fn=None*, *serializer=None*)
KombuQueueConsumerFactory constructor.

Parameters

- **baseplate** (*Baseplate*) – The Baseplate set up for your consumer.
- **exchange** – The *kombu.Exchange* that you will bind your *Queue* s to.
- **queues** (*Sequence[Queue]*) – List of *Queue* s to consume from.
- **queue_name** – Name for your queue.
- **routing_keys** – List of routing keys that you will create *Queue* s to consume from.
- **handler_fn** (*Callable[[RequestContext, Any, Message], None]*) – A function that will process an individual message from a queue.
- **error_handler_fn** (*Optional[Callable[[RequestContext, Any, Message, Exception], None]]*) – A function that will be called when an error is thrown while executing the *handler_fn*. This function will be responsible for calling *message.ack* or *message.requeue* as it will not be automatically called by *KombuMessageHandler*’s *handle* function.

- **health_check_fn** (`Optional[Callable[[Dict[str, Any]], bool]]`) – A *baseplate.server.queue_consumer.HealthcheckCallback* function that can be used to customize your health check.
- **serializer** (`Optional[KombuSerializer[~T]]`) – A *baseplate.clients.kombu.KombuSerializer* that should be used to decode the messages you are consuming.

Errors

class `baseplate.frameworks.queue_consumer.kombu.FatalMessageHandlerError`

An error that signals that the queue process should exit.

Raising an Exception that is a subclass of `FatalMessageHandlerError` will cause the `KombuMessageHandler` to re-raise the exception rather than swallowing it which will cause the handler thread/process to stop. This, in turn, will gracefully shut down the `QueueConsumerServer` currently running.

Exceptions of this nature should be reserved for errors that are due to problems with the environment rather than the message itself. For example, a node that cannot get its AWS credentials.

`baseplate.frameworks.queue_consumer.deprecated`

Deprecated since version 1.1: This way of creating a Baseplate queue consumer is deprecated in favor of using a *QueueConsumerServer* and will be removed in the next major release. Instructions for upgrading are included below.

Upgrading to a QueueConsumerServer

To start, you will be running your queue consumer as a “server” now, so it will use *baseplate-serve* rather than *baseplate-script* as the entrypoint.

```
- baseplate-script run.ini consumer:run
+ baseplate-serve run.ini --bind 0.0.0.0:8080
```

This also means that you will need to update your config file with similar sections to what you would have for an HTTP or Thrift service.

```
[DEFAULT]
rabbitmq.hostname = amqp://rabbit.local:5672
rabbitmq.exchange_name = my-exchange
rabbitmq.exchange_type = direct

+ [app:main]
+ factory = my_service:make_consumer_factory
+
+ [server:main]
+ factory = baseplate.server.queue_consumer
+ max_concurrency = 1
```

You will also need to change your code to create a `KombuQueueConsumerFactory` with a *make_consumer_factory* function rather than using *queue_consumer.consume* as you did for this.

```
from kombu import Connection, Exchange
-
- from baseplate import queue_consumer
```

(continues on next page)

(continued from previous page)

```

-
- def process_links(context, msg_body, msg):
-     print('processing %s' % msg_body)
-
- def run():
-     queue_consumer.consume(
-         baseplate=make_baseplate(cfg, app_config),
-         exchange=Exchange('reddit_exchange', 'direct'),
-         connection=Connection(
-             hostname='amqp://guest:guest@reddit.local:5672',
-             virtual_host='/',
-         ),
-         queue_name='process_links_q',
-         routing_keys=[
-             'link_created',
-             'link_deleted',
-             'link_updated',
-         ],
-         handler=process_links,
-     )
+
+ from baseplate import Baseplate
+ from baseplate.frameworks.queue_consumer.kombu import (
+     KombuQueueConsumerFactory,
+ )
+
+ def process_links(context, message):
+     body = message.decode()
+     print('processing %s' % body)
+
+ def make_consumer_factory(app_config):
+     baseplate = Baseplate(app_config)
+     exchange = Exchange('reddit_exchange', 'direct')
+     connection = Connection(
+         hostname='amqp://guest:guest@reddit.local:5672',
+         virtual_host='/',
+     )
+     queue_name = 'process_links_q'
+     routing_keys = ['link_created', 'link_deleted', 'link_updated']
+     return KombuQueueConsumerFactory.new(
+         baseplate=baseplate,
+         exchange=exchange,
+         connection=connection,
+         queue_name=queue_name,
+         routing_keys=routing_keys,
+         handler_fn=process_links,
+     )

```

Original docs

To create a long-running process to consume from a queue:

```

from kombu import Connection, Exchange
from baseplate import queue_consumer

```

(continues on next page)

(continued from previous page)

```

def process_links(context, msg_body, msg):
    print('processing %s' % msg_body)

def run():
    queue_consumer.consume(
        baseplate=make_baseplate(cfg, app_config),
        exchange=Exchange('reddit_exchange', 'direct'),
        connection=Connection(
            hostname='amqp://guest:guest@reddit.local:5672',
            virtual_host='/',
        ),
        queue_name='process_links_q',
        routing_keys=[
            'link_created',
            'link_deleted',
            'link_updated',
        ],
        handler=process_links,
    )

```

This will create a queue named 'process_links_q' and bind the routing keys 'link_created', 'link_deleted', and 'link_updated'. It will then register a consumer for 'process_links_q' to read messages and feed them to process_links.

Register and run a queue consumer

baseplate.frameworks.queue_consumer.deprecated.**consume** (*baseplate, exchange, connection, queue_name, routing_keys, handler*)

Create a long-running process to consume messages from a queue.

A queue with name `queue_name` is created and bound to the `routing_keys` so messages published to the `routing_keys` are routed to the queue.

Next, the process registers a consumer that receives messages from the queue and feeds them to the `handler`.

The handler function must take 3 arguments:

- `context`: a baseplate context
- `message_body`: the text body of the message
- `message`: `kombu.message.Message`

The consumer will automatically `ack` each message after the handler method exits. If there is an error in processing and the message must be retried the handler should raise an exception to crash the process. This will prevent the `ack` and the message will be re-queued at the head of the queue.

Parameters

- **baseplate** (*Baseplate*) – A baseplate instance for the service.
- **exchange** (*Exchange*) –
- **connection** (*Connection*) –
- **queue_name** (*str*) – The name of the queue.
- **routing_keys** (*Sequence[str]*) – List of routing keys.

- **handler** (`Callable[[RequestContext, str, Message], None]`) – The handler method.

Return type `_Specialform`

class `baseplate.frameworks.queue_consumer.deprecated.KombuConsumer` (*base_consumer*)
Consumer for use in baseplate.

The `get_message()` and `get_batch()` methods will automatically record diagnostic information.

classmethod `new` (*connection*, *queues*, *queue_size=100*)
Create and initialize a consumer.

Parameters

- **connection** (`Connection`) – The connection
- **queues** (`Sequence[Queue]`) – List of queues.
- **queue_size** (`int`) – The maximum number of messages to cache in the internal `queue.Queue` worker queue. Defaults to 100. For an infinite size (not recommended), use `queue_size=0`.

Return type `KombuConsumer`

get_message (*server_span*)
Return a single message.

Parameters **server_span** (`Span`) – The span.

Return type `Message`

get_batch (*server_span*, *max_items*, *timeout*)
Return a batch of messages.

Parameters

- **server_span** (`Span`) – The span.
- **max_items** (`int`) – The maximum batch size.
- **timeout** (`Optional[float]`) – The maximum time to wait in seconds, or `None` for no timeout.

Return type `Sequence[Message]`

If you require more direct control

class `baseplate.frameworks.queue_consumer.deprecated.BaseKombuConsumer` (*worker*,
worker_thread,
work_queue)

Base object for consuming messages from a queue.

A worker process accepts messages from the queue and puts them in a local work queue. The “real” consumer can then get messages with `get_message()` or `get_batch()`. It is that consumer’s responsibility to ack or reject messages.

Can be used directly, outside of standard baseplate context.

classmethod `new` (*connection*, *queues*, *queue_size=100*)
Create and initialize a consumer.

Parameters

- **connection** (`Connection`) – The connection

- **queues** (*Sequence*[*Queue*]) – List of queues.
- **queue_size** (*int*) – The maximum number of messages to cache in the internal *queue.Queue* worker queue. Defaults to 100. For an infinite size (not recommended), use *queue_size=0*.

Return type *BaseKombuConsumer*

get_message (*timeout=None*)
Return a single message.

Return type *Message*

get_batch (*max_items, timeout*)
Return a batch of messages.

Parameters

- **max_items** (*int*) – The maximum batch size.
- **timeout** (*Optional[float]*) – The maximum time to wait in seconds, or *None* for no timeout.

Return type *Sequence*[*Message*]

baseplate.observers

Observers watch Baseplate for events that happen during requests, such as requests starting and ending and service calls being made. Observers can also add attributes to the *RequestContext* for your application to use during the request.

To enable observers, call *configure_observers()* on your *Baseplate* object during application startup and supply the application configuration. See each observer below for what configuration options are available.

```
def make_wsgi_app(app_config):  
    baseplate = Baseplate(app_config)  
    baseplate.configure_observers()  
  
    ...
```

Logging

The logging observer adds request-specific metadata to log lines coming out of your application.

Configuration

No configuration is necessary, this observer is always enabled when you call *configure_observers()*.

If your application is run with **baseplate-serve**, logging can be controlled with Python's standard logging configuration. See *Logging* for more information.

Outputs

When used with **baseplate-serve**, log lines look like:

For each span representing a call to a remote service or database, the timer has a name like `{namespace}.clients.{context_name}.{method}` and the counter `{namespace}.clients.{context_name}.{method}.{success,failure}` where `context_name` is the name of the client in the context configuration.

Calls to `incr_tag()` will increment a counter like `{namespace}.{tag_name}` by the amount specified.

When using **baseplate-serve**, various process-level runtime metrics will also be emitted. These are not tied to individual requests but instead give insight into how the whole application is functioning. See *Process-level metrics* for more information.

Direct Use

When enabled, the metrics observer also adds a *Client* object as an attribute named `metrics` to the *RequestContext*:

```
def my_handler(request):
    request.metrics.counter("foo").increment()
```

To keep your application more generic, it's better to use local spans for custom local timers and `incr_tag()` for custom counters.

Sentry (Crash Reporting)

The Sentry observer integrates *Raven* with your application to record tracebacks for crashes to *Sentry*.

Configuration

Make sure your service calls `configure_observers()` during application startup and then add the following to your configuration file to enable and configure the Sentry observer.

```
[app:main]

...

# required to enable the observer
# the DSN provided by Sentry for your project
sentry.dsn = https://decaf:face@sentry.local/123

# optional: string to identify this client installation
sentry.site = my site

# optional: the environment this application is running in
sentry.environment = staging

# optional: comma-delimited list of module prefixes to ignore when
# determining where an error came from
sentry.exclude_paths = foo, bar

# optional: comma-delimited list of module prefixes to include for
# consideration when drilling down into an exception
sentry.include_paths = foo, bar

# optional: comma-delimited list of fully qualified names of exception
# classes (potentially with * globs) to not report.
```

(continues on next page)

(continued from previous page)

```
sentry.ignore_exceptions = my_service.UninterestingException

# optional: percent chance that a given error will be reported
# (defaults to 100%)
sentry.sample_rate = 37%

# optional: comma-delimited list of fully qualified names of processor
# classes to apply to events before sending to Sentry. defaults to
# raven.processors.SanitizePasswordsProcessor
sentry.processors = my_service.SanitizeTokenProcessor

...
```

Outputs

Any unexpected exceptions that cause the request to crash (including outside request context) will be reported to Sentry. The Trace ID of the current request will be included in the context reported to Sentry.

Direct Use

When enabled, the error reporting observer also adds a `raven.Client` object as an attribute named `sentry` to the `RequestContext`:

```
def my_handler(request):
    try:
        ...
    except Exception:
        request.sentry.captureException()
```

Tracing

The tracing observer reports span information to a Zipkin-compatible distributed trace aggregator. This can be used to build up cross-service views of request processing across many services.

See [the OpenTracing overview](#) for more info on what tracing is and how it is helpful to you.

Configuration

Make sure your service calls `configure_observers()` during application startup and then add the following to your configuration file to enable and configure the tracing observer.

```
[app:main]

...

# required to enable observer
# the name of the service reporting traces
tracing.service_name = my_service

# optional: traces won't be reported if not set
```

(continues on next page)

(continued from previous page)

```
# the name of the POSIX queue the trace publisher sidecar
# is listening on
tracing.queue_name = some-queue

# optional: what percent of requests to report spans for
# (defaults to 10%). note: this only kicks in if sampling
# has not already been determined by an upstream service.
tracing.sample_rate = 10%

...
```

Server Timeouts

The timeout observer ends processing of requests in your service if they take too long. This is particularly important when an upstream service times out on its end and retries requests to your services which will cause a pileup.

This is entirely configured in-service at the moment and no headers from upstream services are yet taken into account.

Warning: The timeout mechanism is entirely cooperative. If request processing is taking a long time because it is doing compute-heavy actions and not yielding to the event loop it might go on longer than the allotted timeout.

Configuration

Make sure your service calls `configure_observers()` during application startup. By default, requests will time out after 10 seconds. The following configuration settings allow you to customize this.

```
[app:main]

...

# optional: defaults to 10 seconds if not specified. this timeout
# is used for any endpoint not specified in the by_endpoint
# section below.
server_timeout.default = 200 milliseconds

# optional: defaults to false. if enabled, tracebacks will be
# printed to the logs when timeouts occur.
server_timeout.debug = true

# optional: timeout values for specific endpoints. the name
# used must match the name of the server span generated.
# this overrides the default timeout.
# - thrift services: the name of the thrift RPC method
# - pyramid services: the name of the route (config.add_route)
server_timeout.by_endpoint.is_healthy = 300 milliseconds
server_timeout.by_endpoint.my_method = 12 seconds

...
```

Outputs

When a request times out, Baseplate.py will end the greenlet processing that request and emit some diagnostics:

- A log entry like `Server timed out processing for 'is_healthy' after 0.30 seconds`. If `server_timeout.debug` was configured to `True`, the full stack trace of the place the greenlet timed out will also be included.
- A *counter metric* named `{namespace}.server.{route_or_method_name}.timed_out`.
- A tag on the *server span sent to distributed tracing* indicating `timed_out=True`.

