
Gabbi Documentation

Chris Dent

Sep 23, 2022

Contents

1	Test Structure	3
2	Fixtures	7
3	Response Handlers	9
4	Substitution	11
5	Data	15
6	Loading and Running Tests	17
7	Example Tests	21
8	JSONPath	27
9	Target Host	31
10	Fixtures	33
11	Inner Fixtures	35
12	Content Handlers	37
13	YAML Runner	41
14	Release Notes	43
15	Contributors	51
16	Frequently Asked Questions	53
17	gabbi Package	55
18	Gabbi	61
	Python Module Index	63
	Index	65

Gabbi tests are expressed in YAML as a series of HTTP requests with their expected response:

```
tests:  
  - name: retrieve root  
    GET: /  
    status: 200
```

This will trigger a GET request to / on the configured *Target Host*. The test will pass if the response's status code is 200.

CHAPTER 1

Test Structure

The top-level `tests` category contains an ordered sequence of test declarations, each describing the expected response to a given request:

1.1 Metadata

Key	Description	Notes
name	The test's name. Must be unique within a file.	re- quired
desc	An arbitrary string describing the test.	
verbose	If <code>True</code> or <code>all</code> (synonymous), prints a representation of the current request and response to <code>stdout</code> , including both headers and body. If set to <code>headers</code> or <code>body</code> , only the corresponding part of the request and response will be printed. If the output is a TTY, colors will be used. If the body content-type is JSON it will be formatted for improved readability. See <code>VerboseHttp</code> for details.	de- faults to False
skip	A string message which if set will cause the test to be skipped with the provided message.	de- faults to False
xfail	Determines whether to expect this test to fail. Note that the test will be run anyway.	de- faults to False
use_prior	Determines if this test will be run in sequence (after) the test prior to it in the list of tests within a file. To be concrete, when this is <code>True</code> the test is dependent on the prior test and if that prior has not yet run, it will be run, even if only the current test has been selected. Set this to <code>False</code> to allow selecting a test without dependencies.	de- faults to True
cert_val	States whether the underlying HTTP client should attempt to validate SSL certificates. In many test environment certificates will be self-signed so changing this may be required. It can also be changed when <i>Loading and Running Tests</i> or using <i>gabbi-run</i> .	de- faults to True
timeout	Sets a timeout (in seconds) for the HTTP request.	de- faults to 30 sec- onds
disable_response_handler	If <code>True</code> , means that the response body will not be processed to Python data. This can be necessary if a response claims a <code>content-type</code> but the body is not actually that type but it is still necessary to run tests against the response. In that situation, if <code>disable_response_handler</code> is <code>False</code> the test will be treated as a failure.	de- faults to False

Note: When tests are generated dynamically, the `TestCase` name will include the respective test's name, lowercased with spaces transformed to `_`. In at least some test runners this will allow you to select and filter on test name.

1.2 Request Parameters

Key	Description	Notes
any uppercase string	Any such key is considered an HTTP method, with the corresponding value expressing the URL. This is a shortcut combining <code>method</code> and <code>url</code> into a single statement: <code>GET: /index</code> corresponds to: <code>method: GET</code> <code>url: /index</code>	
<code>method</code>	The HTTP request method.	defaults to <code>GET</code>
<code>url</code>	The URL to request. This can either be a full path (e.g. <code>/index</code>) or a fully qualified URL (i.e. including host and scheme, e.g. <code>http://example.org/index</code>) — see <i>Target Host</i> for details.	Either this or the shortcut above is required
<code>request_headers</code>	A dictionary of key-value pairs representing request header names and values. These will be added to the constructed request.	
<code>query_parameters</code>	A dictionary of query parameters that will be added to the <code>url</code> as query string. If that URL already contains a set of query parameters, those will be extended. See <i>Example Tests</i> for a demonstration of how the data is structured.	
<code>data</code>	A representation to pass as the body of a request. Note that <code>content-type</code> in <code>request_headers</code> should also be set — see <i>Data</i> for details.	
<code>redirects</code>	If <code>True</code> , redirects will automatically be followed.	defaults to <code>False</code>
<code>ssl</code>	Determines whether the request uses SSL (i.e. HTTPS). Note that the <code>url</code> 's scheme takes precedence if present — see <i>Target Host</i> for details.	defaults to <code>False</code>

1.3 Response Expectations

Key	Description	Notes
status	The expected response status code. Multiple acceptable response codes may be provided, separated by <code> </code> (e.g. <code>302 301</code> — note, however, that this indicates ambiguity, which is generally undesirable).	defaults to 200
response_headers	A dictionary of key-value pairs representing expected response header names and values. If a header's value is wrapped in <code>/.../</code> , it will be treated as a regular expression to search for in the response header.	
response_forbidden_headers	A list of headers which must <i>not</i> be present.	
response_strings	A list of string fragments expected to be present in the response body.	
response_json_paths	A dictionary of JSONPath rules paired with expected matches. Using this rule requires that the content being sent from the server is JSON (i.e. a content type of <code>application/json</code> or containing <code>+json</code>) If the value is wrapped in <code>/.../</code> the result of the JSONPath query will be searched for the value as a regular expression.	
poll	A dictionary of two keys: <ul style="list-style-type: none"> <code>count</code>: An integer stating the number of times to attempt this test before giving up. <code>delay</code>: A floating point number of seconds to delay between attempts. This makes it possible to poll for a resource created via an asynchronous request. Use with caution.	

Note that many of these items allow *substitutions*.

Default values for a file's `tests` may be provided via the top-level `defaults` category. These take precedence over the global defaults (explained below).

For examples see the [gabbi tests](#), [Example Tests](#) and the [gabbi-demo](#) tutorial.

CHAPTER 2

Fixtures

The top-level `fixtures` category contains a sequence of named *Fixtures*.

CHAPTER 3

Response Handlers

`response_*` keys are examples of Response Handlers. Custom handlers may be created by test authors for specific use cases. See *Content Handlers* for more information.

There are a number of magical variables that can be used to make reference to the state of a current test, the one just prior or any test prior to the current one. The variables are replaced with real values during test processing.

4.1 Global

- `$_ENVIRON['<environment variable>']`: The name of an environment variable. Its value will replace the magical variable. If the string value of the environment variable is "True" or "False" then the resulting value will be the corresponding boolean, not a string.

4.2 Current Test

- `$_SCHEME`: The current scheme/protocol (usually `http` or `https`).
- `$_NETLOC`: The host and potentially port of the request.

4.3 Immediately Prior Test

- `$_COOKIE`: All the cookies set by any `Set-Cookie` headers in the prior response, including only the cookie key and value pairs and no metadata (e.g. `expires` or `domain`).
- `$_URL`: The URL defined in the prior request, after substitutions have been made. For backwards compatibility with earlier releases `$_LAST_URL` may also be used, but if `$_HISTORY` (see below) is being used, `$_URL` must be used.
- `$_LOCATION`: The location header returned in the prior response.
- `$_HEADERS['<header>']`: The value of any header from the prior response.
- `$_RESPONSE['<json path>']`: A `JSONPath` query into the prior response. See [JSONPath](#) for more on formatting.

4.4 Any Previous Test

- `$HISTORY['<test name>'].<magical variable expression>`: Any variable which refers to a prior test may be used in an expression that refers to any earlier test in the same file by identifying the target test by its name in a `$HISTORY` dictionary. For example, to refer to a value in a JSON object in the response of a test named `post_json`:

```
$HISTORY['post_json'].$RESPONSE['$.key']
```

This is a very powerful feature that could lead to test that are difficult for humans to read. Take care to optimize for the maintainers that will come after you, not yourself.

