hyper Documentation

Release 0.2.0

Cory Benfield

March 06, 2015

Contents

Release v0.2.0.

HTTP is changing under our feet. HTTP/1.1, our old friend, is being supplemented by the brand new HTTP/2 standard. HTTP/2 provides many benefits: improved speed, lower bandwidth usage, better connection management, and more.

`hyper` provides these benefits to your Python code. How? Like this:

```
from hyper import HTTP20Connection

conn = HTTP20Connection('twitter.com:443')
conn.request('GET', '/')
resp = conn.getresponse()

print(resp.read())
```

Simple. `hyper` is written in 100% pure Python, which means no C extensions. For recent versions of Python (3.4 and onward, and 2.7.9 and onward) it's entirely self-contained with no external dependencies.

`hyper` supports Python 3.4 and Python 2.7.9.

CHAPTER 1

Caveat Emptor!

Please be warned: `hyper` is in a very early alpha. You *will* encounter bugs when using it. In addition, there are very many rough edges. With that said, please try it out in your applications: I need your feedback to fix the bugs and file down the rough edges.

CHAPTER 2

Get Stuck In

The quickstart documentation will help get you going with `hyper`.

2.1 Quickstart Guide

First, congratulations on picking `hyper` for your HTTP/2 needs. `hyper` is the premier (and, as far as we're aware, the only) Python HTTP/2 library. In this section, we'll walk you through using `hyper`.

2.1.1 Installing hyper

To begin, you will need to install `hyper`. This can be done like so:

```
$ pip install hyper
```

If `pip` is not available to you, you can try:

```
$ easy_install hyper
```

If that fails, download the library from its GitHub page and install it using:

```
$ python setup.py install
```

Installation Requirements

The HTTP/2 specification requires very modern TLS support from any compliant implementation. When using Python 3.4 and later this is automatically provided by the standard library. For earlier releases of Python, we use PyOpenSSL to provide the TLS support we need.

Unfortunately, this is not always totally trivial. You will need to build PyOpenSSL against a version of OpenSSL that is at least 1.0.1, and to do that you'll actually need to obtain that version of OpenSSL.

To install against the relevant version of OpenSSL for your system, follow the instructions from the cryptography project, replacing references to `cryptography` with `hyper`.

2.1.2 Making Your First Request

With `hyper` installed, you can start making HTTP/2 requests. At this stage, `hyper` can only be used with services that *definitely* support HTTP/2. Before you begin, ensure that whichever service you're contacting definitely supports HTTP/2. For the rest of these examples, we'll use Twitter.

Begin by getting the Twitter homepage:

```
>>> from hyper import HTTP20Connection
>>> c = HTTP20Connection('twitter.com:443')
>>> c.request('GET', '/')
1
>>> resp = c.getresponse()
```

Used in this way, `hyper` behaves exactly like `http.client`. You can make sequential requests using the exact same API you're accustomed to. The only difference is that `HTTP20Connection.request()` returns a value, unlike the equivalent `http.client` function. The return value is the HTTP/2 *stream identifier*. If you're planning to use `hyper` in this very simple way, you can choose to ignore it, but it's potentially useful. We'll come back to it.

Once you've got the data, things continue to behave exactly like `http.client`:

```
>>> resp.getheader('content-encoding')
'deflate'
>>> resp.getheaders()
[('x-xss-protection', '1; mode=block')...
>>> resp.status
200
```

We know that Twitter has compressed the response body. `hyper` will automatically decompress that body for you, no input required:

```
>>> body = resp.read()
b'<!DOCTYPE html>\n<!--[if IE 8]><html clas ....
```

That's all it takes.

2.1.3 Streams

In HTTP/2, connections are divided into multiple streams. Each stream carries a single request-response pair. You may start multiple requests before reading the response from any of them, and switch between them using their stream IDs.

For example:

```
>>> from hyper import HTTP20Connection
>>> c = HTTP20Connection('twitter.com:443')
>>> first = c.request('GET', '/')
>>> second = c.request('GET', '/lukasaoz')
>>> third = c.request('GET', '/about')
>>> second_response = c.getresponse(second)
>>> first_response = c.getresponse(first)
>>> third_response = c.getresponse(third)
```

`hyper` will ensure that each response is matched to the correct request.

2.1.4 Requests Integration

Do you like requests? Of course you do, everyone does! It's a shame that requests doesn't support HTTP/2 though. To rectify that oversight, `hyper` provides a transport adapter that can be plugged directly into Requests, giving it instant HTTP/2 support.