1.2.2 The Library

Baseplate also provides a collection of “extra batteries”. These independent modules provide commonly needed functionality to applications. They can be used separately from the rest of Baseplate.

`baseplate.lib.config`

Configuration parsing and validation.

This module provides `parse_config` which turns a dictionary of stringy keys and values into a structured and typed configuration object.

For example, an INI file like the following:

```
[app:main]
simple = true
cards = clubs, spades, diamonds
nested.once = 1
nested.really.deep = 3 seconds
some_file = /var/lib/whatever.txt
sample_rate = 37.1%
interval = 30 seconds
```

Might be parsed like the following. Note: when running under the baseplate server, The `config_parser.items(...)` step is taken care of for you and `raw_config` is passed as the only argument to your factory function.

```
>>> raw_config = dict(config_parser.items("app:main"))

>>> CARDS = config.OneOf(clubs=1, spades=2, diamonds=3, hearts=4)
>>> cfg = config.parse_config(raw_config, {
...     "simple": config.Boolean,
...     "cards": config.TupleOf(CARDS),
...     "nested": {
...         "once": config.Integer,
...         "really": {
...             "deep": config.Timespan,
...         },
...     },
...     "some_file": config.File(mode="r"),
...     "optional": config.Optional(config.Integer, default=9001),
...     "sample_rate": config.Percent,
...     "interval": config.Fallback(config.Timespan, config.Integer),
... })
```

(continues on next page)

(continued from previous page)

```
>>> print(cfg.simple)
True

>>> print(cfg.cards)
[1, 2, 3]

>>> print(cfg.nested.really.deep)
0:00:03

>>> cfg.some_file.read()
'cool'
>>> cfg.some_file.close()

>>> cfg.sample_rate
0.371

>>> print(cfg.interval)
0:00:30
```

Parser

`baseplate.lib.config.parse_config` (*config*, *spec*)
Parse options against a spec and return a structured representation.

Parameters

- **config** (`Dict[str, str]`) – The raw stringy configuration dictionary.
- **spec** (`Dict[str, Union[Parser[~T], Dict[str, Any], Callable[[str], ~T]]]`) – A specification of what the configuration should look like.

Raises `ConfigurationError` The configuration violated the spec.

Return type `ConfigNamespace`

Returns A structured configuration object.

Value Types

Each option can have a type specified. Some types compose with other types to make complicated expressions.

`baseplate.lib.config.String` (*text*)
A raw string.

Return type `str`

`baseplate.lib.config.Float` (*text*)
A floating-point number.

Return type `float`

`baseplate.lib.config.Integer` (*text=None*, *base=10*)
An integer.

To prevent mistakes, this will raise an error if the user attempts to configure a non-whole number.

Parameters **base** (`int`) – (Optional) If specified, the base of the integer to parse.

Return type `Union[int, Callable[[str], int]]`

`baseplate.lib.config.Boolean(text)`

True or False, case insensitive.

Return type `bool`

`baseplate.lib.config.Endpoint(text)`

A remote endpoint to connect to.

Returns an `EndpointConfiguration`.

If the endpoint is a hostname:port pair, the family will be `socket.AF_INET` and address will be a two-tuple of host and port, as expected by `socket`.

If the endpoint contains a slash (/), it will be interpreted as a path to a UNIX domain socket. The family will be `socket.AF_UNIX` and address will be the path as a string.

Return type `EndpointConfiguration`

`baseplate.lib.config.Timespan(text)`

A span of time.

This takes a string of the form “1 second” or “3 days” and returns a `datetime.timedelta` representing that span of time.

Units supported are: milliseconds, seconds, minutes, hours, days.

Return type `timedelta`

`baseplate.lib.config.Base64(text)`

A base64 encoded block of data.

This is useful for arbitrary binary blobs.

Return type `bytes`

`baseplate.lib.config.File(mode='r')`

A path to a file.

This takes a path to a file and returns an open file object, like returned by `open()`.

Parameters `mode(str)` – an optional string that specifies the mode in which the file is opened.

Return type `Callable[[str], Io[AnyStr]]`

`baseplate.lib.config.Percent(text)`

A percentage.

This takes a string of the form “37.2%” or “44%” and returns a float in the range [0.0, 1.0].

Return type `float`

`baseplate.lib.config.UnixUser(text)`

A Unix user name.

The parsed value will be the integer user ID.

Return type `int`

`baseplate.lib.config.UnixGroup(text)`

A Unix group name.

The parsed value will be the integer group ID.

Return type `int`

`baseplate.lib.config.OneOf(**options)`

One of several choices.

For each option, the name is what should be in the configuration file and the value is what it is mapped to.

For example:

```
OneOf(hearts="H", spades="S")
```

would parse:

```
"hearts"
```

into:

```
"H"
```

Return type `Callable[[str], ~T]`

`baseplate.lib.config.TupleOf(item_parser)`

A comma-delimited list of type T.

At least one value must be provided. If you want an empty list to be a valid choice, wrap with `Optional()`.

Return type `Callable[[str], Sequence[~T]]`

If you need something custom or fancy for your application, just use a callable which takes a string and returns the parsed value or raises `ValueError`.

Combining Types

These options are used in combination with other types to form more complex configurations.

`baseplate.lib.config.Optional(item_parser, default=None)`

An option of type T, or default if not configured.

Return type `Callable[[str], Optional[~T]]`

`baseplate.lib.config.Fallback(primary_parser, fallback_parser)`

An option of type T1, or if that fails to parse, of type T2.

This is useful for backwards-compatible configuration changes.

Return type `Callable[[str], ~T]`

`baseplate.lib.config.DictOf(*args, **kws)`

A group of options of a given type.

This is useful for providing data to the application without the application having to know ahead of time all of the possible keys.

```
[app:main]
population.cn = 1383890000
population.in = 1317610000
population.us = 325165000
population.id = 263447000
population.br = 207645000
```

```
>>> cfg = config.parse_config(raw_config, {
...     "population": config.DictOf(config.Integer),
... })

>>> len(cfg.population)
5

>>> cfg.population["br"]
207645000
```

It can also be combined with other configuration specs or parsers to parse more complicated structures:

```
[app:main]
countries.cn.population = 1383890000
countries.cn.capital = Beijing
countries.in.population = 1317610000
countries.in.capital = New Delhi
countries.us.population = 325165000
countries.us.capital = Washington D.C.
countries.id.population = 263447000
countries.id.capital = Jakarta
countries.br.population = 207645000
countries.br.capital = Brasília
```

```
>>> cfg = config.parse_config(raw_config, {
...     "countries": config.DictOf({
...         "population": config.Integer,
...         "capital": config.String,
...     }),
... })

>>> len(cfg.countries)
5

>>> cfg.countries["cn"].capital
'Beijing'

>>> cfg.countries["id"].population
263447000
```

Data Types

class `baseplate.lib.config.EndpointConfiguration`

A description of a remote endpoint.

This is a 2-tuple of (family and address).

family One of `socket.AF_INET` or `socket.AF_UNIX`.

address An address appropriate for the family.

See also:

`baseplate.lib.config.Endpoint()`

Add a new parser

class `baseplate.lib.config.Parser`
Base class for configuration parsers.

Exceptions

exception `baseplate.lib.config.ConfigurationError` (*key, error*)
Raised when the configuration violates the spec.

`baseplate.lib.crypto`

Utilities for common cryptographic operations.

```
message = "Hello, world!"

secret = secrets.get_versioned("some_signing_key")
signature = make_signature(
    secret, message, max_age=datetime.timedelta(days=1))

try:
    validate_signature(secret, message, signature)
except SignatureError:
    print("Oh no, it was invalid!")
else:
    print("Message was valid!")
```

```
Message was valid!
```

Message Signing

`baseplate.lib.crypto.make_signature` (*secret, message, max_age*)
Return a signature for the given message.

To ensure that key rotation works automatically, always fetch the secret token from the secret store immediately before use and do not cache / save the token anywhere. The `current` version of the secret will be used to sign the token.

Parameters

- **secret** (*VersionedSecret*) – The secret signing key from the secret store.
- **message** (*str*) – The message to sign.
- **max_age** (*timedelta*) – The amount of time in the future the signature will be valid for.

Return type `bytes`

Returns An encoded signature.

`baseplate.lib.crypto.validate_signature` (*secret, message, signature*)
Validate and assert a message's signature is correct.

If the signature is valid, the function will return normally with a *SignatureInfo* with some details about the signature. Otherwise, an exception will be raised.

To ensure that key rotation works automatically, always fetch the secret token from the secret store immediately before use and do not cache / save the token anywhere. All active versions of the secret will be checked when validating the signature.

Parameters

- **secret** (*VersionedSecret*) – The secret signing key from the secret store.
- **message** (*str*) – The message payload to validate.
- **signature** (*bytes*) – The signature supplied with the message.

Raises *UnreadableSignatureError* The signature is corrupt.

Raises *IncorrectSignatureError* The digest is incorrect.

Raises *ExpiredSignatureError* The signature expired.

Return type *SignatureInfo*

class `baseplate.lib.crypto.SignatureInfo`
Information about a valid signature.

Variables

- **version** – The version of the packed signature format.
- **expiration** – The time, in seconds since the UNIX epoch, at which the signature will expire.

Exceptions

exception `baseplate.lib.crypto.SignatureError`
Base class for all message signing related errors.

exception `baseplate.lib.crypto.UnreadableSignatureError`
Raised when the signature is corrupt or wrongly formatted.

exception `baseplate.lib.crypto.IncorrectSignatureError`
Raised when the signature is readable but does not match the message.

exception `baseplate.lib.crypto.ExpiredSignatureError` (*expiration*)
Raised when the signature is valid but has expired.

The *expiration* attribute is the time (as seconds since the UNIX epoch) at which the signature expired.

`baseplate.lib.datetime`

Extensions to the standard library *datetime* module.

`baseplate.lib.datetime.datetime_to_epoch_milliseconds` (*dt*)
Convert datetime object to epoch milliseconds.

Return type `int`

`baseplate.lib.datetime.datetime_to_epoch_seconds` (*dt*)
Convert datetime object to epoch seconds.

Return type `int`

`baseplate.lib.datetime.epoch_seconds_to_datetime` (*sec*)
Convert epoch seconds to UTC datetime.

Return type `datetime`

`baseplate.lib.datetime.get_utc_now()`

Get current offset-aware datetime which has timezone information.

Return type `datetime`

`baseplate.lib.edge_context`

The *EdgeRequestContext* provides an interface into both authentication and context information about the original request from a user. For edge services, it provides helpers to create the initial object and serialize the context information into the appropriate headers. Once this object is created and attached to the context, Baseplate will automatically forward the headers to downstream services so they can access the authentication and context data as well.

class `baseplate.lib.edge_context.EdgeRequestContextFactory` (*secrets*)

Factory for creating *EdgeRequestContext* objects.

Every application should set one of these up. Edge services that talk directly with clients should use *new()* directly. For internal services, pass the object off to Baseplate's framework integration (Thrift/Pyramid) for automatic use.

Parameters *secrets* (`baseplate.lib.secrets.SecretsStore`) – A configured secrets store.

new (*authentication_token=None*, *loid_id=None*, *loid_created_ms=None*, *session_id=None*, *device_id=None*, *origin_service_name=None*, *country_code=None*)

Return a new *EdgeRequestContext* object made from scratch.

Services at the edge that communicate directly with clients should use this to pass on the information they get to downstream services. They can then use this information to check authentication, run experiments, etc.

To use this, create and attach the context early in your request flow:

```
auth_cookie = request.cookies["authentication"]
token = request.authentication_service.authenticate_cookie(cookie)
loid = parse_loid(request.cookies["loid"])
session = parse_session(request.cookies["session"])
device_id = request.headers["x-device-id"]

edge_context = self.edgecontext_factory.new(
    authentication_token=token,
    loid_id=loid.id,
    loid_created_ms=loid.created,
    session_id=session.id,
    device_id=device_id,
)
edge_context.attach_context(request)
```

Parameters

- **authentication_token** (`Optional[bytes]`) – A raw authentication token as returned by the authentication service.
- **loid_id** (`Optional[str]`) – ID for the current LoID in fullname format.
- **loid_created_ms** (`Optional[int]`) – Epoch milliseconds when the current LoID cookie was created.

- **session_id** (*Optional[str]*) – ID for the current session cookie.
- **device_id** (*Optional[str]*) – ID for the device where the request originated from.
- **origin_service_name** (*Optional[str]*) – Name for the “origin” service handling the request from the client.
- **country_code** (*Optional[str]*) – two-character ISO 3166-1 country code where the request originated from.

Return type *EdgeRequestContext*

from_upstream (*edge_header*)

Create and return an *EdgeRequestContext* from an upstream header.

This is generally used internally to Baseplate by framework integrations that automatically pick up context from inbound requests.

Parameters **edge_header** (*Optional[bytes]*) – Raw payload of Edge-Request header from upstream service.

Return type *EdgeRequestContext*

class `baseplate.lib.edge_context.EdgeRequestContext` (*authn_token_validator, header*)

Contextual information about the initial request to an edge service.

Construct this using an *EdgeRequestContextFactory*.

attach_context (*context*)

Attach this to the provided *RequestContext*.

Parameters **context** (*RequestContext*) – request context to attach this to

Return type *None*

event_fields ()

Return fields to be added to events.

Return type *Dict[str, Any]*

user

User object for the current context.

oauth_client

OAuthClient object for the current context.

device

Device object for the current context.

session

Session object for the current context.

service

Service object for the current context.

origin_service

Origin object for the current context.

geolocation

Geolocation object for the current context.

class `baseplate.lib.edge_context.User`

Wrapper for the user values in *AuthenticationToken* and the *LoId* cookie.

authentication_token

Alias for field number 0

loid

Alias for field number 1

cookie_created_ms

Alias for field number 2

id

Return the authenticated account_id for the current User.