4.5 Casting

For `$ENVIRON` and `$RESPONSE` it is possible to attempt to cast the value to another type: `int`, `float`, `str`, or `bool`. If the cast fails an exception will be raised and the test will fail.

This functionality only works when the magical variable is the whole value of a YAML entry. If the variable is intermixed with other data, an exception will be raised and the test will fail.

The format for a cast is to append a `:` and the cast type after the type of the magical variable. For example:

```
$RESPONSE:int['$.some_string_value']
```

Warning: Prior to the introduction of this feature, `$ENVIRON` would already do some automatic casting of numbers to ints and floats and the strings `True` and `False` to booleans. This continues to be the case, but only if no cast is provided.

Note: Where a single-quote character, `'`, is shown in the variables above you may also use a double-quote character, `"`, but in any given expression the same character must be used at both ends.

All of these variables may be used in all of the following fields:

- `skip`
- `url`
- `query_parameters`
- `data`
- `request_headers` (in both the key and value)
- `response_strings`
- `response_json_paths` (in both the key and value, see [json path substitution](#) for more info)
- `response_headers` (in both the key and value)
- `response_forbidden_headers`
- `count` and `delay` fields of `poll`

With these variables it ought to be possible to traverse an API without any explicit statements about the URLs being used. If you need a replacement on a field that is not currently supported please raise an issue or provide a patch.

As all of these features needed to be tested in the development of gabbi itself, the [gabbi tests](#) are a good source of examples on how to use the functionality. See also [Example Tests](#) for a collection of examples and the [gabbi-demo tutorial](#).

The `data` key has some special handing to allow for a bit more flexibility when doing a `POST` or `PUT`:

- If the value is not a string (that is, it is a sequence or structure) it is treated as a data structure that will be turned into a string by the `dumps` method on the relevant *content handler*. For example if the content-type of the body is `application/json` the data structure will be turned into a JSON string.
- If the value is a string that begins with `<@` then the rest of the string is treated as a filepath to be loaded. The path is relative to the test directory and may not traverse up into parent directories.
- If the value is an undecorated string, that's the value.

Note: When reading from a file care should be taken to ensure that a reasonable content-type is set for the data as this will control if any encoding is done of the resulting string value. If it is text, json, xml or javascript it will be encoded to UTF-8.

Loading and Running Tests

To run gabbi tests with a test harness they must be generated in some fashion and then run. This is accomplished by a test loader. Initially gabbi only supported those test harnesses that supported the `load_tests` protocol in `UnitTest`. It is now possible to also build and run tests with `pytest` with some limitations described below.

Note: It is also possible to run gabbi tests from the command line. See *YAML Runner*.

Note: By default gabbi will load YAML files using the `safe_load` function. This means only basic YAML types are allowed in the file. For most use cases this is fine. If you need custom types (for example, to match NaN) it is possible to set the `safe_yaml` parameter of `build_tests()` to `False`. If custom types are used, please keep in mind that this can limit the portability of the YAML files to other contexts.

Warning: If tests are being run with a runner that supports concurrency (such as `testrepository`) it is critical that the test runner is informed of how to group the tests into their respective suites. The usual way to do this is to use a regular expression that groups based on the name of the yml files. For example, when using `testrepository` the `.testr.conf` file needs an entry similar to the following:

```
group_regex=gabbi\.suitemaker\.(test_[^_]+\.[^_]+)
```

6.1 UnitTest Style Loader

To run the tests with a `load_tests` style loader a test file containing a `load_tests` method is required. That will look a bit like:

```
"""A sample test module."""
# For pathname munging
```

(continues on next page)

```
import os

# The module that build_tests comes from.
from gabbi import driver

# We need access to the WSGI application that hosts our service
from myapp import wsgiapp

# We're using fixtures in the YAML files, we need to know where to
# load them from.
from myapp.test import fixtures

# By convention the YAML files are put in a directory named
# "gabbits" that is in the same directory as the Python test file.
TESTS_DIR = 'gabbits'

def load_tests(loader, tests, pattern):
    """Provide a TestSuite to the discovery process."""
    test_dir = os.path.join(os.path.dirname(__file__), TESTS_DIR)
    # Pass "require_ssl=True" as an argument to force all tests
    # to use SSL in requests.
    return driver.build_tests(test_dir, loader,
                              intercept=wsgiapp.app,
                              fixture_module=fixtures)
```

For details on the arguments available when building tests see `build_tests()`.

Once the test loader has been created, it needs to be run. There are *many* options. Which is appropriate depends very much on your environment. Here are some examples using `unittest` or `testtools` that require minimal knowledge to get started.

By file:

```
python -m testtools.run -v test/test_loader.py
```

By module:

```
python -m testtools.run -v test.test_loader
python -m unittest -v test.test_loader
```

Using test discovery to locate all tests in a directory tree:

```
python -m testtools.run discover
python -m unittest discover test
```

See the [source distribution](#) and the [tutorial repo](#) for more advanced options, including using `testrepository` and `subunit`.

6.2 pytest

Since `pytest` does not support the `load_tests` system, a different way of generating tests is required. Two techniques are supported.

The original method (described below) used yield statements to generate tests which pytest would collect. This style of tests is deprecated as of `pytest>=3.0` so a new style using `pytest` fixtures has been developed.

Warning: The `pytest` loader now requires that `test_loader_name` be set when `gabbi.driver.py_test_generator()` is called.

6.2.1 `pytest >= 3.0`

In the newer technique, a test file is created that uses the `pytest_generate_tests` hook. Special care must be taken to always import the `test_pytest` method which is the base test that the `pytest` hook parametrizes to generate the tests from the YAML files. Without the method, the hook will not be called and no tests generated.

Here is a simple example file:

```
"""A sample pytest module for pytest >= 3.0."""

# For pathname munging
import os

# The module that py_test_generator comes from.
from gabbi import driver

# We need test_pytest so that pytest test collection works properly.
# Without this, the pytest_generate_tests method below will not be
# called.
from gabbi.driver import test_pytest # noqa

# We need access to the WSGI application that hosts our service
from myapp import wsgiapp

# We're using fixtures in the YAML files, we need to know where to
# load them from.
from myapp.test import fixtures

# By convention the YAML files are put in a directory named
# "gabbits" that is in the same directory as the Python test file.
TESTS_DIR = 'gabbits'

def pytest_generate_tests(metafunc):
    test_dir = os.path.join(os.path.dirname(__file__), TESTS_DIR)
    driver.py_test_generator(
        test_dir, intercept=wsgiapp.app,
        test_loader_name=__name__,
        fixture_module=fixtures, metafunc=metafunc)
```

This can then be run with the usual `pytest` commands. For example:

```
py.test -svx pytest3.0-example.py
```

6.2.2 `pytest < 3.0`

When using the older technique, test file must be created that calls `py_test_generator()` and yields the generated tests. That will look a bit like this:

```
"""A sample pytest module."""

# For pathname munging
import os

# The module that build_tests comes from.
from gabbi import driver

# We need access to the WSGI application that hosts our service
from myapp import wsgiapp

# We're using fixtures in the YAML files, we need to know where to
# load them from.
from myapp.test import fixtures

# By convention the YAML files are put in a directory named
# "gabbits" that is in the same directory as the Python test file.
TESTS_DIR = 'gabbits'

def test_gabbits():
    test_dir = os.path.join(os.path.dirname(__file__), TESTS_DIR)
    # Pass "require_ssl=True" as an argument to force all tests
    # to use SSL in requests.
    test_generator = driver.py_test_generator(
        test_loader_name=__name__,
        test_dir, intercept=wsgiapp.app,
        fixture_module=fixtures)

    for test in test_generator:
        yield test
```

This can then be run with the usual pytest commands. For example:

```
py.test -svx pytest-example.py
```

The older technique will continue to work with all versions of `pytest<4.0` but `>=3.0` will produce warnings. If you want to use the older technique but not see the warnings add `--disable-pytest-warnings` parameter to the invocation of `py.test`.