All you have to do is identify a host that you'd like to communicate with over HTTP/2. Once you've worked that out, you can get started straight away:

```
>>> import requests
>>> from hyper.contrib import HTTP20Adapter
>>> s = requests.Session()
>>> s.mount('https://twitter.com', HTTP20Adapter())
>>> r = s.get('https://twitter.com')
>>> print(r.status_code)
200
```

This transport adapter is subject to all of the limitations that apply to `hyper`, and provides all of the goodness of requests.

A quick warning: some hosts will redirect to new hostnames, which may redirect you away from HTTP/2. Make sure you install the adapter for all the hostnames you're interested in:

```
>>> a = HTTP20Adapter()
>>> s.mount('https://twitter.com', a)
>>> s.mount('https://www.twitter.com', a)
```

CHAPTER 3

Advanced Documentation

More advanced topics are covered here.

3.1 Advanced Usage

This section of the documentation covers more advanced use-cases for `hyper`.

3.1.1 Responses as Context Managers

If you're concerned about having too many TCP sockets open at any one time, you may want to keep your connections alive only as long as you know you'll need them. In HTTP/2 this is generally not something you should do unless you're very confident you won't need the connection again anytime soon. However, if you decide you want to avoid keeping the connection open, you can use the `HTTP20Connection` as a context manager:

```
with HTTP20Connection('twitter.com:443') as conn:
    conn.request('GET', '/')
    data = conn.getresponse().read()

analyse(data)
```

You may not use any `HTTP20Response` objects obtained from a connection after that connection is closed. Interacting with these objects when a connection has been closed is considered undefined behaviour.

3.1.2 Multithreading

Currently, `hyper`'s `HTTP20Connection` class is **not** thread-safe. Thread-safety is planned for `hyper`'s core objects, but in this early alpha it is not a high priority.

To use `hyper` in a multithreaded context the recommended thing to do is to place each connection in its own thread. Each thread should then have a request queue and a response queue, and the thread should be able to spin over both, sending requests and returning responses. The stream identifiers provided by `hyper` can be used to match the two together.

3.1.3 SSL/TLS Certificate Verification

By default, all HTTP/2 connections are made over TLS, and `hyper` bundles certificate authorities that it uses to verify the offered TLS certificates. Currently certificate verification cannot be disabled.

3.1.4 Streaming Uploads

Just like the ever-popular `requests` module, `hyper` allows you to perform a 'streaming' upload by providing a file-like object to the 'data' parameter. This will cause `hyper` to read the data in 1kB at a time and send it to the remote server. You *must* set an accurate Content-Length header when you do this, as `hyper` won't set it for you.

3.1.5 Content Decompression

In HTTP/2 it's mandatory that user-agents support receiving responses that have their bodies compressed. As demonstrated in the quickstart guide, `hyper` transparently implements this decompression, meaning that responses are automatically decompressed for you. If you don't want this to happen, you can turn it off by passing the `decode_content` parameter to `read()`, like this:

```
>>> resp.read(decode_content=False)
b'\xc9...'
```

3.1.6 Flow Control & Window Managers

HTTP/2 provides a facility for performing 'flow control', enabling both ends of a HTTP/2 connection to influence the rate at which data is received. When used correctly flow control can be a powerful tool for maximising the efficiency of a connection. However, when used poorly, flow control leads to severe inefficiency and can adversely affect the throughput of the connection.

By default `hyper` does its best to manage the flow control window for you, trying to avoid severe inefficiencies. In general, though, the user has a much better idea of how to manage the flow control window than `hyper` will: you know your use case better than `hyper` possibly can.

For that reason, `hyper` provides a facility for using pluggable *window managers*. A *window manager* is an object that is in control of resizing the flow control window. This object gets informed about every frame received on the connection, and can make decisions about when to increase the size of the receive window. This object can take advantage of knowledge from layers above `hyper`, in the user's code, as well as knowledge from `hyper`'s layer.

To implement one of these objects, you will want to subclass the `BaseFlowControlManager` class and implement the `increase_window_size()` method. As a simple example, we can implement a very stupid flow control manager that always resizes the window in response to incoming data like this:

```
class StupidFlowControlManager(BaseFlowControlManager):
    def increase_window_size(self, frame_size):
        return frame_size
```

The *class* can then be plugged straight into a connection object:

```
HTTP20Connection('twitter.com:443', window_manager=StupidFlowControlManager)
```

Note that we don't plug an instance of the class in, we plug the class itself in. We do this because the connection object will spawn instances of the class in order to manage the flow control windows of streams in addition to managing the window of the connection itself.

3.1.7 Server Push

HTTP/2 provides servers with the ability to "push" additional resources to clients in response to a request, as if the client had requested the resources themselves. When minimizing the number of round trips is more critical than maximizing bandwidth usage, this can be a significant performance improvement.