Raises *NoAuthenticationError* if there was no authentication token, it was invalid, or the subject is not an account.

Return type *Optional[str]*

is_logged_in

Return if the User has a valid, authenticated id.

Return type *bool*

roles

Return the authenticated roles for the current User.

Raises *NoAuthenticationError* if there was no authentication token or it was invalid

Return type *Set[str]*

has_role (*role*)

Return if the authenticated user has the specified role.

Parameters *client_types* – Case-insensitive sequence role name to check.

Raises *NoAuthenticationError* if there was no authentication token defined for the current context

Return type *bool*

event_fields ()

Return fields to be added to events.

Return type *Dict[str, Any]*

class *baseplate.lib.edge_context.OAuthClient*

Wrapper for the OAuth2 client values in AuthenticationToken.

authentication_token

Alias for field number 0

id

Return the authenticated id for the current client.

Raises *NoAuthenticationError* if there was no authentication token defined for the current context

Return type *Optional[str]*

is_type (**client_types*)

Return if the authenticated client type is one of the given types.

When checking the type of the current OAuthClient, you should check that the type “is” one of the allowed types rather than checking that it “is not” a disallowed type.

For example:


```
if oauth_client.is_type("third_party"):
    ...
```

not:

```
if not oauth_client.is_type("first_party"):
    ...
```

Parameters `client_types` (`str`) – Case-insensitive sequence of client type names that you want to check.

Raises `NoAuthenticationError` if there was no authentication token defined for the current context

Return type `bool`

event_fields ()

Return fields to be added to events.

Return type `Dict[str, Any]`

class `baseplate.lib.edge_context.Session`

Wrapper for the session values in the `EdgeRequestContext`.

id

Alias for field number 0

class `baseplate.lib.edge_context.Service`

Wrapper for the Service values in `AuthenticationToken`.

authentication_token

Alias for field number 0

name

Return the authenticated service name.

Type name string or None if context authentication is invalid

Raises `NoAuthenticationError` if there was no authentication token, it was invalid, or the subject is not a service.

Return type `str`

class `baseplate.lib.edge_context.AuthenticationToken`

Information about the authenticated user.

`EdgeRequestContext` provides high-level helpers for extracting data from authentication tokens. Use those instead of direct access through this class.

subject

Return the raw *subject* that is authenticated.

Return type `Optional[str]`

exception `baseplate.lib.edge_context.NoAuthenticationError`

Raised when trying to use an invalid or missing authentication token.

baseplate.lib.events

Client library for sending events to the data processing system.

This is for use with the event collector system. Events generally track something that happens in production that we want to instrument for planning and analytical purposes.

Events are serialized and put onto a message queue on the same server. These serialized events are then consumed and published to the remote event collector by a separate daemon.

Building Events

Thrift Schema v2 Events

For modern Thrift-based events: import the event schemas into your project, instantiate and fill out an event object, and pass it into the queue:

```
import time
import uuid

from baseplate import Baseplate
from baseplate.lib.events import EventQueue
from baseplate.lib.events import serialize_v2_event

from event_schemas.event.ttypes import Event

def my_handler(request):
    event = Event(
        source="baseplate",
        action="test",
        noun="baseplate",
        client_timestamp=time.time() * 1000,
        uuid=str(uuid.uuid4()),
    )
    request.events_v2.put(ev2)

def make_wsgi_app(app_config):
    ...

    baseplate = Baseplate(app_config)
    baseplate.configure_context(
        {
            ...

            "events_v2": EventQueue("v2", serialize_v2_events),

            ...
        }
    )

    ...
```

Queuing Events

class baseplate.lib.events.**EventQueue** (*name, event_serializer*)
A queue to transfer events to the publisher.

Parameters

- **name** (`str`) – The name of the event queue to send to. This specifies which publisher should send the events which can be useful for routing to different event pipelines (prod/test/v2 etc.).
- **event_serializer** (`Callable[[~T], bytes]`) – A callable that takes an event object and returns serialized bytes ready to send on the wire. See below for options.

put (`event`)

Add an event to the queue.

The queue is local to the server this code is run on. The event publisher on the server will take these events and send them to the collector.

Parameters **event** (`~T`) – The event to send. The type of event object passed in depends on the selected `event_serializer`.

Raises `EventTooLargeError` The serialized event is too large.

Raises `EventQueueFullError` The queue is full. Events are not being published fast enough.

Return type `None`

The `EventQueue` also implements `ContextFactory` so it can be used with `add_to_context()`:

```
event_queue = EventQueue("production", serialize_v2_event)
baseplate.add_to_context("events_production", event_queue)
```

It can then be used from the `RequestContext` during requests:

```
def some_service_method(self, context):
    event = Event(...)
    context.events_production.put(event)
```

Serializers

The `event_serializer` parameter to `EventQueue` is a callable which serializes a given event object. Baseplate comes with a serializer for the Thrift schema based V2 event system:

`baseplate.lib.events.serialize_v2_event(event)`
Serialize a Thrift struct to bytes for the V2 event protocol.

Parameters **event** (`Any`) – A Thrift struct from the event schemas.

Return type `bytes`

Exceptions

exception `baseplate.lib.events.EventError`
Base class for event related exceptions.

exception `baseplate.lib.events.EventTooLargeError(size)`
Raised when a serialized event is too large to send.

exception `baseplate.lib.events.EventQueueFullError`
Raised when the queue of events is full.

This usually indicates that the event publisher is having trouble talking to the event collector.

Publishing Events

Events that are put onto an *EventQueue* are consumed by a separate process and published to the remote event collector service. The publisher is in baseplate and can be run as follows:

```
$ python -m baseplate.sidecars.event_publisher --queue-name something config_file.ini
```

The publisher will look at the specified INI file to find its configuration. Given a queue name of something (as in the example above), it will expect a section in the INI file called `[event-publisher:something]` with content like below:

```
[event-publisher:something]
collector.hostname = some-domain.example.com

key.name = NameOfASecretKey
key.secret = Base64-encoded-blob-of-randomness

metrics.namespace = a.name.to.put.metrics.under
metrics.endpoint = the-statsd-host:1234
```

baseplate.lib.experiments

Configuration Parsing

```
baseplate.lib.experiments.experiments_client_from_config(app_config,
                                                         event_logger,    pre-
                                                         fix='experiments.')
```

Configure and return an *ExperimentsContextFactory* object.

The keys useful to *experiments_client_from_config()* should be prefixed, e.g. `experiments.path`, etc.

Supported keys:

path: the path to the experiment configuration file generated by the experiment configuration fetcher daemon.

timeout (optional): the time that we should wait for the file specified by `path` to exist. Defaults to *None* which is *infinite*.

Parameters

- **raw_config** – The application configuration which should have settings for the experiments client.
- **event_logger** (*EventLogger*) – The *EventLogger* to be used to log bucketing events.
- **prefix** (*str*) – the prefix used to filter keys (defaults to “experiments.”).
- **backoff** – retry backoff time for experiments file watcher. Defaults to *None*, which is mapped to `DEFAULT_FILEWATCHER_BACKOFF`.

Return type *ExperimentsContextFactory*

Classes

```
class baseplate.lib.experiments.ExperimentsContextFactory (path,  
                                                         event_logger=None,  
                                                         timeout=None, back-  
                                                         off=None)
```

Experiment client context factory.

This factory will attach a new `baseplate.lib.experiments.Experiments` to an attribute on the `RequestContext`.

Parameters

- **path** (`str`) – Path to the experiment configuration file.
- **event_logger** (`Optional[EventLogger]`) – The logger to use to log experiment eligibility events. If not provided, a `DebugLogger` will be created and used.
- **timeout** (`Optional[float]`) – How long, in seconds, to block instantiation waiting for the watched experiments file to become available (defaults to not blocking).
- **backoff** (`Optional[float]`) – retry backoff time for experiments file watcher. Defaults to `None`, which is mapped to `DEFAULT_FILEWATCHER_BACKOFF`.

```
class baseplate.lib.experiments.Experiments (config_watcher, server_span, context_name,  
                                                         event_logger=None)
```

Access to experiments with automatic refresh when changed.

This experiments client allows access to the experiments cached on disk by the experiment configuration fetcher daemon. It will automatically reload the cache when changed. This client also handles logging bucketing events to the event pipeline when it is determined that the request is part of an active variant.

get_all_experiment_names ()

Return a list of all valid experiment names from the configuration file.

Return type `Sequence[str]`

Returns List of all valid experiment names.

is_valid_experiment (*name*)

Return true if the provided experiment name is a valid experiment.

Parameters **name** (`str`) – Name of the experiment you want to check.

Return type `bool`

Returns Whether or not a particular experiment is valid.

variant (*name, user=None, bucketing_event_override=None, **kwargs*)

Return which variant, if any, is active.

If a variant is active, a bucketing event will be logged to the event pipeline unless any one of the following conditions are met:

1. `bucketing_event_override` is set to `False`.
2. The experiment specified by “name” explicitly disables bucketing events.
3. We have already logged a bucketing event for the value specified by `experiment.get_unique_id(**kwargs)` within the current request.

Since checking the status an experiment will fire a bucketing event, it is best to only check the variant when you are making the decision that will expose the experiment to the user. If you absolutely must check the

status of an experiment before you are sure that the experiment will be exposed to the user, you can use *bucketing_event_override* to disabled bucketing events for that check.

Parameters

- **name** (*str*) – Name of the experiment you want to run.
- **user** (*Optional[User]*) – User object for the user you want to check the experiment variant for. If you set user, the experiment parameters for that user (“user_id”, “logged_in”, and “user_roles”) will be extracted and added to the inputs to the call to `Experiment.variant`. The user’s event_fields will also be extracted and added to the bucketing event if one is logged. It is recommended that you provide a value for user rather than setting the user parameters manually in `kwargs`.
- **bucketing_event_override** (*Optional[bool]*) – Set if you need to override the default behavior for sending bucketing events. This parameter should be set sparingly as it breaks the assumption that you will fire a bucketing event when you first check the state of an experiment. If set to `False`, will never send a bucketing event. If set to `None`, no override will be applied. Set to `None` by default. Note that setting `bucketing_event_override` to `True` has no effect, it will behave the same as when it is set to `None`.
- **kwargs** (*str*) – Arguments that will be passed to `experiment.variant` to determine bucketing, targeting, and overrides. These values will also be passed to the logger.

Return type *Optional[str]*

Returns Variant name if a variant is active, `None` otherwise.

expose (*experiment_name, variant_name, user=None, **kwargs*)

Log an event to indicate that a user has been exposed to an experimental treatment.

Parameters

- **experiment_name** (*str*) – Name of the experiment that was exposed.
- **variant_name** (*str*) – Name of the variant that was exposed.
- **user** (*Optional[User]*) – User object for the user you want to check the experiment variant for. If unset, it is expected that `user_id` and `logged_in` values will be set in the keyword arguments.
- **kwargs** (*str*) – Additional arguments that will be passed to logger.

Return type `None`

`baseplate.lib.file_watcher`

Watch a file and keep a parsed copy in memory that’s updated on changes.

The contents of the file are re-loaded and parsed only when necessary.

For example, a JSON file like the following:

```
{
  "one": 1,
  "two": 2
}
```

might be watched and parsed like this:

```
>>> watcher = FileWatcher(path, parser=json.load)
>>> watcher.get_data() == {u"one": 1, u"two": 2}
True
```

The return value of `get_data()` would change whenever the underlying file changes.

```
class baseplate.lib.file_watcher.FileWatcher (path, parser, timeout=None, binary=False,
                                              encoding=None, newline=None, back-
                                              off=None)
```

Watch a file and load its data when it changes.

Parameters

- **path** (`str`) – Full path to a file to watch.
- **parser** (`Callable[[io[AnyStr]], ~T]`) – A callable that takes an open file object, parses or otherwise interprets the file, and returns whatever data is meaningful.
- **timeout** (`Optional[float]`) – How long, in seconds, to block instantiation waiting for the watched file to become available (defaults to not blocking).
- **binary** (`bool`) – Should the file be opened in binary mode. If `True` the file will be opened with the mode “`rb`”, otherwise it will be opened with the mode “`r`”. (defaults to “`r`”)
- **encoding** (`Optional[str]`) – The name of the encoding used to decode the file. The default encoding is platform dependent (whatever `locale.getpreferredencoding()` returns), but any text encoding supported by Python can be used. This is not supported if `binary` is set to `True`.
- **newline** (`Optional[str]`) – Controls how universal newlines mode works (it only applies to text mode). It can be `None`, “”, “`\n`”, “`\r`”, and “`\r\n`”. This is not supported if `binary` is set to `True`.
- **backoff** (`Optional[float]`) – retry backoff time for the file watcher. Defaults to `None`, which is mapped to `DEFAULT_FILEWATCHER_BACKOFF`.

`get_data()`

Return the current contents of the file, parsed.

The watcher ensures that the file is re-loaded and parsed whenever its contents change. Parsing only occurs when necessary, not on each call to this method. This method returns whatever the most recent call to the parser returned.

Make sure to call this method each time you need data from the file rather than saving its results elsewhere. This ensures you always have the freshest data.

Return type `~T`

Exceptions

```
exception baseplate.lib.file_watcher.WatchedFileNotAvailableError (path, inner)
```

Raised when the watched file could not be loaded.

`baseplate.lib.live_data`

This component of Baseplate provides real-time synchronization of data across a cluster of servers. It is intended for situations where data is read frequently, does not change super often, and when it does change needs to change

everywhere at once. In most cases, this will be an underlying feature of some other system (e.g. an experiments framework.)

There are four main components of the live data system:

- *ZooKeeper*, a highly available data store that can push change notifications.
- The watcher, a sidecar daemon that watches nodes in ZooKeeper and syncs their contents to disk.
- *FileWatcher* instances in your application that load the synchronized data into memory.
- Something that writes to ZooKeeper (potentially the writer tool).

The watcher daemon and tools for writing data to ZooKeeper are covered on this page.