CHAPTER 7

Example Tests

What follows is a commented example of some tests in a single file demonstrating many of the *Test Format* features. See *Loading and Running Tests* for the Python needed to integrate with a testing harness.

```
# Fixtures can be used to set any necessary configuration, such as a
# persistence layer, and establish sample data. They operate per
# file. They are context managers, each one wrapping the next in the
# sequence.

fixtures:
    - ConfigFixture
    - SampleDataFixture

# There is an included fixture named "SkipAllFixture" which can be
# used to declare that all the tests in the given file are to be
# skipped.

# Each test file can specify a set of defaults that will be used for
# every request. This is useful for always specifying a particular
# header or always requiring SSL. These values will be used on every
# test in the file unless overridden. Lists and dicts are merged one
# level deep, except for "data" which is copied verbatim whether it
# is a string, list or dict (it can be all three).

defaults:
    ssl: True
    request_headers:
        x-my-token: zoom

# The tests themselves are a list under a "tests" key. It's useful
# to use plenty of whitespace to help readability.

tests:

# Each request *must* have a name which is unique to the file. When it
```

(continues on next page)

```
# becomes a TestCase the name will be lowercased and spaces will
# become "_". Use that generated name when limiting test runs.

- name: a test for root
  desc: Some explanatory text that could be used by other tooling

# The URL can either be relative to a host specified elsewhere or
# be a fully qualified "http" or "https" URL. *You* are responsible
# for url-encoding the URL.

  url: /
  method: GET

# If no status or method are provided they default to "200" and
# "GET".

# Instead of explicitly stating "url" and "method" you can join
# those two keys into one key representing the method. The method
# *must* be uppercase.

- name: another test for root
  desc: Same test as above but with GET key
  GET: /

# A single test can override settings in defaults (set above).

- name: root without ssl redirects
  ssl: False
  GET: /
  status: 302

# When evaluating response headers it is possible to use a regular
# expression to not have to test the whole value. Regular expressions match
# anywhere in the output, not just at the beginning.

  response_headers:
    location: /^https/

# By default redirects will not be followed. This can be changed.

- name: follow root without ssl redirect
  ssl: False
  redirects: True
  GET: /
  status: 200 # This is the response code after the redirect.

# URLs can express query parameters in two ways: either in the url
# value directly, or as query_parameters. If both are used then
# query_parameters are appended. In this example the resulting URL
# will be equivalent to
# /foo?section=news&article=1&article=2&date=yesterday
# but not necessarily in that order.

- name: create a url with parameters
  GET: /foo?section=news
  query_parameters:
    article:
```

(continues on next page)

(continued from previous page)

```
    - 1
    - 2
    date: yesterday

# Request headers can be used to declare media-type choices and
# experiment with authorization handling (amongst other things).
# Response headers allow evaluating headers in the response. These
# two together form the core value of gabbi.

- name: test accept
  GET: /resource
  request_headers:
    accept: application/json
  response_headers:
    content-type: /application/json/

# If a header must not be present in a response at all that can be
# expressed in a test as follows.

- name: test forbidden headers
  GET: /resource
  response_forbidden_headers:
    - x-special-header

# All of the above requests have defaulted to a "GET" method. When
# using "POST", "PUT" or "PATCH", the "data" key provides the
# request body.

- name: post some text
  POST: /text_repo
  request_headers:
    content-type: text/plain
  data: "I'm storing this"
  status: 201

# If the data is not a string, it will be transformed into JSON.
# You must supply an appropriate content-type request header.

- name: post some json
  POST: /json_repo
  request_headers:
    content-type: application/json
  data:
    name: smith
    abode: castle
  status: 201

# If the data is a string prepended with "<@" the value will be
# treated as the name of a file in the same directory as the YAML
# file. Again, you must supply an appropriate content-type. If the
# content-type is one of several "text-like" types, the content will
# be assumed to be UTF-8 encoded.

- name: post an image
  POST: /image_repo
  request_headers:
    content-type: image/png
```

(continues on next page)

```
    data: <@kittens.png

# A single request can be marked to be skipped.

- name: patch an image
  skip: patching images not yet implemented
  PATCH: /image_repo/12d96fb8-e78c-11e4-8c03-685b35afa334

# Or a single request can be marked that it is expected to fail.

- name: check allow headers
  desc: the framework doesn't do allow yet
  xfail: True
  PUT: /post_only_url
  status: 405
  response_headers:
    allow: POST

# The body of a response can be evaluated with response handlers.
# The most simple checks for a set of strings anywhere in the
# response. Note that the strings are members of a list.

- name: check for css file
  GET: /blog/posts/12
  response_strings:
    - normalize.css

# For JSON responses, JSONPath rules can be used.

- name: post some json get back json
  POST: /json_repo
  request_headers:
    content-type: application/json
  data:
    name: smith
    abode: castle
  status: 201
  response_json_paths:
    $.name: smith
    $.abode: castle

# Requests run in sequence. One test can make reference to the test
# immediately prior using some special variables.
# "$LOCATION" contains the "location" header in the previous
# response.
# "$HEADERS" is a pseudo dictionary containing all the headers of
# the previous response.
# "$ENVIRON" is a pseudo dictionary providing access to the current
# environment.
# "$RESPONSE" provides access to the JSON in the prior response, via
# JSONPath. See http://jsonpath-rw.readthedocs.io/ for
# jsonpath-rw formatting.
# $SCHEME and $NETLOC provide access to the current protocol and
# location (host and port).

- name: get the thing we just posted
  GET: $LOCATION
```

(continues on next page)

(continued from previous page)

```
request_headers:
  x-magic-exchange: $HEADERS['x-magic-exchange']
  x-token: $ENVIRON['OS_TOKEN']
response_json_paths:
  $.name: $RESPONSE['$.name']
  $.abode: $RESPONSE['$.abode']
response_headers:
  content-location: /$SCHEME://$NETLOC/

# For APIs where resource creation is asynchronous it can be
# necessary to poll for the resulting resource. First we create the
# resource in one test. The next test uses the "poll" key to loop
# with a delay for a set number of times.

- name: create async
  POST: /async_creator
  request_headers:
    content-type: application/json
  data:
    name: jones
    abode: bungalow
  status: 202

- name: poll for created resource
  GET: $LOCATION
  poll:
    count: 10 # try up to ten times
    delay: .5 # wait .5 seconds between each try
  response_json_paths:
    $.name: $RESPONSE['$.name']
    $.abode: $RESPONSE['$.abode']
```


Gabbi supports JSONPath both for validating JSON response bodies and within *substitutions*.

JSONPath expressions are provided by `jsonpath_rw`, with `jsonpath_rw_ext` custom extensions to address common requirements:

1. Sorting via `sorted` and `[/property]`.
2. Filtering via `[?property = value]`.
3. Returning the respective length via `len`.

(These apply both to arrays and key-value pairs.)

Here is a JSONPath example demonstrating some of these features. Given JSON data as follows:

```
{
  "pets": [
    {"type": "cat", "sound": "meow"},
    {"type": "dog", "sound": "woof"}
  ]
}
```

If the ordering of the list in `pets` is predictable and reliable it is relatively straightforward to test values:

```
response_json_paths:
  # length of list is two
  $.pets.`len`: 2
  # sound of second item in list is woof
  $.pets[1].sound: woof
```

If the ordering is *not* predictable additional effort is required:

```
response_json_paths:
  # sort by type
  $.pets[/type][0].sound: meow
  # sort by type, reversed
```

(continues on next page)

(continued from previous page)

```
$.pets[\\type][0].sound: woof
# all the sounds
$.pets[/type]..sound: ['meow', 'woof']
# filter by type = dog
$.pets[?type = "dog"].sound: woof
```

If it is necessary to validate the entire JSON response use a JSONPath of \$:

```
response_json_paths:
  $:
    pets:
      - type: cat
        sound: meow
      - type: dog
        sound: woof
```

This is not a technique that should be used frequently as it can lead to difficult to read tests and it also indicates that your gabbi tests are being used to test your serializers and data models, not just your API interactions.

It is also possible to read raw JSON from disk for either all or some of a JSON response:

```
response_json_paths:
  $: @<data.json
```

or:

```
response_json_paths:
  $.pets: <@pets.json
  $.pets[0]: <@cat.json
```

Examples like this can be found in one of gabbi's own tests.

If it is desired to load YAML files like the JSON ones above, two things must be done:

1. The `YAMLDiskLoadingJSONHandler` custom content handler must be passed to the driver through the `content_handlers` argument. See [Extensions](#) on how to do this.
2. The YAML files to load must be placed in a subdirectory to prevent the test runner from consuming them as test files to run:

```
response_json_paths:
  $: @<subdir/values.yaml
```

When reading from disk you can apply the same JSONPath by adding a ':' to the end of your file name. This allows you to store multiple API responses into a single file to reduce file management when constructing your tests.

Given JSON data as follows:

```
{
  "values": [{
    "pets": [{
      "type": "cat",
      "sound": "meow"
    }, {
      "type": "dog",
      "sound": "woof"
    }
  ]
}, {
```

(continues on next page)

(continued from previous page)

```

    "people": [{
      "name": "chris",
      "id": 1
    }, {
      "name": "justin",
      "id": 2
    }]
  }]
}

```

You can write your tests like the following:

```

response_json_paths:
  $.pets: <@pets.json
  $.pets[?type = "cat"].sound: <@values.json:$.values[0].pets[?type = "cat"].sound

```

Although placing more than one API response into a single JSON file may seem convenient, keep in mind there is a tradeoff in readability that should not be overlooked before implementing this technique.

Examples like this can be found in one of gabbi's [yaml-from-disk tests](#).

There are more JSONPath examples in [Example Tests](#) and in the [jsonpath_rw](#) and [jsonpath_rw_ext](#) documentation.

8.1 Substitution

Substitutions can be made in both the left (query) and right (expected) hand sides of the json path expression. When substitutions are used in the query, care must be taken to ensure proper quoting of the resulting value. For example if there is a uuid (with hyphens) at `$RESPONSE['$.id']` then this expression may fail:

```
$.nested.structure.$RESPONSE['$.id'].name: foobar
```

as it will evaluate to something like:

```
$.nested.structure.ADC8AAFC-D564-40D1-9724-7680D3C010C2.name: foobar
```

which may be treated as an arithmetic expression by the json path parser. The test author should write:

```
$.nested.structure["$RESPONSE['$.id']"].name: foobar
```

to quote the result of the substitution.

Target Host

The target host is the host on which the API to be tested can be found. Gabbi intends to preserve the flow and semantics of HTTP interactions as much as possible, and every HTTP request needs to be directed at a host of some form. Gabbi provides three ways to control this:

- Using `wsgi-intercept` to provide a fake socket and WSGI environment on an arbitrary host and port attached to a WSGI application (see [intercept examples](#)).
- Using fully qualified `url` values in the YAML defined tests (see [full examples](#)).
- Using a host and (optionally) port defined at test build time (see [live examples](#)).

The intercept and live methods are mutually exclusive per test builder, but either kind of test can freely intermix fully qualified URLs into the sequence of tests in a YAML file.

For test driven development and local tests the intercept style of testing lowers test requirements (no web server required) and is fast. Interception is performed as part of *Fixtures* processing as the most deeply nested fixture. This allows any configuration or database setup to be performed prior to the WSGI application being created.

For the implementation of the above see `build_tests()`.

Each suite of tests is represented by a single YAML file, and may optionally use one or more fixtures to provide the necessary environment required by the tests in that file.

Fixtures are implemented as nested context managers. Subclasses of *GabbiFixture* must implement `start_fixture` and `stop_fixture` methods for creating and destroying, respectively, any resources managed by the fixture. While the subclass may choose to implement `__init__` it is important that no exceptions are thrown in that method, otherwise the stack of context managers will fail in unexpected ways. Instead initialization of real resources should happen in `start_fixture`.

At this time there is no mechanism for the individual tests to have any direct awareness of the fixtures. The fixtures exist, conceptually, on the server side of the API being tested.

Fixtures may do whatever is required by the testing environment, however there are two common scenarios:

- Establishing (and then resetting when a test suite has finished) any baseline configuration settings and persistence systems required for the tests.
- Creating sample data for use by the tests.

If a fixture raises `unittest.case.SkipTest` during `start_fixture` all the tests in the current file will be skipped. This makes it possible to skip the tests if some optional configuration (such as a particular type of database) is not available.

If an exception is raised while a fixture is being used, information about the exception will be stored on the fixture so that the `stop_fixture` method can decide if the exception should change how the fixture should clean up. The exception information can be found on `exc_type`, `exc_value` and `traceback` method attributes.

If an exception is raised when a fixture is started (in `start_fixture`) the first test in the suite using the fixture will be marked with an error using the traceback from the exception and all the tests in the suite will be skipped. This ensures that fixture failure is adequately captured and reported by test runners.

Inner Fixtures

In some contexts (for example CI environments with a large number of tests being run in a broadly concurrent environment where output is logged to a single file) it can be important to capture and consolidate stray output that is produced during the tests and display it associated with an individual test. This can help debugging and avoids unusable output that is the result of multiple streams being interleaved.

Inner fixtures have been added to support this. These are fixtures more in line with the traditional `unittest` concept of fixtures: a class on which `setUp` and `tearDown` is automatically called.

`build_tests()` accepts a named parameter arguments of `inner_fixtures`. The value of that argument may be an ordered list of `fixtures.Fixture` classes that will be called when each individual test is set up.

An example fixture that could be useful is the `FakeLogger`.

Note: At this time `inner_fixtures` are not supported when using the pytest *loader*.

Content Handlers

Content handlers are responsible for preparing request data and evaluating response data based on the content-type of the request and response. A content handler operates as follows:

- Structured YAML data provided via the `data` attribute is converted to a string or bytes sequence and used as request body.
- The response body (a string or sequence of bytes) is transformed into a content-type dependent structure and stored in an internal attribute named `response_data` that is:
 - used when evaluating the response body
 - used in `RESPONSE [] substitutions`

By default, `gabbi` provides content handlers for JSON. In that content handler the `data` test key is converted from structured YAML into a JSON string. Response bodies are converted from a JSON string into a data structure in `response_data` that is used when evaluating `response_json_paths` entries in a test or doing JSONPath-based `RESPONSE [] substitutions`.

A `YAMLDiskLoadingJSONHandler` has been added to extend the JSON handler. It works the same way as the JSON handler except for when evaluating the `response_json_paths` handle, data that is read from disk can be either in JSON or YAML format. The `YAMLDiskLoadingJSONHandler` is not enabled by default and must be added as shown in the *Extensions* section in order to be used in the tests.