Watcher Daemon

The watcher daemon is a sidecar that watches nodes in ZooKeeper and syncs their contents to local files on change. It is entirely configured via INI file and is run like so:

```
$ python -m baseplate.sidecars.live_data_watcher some_config.ini
```

Where `some_config.ini` might look like:

```
[live-data]
zookeeper.hosts = zk01:2181,zk02:2181
zookeeper.credentials = secret/myservice/zookeeper_credentials

nodes.a.source = /a/node/in/zookeeper
nodes.a.dest = /var/local/file-on-disk

nodes.b.source = /another/node/in/zookeeper
nodes.b.dest = /var/local/another-file
nodes.b.owner = www-data
nodes.b.group = www-data
nodes.b.mode = 0400
```

Each of the defined nodes will be watched by the daemon.

The watcher daemon will touch the `mtime` of the local files periodically to indicate liveliness to monitoring tools.

The Writer Tool

For simple cases where you just want to put the contents of a file into ZooKeeper (perhaps in a CI task) you can use the live data writer. It expects a configuration file with ZooKeeper connection information, like the watcher, and takes some additional parameters on the command line.

```
$ python -m baseplate.lib.live_data.writer some_config.ini \
    input.json /some/node/in/zookeeper
Writing input.json to ZooKeeper /some/node/in/zookeeper...
---
+++
@@ -1,4 +1,4 @@
{
```

(continues on next page)

(continued from previous page)

```

-     "key": "one"
+     "key": "two"
}
Wrote data to Zookeeper.

```

The ZooKeeper node must be created before this tool can be used so that appropriate ACLs can be configured.

Direct access to ZooKeeper

If you're doing something more complicated with your data that the above tools don't cover, you'll want to connect directly to ZooKeeper.

```
baseplate.lib.live_data.zookeeper_client_from_config(secrets, app_config,
                                                    read_only=None)
```

Configure and return a ZooKeeper client.

There are several configuration options:

zookeeper.hosts A comma-delimited list of hosts with optional `chroot` at the end. For example `zk01:2181, zk02:2181` or `zk01:2181, zk02:2181/some/root`.

zookeeper.credentials (Optional) A comma-delimited list of paths to secrets in the secrets store that contain ZooKeeper authentication credentials. Secrets should be of the "simple" type and contain `username:password`.

zookeeper.timeout (Optional) A time span of how long to wait for each connection attempt.

The client will attempt forever to reconnect on connection loss.

Parameters

- **secrets** (*SecretsStore*) – A secrets store object
- **raw_config** – The application configuration which should have settings for the ZooKeeper client.
- **read_only** (*Optional[bool]*) – Whether or not to allow connections to read-only ZooKeeper servers.

Return type *KazooClient*

baseplate.lib.message_queue

This module provides a thin wrapper around POSIX Message queues.

Note: This implementation uses [POSIX Message queues](#) and is not portable to all operating systems.

There are also various limits on the sizes of queues:

- The `msgqueue.rlimit` limits the amount of space the user can use on message queues.
 - The `fs.mqueue.msg_max` and `fs.mqueue.msgsize_max` sysctls limit the maximum number of messages and the maximum size of each message which a queue can be configured to have.
-

Minimal Example

Here's a minimal, artificial example of a separate producer and consumer process pair (run the producer then the consumer):

```
# producer.py
from baseplate.lib.message_queue import MessageQueue

# If the queue doesn't already exist, we'll create it.
mq = MessageQueue(
    "/baseplate-testing", max_messages=1, max_message_size=1)
message = "1"
mq.put(message)
print("Put Message: %s" % message)
```

You should see:

```
Put Message: 1
```

After running the producer once, we have a single message pushed on to our POSIX message queue. Next up, run the consumer:

```
# consumer.py
from baseplate.lib.message_queue import MessageQueue

mq = MessageQueue(
    "/baseplate-testing", max_messages=1, max_message_size=1)
# Unless a `timeout` kwarg is passed, this will block until
# we can pop a message from the queue.
message = mq.get()
print("Get Message: %s" % message.decode())
```

You'll end up seeing:

```
Get Message: 1
```

The `/baseplate-testing` value is the name of the queue. Queues names should start with a forward slash, followed by one or more characters (but no additional slashes).

Multiple processes can bind to the same queue by specifying the same queue name.

Message Queue Default Limits

Most operating systems with POSIX queues include very low defaults for the maximum message size and maximum queue depths. On Linux 2.6+, you can list and check the values for these by running:

```
$ ls /proc/sys/fs/mqueue/
msg_default  msg_max  msgsize_default  msgsize_max  queues_max
$ cat /proc/sys/fs/mqueue/msgsize_max
8192
```

Explaining these in detail is outside the scope of this document, so we'll refer you to [POSIX Message queues](#) (or `man 7 mq_overview`) for detailed instructions on what these mean.

Gotchas

If you attempt to create a POSIX Queue where one of your provided values is over the limits defined under `/proc/sys/fs/mqueue/`, you'll probably end up seeing a vague `ValueError` exception. Here's an example:

```
>>> from baseplate.lib.message_queue import MessageQueue
>>> mq = MessageQueue(
    "/over-the-limit", max_messages=11, max_message_size=8096)
Traceback (most recent call last):
  File "<input>", line 2, in <module>
  File "/home/myuser/baseplate/baseplate/lib/message_queue.py", line 83, in __init__
    max_message_size=max_message_size,
ValueError: Invalid parameter(s)
```

Since the default value for `/proc/sys/fs/mqueue/msg_max` on Linux is 10, our `max_messages=11` is invalid. You can raise these limits by doing something like this as a privileged user:

```
$ echo "50" > /proc/sys/fs/mqueue/msg_max
```

CLI Usage

The `message_queue` module can also be run as a command-line tool to consume, log, and discard messages from a given queue:

```
$ python -m baseplate.lib.message_queue --read /queue
```

or to write arbitrary messages to the queue:

```
$ echo hello! | python -m baseplate.lib.message_queue --write /queue
```

See `--help` for more info.

`baseplate.lib.message_queue`

A Gevent-friendly POSIX message queue.

```
class baseplate.lib.message_queue.MessageQueue (name, max_messages,
                                                max_message_size)
```

A Gevent-friendly (but not required) inter process message queue.

`name` should be a string of up to 255 characters consisting of an initial slash, followed by one or more characters, none of which are slashes.

Note: This relies on POSIX message queues being available and `select(2)`-able like other file descriptors. Not all operating systems support this.

get (*timeout=None*)

Read a message from the queue.

Parameters `timeout` (`Optional[float]`) – If the queue is empty, the call will block up to `timeout` seconds or forever if `None`.

Raises `TimedOutError` The queue was empty for the allowed duration of the call.

Return type `bytes`

put (*message*, *timeout=None*)

Add a message to the queue.

Parameters *timeout* (`Optional[float]`) – If the queue is full, the call will block up to *timeout* seconds or forever if `None`.

Raises `TimeoutError` The queue was full for the allowed duration of the call.

Return type `None`

unlink ()

Remove the queue from the system.

The queue will not leave until the last active user closes it.

Return type `None`

close ()

Close the queue, freeing related resources.

This must be called explicitly if queues are created/destroyed on the fly. It is not automatically called when the object is reclaimed by Python.

Return type `None`

Exceptions

exception `baseplate.lib.message_queue.MessageQueueError`

Base exception for message queue related errors.

exception `baseplate.lib.message_queue.TimeoutError`

Raised when a message queue operation times out.

`baseplate.lib.metrics`

Application metrics via StatsD.

A client for the application metrics aggregator `StatsD`. Metrics sent to StatsD are aggregated and written to graphite. StatsD is generally used for whole-system health monitoring and insight into usage patterns.

Basic example usage:

```
from baseplate.lib.metrics import metrics_client_from_config

client = metrics_client_from_config(app_config)
client.counter("events.connect").increment()
client.gauge("workers").replace(4)

with client.timer("something.todo"):
    do_something()
do_something_else()
```

If you have multiple metrics to send, you can batch them up for efficiency:

```
with client.batch() as batch:
    batch.counter("froozles").increment()
    batch.counter("blargs").decrement(delta=3)

    with batch.timer("something"):
        do_another_thing()
```

and the batch will be sent in as few packets as possible when the *with* block ends.

Clients

`baseplate.lib.metrics.metrics_client_from_config(raw_config)`

Configure and return a metrics client.

This expects two configuration options:

metrics.namespace The root key to prefix all metrics in this application with.

metrics.endpoint A `host:port` pair, e.g. `localhost:2014`. If an empty string, a client that discards all metrics will be returned.

Parameters `raw_config` (`Dict[str, str]`) – The application configuration which should have settings for the metrics client.

Return type `Client`

Returns A configured client.

class `baseplate.lib.metrics.Client` (`transport, namespace`)

A client for StatsD.

batch ()

Return a client-like object which batches up metrics.

Batching metrics can reduce the number of packets that are sent to the stats aggregator.

Return type `Batch`

counter (`name`)

Return a Counter with the given name.

The sample rate is currently up to your application to enforce.

Parameters `name` (`str`) – The name the counter should have.

Return type `Counter`

gauge (`name`)

Return a Gauge with the given name.

Parameters `name` (`str`) – The name the gauge should have.

Return type `Gauge`

histogram (`name`)

Return a Histogram with the given name.

Parameters `name` (`str`) – The name the histogram should have.

Return type `Histogram`

timer (`name`)

Return a Timer with the given name.

Parameters `name` (`str`) – The name the timer should have.

Return type `Timer`

class baseplate.lib.metrics.**Batch** (*transport, namespace*)

A batch of metrics to send to StatsD.

The batch also supports the [context manager protocol](#), for use with Python's `with` statement. When the context is exited, the batch will automatically `flush()`.

flush ()

Immediately send the batched metrics.

Return type `None`

counter (*name*)

Return a BatchCounter with the given name.

The sample rate is currently up to your application to enforce.

Parameters **name** (`str`) – The name the counter should have.

Return type `Counter`

gauge (*name*)

Return a Gauge with the given name.

Parameters **name** (`str`) – The name the gauge should have.

Return type `Gauge`

histogram (*name*)

Return a Histogram with the given name.

Parameters **name** (`str`) – The name the histogram should have.

Return type `Histogram`

timer (*name*)

Return a Timer with the given name.

Parameters **name** (`str`) – The name the timer should have.

Return type `Timer`

Metrics

class baseplate.lib.metrics.**Counter** (*transport, name*)

A counter for counting events over time.

increment (*delta=1.0, sample_rate=1.0*)

Increment the counter.

Parameters

- **delta** (`float`) – The amount to increase the counter by.
- **sample_rate** (`float`) – What rate this counter is sampled at. [0-1].

Return type `None`

decrement (*delta=1.0, sample_rate=1.0*)

Decrement the counter.

This is equivalent to `increment()` with delta negated.

Return type `None`

send (*delta*, *sample_rate*)

Send the counter to the backend.

Parameters

- **delta** (*float*) – The amount to increase the counter by.
- **sample_rate** (*float*) – What rate this counter is sampled at. [0-1].

Return type None

class baseplate.lib.metrics.**Timer** (*transport*, *name*)

A timer for recording elapsed times.

The timer also supports the [context manager protocol](#), for use with Python’s `with` statement. When the context is entered the timer will `start()` and when exited, the timer will automatically `stop()`.

start (*sample_rate=1.0*)

Record the current time as the start of the timer.

Return type None

stop ()

Stop the timer and record the total elapsed time.

Return type None

send (*elapsed*, *sample_rate=1.0*)

Directly send a timer value without having to stop/start.

This can be useful when the timing was managed elsewhere and we just want to report the result.

Parameters **elapsed** (*float*) – The elapsed time in seconds to report.

Return type None

class baseplate.lib.metrics.**Gauge** (*transport*, *name*)

A gauge representing an arbitrary value.

Note: The StatsD protocol supports incrementing/decrementing gauges from their current value. We do not support that here because this feature is unpredictable in face of the StatsD server restarting and the “current value” being lost.

replace (*new_value*)

Replace the value held by the gauge.

This will replace the value held by the gauge with no concern for its previous value.

Note: Due to the way the protocol works, it is not possible to replace gauge values with negative numbers.

Parameters **new_value** (*float*) – The new value to store in the gauge.

Return type None

class baseplate.lib.metrics.**Histogram** (*transport*, *name*)

A bucketed distribution of integer values across a specific range.

Records data value counts across a configurable integer value range with configurable buckets of value precision within that range.

Configuration of each histogram is managed by the backend service, not by this interface. This implementation also depends on histograms being supported by the StatsD backend. Specifically, the StatsD backend must support the `h` key, e.g. `metric_name:320|h`.

add_sample (*value*)

Add a new value to the histogram.

This records a new value to the histogram; the bucket it goes in is determined by the backend service configurations.

Return type `None`

baseplate.lib.random

Extensions to the standard library *random* module.

class `baseplate.lib.random.WeightedLottery` (*items*, *weight_key*)

A lottery where items can have different chances of selection.

Items will be picked with chance proportional to their weight relative to the sum of all weights, so the higher the weight, the higher the chance of being picked.

Parameters

- **items** (`Iterable[~T]`) – Items to choose from.
- **weight_key** (`Callable[[~T], int]`) – A function that takes an item in *items* and returns a non-negative integer weight for that item.

Raises `ValueError` if any weights are negative or there are no items.

An example of usage:

```
>>> words = ["apple", "banana", "canelope"]
>>> lottery = WeightedLottery(words, weight_key=len)
>>> lottery.pick()
'banana'
>>> lottery.sample(2)
['apple', 'canelope']
```

pick ()

Pick a random element from the lottery.

Return type `~T`

sample (*sample_size*)

Sample elements from the lottery without replacement.

Parameters **sample_size** (`int`) – The number of items to sample from the lottery.

Return type `Iterable[~T]`

baseplate.lib.ratelimit

Configuring a rate limiter for your request context requires a context factory for the backend and a factory for the rate limiter itself:


```
redis_pool = pool_from_config(app_config)
backend_factory = RedisRateLimitBackendContextFactory(redis_pool)
ratelimiter_factory = RateLimiterContextFactory(backend_factory, allowance, interval)
baseplate.add_to_context('ratelimiter', ratelimiter_factory)
```

The rate limiter can then be used during a request with:

```
try:
    context.ratelimiter.consume(context.request_context.user.id)
    print('RateLimit passed')
except RateLimitExceededException:
    print('Too many requests')
```

Classes

class `baseplate.lib.ratelimit.RateLimiter` (*backend, allowance, interval*)

A class for rate limiting actions.