Further content handlers can be added as extensions. Test authors may need these extensions for their own suites, or enterprising developers may wish to create and distribute extensions for others to use.

Note: One extension that is likely to be useful is a content handler that turns `data` into url-encoded form data suitable for POST and turns an HTML response into a DOM object.

12.1 Extensions

Content handlers are an evolution of the response handler concept in earlier versions `gabbi`. To preserve backwards compatibility with existing response handlers, old style response handlers are still allowed, but new handlers should

implement the content handler interface (described below).

Registering additional custom handlers is done by passing a subclass of `ContentHandler` to `build_tests()`:

```
driver.build_tests(test_dir, loader, host=None,
                  intercept=simple_wsgi.SimpleWsgi,
                  content_handlers=[MyContentHandler])
```

If `pytest` is being used:

```
driver.py_test_generator(test_dir, intercept=simple_wsgi.SimpleWsgi,
                        test_loader_name=__name__,
                        content_handlers=[MyContentHandler])
```

Gabbi provides an additional custom handler named `YAMLDiskLoadingJSONHandler`. This can be used for loading JSON and YAML files from disk when evaluating the `response_json_paths` handle.

Warning: `YAMLDiskLoadingJSONHandler` shares the same content-type as the default `JSONHandler`. When there are multiple handlers listed that accept the same content-type, the one that is earliest in the list will be used.

With `gabbi-run`, custom handlers can be loaded via the `--response-handler` option – see `load_response_handlers()` for details.

Note: The use of the `--response-handler` argument is done to preserve backwards compatibility and avoid excessive arguments. Both types of handler may be passed to the argument.

12.1.1 Implementation Details

Creating a content handler requires subclassing `ContentHandler` and implementing several methods. These methods are described below, but inspecting `JSONHandler` will be instructive in highlighting required arguments and techniques.

To provide a `response_<something>` response-body evaluator a subclass must define:

- `test_key_suffix`: This, along with the prefix `response_`, forms the key used in the test structure. It is a class level string.
- `test_key_value`: The key's default value, either an empty list (`[]`) or empty dict (`{}`). It is a class level value.
- `action`: An instance method which tests the expected values against the HTTP response - it is invoked for each entry, with the parameters depending on the default value. The arguments to `action` are (in order):
 - `self`: The current instance.
 - `test`: The currently active `HTTPTestCase`
 - `item`: The current entry if `test_key_value` is a list, otherwise the key half of the key/value pair at this entry.
 - `value`: None if `test_key_value` is a list, otherwise the value half of the key/value pair at this entry.

To translate request or response bodies to or from structured data a subclass must define an `accepts` method. This should return `True` if this class is willing to translate the provided content-type. During request processing it is given the value of the content-type header that will be sent in the request. During response processing it is given the value of the content-type header of the response. This makes it possible to handle different request and response bodies in the

same handler, if desired. For example a handler might accept `application/x-www-form-urlencoded` and `text/html`.

If `accepts` is defined two additional static methods should be defined:

- `dumps`: Turn structured Python data from the `data` key in a test into a string or byte stream. The optional `test` param allows you to access the current test case which may help with manipulations for custom content handlers, e.g. `multipart/form-data` needs to add a boundary to the `Content-Type` header in order to mark the appropriate sections of the body.
- `loads`: Turn a string or byte stream in a response into a Python data structure. Gabbi will put this data on the `response_data` attribute on the test, where it can be used in the evaluations described above (in the `action` method) or in `$RESPONSE` handling. An example usage here would be to turn HTML into a DOM.
- `load_data_file`: Load data from disk into a Python data structure. Gabbi will call this method when `response_<something>` contains an item where the right hand side value starts with `<@`. The `test` param allows you to access the current test case and provides a `load_data_file` method which should be used because it verifies the data is loaded within the test dictionary and returns the file source as a string. The `load_data_file` method was introduced to re-use the `JSONHandler` in order to support loading YAML files from disk through the implementation of an additional custom handler, see `YAMLDiskLoadingJSONHandler` for details.

Finally if a `replacer` class method is defined, then when a `$RESPONSE` substitution is encountered, `replacer` will be passed the `response_data` of the prior test and the argument within the `$RESPONSE`.

Please see the [JSONHandler source](#) for additional detail.

CHAPTER 13

YAML Runner

If there is a running web service that needs to be tested and creating a test loader with `build_tests()` is either inconvenient or overkill it is possible to run YAML test files directly from the command line with the console-script `gabbi-run`. It accepts YAML on `stdin` or as multiple file arguments, and generates and runs tests and outputs a summary of the results.

The provided YAML may not use custom *Fixtures* but otherwise uses the default *Test Format*. *Target Host* information is either expressed directly in the YAML file or provided on the command line:

```
gabbi-run [host[:port]] < /my/test.yaml
```

or:

```
gabbi-run http://host:port < /my/test.yaml
```

To test with one or more files the following command syntax may be used:

```
gabbi-run http://host:port -- /my/test.yaml /my/other.yaml
```

Note: The filename arguments must come after a `--` and all other arguments (host, port, prefix, failfast) must come before the `--`.

Note: If files are provided, test output will use names including the name of the file. If any single file includes an error, the name of the file will be included in a summary of failed files at the end of the test report.

To facilitate using the same tests against the same application mounted in different locations in a WSGI server, a `prefix` may be provided as a second argument:

```
gabbi-run host[:port] [prefix] < /my/test.yaml
```

or in the target URL:

```
gabbi-run http://host:port/prefix < /my/test.yaml
```

The value of `prefix` will be prepended to the path portion of URLs that are not fully qualified.

Anywhere `host` is used, if it is a raw IPV6 address it should be wrapped in `[and]`.

If `https` is used in the target, then the tests in the provided YAML will default to `ssl: True`.

Use `-k` or `--insecure` to **not** validate certificates when making `https` connections.

If a `-x` or `--failfast` argument is provided then `gabbi-run` will exit after the first test failure.

Use `-v` or `--verbose` with a value of `all`, `headers` or `body` to turn on *verbosity* for all tests being run.

Use `-q` or `--quiet` to silence test runner output.

Use `-r` or `--response-handler` to load a custom response or content handler for use with tests.

Use `-l` to load response handlers relative to the current working directory.

For example, to load a handler named `HTMLHandler` from the `handlers.html` module relative to the current directory:

```
gabbi-run -l -r handlers.html:HTMLHandler http://example.com < my.yaml
```

These are informal release notes for gabbi since version 1.0.0, highlighting major features and changes. For more detail see the [commit logs](#) on GitHub.

14.1 2.8.0

- Add timeout parameter per test, defaulting to 30 seconds.
- Ensure `stestr > 4.0.0` in `test-requirements.txt`

14.2 2.7.2

- Handle SNI when `host` request header is any case.

14.3 2.7.1

- Clear the PoolManager after each test to avoid running out of socket file descriptors.

14.4 2.7.0

- Add `-l` argument to `gabbi-run` to allow loading content and response handlers relative to the local directory.

14.5 2.6.0

- Fix naming of tests when using pytest. This changes means that setting `test_loader_name` is now **required** when using pytest.

14.6 2.5.0

- Support SNI in https requests when `host` request header is set.

14.7 2.4.0

- Publish a docker container that runs `gabbi-run` with each tagged release.

14.8 2.3.0

- For the `$ENVIRON` and `$RESPONSE` *substitutions* it is now possible to *cast* the value to a type of `int`, `float`, `str`, or `bool`.
- The `JSONHandler` is now more strict about how it detects that a body content is JSON, avoiding some errors where the content-type header suggests JSON but the content cannot be decoded as such.
- Better error message when content cannot be decoded.
- Addition of the `disable_response_handler` test setting for those cases when the test author has no control over the content-type header and it is wrong.

14.9 2.2.0

- Allow substitutions in *skip* for more flexible skipping, notably with `$ENVIRON`.
- Better test name output when running with `unittest` and `verbose`.

14.10 2.1.0

- Remove support for Python 3.5. Ensure testing support for 3.8, 3.9, and `pypy3`.
- Adapt to new behavior in `urllib.parse`.
- Correct fixture start when using `pytest`.

14.11 2.0.4

- If no content-type is provided with a response and `verbose` is on for responses, display the response body as if it were text.

14.12 2.0.3

- Properly declare that `gabbi 2.x` is Python 3 only.

14.13 2.0.1

- Correct management of response handler default fields.

14.14 2.0.0

- Drop support for Python 2. If you need Python 2 support, use an older version.
- Stop using `testtools` and `fixtures`. These two modules present several difficulties and their maintenance situation suggests those difficulties will not be resolved. Since Python 2 support is being removed, the need for the modules can be removed as well without losing functionality. “Inner fixtures” that use the `fixtures.Fixture` interface should continue to work.

14.15 1.49.0

- Add support for not validating certificates in `https` requests. Controlled by the `cert_validate` attribute in individual tests and `build_tests()` and the `-k` or `--insecure` argument to `gabbi-run`.

14.16 1.48.0

- Support `pytest 5.0.0` in Python `>=3.5`. For earlier versions of Python, `pytest<5.0.0` will be used; the `pytest` project is dropping support for older versions of Python.

14.17 1.47.0

- Use `pytest<5.0.0` until `gabbi` has solutions for the changes in `5.0.0`.

14.18 1.46.0

- A `-q` argument is added to `gabbi-run` to suppress output from the test runner.

14.19 1.45.0

- Adjust loading of `YAML` to be ready for new release of `PyYAML`.

14.20 1.44.0

- Provide the `YAMLDiskLoadingJSONHandler` class that allows test result data for `response_json_path` checks to be loaded from `YAML-on-disk`.

14.21 1.43.0

- Use *JSONPath* to select a portion of data-on-disk in `response_json_path` checks.
- Restrict PyYAML to <4.0.

14.22 1.42.0

- Allow listing of tests with no host configured. When host is an empty string, tests can be listed (for discovery), but will be skipped on run.

14.23 1.41.0

- JSON `$RESPONSE` *substitutions* in the `data` field may be complex types (lists and dicts), not solely strings.

14.24 1.40.0

- When the HTTP response begins with a bad status line, have `BadStatusLine` be raised from `urllib3`.

14.25 1.39.0

- Allow *substitutions* in the key portion of request and response headers, not just the value.

14.26 1.38.0

- Remove support for Python 3.3.
- Make handling of fixture-level skips in `pytest` actually work.

14.27 1.37.0

- Add `safe_yaml` parameter to `build_tests()`.

14.28 1.36.0

- `use_prior_test` is added to test *Metadata*.
- Extensive cleanups in regular expression handling when constructing tests from YAML.

14.29 1.35.0

JSONPath handling gets two improvements:

- The value side of a `response_json_paths` entry can be loaded from a file using the `<@file.json` syntax also used in *Data*.
- The key side of a `response_json_paths` entry can use *substitutions*. This was already true for the value side.

14.30 1.34.0

Substitutions in `$RESPONSE` handling now preserve numeric types instead of casting to a string. This is useful when servers are expecting strong types and tests want to send response data back to the server.

14.31 1.33.0

`count` and `delay` test keys allow *substitutions*. `gabbi.driver.build_tests()` accepts a `verbose` parameter to set test *verbosity* for an entire session.

14.32 1.32.0

Better failure reporting when using *gabbi-run* with multiple files. Test names are based on the files and a summary of failed files is provided at the end of the report.

14.33 1.31.0

Effectively capture a failure in a *fixture* and report the traceback. Without this some test runners swallow the error and discovering problems when developing fixtures can be quite challenging.

14.34 1.30.0

Thanks to Samuel Fekete, tests can use the `$HISTORY` dictionary to refer to any prior test in the same file, not just the one immediately prior, when doing *substitutions*.

14.35 1.29.0

Filenames used to read data into tests using the `<@` syntax may now use pathnames relative to the YAML file. See *Data*.

gabbi-run gains a `--verbose` parameter to force all tests run in a session to run with *verbose* set.

When using *pytest* to load tests, a new mechanism is available which avoids warnings produced in when using a version of *pytest* greater than 3.0.

14.36 1.28.0

When verbosely displaying request and response bodies that are JSON, pretty print for improved readability.

14.37 1.27.0

Allow *gabbi-run* to accept multiple filenames as command line arguments instead of reading tests from stdin.

14.38 1.26.0

Switch from response handlers to *Content Handlers* to allow more flexible processing of both response `_and_` request bodies.

Add *inner fixtures* for per test fixtures, useful for output capturing.

14.39 1.25.0

Allow the `test_loader_name` arg to `gabbi.driver.build_tests()` to override the prefix of the pretty printed name of generated tests.

14.40 1.24.0

String values in JSONPath matches may be wrapped in `/.../`` to be treated as regular expressions.

14.41 1.23.0

Better *documentation* of how to run gabbi in a concurrent environment. Improved handling of pytest fixtures and test counts.

14.42 1.22.0

Add `url` to `gabbi.driver.build_tests()` to use instead of `host`, `port` and `prefix`.

14.43 1.21.0

Add `require_ssl` to `gabbi.driver.build_tests()` to force use of SSL.

14.44 1.20.0

Add `$COOKIE substitution`.

14.45 1.19.1

Correctly support IPV6 hosts.

14.46 1.19.0

Add `$_LAST_URL` *substitution*.

14.47 1.17.0

Introduce support for loading and running tests with pytest.

14.48 1.16.0

Use urllib3 instead of httplib2 for driving HTTP requests.

14.49 1.13.0

Add sorting and filtering to *JSONPath* handling.

14.50 1.11.0

Add the `response_forbidden_headers` to *response expectations*.

14.51 1.7.0

Instead of:

```
tests:  
- name: a simple get  
  url: /some/path  
  method: get
```

1.7.0 also makes it possible to:

```
tests:  
- name: a simple get  
  GET: /some/path
```

Any upper case key is treated as a method.

14.52 1.4.0 and 1.5.0

Enhanced flexibility and colorization when setting tests to be *verbose*.

14.53 1.3.0

Adds the `query_parameters` key to *request parameters*.

14.54 1.2.0

The start of improvements and extensions to *JSONPath* handling. In this case the addition of the `len` function.

14.55 1.1.0

Vastly improved output and behavior in *gabbi-run*.

14.56 1.0.0

Version 1 was the first release with a commitment to a stable *Test Format*. Since then new fields have been added but have not been taken away.

CHAPTER 15

Contributors

The following people have contributed code to gabbi. Thanks to them. Thanks also to all the people who have made gabbi better by reporting [issues](#) and their successes and failures with using gabbi.

- Chris Dent
- FND
- Mehdi Abaakouk
- Tom Viner
- Jason Myers
- Josh Leeb-du Toit
- Duc Truong
- Zane Bitter
- Ryan Spencer
- Kim Raymoure
- Travis Truman
- Samuel Fekete
- Michael McCune
- Imran Hayder
- Julien Danjou
- Trevor McCasland
- Danek Duvall
- Marc Abramowitz
- Scott Wallace

Frequently Asked Questions

Note: This section provides a collection of questions with answers that don't otherwise fit in the rest of the documentation. If something is missing, please create an [issue](#).