Parameters

- **backend** (*RateLimitBackend*) – The backend to use for storing rate limit counters.
- **allowance** (*int*) – The maximum allowance allowed per key.
- **interval** (*int*) – The interval (in seconds) to reset allowances.

consume (*key, amount=1*)

Consume the given *amount* from the allowance for the given *key*.

This will raise `baseplate.lib.ratelimit.RateLimitExceededException` if the allowance for *key* is exhausted.

Parameters

- **key** (*str*) – The name of the rate limit bucket to consume from.
- **amount** (*int*) – The amount to consume from the rate limit bucket.

Return type `None`

class `baseplate.lib.ratelimit.RateLimiterContextFactory` (*backend_factory, allowance, interval*)

RateLimiter context factory.

Parameters

- **backend_factory** (*ContextFactory*) – An instance of `baseplate.clients.ContextFactory`. The context factory must return an instance of `baseplate.lib.ratelimit.backends.RateLimitBackend`
- **allowance** (*int*) – The maximum allowance allowed per key.
- **interval** (*int*) – The interval (in seconds) to reset allowances.

make_object_for_context (*name, span*)

Return an object that can be added to the context object.

Parameters

- **name** (*str*) – The name assigned to this object on the context.
- **span** (*Span*) – The current span this object is being made for.

Return type `RateLimiter`

class `baseplate.lib.ratelimit.RateLimitExceededException`
This exception gets raised whenever a rate limit is exceeded.

Backends

class `baseplate.lib.ratelimit.backends.RateLimitBackend`
An interface for rate limit backends to implement.

consume (*key, amount, allowance, interval*)

Consume the given *amount* from the allowance for the given *key*.

This will return true if the *key* remains below the *allowance* after consuming the given *amount*.

Parameters

- **key** (`str`) – The name of the rate limit bucket to consume from.
- **amount** (`int`) – The amount to consume from the rate limit bucket.
- **allowance** (`int`) – The maximum allowance for the rate limit bucket.
- **interval** (`int`) – The interval to reset the allowance.

Return type `bool`

Memcache

class `baseplate.lib.ratelimit.backends.memcache.MemcacheRateLimitBackendContextFactory` (*memcache, prefix*)
MemcacheRateLimitBackend context factory.

Parameters

- **memcache_pool** (`PooledClient`) – The memcache pool to back this ratelimit.
- **prefix** (`str`) – A prefix to add to keys during rate limiting. This is useful if you will have two different rate limiters that will receive the same keys.

make_object_for_context (*name, span*)

Return an object that can be added to the context object.

Parameters

- **name** (`str`) – The name assigned to this object on the context.
- **span** (`Span`) – The current span this object is being made for.

Return type `MemcacheRateLimitBackend`

class `baseplate.lib.ratelimit.backends.memcache.MemcacheRateLimitBackend` (*memcache, prefix*)
A Memcache backend for rate limiting.

Parameters

- **memcache** (`MonitoredMemcacheConnection`) – A memcached connection.
- **prefix** (`str`) – A prefix to add to keys during rate limiting. This is useful if you will have two different rate limiters that will receive the same keys.

consume (*key, amount, allowance, interval*)

Consume the given *amount* from the allowance for the given *key*.

This will return true if the *key* remains below the *allowance* after consuming the given *amount*.

Parameters

- **key** (*str*) – The name of the rate limit bucket to consume from.
- **amount** (*int*) – The amount to consume from the rate limit bucket.
- **allowance** (*int*) – The maximum allowance for the rate limit bucket.
- **interval** (*int*) – The interval to reset the allowance.

Return type *bool*

Redis

class baseplate.lib.ratelimit.backends.redis.**RedisRateLimitBackendContextFactory** (*redis_pool, pre-fix='rl:'*)

RedisRateLimitBackend context factory.

Parameters

- **redis_pool** (*ConnectionPool*) – The redis pool to back this ratelimit.
- **prefix** (*str*) – A prefix to add to keys during rate limiting. This is useful if you will have two different rate limiters that will receive the same keys.

make_object_for_context (*name, span*)

Return an object that can be added to the context object.

Parameters

- **name** (*str*) – The name assigned to this object on the context.
- **span** (*Span*) – The current span this object is being made for.

Return type *RedisRateLimitBackend*

class baseplate.lib.ratelimit.backends.redis.**RedisRateLimitBackend** (*redis, pre-fix='rl:'*)

A Redis backend for rate limiting.

Parameters

- **redis** (*MonitoredRedisConnection*) – An instance of *baseplate.clients.redis.MonitoredRedisConnection*.
- **prefix** (*str*) – A prefix to add to keys during rate limiting. This is useful if you will have two different rate limiters that will receive the same keys.

consume (*key, amount, allowance, interval*)

Consume the given *amount* from the allowance for the given *key*.

This will return true if the *key* remains below the *allowance* after consuming the given *amount*.

Parameters

- **key** (*str*) – The name of the rate limit bucket to consume from.
- **amount** (*int*) – The amount to consume from the rate limit bucket.
- **allowance** (*int*) – The maximum allowance for the rate limit bucket.

- **interval** (*int*) – The interval to reset the allowance.

Return type *bool*

`baseplate.lib.retry`

Note: This module is a low-level helper, many client libraries have protocol-aware retry logic built in. Check your library before using this.

Policies for retrying an operation safely.

class `baseplate.lib.retry.RetryPolicy`

A policy for retrying operations.

Policies are meant to be used as an iterable:

```
for time_remaining in RetryPolicy.new(attempts=3):
    try:
        some_operation.do(timeout=time_remaining)
        break
    except SomeError:
        pass
else:
    raise MaxRetriesError
```

yield_attempts ()

Return an iterator which controls attempts.

On each iteration, the iterator will yield the number of seconds left to retry, this should be used to set the timeout on the operation being carried out. If there is no maximum time remaining, *None* is yielded instead.

The iterable will raise *StopIteration* once the operation should not be retried any further.

Return type *Iterator[Optional[float]]*

__iter__ ()

Return the result of *yield_attempts* ().

This allows policies to be directly iterated over.

Return type *Iterator[Optional[float]]*

static new (*attempts=None, budget=None, backoff=None*)

Create a new retry policy with the given constraints.

Parameters

- **attempts** (*Optional[int]*) – The maximum number of times the operation can be attempted.
- **budget** (*Optional[float]*) – The maximum amount of time, in seconds, that the local service will wait for the operation to succeed.
- **backoff** (*Optional[float]*) – The base amount of time, in seconds, for exponential back-off between attempts. *N* in $(N * 2^{**attempts})$.

Return type *RetryPolicy*

baseplate.lib.secrets

Application integration with the secret fetcher daemon.

Fetcher Daemon

The secret fetcher is a sidecar that is run as a single daemon on each server. It can authenticate to Vault either as the server itself (through an AWS-signed instance identity document) or through a mounted JWT when running within a Kubernetes pod. It then gets access to secrets based upon the policies mapped to the role it authenticated as. Once authenticated, it fetches a given list of secrets from Vault and stores all of the data in a local file. It will automatically re-fetch secrets as their leases expire, ensuring that key rotation happens on schedule.

Because this is a sidecar, individual application processes don't need to talk directly to Vault for simple secret tokens (but can do so if needed for more complex operations like using the Transit backend). This reduces the load on Vault and adds a safety net if Vault becomes unavailable.

Secret Store

The secret store is the in-application integration with the file output of the fetcher daemon.

```
baseplate.lib.secrets.secrets_store_from_config(app_config, timeout=None, prefix='secrets.')
```

Configure and return a secrets store.

The keys useful to `secrets_store_from_config()` should be prefixed, e.g. `secrets.url`, etc.

Supported keys:

`path`: the path to the secrets file generated by the secrets fetcher daemon.

Parameters

- **app_config** (`Dict[str, str]`) – The application configuration which should have settings for the secrets store.
- **timeout** (`Optional[int]`) – How long, in seconds, to block instantiation waiting for the secrets data to become available (defaults to not blocking).
- **prefix** (`str`) – Specifies the prefix used to filter keys. Defaults to “secrets.”
- **backoff** – retry backoff time for secrets file watcher. Defaults to None, which is mapped to `DEFAULT_FILEWATCHER_BACKOFF`.

Return type `SecretsStore`

```
class baseplate.lib.secrets.SecretsStore(path, timeout=None, backoff=None)
```

Access to secret tokens with automatic refresh when changed.

This local vault allows access to the secrets cached on disk by the fetcher daemon. It will automatically reload the cache when it is changed. Do not cache or store the values returned by this class's methods but rather get them from this class each time you need them. The secrets are served from memory so there's little performance impact to doing so and you will be sure to always have the current version in the face of key rotation etc.

get_raw (`path`)

Return a dictionary of key/value pairs for the given secret path.

This is the raw representation of the secret in the underlying store.

Return type `Dict[str, str]`

get_credentials (*path*)

Decode and return a credential secret.

Credential secrets are a convention of username/password pairs stored as separate values in the raw secret payload.

The following keys are significant:

type This must always be `credential` for this method.

encoding This must be unset or set to `identity`.

username This contains the raw username.

password This contains the raw password.

Return type `CredentialSecret`

get_simple (*path*)

Decode and return a simple secret.

Simple secrets are a convention of key/value pairs in the raw secret payload. The following keys are significant:

type This must always be `simple` for this method.

value This contains the raw value of the secret token.

encoding (Optional) If present, how to decode the value from how it's encoded at rest (only `base64` currently supported).

Return type `bytes`

get_versioned (*path*)

Decode and return a versioned secret.

Versioned secrets are a convention of key/value pairs in the raw secret payload. The following keys are significant:

type This must always be `versioned` for this method.

current, next, and previous The raw secret value's versions. `current` is the "active" version, which is used for new creation/signing operations. `previous` and `next` are only used for validation (e.g. checking signatures) to ensure continuity when keys rotate. Both `previous` and `next` are optional.

encoding (Optional) If present, how to decode the values from how they are encoded at rest (only `base64` currently supported).

Return type `VersionedSecret`

get_vault_url ()

Return the URL for accessing Vault directly.

See also:

The `baseplate.clients.hvac` module provides integration with HVAC, a Vault client.

Return type `str`

get_vault_token()

Return a Vault authentication token.

The token will have policies attached based on the current EC2 server's Vault role. This is only necessary if talking directly to Vault.

See also:

The `baseplate.clients.hvac` module provides integration with HVAC, a Vault client.

Return type `str`

make_object_for_context(name, span)

Return an object that can be added to the context object.

This allows the secret store to be used with `add_to_context()`:

```
secrets = SecretsStore("/var/local/secrets.json")
baseplate.add_to_context("secrets", secrets)
```

Return type `SecretsStore`

class baseplate.lib.secrets.VersionedSecret

A versioned secret.

Versioned secrets allow for seamless rotation of keys. When using the secret to generate tokens (e.g. signing a message) always use the `current` value. When validating tokens, check against all the versions in `all_versions`. This will allow keys to rotate smoothly even if not done instantly across all users of the secret.

previous

Alias for field number 0

current

Alias for field number 1

next

Alias for field number 2

all_versions

Return an iterator over the available versions of this secret.

Return type `Iterator[bytes]`

classmethod from_simple_secret(value)

Make a fake versioned secret from a single value.

This is a backwards compatibility shim for use with APIs that take versioned secrets. Try to use proper versioned secrets fetched from the secrets store instead.

Return type `VersionedSecret`

class baseplate.lib.secrets.CredentialSecret

A secret for storing username/password pairs.

Credential secrets allow us to store usernames and passwords together in a single secret. Note that they are not versioned since the general pattern for rotating credentials like this would be to generate a new username/password pair. This object has two properties:

username

Alias for field number 0

password

Alias for field number 1

Exceptions

exception `baseplate.lib.secrets.CorruptSecretError` (*path, message*)
Raised when the requested secret does not match the expected format.

exception `baseplate.lib.secrets.SecretNotFoundError` (*name*)
Raised when the requested secret is not in the local vault.

exception `baseplate.lib.secrets.SecretsNotAvailableError` (*inner*)
Raised when the secrets store was not accessible.

baseplate.lib.thrift_pool

A Thrift client connection pool.

Note: See `baseplate.clients.thrift.ThriftContextFactory` for a convenient way to integrate the pool with your application.

The pool lazily creates connections and maintains them in a pool. Individual connections have a maximum lifetime, after which they will be recycled.

A basic example of usage:

```
pool = thrift_pool_from_config(app_config, "example_service.")
with pool.connection() as protocol:
    client = ExampleService.Client(protocol)
    client.do_example_thing()
```

Configuration Parsing

`baseplate.lib.thrift_pool.thrift_pool_from_config` (*app_config, prefix, **kwargs*)
Make a ThriftConnectionPool from a configuration dictionary.

The keys useful to `thrift_pool_from_config()` should be prefixed, e.g. `example_service.endpoint` etc. The `prefix` argument specifies the prefix used to filter keys. Each key is mapped to a corresponding keyword argument on the `ThriftConnectionPool` constructor. Any keyword arguments given to this function will be also be passed through to the constructor. Keyword arguments take precedence over the configuration file.

Supported keys:

- **endpoint (required):** A `host:port` pair, e.g. `localhost:2014`, where the Thrift server can be found.
- **size:** The size of the connection pool.
- **max_age:** The oldest a connection can be before it's recycled and replaced with a new one. Written as a `Timespan()` e.g. `1 minute`.
- **timeout:** The maximum amount of time a connection attempt or RPC call can take before a `TimeoutError` is raised. (`Timespan()`)

- **max_connection_attempts:** The maximum number of times the pool will attempt to open a connection.

Changed in version 1.2: `max_retries` was renamed `max_connection_attempts`.

Return type `ThriftConnectionPool`

Classes

```
class baseplate.lib.thrift_pool.ThriftConnectionPool (endpoint, size=10,
                                                    max_age=120, timeout=1,
                                                    max_connection_attempts=None,
                                                    max_retries=None, proto-
                                                    col_factory=<thrift.protocol.THeaderProtocol.THeader
                                                    object>)
```

A pool that maintains a queue of open Thrift connections.