As this document grows it will gain a more refined structure.

16.1 General

16.1.1 Is gabbi only for testing Python-based APIs?

No, you can use *gabbi-run* to test an HTTP service built in any programming language.

16.1.2 How do I run just one test?

Each YAML file contains a sequence of tests, each test within each file has a name. That name is translated to the name of the test by replacing spaces with an `_`.

When running tests that are *generated dynamically*, filtering based on the test name prior to the test being collected will not work in some test runners. Test runners that use a `--load-list` functionality can be convinced to filter after discovery.

pytest does this directly with the `-k` keyword flag.

When using *testrepository* with *tox* as used in *gabbi*'s own tests it is possible to pass a filter in the *tox* command:

```
tox -epy27 -- get_the_widget
```

When using `testtools.run` and similar test runners it's a bit more complicated. It is necessary to provide the full name of the test as a list to `--load-list`:

```
python -m testtools.run --load-list \  
    <(echo package.tests.test_api.yamlfile_get_the_widge.test_request)
```

16.1.3 How do I run just one test, without running prior tests in a sequence?

By default, when you select a single test to run, all tests prior to that one in a file will be run as well: the file is treated as a sequence of dependent tests. If you do not want this you can adjust the `use_prior_test` test *metadata* in one of three ways:

- Set it in the YAML file for the one test you are concerned with.
- Set the defaults for all tests in that file.
- set `use_prior_test` to `false` when calling `build_tests()`

Be aware that doing this breaks a fundamental assumption that gabbi makes about how tests work. Any *substitutions* will fail.

16.2 Testing Style

16.2.1 Can I have variables in my YAML file?

Gabbi provides the `$ENVIRON` *substitution* which can operate a bit like variables that are set elsewhere and then used in the tests defined by the YAML.

If you find it necessary to have variables within a single YAML file you take advantage of YAML *alias nodes* list this:

```
vars:  
  - &uuid_1 5613AABF-BAED-4BBA-887A-252B2D3543F8  
  
tests:  
- name: send a uuid to a post  
  POST: /resource  
  request_headers:  
    content-type: application/json  
  data:  
    uuid: *uuid_1
```

You can alias all sorts of nodes, not just single items. Be aware that the replacement of an alias node happens while the YAML is being loaded, before gabbi does any processing.

16.2.2 How many tests should be put in one YAML file?

For the sake of readability it is best to keep each YAML file relatively short. Since each YAML file represents a sequence of requests, it usually makes sense to create a new file when a test is not dependent on any before it.

It's tempting to put all the tests for any resource or URL in the same file, but this eventually leads to files that are too long and are thus difficult to read.

17.1 case Module

17.2 driver Module

17.3 suitemaker Module

17.4 fixture Module

Manage fixtures for gabbi at the test suite level.

```
class gabbi.fixture.GabbiFixture
    Bases: object
```

A context manager that operates as a fixture.

Subclasses must implement `start_fixture` and `stop_fixture`, each of which contain the logic for stopping and starting whatever the fixture is. What a fixture is is left as an exercise for the implementor.

These context managers will be nested so any actual work needs to happen in `start_fixture` and `stop_fixture` and not in `__init__`. Otherwise exception handling will not work properly.

```
start_fixture()
```

Implement the actual workings of starting the fixture here.

```
stop_fixture()
```

Implement the actual workings of stopping the fixture here.

```
exception gabbi.fixture.GabbiFixtureError
    Bases: exceptions.Exception
```

Generic exception for GabbiFixture.

class `gabbi.fixture.SkipAllFixture`

Bases: `gabbi.fixture.GabbiFixture`

A fixture that skips all the tests in the current suite.

start_fixture ()

Implement the actual workings of starting the fixture here.

`gabbi.fixture.nest` (*args, **kws)

Nest a series of fixtures.

This is duplicated from `nested` in the `stdlib`, which has been deprecated because of issues with how exceptions are difficult to handle during `__init__`. Gabbi needs to nest an unknown number of fixtures dynamically, so the `with` syntax that replaces `nested` will not work.

17.5 handlers Module

17.5.1 handlers.base Module

17.5.2 handlers.core Module

17.5.3 handlers.jsonhandler Module

17.5.4 handlers.yaml_disk_loading_jsonhandler Module

17.6 suite Module

A `TestSuite` for containing gabbi tests.

This suite has two features: the contained tests are ordered and there are suite-level fixtures that operate as context managers.

class `gabbi.suite.GabbiSuite` (tests=())

Bases: `unittest.suite.TestSuite`

A `TestSuite` with fixtures.

The suite wraps the tests with a set of nested context managers that operate as fixtures.

If a fixture raises `unittest.case.SkipTest` during setup, all the tests in this suite will be skipped.

run (result, debug=False)

Override `TestSuite` run to start suite-level fixtures.

To avoid exception confusion, use a null `Fixture` when there are no fixtures.

start (result, tests=None)

Start fixtures when using `pytest`.

stop ()

Stop fixtures when using `pytest`.

`gabbi.suite.noop` (*args)

A `noop` method used to disable collected tests.

17.7 runner Module

17.8 reporter Module

TestRunner and TestResult for gabbi-run.

class `gabbi.reporter.ConciseTestResult` (*stream, descriptions, verbosity*)

Bases: `unittest.runner.TextTestResult`

A `TextTestResult` with simple but useful output.

If the output is a tty or `GABBI_FORCE_COLOR` is set in the environment, output will be colorized.

addError (*test, err*)

Called when an error has occurred. 'err' is a tuple of values as returned by `sys.exc_info()`.

addExpectedFailure (*test, err*)

Called when an expected failure/error occurred.

addFailure (*test, err*)

Called when an error has occurred. 'err' is a tuple of values as returned by `sys.exc_info()`.

addSkip (*test, reason*)

Called when a test is skipped.

addSuccess (*test*)

Called when a test has completed successfully

addUnexpectedSuccess (*test*)

Called when a test was expected to fail, but succeed.

getDescription (*test*)

printErrorList (*flavor, errors*)

startTest (*test*)

Called when the given test is about to be run

class `gabbi.reporter.ConciseTestRunner` (*stream=<open file '<stderr>', mode 'w'>, descriptions=True, verbosity=1, failfast=False, buffer=False, resultclass=None*)

Bases: `unittest.runner.TextTestRunner`

A `TextTestRunner` that uses `ConciseTestResult` for reporting results.

resultclass

alias of `ConciseTestResult`

class `gabbi.reporter.PyTestResult` (*stream=None, descriptions=None, verbosity=None*)

Bases: `unittest.result.TestResult`

Wrap a test result to allow it to work with pytest.

The main behaviors here are:

- to turn what had been exceptions back into exceptions
- use pytest's skip and xfail methods

addError (*test, err*)

Called when an error has occurred. 'err' is a tuple of values as returned by `sys.exc_info()`.

addExpectedFailure (*test, err*)

Called when an expected failure/error occurred.

addFailure (*test, err*)

Called when an error has occurred. 'err' is a tuple of values as returned by `sys.exc_info()`.

addSkip (*test, reason*)

Called when a test is skipped.

17.9 `utils` Module

Utility functions grab bag.

`gabbi.utils.create_url` (*base_url, host, port=None, prefix="", ssl=False*)

Given pieces of a path-based url, return a fully qualified url.

`gabbi.utils.decode_response_content` (*header_dict, content*)

Decode content to a proper string.

`gabbi.utils.extract_content_type` (*header_dict, default='application/binary'*)

Extract parsed content-type from headers.