Parameters

- **endpoint** (`EndpointConfiguration`) – The remote address of the Thrift service.
- **size** (`int`) – The maximum number of connections that can be open before new attempts to open block.
- **max_age** (`int`) – The maximum number of seconds a connection should be kept alive. Connections older than this will be reaped.
- **timeout** (`int`) – The maximum number of seconds a connection attempt or RPC call can take before a `TimeoutError` is raised.
- **max_connection_attempts** (`Optional[int]`) – The maximum number of times the pool will attempt to open a connection.
- **protocol_factory** (`TProtocolFactory`) – The factory to use for creating protocols from transports. This is useful for talking to services that don't support `THeaderProtocol`.

All exceptions raised by this class derive from `TTransportException`.

Changed in version 1.2: `max_retries` was renamed `max_connection_attempts`.

connection()

Acquire a connection from the pool.

This method is to be used with a context manager. It returns a connection from the pool, or blocks up to `timeout` seconds waiting for one if the pool is full and all connections are in use.

When the context is exited, the connection is returned to the pool. However, if it was exited via an unexpected Thrift exception, the connection is closed instead because the state of the connection is unknown.

Return type `Generator[TProtocolBase, None, None]`

baseplate.lib.service_discovery

Integration with Synapse's `file_output` service discovery method.

Note: Production Baseplate services have Synapse hooked up to a local HAProxy instance which will automatically route connections to services for you if you connect to the correct address/port on localhost. That is the preferred method of connecting to services.

The contents of this module are useful for inspecting the service inventory directly for cases where a blind TCP connection is insufficient (e.g. to give service addresses to a client, or for topology-aware clients like Cassandra).

A basic example of usage:

```
inventory = ServiceInventory("/var/lib/synapse/example.json")
backend = inventory.get_backend()
print(backend.endpoint.address)
```

class `baseplate.lib.service_discovery.ServiceInventory` (*filename*)

The inventory enumerates available backends for a single service.

Parameters `filename` (`str`) – The absolute path to the Synapse-generated inventory file in JSON format.

get_backends ()

Return a list of all available backends in the inventory.

If the inventory file becomes unavailable, the previously seen inventory is returned.

Return type `Sequence[Backend]`

get_backend ()

Return a randomly chosen backend from the available backends.

If weights are specified in the inventory, they will be respected when making the random selection.

Raises `NoBackendsAvailableError` if the inventory has no available endpoints.

Return type `Backend`

class `baseplate.lib.service_discovery.Backend`

A description of a service backend.

This is a tuple of several values:

id A unique integer ID identifying the backend.

name The name of the backend.

endpoint An `EndpointConfiguration` object describing the network address of the backend.

weight An integer weight indicating how much to prefer this backend when choosing whom to connect to.

Exceptions

exception `baseplate.lib.service_discovery.NoBackendsAvailableError`

Raised when no backends are available for this service.

1.3 CLI Tools

1.3.1 baseplate-healthcheck

Baseplate services have well-defined health-check endpoints. The `baseplate-healthcheck` tool connects to a given service and checks these endpoints to see if they're alive.

Command Line

There are two required arguments on the command line: the protocol of the service to check (`thrift` or `wsgi`) and the endpoint to connect to.

For example, to check a Thrift-based service listening on port 9090:

```
$ baseplate-healthcheck thrift 127.0.0.1:9090
```

or a WSGI (HTTP) service listening on a UNIX domain socket:

```
$ baseplate-healthcheck wsgi /run/myservice.sock
```

Results

If the service is healthy, the tool will exit with a status code indicating success (0) and print “OK!”. If the service is unhealthy, the tool will exit with a status code indicating failure (1) and print an error message explaining what went wrong.

Usage

This script can be used as part of a process to validate a server after creation, or to check service liveliness for a service discovery system.

1.3.2 baseplate-serve

Baseplate comes with a simple Gevent-based server for both Thrift and WSGI applications called `baseplate-serve`.

Configuration

There is one required parameter on the command line, the path to an INI-format configuration file. There should be two sections in the file: the `server` section and the `app` section. The section headers look like `server:main` or `app:main` where the part before the `:` is the type of section and the part after is the “name”. Baseplate looks for sections named `main` by default but can be overridden with the `--server-name` and `--app-name` options.

The Server

Here’s an example of a `server` section:

```
[server:main]
factory = baseplate.server.thrift
stop_timeout = 30
```

The `factory` tells baseplate what code to use to run the server. Baseplate comes with two servers built in:

`baseplate.server.thrift` A Gevent Thrift server.

`baseplate.server.wsgi` A Gevent WSGI server.

Both take two configuration values as well:

max_concurrency The maximum number of simultaneous clients the server will handle. Note that this is how many connections will be accepted, but some of those connections may be idle at any given time.

stop_timeout (Optional) How long, in seconds, to wait for active connections to finish up gracefully when shutting down. By default, the server will shut down immediately.

The WSGI server takes an additional optional parameter:

handler A full name of a class which subclasses `gevent.pywsgi.WSGIHandler` for extra functionality.

There are some additional configuration settings in this section that start with a `monitoring` prefix. For more information on those, see [Process-level metrics](#).

The Application

And now the real bread and butter, your `app` section:

```
[app:main]
factory = my_app.processor:make_processor
foo = 3
bar = 22
noodles.blah = one, two, three
```

The `app` section also takes a `factory`. This should be the name of a callable in your code which builds and returns your application. The part before the `:` is a Python module. The part after the `:` is the name of a callable object within that module.

The rest of the options in the `app` section of the configuration file get passed as a dictionary to your application callable. You can parse these options with [baseplate.lib.config](#).

The application factory should return an appropriate object for your server:

Thrift A `TProcessor`.

WSGI A WSGI callable.

Logging

The baseplate server provides a default configuration for the Python standard `logging` system. The root logger will print to `stdout` with a format that includes trace information. The default log level is `INFO` or `DEBUG` if the `--debug` flag is passed to `baseplate-serve`.

If more complex logging configuration is necessary, the configuration file will override the default setup. The [Configuration file format](#) is documented in the standard library.

Automatic reload on source changes

In development, it's useful for the server to restart itself when you change code. You can do this by passing the `--reload` flag to `baseplate-serve`.

This should not be used in production environments.

Einhorn

`baseplate-serve` can run as a worker in [Stripe's Einhorn socket manager](#). This allows Einhorn to handle binding the socket, worker management, rolling restarts, and worker health checks.

Baseplate supports Einhorn’s “manual ACK” protocol. Once the application is loaded and ready to serve, Baseplate notifies the Einhorn master process via its command socket.

An example command line:

```
einhorn -m manual -n 4 --bind localhost:9190 \
  baseplate-serve myapp.ini
```

Debug Signal

Applications running under `baseplate-serve` will respond to `SIGUSR1` by printing a stack trace to the logger. This can be useful for debugging deadlocks and other issues.

Note that Einhorn will exit if you send it a `SIGUSR1`. You can instead open up `einhornsh` and instruct the master to send the signal to all workers:

```
$ einhornsh
> signal SIGUSR1
Successfully sent USRs to 4 processes: [...]
```

Process-level metrics

If your application has registered a metrics client with `configure_metrics()`, `baseplate-serve` will automatically send process-level metrics every 10 seconds. Which metrics are sent depends on your server configuration, for example:

```
[server:main]
factory = baseplate.server.thrift

monitoring.blocked_hub = 100 milliseconds
monitoring.concurrency = true
```

will enable the `blocked_hub` reporter (configuring it to trigger at a 100ms threshold) and the `concurrency` reporter (which has no special configuration).

The following reporters are available:

monitoring.blocked_hub Enabled if a valid `Timespan()` is set, defaults to disabled.

This will turn on `Gevent’s monitoring thread` and report events indicating that `Gevent` detects the main event loop was blocked by a `greenlet` for longer than the given time span. This can indicate excessive CPU usage causing event loop starvation or the use of non-patched blocking IO calls. More detailed information, including stack traces, is also printed to the logging system.

Each instance of the hub being blocked will be reported as a `Timer` measuring the duration of the blockage.

Note: the performance impact of this reporter is not currently understood. Watch your metrics closely if you turn this on.

monitoring.concurrency Enabled if `true`, disabled if `false`. Defaults to enabled.

This will track the number of in-flight requests being processed concurrently by this server process.

At each report interval, this will update two `Gauge` metrics with the current number of open connections (`open_connections`) and current number of in-flight requests being processed concurrently (`active_requests`).

monitoring.connection_pool Enabled if `true`, disabled if `false`. Defaults to disabled.

This will track the usage of connection pools for various clients in the application. The metrics generated will depend on which clients are used.

monitoring.gc.stats Enabled if `true`, disabled if `false`. Defaults to enabled.

This will report the Python garbage collector's statistics to the metrics system.

At each report interval, this will update gauges with the current values returned by `gc.get_stats()`.

monitoring.gc.timing Enabled if `true`, disabled if `false`. Defaults to disabled.

This will track the duration of time taken by Python's garbage collector doing a collection sweep.

The duration of each pass of the garbage collector will be reported as a timer.

Note: the performance impact of this reporter is not currently understood. Watch your metrics closely if you turn this on.

monitoring.gc.refcycle Enabled if a path to a writable directory is set, defaults to disabled.

This should only be used in debugging, it will certainly have a negative performance impact.

This will turn off automatic garbage collection and instead run a sweep every reporting interval. Any objects found by the collector will be graphed using `objgraph` to help find reference cycles. The resulting graphs images will be written to the directory specified.

The `objgraph` library and `graphviz` package must be installed for this to work properly.

All metrics generated by `baseplate-serve` are prefixed with your application's configured namespace, followed by `runtime.{hostname}.PID{process_id}`.

1.3.3 baseplate-script

This command allows you to run a piece of Python code with the application configuration loaded similarly to `baseplate-serve`. The command is `baseplate-script`.

Command Line

There are two required arguments on the command line: the path to an INI-format configuration file, and the fully qualified name of a Python function to run.

The function should be specified as a module path, a colon, and a function name. For example, `my_service.models:create_schema`. The function should take a single argument which will be the application's configuration as a dictionary. This is the same as the application factory used by the server.

Just like with `baseplate-serve`, the `app:main` section will be loaded by default. This can be overridden with the `--app-name` option.

Example

Given a configuration file, `printer.ini`:

```
[app:main]
message = Hello!

[app:bizarro]
message = !olleH
```

and a small script, `printer.py`:

```
def run(app_config, args):
    parser = argparse.ArgumentParser()
    parser.add_argument("name")
    args = parser.parse_args(args)

    print(f"{app_config['message']} {args.name}")
```

You can run the script with various configurations:

```
$ baseplate-script printer.ini printer:run Goodbye.
Hello! Goodbye.

$ baseplate-script printer.ini --app-name=bizarro printer:run Goodbye.
!olleH Goodbye.
```

1.3.4 baseplate-shell

This command allows you to run an interactive Python shell for Baseplate.py services with the application configuration and context loaded. The command is `baseplate-shell`.

This shell can be used for any kind of Baseplate.py service: Thrift, HTTP, etc.

Command Line

This command requires the path to an INI-format configuration file to run.

Just like with `baseplate-serve`, the `app:main` section will be loaded by default. This can be overridden with the `--app-name` option.

By default, the shell will have variables containing the application and the context exposed. Additional variables can be exposed by providing a `setup` function in the `shell` (or `tshell` for backwards compatibility) section of the configuration file.

Example

Given a configuration file, `example.ini`:

```
[app:main]
factory = baseplate.server.thrift

[shell]
setup = my_service:shell_setup
```

and a small setup function, `my_service.py`:

```
def shell_setup(env, env_banner):
    from my_service import models
    env['models'] = models
    env_banner['models'] = 'Models module'
```

You can begin a shell with the `models` module exposed:

```
$ baseplate-shell example.ini
Baseplate Interactive Shell
Python 2.7.6 (default, Nov 23 2017, 15:49:48)
[GCC 4.8.4]

Available Objects:

  app          This project's app instance
  context      The context for this shell instance's span
  models       Models module
>>>
```

1.4 Linters

Incorporating linters into your service will enforce a coding standard and prevent errors from getting merged into your codebase. The `baseplate.lint` module consists of custom [Pylint](#) checkers which add more lint to Pylint. These lints are based on bugs found at [Reddit](#).

1.4.1 Configuration

Getting Started

Install [Pylint](#) and ensure you have it and its dependencies added to your `requirements-dev.txt` file.

Follow the [Pylint user guide](#) for instructions to generate a default `pylintrc` configuration file and run Pylint.

Adding Custom Checkers

In your `pylintrc` file, add `baseplate.lint` to the `[MASTER]` `load-plugins` configuration.

```
# List of plugins (as comma separated values of python modules names) to load,
# usually to register additional checkers.
load-plugins=baseplate.lint
```

This will allow you to use all the custom checkers in the `baseplate.lint` module when you run Pylint.

Custom Checkers List

- W9000: no-database-query-string-format

Creating Custom Checkers

If there is something you want to lint and a checker does not already exist, you can add a new one to `baseplate.lint`.

The following is an example checker you can reference to create your own.


```

# Pylint documentation for writing a checker: http://pylint.pycqa.org/en/latest/how\_tos/custom\_checkers.html
# This is an example of a Pylint AST checker and should not be registered to use
# In an AST (abstract syntax tree) checker, the code will be represented as nodes of
# a tree
# We will use the astroid library: https://astroid.readthedocs.io/en/latest/api/general.html to visit and leave nodes
# Libraries needed for an AST checker
from astroid import nodes
from pylint.checkers import BaseChecker
from pylint.interfaces import IAstroidChecker
from pylint.lint import PyLinter

# Basic example of a Pylint AST (abstract syntax tree) checker
# Checks for variables that have been reassigned in a function. If it finds a
# reassigned variable, it will throw an error
class NoReassignmentChecker(BaseChecker):
    __implements__ = IAstroidChecker

    # Checker name
    name = "no-reassigned-variable"
    # Set priority to -1
    priority = -1
    # Message dictionary
    msgs = {
        # message-id, consists of a letter and numbers
        # Letter will be one of following letters (C=Convention, W=Warning, E=Error,
        # F=Fatal, R=Refactoring)
        # Numbers need to be unique and in-between 9000-9999
        # Check https://baseplate.readthedocs.io/en/stable/linters/index.html#custom-checkers-list
        # for numbers that are already in use
        "W9001": (
            # displayed-message shown to user
            "Reassigned variable found.",
            # message-symbol used as alias for message-id
            "reassigned-variable",
            # message-help shown to user when calling pylint --help-msg
            "Ensure variables are not reassigned.",
        )
    }

    def __init__(self, linter: PyLinter = None):
        super().__init__(linter)
        self.variables: set = set()

    # The following two methods are called for us by pylint/astroid
    # The linter walks through the tree, visiting and leaving desired nodes
    # Methods should start with visit_ or leave_ followed by lowercase class name of
    # nodes
    # List of available nodes: https://astroid.readthedocs.io/en/latest/api/astroid/nodes.html

    # Visit the Assign node: https://astroid.readthedocs.io/en/latest/api/astroid/nodes.html#astroid.nodes.Assign
    def visit_assign(self, node: nodes) -> None:

```

(continues on next page)

(continued from previous page)

```

    for variable in node.targets:
        if variable.name not in self.variables:
            self.variables.add(variable.name)
        else:
            self.add_message("non-unique-variable", node=node)

    # Leave the FunctionDef node: https://astroid.readthedocs.io/en/latest/api/
    ↪ astroid.nodes.html#astroid.nodes.FunctionDef
    def leave_functiondef(self, node: nodes) -> nodes:
        self.variables = set()
        return node

```

Add a test to the baseplate test suite following this example checker test.

```

# Libraries needed for tests
import astroid
import pylint.testutils

from baseplate.lint import example_plugin

# CheckerTestCase creates a linter that will traverse the AST tree
class TestNoReassignmentChecker(pylint.testutils.CheckerTestCase):
    CHECKER_CLASS = example_plugin.NoReassignmentChecker

    # Use astroid.extract_node() to create a test case
    # Where you put #@ is where the variable gets assigned
    # example, assign_node_a = test = 1, assign_node_b = test = 2
    def test_finds_reassigned_variable(self):
        assign_node_a, assign_node_b = astroid.extract_node(
            """
            test = 1 #@
            test = 2 #@
            """
        )

        self.checker.visit_assign(assign_node_a)
        self.checker.visit_assign(assign_node_b)
        self.assertAddsMessages(
            pylint.testutils.Message(msg_id="reassigned-variable", node=assign_node_a)
        )

    def test_ignores_no_reassigned_variable(self):
        assign_node_a, assign_node_b = astroid.extract_node(
            """
            test1 = 1 #@
            test2 = 2 #@
            """
        )

        with self.assertNoMessages():
            self.checker.visit_assign(assign_node_a)
            self.checker.visit_assign(assign_node_b)

    def test_ignores_variable_outside_function(self):
        func_node, assign_node_a, assign_node_b = astroid.extract_node(
            """

```

(continues on next page)

(continued from previous page)

```
def test1(): #@
    test = 1 #@

def test2():
    test = 2 #@
    """

)

with self.assertNoMessages():
    self.checker.visit_assign(assign_node_a)
    self.checker.leave_functiondef(func_node)
    self.checker.visit_assign(assign_node_b)
```

Register your checker by adding it to the register() function:

```
from pylint.lint import PyLinter

from baseplate.lint.db_query_string_format_plugin import NoDbQueryStringFormatChecker

def register(linter: PyLinter) -> None:
    checker = NoDbQueryStringFormatChecker(linter)
    linter.register_checker(checker)
```

Lastly, add your checker message-id and name to *Custom Checkers List*.

CHAPTER 2

Appendix

- `genindex`
- `modindex`

b

`baseplate`, 13

`baseplate.clients`, 22

`baseplate.clients.cassandra`, 22

`baseplate.clients.hvac`, 24

`baseplate.clients.kombu`, 26

`baseplate.clients.memcache`, 29

`baseplate.clients.memcache.lib`, 31

`baseplate.clients.redis`, 32

`baseplate.clients.sqlalchemy`, 35

`baseplate.clients.thrift`, 38

`baseplate.frameworks`, 41

`baseplate.frameworks.pyramid`, 42

`baseplate.frameworks.queue_consumer`, 43

`baseplate.frameworks.queue_consumer.deprecated`, 51

`baseplate.frameworks.queue_consumer.kafka`, 44

`baseplate.frameworks.queue_consumer.kombu`, 48

`baseplate.frameworks.thrift`, 41

`baseplate.lib.config`, 59

`baseplate.lib.crypto`, 64

`baseplate.lib.datetime`, 65

`baseplate.lib.edge_context`, 66

`baseplate.lib.events`, 69

`baseplate.lib.experiments`, 72

`baseplate.lib.file_watcher`, 74

`baseplate.lib.message_queue`, 79

`baseplate.lib.metrics`, 80

`baseplate.lib.random`, 84

`baseplate.lib.ratelimit`, 84

`baseplate.lib.ratelimit.backends.memcache`, 86

`baseplate.lib.ratelimit.backends.redis`, 87

`baseplate.lib.retry`, 88

`baseplate.lib.secrets`, 89

`baseplate.lib.service_discovery`, 93

`baseplate.lib.thrift_pool`, 92

`baseplate.observers`, 54

Symbols

- `__init__()` (*baseplate.Baseplate* method), 13
 - `__init__()` (*baseplate.frameworks.queue_consumer.kafka.FastConsumerFactory* method), 47
 - `__init__()` (*baseplate.frameworks.queue_consumer.kafka.InOrderConsumerFactory* method), 45
 - `__init__()` (*baseplate.frameworks.queue_consumer.kombu.KombuQueueConsumerFactory* method), 49
 - `__iter__()` (*baseplate.lib.retry.RetryPolicy* method), 88
- ## A
- `add_sample()` (*baseplate.lib.metrics.Histogram* method), 84
 - `add_to_context()` (*baseplate.Baseplate* method), 15
 - `all_versions` (*baseplate.lib.secrets.VersionedSecret* attribute), 91
 - `attach_context()` (*baseplate.lib.edge_context.EdgeRequestContext* method), 67
 - `authentication_token` (*baseplate.lib.edge_context.OAuthClient* attribute), 68
 - `authentication_token` (*baseplate.lib.edge_context.Service* attribute), 69
 - `authentication_token` (*baseplate.lib.edge_context.User* attribute), 67
 - `AuthenticationToken` (class in *baseplate.lib.edge_context*), 69
- ## B
- `Backend` (class in *baseplate.lib.service_discovery*), 94
 - `Base64()` (in module *baseplate.lib.config*), 61
 - `BaseKombuConsumer` (class in *baseplate.frameworks.queue_consumer.deprecated*), 53
 - `Baseplate` (class in *baseplate*), 13
 - baseplate* (module), 13
 - baseplate.clients* (module), 22
 - baseplate.clients.cassandra* (module), 22
 - baseplate.clients.hvac* (module), 24
 - baseplate.clients.kombu* (module), 26
 - baseplate.clients.memcache* (module), 29
 - baseplate.clients.memcache.lib* (module), 31
 - baseplate.clients.redis* (module), 32
 - baseplate.clients.sqlalchemy* (module), 35
 - baseplate.clients.thrift* (module), 38
 - baseplate.frameworks* (module), 41
 - baseplate.frameworks.pyramid* (module), 42
 - baseplate.frameworks.queue_consumer* (module), 43
 - baseplate.frameworks.queue_consumer.deprecated* (module), 51
 - baseplate.frameworks.queue_consumer.kafka* (module), 44
 - baseplate.frameworks.queue_consumer.kombu* (module), 48
 - baseplate.frameworks.thrift* (module), 41
 - baseplate.lib.config* (module), 59
 - baseplate.lib.crypto* (module), 64
 - baseplate.lib.datetime* (module), 65
 - baseplate.lib.edge_context* (module), 66
 - baseplate.lib.events* (module), 69
 - baseplate.lib.experiments* (module), 72
 - baseplate.lib.file_watcher* (module), 74
 - baseplate.lib.message_queue* (module), 79
 - baseplate.lib.metrics* (module), 80
 - baseplate.lib.random* (module), 84
 - baseplate.lib.ratelimit* (module), 84
 - baseplate.lib.ratelimit.backends.memcache* (module), 86
 - baseplate.lib.ratelimit.backends.redis* (module), 87
 - baseplate.lib.retry* (module), 88
 - baseplate.lib.secrets* (module), 89

- baseplate.lib.service_discovery (module), 93
- baseplate.lib.thrift_pool (module), 92
- baseplate.observers (module), 54
- BaseplateConfigurator (class in baseplate.frameworks.pyramid), 42
- baseplateify_processor() (in module baseplate.frameworks.thrift), 41
- BaseplateObserver (class in baseplate), 20
- Batch (class in baseplate.lib.metrics), 81
- batch() (baseplate.lib.metrics.Client method), 81
- Boolean() (in module baseplate.lib.config), 61
- ## C
- CassandraClient (class in baseplate.clients.cassandra), 23
- CassandraContextFactory (class in baseplate.clients.cassandra), 24
- Client (class in baseplate.lib.metrics), 81
- close() (baseplate.clients.redis.MessageQueue method), 35
- close() (baseplate.lib.message_queue.MessageQueue method), 80
- cluster_from_config() (in module baseplate.clients.cassandra), 23
- ConfigurationError, 64
- configure_context() (baseplate.Baseplate method), 14
- configure_error_reporting() (baseplate.Baseplate method), 22
- configure_logging() (baseplate.Baseplate method), 21
- configure_metrics() (baseplate.Baseplate method), 21
- configure_observers() (baseplate.Baseplate method), 14
- configure_tracing() (baseplate.Baseplate method), 21
- connection() (baseplate.lib.thrift_pool.ThriftConnectionPool method), 93
- connection_from_config() (in module baseplate.clients.kombu), 28
- consume() (baseplate.lib.ratelimit.backends.memcache.MemcacheRateLimitBackend method), 86
- consume() (baseplate.lib.ratelimit.backends.RateLimitBackend method), 86
- consume() (baseplate.lib.ratelimit.backends.redis.RedisRateLimitBackend method), 87
- consume() (baseplate.lib.ratelimit.RateLimiter method), 85
- consume() (in module baseplate.frameworks.queue_consumer.deprecated), 52
- ContextFactory (class in baseplate.clients), 40
- cookie_created_ms (baseplate.lib.edge_context.User attribute), 68
- CorruptSecretError, 92
- Counter (class in baseplate.lib.metrics), 82
- counter() (baseplate.lib.metrics.Batch method), 82
- counter() (baseplate.lib.metrics.Client method), 81
- CQLMapperClient (class in baseplate.clients.cassandra), 23
- CQLMapperContextFactory (class in baseplate.clients.cassandra), 24
- CredentialSecret (class in baseplate.lib.secrets), 91
- current (baseplate.lib.secrets.VersionedSecret attribute), 91
- ## D
- datetime_to_epoch_milliseconds() (in module baseplate.lib.datetime), 65
- datetime_to_epoch_seconds() (in module baseplate.lib.datetime), 65
- decompress_and_load() (in module baseplate.clients.memcache.lib), 31
- decompress_and_unpickle() (in module baseplate.clients.memcache.lib), 31
- decrement() (baseplate.lib.metrics.Counter method), 82
- device (baseplate.lib.edge_context.EdgeRequestContext attribute), 67
- DictOf() (in module baseplate.lib.config), 62
- ## E
- EdgeRequestContext (class in baseplate.lib.edge_context), 67
- EdgeRequestContextFactory (class in baseplate.lib.edge_context), 66
- Endpoint() (in module baseplate.lib.config), 61
- EndpointConfiguration (class in baseplate.lib.config), 63
- engine_from_config() (in module baseplate.clients.sqlalchemy), 36
- epoch_seconds_to_datetime() (in module baseplate.lib.datetime), 65
- event_fields() (baseplate.lib.edge_context.EdgeRequestContext method), 67
- event_fields() (baseplate.lib.edge_context.OAuthClient method), 69
- event_fields() (baseplate.lib.edge_context.User method), 68
- EventError, 71
- EventQueue (class in baseplate.lib.events), 70
- EventQueueFullError, 71

EventTooLargeError, 71
 exchange_from_config() (in module baseplate.clients.kombu), 28
 execute_command() (baseplate.clients.redis.MonitoredRedisConnection method), 34
 Experiments (class in baseplate.lib.experiments), 73
 experiments_client_from_config() (in module baseplate.lib.experiments), 72
 ExperimentsContextFactory (class in baseplate.lib.experiments), 73
 ExpiredSignatureError, 65
 expose() (baseplate.lib.experiments.Experiments method), 74

F

Fallback() (in module baseplate.lib.config), 62
 FastConsumerFactory (class in baseplate.frameworks.queue_consumer.kafka), 46
 FatalMessageHandlerError (class in baseplate.frameworks.queue_consumer.kombu), 50
 File() (in module baseplate.lib.config), 61
 FileWatcher (class in baseplate.lib.file_watcher), 75
 finish() (baseplate.ServerSpan method), 16
 finish() (baseplate.Span method), 18
 flags (baseplate.TraceInfo attribute), 19
 Float() (in module baseplate.lib.config), 60
 flush() (baseplate.lib.metrics.Batch method), 82
 from_simple_secret() (baseplate.lib.secrets.VersionedSecret class method), 91
 from_upstream() (baseplate.lib.edge_context.EdgeRequestContextFactory method), 67
 from_upstream() (baseplate.TraceInfo class method), 19