`gabbi.utils.get_colorizer` (*stream*)

Return a function to colorize a string.

Only if stream is a tty .

`gabbi.utils.host_info_from_target` (*target, prefix=None*)

Turn url or host:port and target into test destination.

`gabbi.utils.load_yaml` (*handle=None, yaml_file=None, safe=True*)

Read and parse any YAML file or filehandle.

Let exceptions flow where they may.

If no file or handle is provided, read from STDIN.

`gabbi.utils.not_binary` (*content_type*)

Decide if something is content we'd like to treat as a string.

`gabbi.utils.parse_content_type` (*content_type, default_charset='utf-8'*)

Parse content type value for media type and charset.

17.10 `exception` Module

Gabbi specific exceptions.

exception `gabbi.exception.GabbiDataLoadError`

Bases: `exceptions.ValueError`

An exception to alert when data streams cannot be loaded.

exception `gabbi.exception.GabbiFormatError`

Bases: `exceptions.ValueError`

An exception to encapsulate poorly formed test data.

exception `gabbi.exception.GabbiSyntaxWarning`

Bases: `exceptions.SyntaxWarning`

A warning about syntax that is not desirable.

17.11 `httpclient` Module

17.12 `json_parser` Module

Keep one single global jsonpath parser.

`gabbi.json_parser.parse` (*path*)

Parse a JSONPath expression use the global parser.

Gabbi is a tool for running HTTP tests where requests and responses are expressed as declarations in YAML files:

```
tests:
- name: retrieve items
  GET: /items
```

See the rest of these docs for more details on the many features and formats for setting request headers and bodies and evaluating responses.

Tests can be run from the command line with *gabbi-run* or programmatically using either *py.test* or *unittest*-style test runners. See *installation instructions* below.

The name is derived from “gabby”: excessively talkative. In a test environment having visibility of what a test is actually doing is a good thing. This is especially true when the goal of a test is to test the HTTP, not the testing infrastructure. Gabbi tries to put the HTTP interaction in the foreground of testing.

If you want to get straight to creating tests look at *Example Tests*, the test files in the *source distribution* and *Test Format*. A *gabbi-demo* repository provides a tutorial of using gabbi to build an API, via the commit history of the repo.

18.1 Purpose

Gabbi works to bridge the gap between human readable YAML files (see *Test Format* for details) that represent HTTP requests and expected responses and the rather complex world of automated testing.

Each YAML file represents an ordered list of HTTP requests along with the expected responses. This allows a single file to represent a process in the API being tested. For example:

- Create a resource.
- Retrieve a resource.
- Delete a resource.
- Retrieve a resource again to confirm it is gone.

At the same time it is still possible to ask gabbi to run just one request. If it is in a sequence of tests, those tests prior to it in the YAML file will be run (in order). In any single process any test will only be run once. Concurrency is handled such that one file runs in one process.

These features mean that it is possible to create tests that are useful for both humans (as tools for learning, improving and developing APIs) and automated CI systems.

Significant flexibility and power is available in the *Test Format* to make it relatively straightforward to test existing complex APIs. This extended functionality includes the use of [JSONPath](#) to query response bodies and templating of test data to allow access to the prior HTTP response in the current request. For APIs which do not use JSON additional *Content Handlers* can be created.

Care should be taken when using this functionality when you are creating a new API. If your API is so complex that it needs complex test files then you may wish to take that as a sign that your API itself too complex. One goal of gabbi is to encourage transparent and comprehensible APIs.

Though gabbi is written in Python and under the covers uses `unittest` data structures and processes, there is no requirement that the *Target Host* be a Python-based service. Anything talking HTTP can be tested. A *YAML Runner* makes it possible to simply create YAML files and point them at a running server.

18.2 Installation

As a Python package, gabbi is typically installed via pip:

```
pip install gabbi
```

You might want to create a virtual environment; an isolated context for Python packages, keeping gabbi cleanly separated from the rest of your system.

Python 3 comes with a built-in tool to create virtual environments:

```
python3 -m venv venv
. venv/bin/activate

pip install gabbi
```

This way we can later use `deactivate` and safely remove the `venv` directory, thus erasing any trace of gabbi from the system.

If you prefer to not install gabbi, or perhaps want to use it in a dynamic fashion in a CI setting, there is an official container image hosted at [docker hub](#) as `cdent/gabbi`. It allows running *gabbi-run* with any arguments you might need, providing tests on STDIN or via a mounted volume.

g

`gabbi.exception`, 58
`gabbi.fixture`, 55
`gabbi.json_parser`, 59
`gabbi.reporter`, 57
`gabbi.suite`, 56
`gabbi.utils`, 58

A

addError() (*gabbi.reporter.ConciseTestResult* method), 57
 addError() (*gabbi.reporter.PyTestResult* method), 57
 addExpectedFailure() (*gabbi.reporter.ConciseTestResult* method), 57
 addExpectedFailure() (*gabbi.reporter.PyTestResult* method), 57
 addFailure() (*gabbi.reporter.ConciseTestResult* method), 57
 addFailure() (*gabbi.reporter.PyTestResult* method), 58
 addSkip() (*gabbi.reporter.ConciseTestResult* method), 57
 addSkip() (*gabbi.reporter.PyTestResult* method), 58
 addSuccess() (*gabbi.reporter.ConciseTestResult* method), 57
 addUnexpectedSuccess() (*gabbi.reporter.ConciseTestResult* method), 57

C

ConciseTestResult (*class in gabbi.reporter*), 57
 ConciseTestRunner (*class in gabbi.reporter*), 57
 create_url() (*in module gabbi.utils*), 58

D

decode_response_content() (*in module gabbi.utils*), 58

E

extract_content_type() (*in module gabbi.utils*), 58

G

gabbi.exception (*module*), 58
 gabbi.fixture (*module*), 55
 gabbi.json_parser (*module*), 59

gabbi.reporter (*module*), 57
 gabbi.suite (*module*), 56
 gabbi.utils (*module*), 58
 GabbiDataLoadError, 58
 GabbiFixture (*class in gabbi.fixture*), 55
 GabbiFixtureError, 55
 GabbiFormatError, 58
 GabbiSuite (*class in gabbi.suite*), 56
 GabbiSyntaxWarning, 58
 get_colorizer() (*in module gabbi.utils*), 58
 getDescription() (*gabbi.reporter.ConciseTestResult* method), 57

H

host_info_from_target() (*in module gabbi.utils*), 58

L

load_yaml() (*in module gabbi.utils*), 58

N

nest() (*in module gabbi.fixture*), 56
 noop() (*in module gabbi.suite*), 56
 not_binary() (*in module gabbi.utils*), 58

P

parse() (*in module gabbi.json_parser*), 59
 parse_content_type() (*in module gabbi.utils*), 58
 printErrorList() (*gabbi.reporter.ConciseTestResult* method), 57
 PyTestResult (*class in gabbi.reporter*), 57

R

resultclass (*gabbi.reporter.ConciseTestRunner* attribute), 57
 run() (*gabbi.suite.GabbiSuite* method), 56

S

SkipAllFixture (*class in gabbi.fixture*), 55

`start()` (*gabbi.suite.GabbiSuite method*), 56
`start_fixture()` (*gabbi.fixture.GabbiFixture method*), 55
`start_fixture()` (*gabbi.fixture.SkipAllFixture method*), 56
`startTest()` (*gabbi.reporter.ConciseTestResult method*), 57
`stop()` (*gabbi.suite.GabbiSuite method*), 56
`stop_fixture()` (*gabbi.fixture.GabbiFixture method*), 55