G

Gauge (class in baseplate.lib.metrics), 83
 gauge() (baseplate.lib.metrics.Batch method), 82
 gauge() (baseplate.lib.metrics.Client method), 81
 geolocation (baseplate.lib.edge_context.EdgeRequestContext attribute), 67
 get() (baseplate.clients.redis.MessageQueue method), 35
 get() (baseplate.lib.message_queue.MessageQueue method), 79
 get_all_experiment_names() (baseplate.lib.experiments.Experiments method), 73

get_backend() (baseplate.lib.service_discovery.ServiceInventory method), 94
 get_backends() (baseplate.lib.service_discovery.ServiceInventory method), 94
 get_batch() (baseplate.frameworks.queue_consumer.deprecated.BaseKombuConsumer method), 54
 get_batch() (baseplate.frameworks.queue_consumer.deprecated.KombuConsumer method), 53
 get_credentials() (baseplate.lib.secrets.SecretsStore method), 89
 get_data() (baseplate.lib.file_watcher.FileWatcher method), 75
 get_message() (baseplate.frameworks.queue_consumer.deprecated.BaseKombuConsumer method), 54
 get_message() (baseplate.frameworks.queue_consumer.deprecated.KombuConsumer method), 53
 get_raw() (baseplate.lib.secrets.SecretsStore method), 89
 get_simple() (baseplate.lib.secrets.SecretsStore method), 90
 get_utc_now() (in module baseplate.lib.datetime), 66
 get_vault_token() (baseplate.lib.secrets.SecretsStore method), 90
 get_vault_url() (baseplate.lib.secrets.SecretsStore method), 90
 get_versioned() (baseplate.lib.secrets.SecretsStore method), 90

H

has_role() (baseplate.lib.edge_context.User method), 68
 HeaderTrustHandler (class in baseplate.frameworks.pyramid), 42
 Histogram (class in baseplate.lib.metrics), 83
 histogram() (baseplate.lib.metrics.Batch method), 82
 histogram() (baseplate.lib.metrics.Client method), 81
 hvac_factory_from_config() (in module baseplate.clients.hvac), 25
 HvacClient (class in baseplate.clients.hvac), 25
 HvacContextFactory (class in baseplate.clients.hvac), 25

I

id (baseplate.lib.edge_context.OAuthClient attribute), 68
 id (baseplate.lib.edge_context.Session attribute), 69
 id (baseplate.lib.edge_context.User attribute), 68

`IncorrectSignatureError`, 65
`incr_tag()` (*baseplate.ServerSpan* method), 16
`incr_tag()` (*baseplate.Span* method), 18
`increment()` (*baseplate.lib.metrics.Counter* method), 82
`InOrderConsumerFactory` (class in *baseplate.frameworks.queue_consumer.kafka*), 44
`Integer()` (in module *baseplate.lib.config*), 60
`is_logged_in` (*baseplate.lib.edge_context.User* attribute), 68
`is_type()` (*baseplate.lib.edge_context.OAuthClient* method), 68
`is_valid_experiment()` (*baseplate.lib.experiments.Experiments* method), 73

K

`KombuConsumer` (class in *baseplate.frameworks.queue_consumer.deprecated*), 53
`KombuProducer` (class in *baseplate.clients.kombu*), 27
`KombuProducerContextFactory` (class in *baseplate.clients.kombu*), 28
`KombuQueueConsumerFactory` (class in *baseplate.frameworks.queue_consumer.kombu*), 48
`KombuSerializer` (class in *baseplate.clients.kombu*), 27
`KombuThriftSerializer` (class in *baseplate.clients.kombu*), 27

L

`lock()` (*baseplate.clients.redis.MonitoredRedisConnection* method), 34
`log()` (*baseplate.ServerSpan* method), 16
`log()` (*baseplate.Span* method), 18
`loid` (*baseplate.lib.edge_context.User* attribute), 68

M

`make_child()` (*baseplate.ServerSpan* method), 16
`make_child()` (*baseplate.Span* method), 18
`make_context_object()` (*baseplate.Baseplate* method), 15
`make_dump_and_compress_fn()` (in module *baseplate.clients.memcache.lib*), 31
`make_object_for_context()` (*baseplate.clients.ContextFactory* method), 40
`make_object_for_context()` (*baseplate.clients.hvac.HvacContextFactory* method), 26
`make_object_for_context()` (*baseplate.clients.memcache.MemcacheContextFactory* method), 30

`make_object_for_context()` (*baseplate.clients.redis.RedisContextFactory* method), 34
`make_object_for_context()` (*baseplate.lib.ratelimit.backends.memcache.MemcacheRateLimitBackend* method), 86
`make_object_for_context()` (*baseplate.lib.ratelimit.backends.redis.RedisRateLimitBackendContextFactory* method), 87
`make_object_for_context()` (*baseplate.lib.ratelimit.RateLimiterContextFactory* method), 85
`make_object_for_context()` (*baseplate.lib.secrets.SecretsStore* method), 91
`make_pickle_and_compress_fn()` (in module *baseplate.clients.memcache.lib*), 32
`make_server_span()` (*baseplate.Baseplate* method), 15
`make_signature()` (in module *baseplate.lib.crypto*), 64
`MemcacheClient` (class in *baseplate.clients.memcache*), 29
`MemcacheContextFactory` (class in *baseplate.clients.memcache*), 30
`MemcacheRateLimitBackend` (class in *baseplate.lib.ratelimit.backends.memcache*), 86
`MemcacheRateLimitBackendContextFactory` (class in *baseplate.lib.ratelimit.backends.memcache*), 86
`MessageQueue` (class in *baseplate.clients.redis*), 35
`MessageQueue` (class in *baseplate.lib.message_queue*), 79
`MessageQueueError`, 80
`metrics_client_from_config()` (in module *baseplate.lib.metrics*), 81
`MonitoredMemcacheConnection` (class in *baseplate.clients.memcache*), 31
`MonitoredRedisConnection` (class in *baseplate.clients.redis*), 34

N

`name` (*baseplate.lib.edge_context.Service* attribute), 69
`new()` (*baseplate.frameworks.queue_consumer.deprecated.BaseKombuConsumer* class method), 53
`new()` (*baseplate.frameworks.queue_consumer.deprecated.KombuConsumer* class method), 53
`new()` (*baseplate.frameworks.queue_consumer.kafka.FastConsumerFactory* class method), 46
`new()` (*baseplate.frameworks.queue_consumer.kafka.InOrderConsumerFactory* class method), 44
`new()` (*baseplate.frameworks.queue_consumer.kombu.KombuQueueConsumer* class method), 48
`new()` (*baseplate.lib.edge_context.EdgeRequestContextFactory* method), 66

- new () (*baseplate.lib.retry.RetryPolicy* static method), 88
 new () (*baseplate.TraceInfo* class method), 19
 next (*baseplate.lib.secrets.VersionedSecret* attribute), 91
 NoAuthenticationError, 69
 NoBackendsAvailableError, 94
- ## O
- oauth_client (*baseplate.lib.edge_context.EdgeRequestContext* attribute), 67
 OAuthClient (class in *baseplate.lib.edge_context*), 68
 on_child_span_created () (*baseplate.ServerSpanObserver* method), 20
 on_child_span_created () (*baseplate.SpanObserver* method), 21
 on_finish () (*baseplate.ServerSpanObserver* method), 20
 on_finish () (*baseplate.SpanObserver* method), 21
 on_incr_tag () (*baseplate.ServerSpanObserver* method), 20
 on_incr_tag () (*baseplate.SpanObserver* method), 21
 on_log () (*baseplate.ServerSpanObserver* method), 20
 on_log () (*baseplate.SpanObserver* method), 21
 on_server_span_created () (*baseplate.BaseplateObserver* method), 20
 on_set_tag () (*baseplate.ServerSpanObserver* method), 20
 on_set_tag () (*baseplate.SpanObserver* method), 21
 on_start () (*baseplate.ServerSpanObserver* method), 20
 on_start () (*baseplate.SpanObserver* method), 21
 OneOf () (in module *baseplate.lib.config*), 61
 Optional () (in module *baseplate.lib.config*), 62
 origin_service (*baseplate.lib.edge_context.EdgeRequestContext* attribute), 67
- ## P
- parent_id (*baseplate.TraceInfo* attribute), 18
 parse_config () (in module *baseplate.lib.config*), 60
 Parser (class in *baseplate.lib.config*), 64
 password (*baseplate.lib.secrets.CredentialSecret* attribute), 91
 Percent () (in module *baseplate.lib.config*), 61
 pick () (*baseplate.lib.random.WeightedLottery* method), 84
 pipeline () (*baseplate.clients.redis.MonitoredRedisConnection* method), 34
 pool_from_config () (in module *baseplate.clients.memcache*), 30
 pool_from_config () (in module *baseplate.clients.redis*), 33
 previous (*baseplate.lib.secrets.VersionedSecret* attribute), 91
 pubsub () (*baseplate.clients.redis.MonitoredRedisConnection* method), 35
 put () (*baseplate.clients.redis.MessageQueue* method), 35
 put () (*baseplate.lib.events.EventQueue* method), 71
 put () (*baseplate.lib.message_queue.MessageQueue* method), 79
- ## R
- RateLimitBackend (class in *baseplate.lib.ratelimit.backends*), 86
 RateLimiter (class in *baseplate.lib.ratelimit*), 85
 RateLimiterContextFactory (class in *baseplate.lib.ratelimit*), 85
 RateLimitExceededException (class in *baseplate.lib.ratelimit*), 86
 RedisClient (class in *baseplate.clients.redis*), 33
 RedisContextFactory (class in *baseplate.clients.redis*), 34
 RedisRateLimitBackend (class in *baseplate.lib.ratelimit.backends.redis*), 87
 RedisRateLimitBackendContextFactory (class in *baseplate.lib.ratelimit.backends.redis*), 87
 register () (*baseplate.Baseplate* method), 20
 register () (*baseplate.ServerSpan* method), 17
 register () (*baseplate.Span* method), 17
 register_serializer () (in module *baseplate.clients.kombu*), 27
 replace () (*baseplate.lib.metrics.Gauge* method), 83
 report_runtime_metrics () (*baseplate.clients.ContextFactory* method), 40
 RequestContext (class in *baseplate*), 15
 RetryPolicy (class in *baseplate.lib.retry*), 88
 roles (*baseplate.lib.edge_context.User* attribute), 68
- ## S
- sample () (*baseplate.lib.random.WeightedLottery* method), 84
 sampled (*baseplate.TraceInfo* attribute), 19
 SecretNotFoundError, 92
 secrets_store_from_config () (in module *baseplate.lib.secrets*), 89
 SecretsNotAvailableError, 92
 SecretsStore (class in *baseplate.lib.secrets*), 89
 send () (*baseplate.lib.metrics.Counter* method), 82
 set () (*baseplate.lib.metrics.Timer* method), 83
 serialize_v2_event () (in module *baseplate.lib.events*), 71
 server_context () (*baseplate.Baseplate* method), 15
 ServerSpan (class in *baseplate*), 16

ServerSpanInitialized (class in baseplate.frameworks.pyramid), 43
ServerSpanObserver (class in baseplate), 20
service (baseplate.lib.edge_context.EdgeRequestContext attribute), 67
Service (class in baseplate.lib.edge_context), 69
ServiceInventory (class in baseplate.lib.service_discovery), 94
session (baseplate.lib.edge_context.EdgeRequestContext attribute), 67
Session (class in baseplate.lib.edge_context), 69
set_tag() (baseplate.ServerSpan method), 17
set_tag() (baseplate.Span method), 18
should_trust_edge_context_payload() (baseplate.frameworks.pyramid.HeaderTrustHandler method), 42
should_trust_edge_context_payload() (baseplate.frameworks.pyramid.StaticTrustHandler method), 43
should_trust_trace_headers() (baseplate.frameworks.pyramid.HeaderTrustHandler method), 42
should_trust_trace_headers() (baseplate.frameworks.pyramid.StaticTrustHandler method), 43
SignatureError, 65
SignatureInfo (class in baseplate.lib.crypto), 65
Span (class in baseplate), 17
span_id (baseplate.TraceInfo attribute), 19
SpanObserver (class in baseplate), 20
SQLAlchemyEngineContextFactory (class in baseplate.clients.sqlalchemy), 37
SQLAlchemySession (class in baseplate.clients.sqlalchemy), 36
SQLAlchemySessionContextFactory (class in baseplate.clients.sqlalchemy), 37
start() (baseplate.lib.metrics.Timer method), 83
start() (baseplate.ServerSpan method), 17
start() (baseplate.Span method), 17
StaticTrustHandler (class in baseplate.frameworks.pyramid), 43
stop() (baseplate.lib.metrics.Timer method), 83
String() (in module baseplate.lib.config), 60
subject (baseplate.lib.edge_context.AuthenticationToken attribute), 69
TimedOutError, 80
Timer (class in baseplate.lib.metrics), 83
timer() (baseplate.lib.metrics.Batch method), 82
timer() (baseplate.lib.metrics.Client method), 81
Timespan() (in module baseplate.lib.config), 61
trace_id (baseplate.TraceInfo attribute), 18
TraceInfo (class in baseplate), 18
transaction() (baseplate.clients.redis.MonitoredRedisConnection method), 34
TupleOf() (in module baseplate.lib.config), 62
U
UnixGroup() (in module baseplate.lib.config), 61
UnixUser() (in module baseplate.lib.config), 61
unlink() (baseplate.clients.redis.MessageQueue method), 35
unlink() (baseplate.lib.message_queue.MessageQueue method), 80
UnreadableSignatureError, 65
user (baseplate.lib.edge_context.EdgeRequestContext attribute), 67
User (class in baseplate.lib.edge_context), 67
username (baseplate.lib.secrets.CredentialSecret attribute), 91
V
validate_signature() (in module baseplate.lib.crypto), 64
variant() (baseplate.lib.experiments.Experiments method), 73
VersionedSecret (class in baseplate.lib.secrets), 91
W
WatchedFileNotAvailableError, 75
WeightedLottery (class in baseplate.lib.random), 84
Y
yield_attempts() (baseplate.lib.retry.RetryPolicy method), 88
Z
zookeeper_client_from_config() (in module baseplate.lib.live_data), 77
T
thrift_pool_from_config() (in module baseplate.lib.thrift_pool), 92
ThriftClient (class in baseplate.clients.thrift), 39
ThriftConnectionPool (class in baseplate.lib.thrift_pool), 93
ThriftContextFactory (class in baseplate.clients.thrift), 39