Servers may declare their intention to push a given resource by sending the headers and other metadata of a request that would return that resource - this is referred to as a "push promise". They may do this before sending the response headers for the original request, after, or in the middle of sending the response body.

In order to receive pushed resources, the `HTTP20Connection` object must be constructed with `enable_push=True`.

You may retrieve the push promises that the server has sent *so far* by calling `getpushes()`, which returns a generator that yields `HTTP20Push` objects. Note that this method is not idempotent; promises returned in one call will not be returned in subsequent calls. If `capture_all=False` is passed (the default), the generator will yield all buffered push promises without blocking. However, if `capture_all=True` is passed, the generator will first yield all buffered push promises, then yield additional ones as they arrive, and terminate when the original stream closes. Using this parameter is only recommended when it is known that all pushed streams, or a specific one, are of higher priority than the original response, or when also processing the original response in a separate thread (N.B. do not do this; `hyper` is not yet thread-safe):

```
conn.request('GET', '/')
response = conn.getheaders()
for push in conn.getpushes(): # all pushes promised before response headers
    print(push.path)
conn.read()
for push in conn.getpushes(): # all other pushes
    print(push.path)
```

To cancel an in-progress pushed stream (for example, if the user already has the given path in cache), call `HTTP20Push.cancel()`.

`hyper` does not currently verify that pushed resources comply with the Same-Origin Policy, so users must take care that they do not treat pushed resources as authoritative without performing this check themselves (since the server push mechanism is only an optimization, and clients are free to issue requests for any pushed resources manually, there is little downside to simply ignoring suspicious ones).

3.1.8 Nghttp2

By default `hyper` uses its built-in pure-Python HPACK encoder and decoder. These are reasonably efficient, and suitable for most use cases. However, they do not produce the best compression ratio possible, and because they're written in pure-Python they incur a cost in memory usage above what is strictly necessary.

nghttp2 is a HTTP/2 library written in C that includes a HPACK encoder and decoder. `nghttp2`'s encoder produces extremely compressed output, and because it is written in C it is also fast and memory efficient. For this reason, performance conscious users may prefer to use `nghttp2`'s HPACK implementation instead of `hyper`'s.

You can do this very easily. If `nghttp2`'s Python bindings are installed, `hyper` will transparently switch to using `nghttp2`'s HPACK implementation instead of its own. No configuration is required.

Instructions for installing `nghttp2` are available here.

CHAPTER 4

Contributing

Want to contribute? Awesome! This guide goes into detail about how to contribute, and provides guidelines for project contributions.

4.1 Contributor's Guide

If you're reading this you're probably interested in contributing to `hyper`. First, I'd like to say: thankyou! Projects like this one live-and-die based on the support they receive from others, and the fact that you're even *considering* supporting `hyper` is incredibly generous of you.

This document lays out guidelines and advice for contributing to `hyper`. If you're thinking of contributing, start by reading this thoroughly and getting a feel for how contributing to the project works. If you've still got questions after reading this, you should go ahead and contact the author: he'll be happy to help.

The guide is split into sections based on the type of contribution you're thinking of making, with a section that covers general guidelines for all contributors.

4.1.1 All Contributions

Be Cordial Or Be On Your Way

`hyper` has one very important guideline governing all forms of contribution, including things like reporting bugs or requesting features. The guideline is be cordial or be on your way. **All contributions are welcome**, but they come with an implicit social contract: everyone must be treated with respect.

This can be a difficult area to judge, so the maintainer will enforce the following policy. If any contributor acts rudely or aggressively towards any other contributor, **regardless of whether they perceive themselves to be acting in retaliation for an earlier breach of this guideline**, they will be subject to the following steps:

1. They must apologise. This apology must be genuine in nature: "I'm sorry you were offended" is not sufficient. The judgement of 'genuine' is at the discretion of the maintainer.
2. If the apology is not offered, any outstanding and future contributions from the violating contributor will be rejected immediately.

Everyone involved in the `hyper` project, the maintainer included, is bound by this policy. Failing to abide by it leads to the offender being kicked off the project.

Get Early Feedback

If you are contributing, do not feel the need to sit on your contribution until it is perfectly polished and complete. It helps everyone involved for you to seek feedback as early as you possibly can. Submitting an early, unfinished version of your contribution for feedback in no way prejudices your chances of getting that contribution accepted, and can save you from putting a lot of work into a contribution that is not suitable for the project.

Contribution Suitability

The project maintainer has the last word on whether or not a contribution is suitable for `hyper`. All contributions will be considered, but from time to time contributions will be rejected because they do not suit the project.

If your contribution is rejected, don't despair! So long as you followed these guidelines, you'll have a much better chance of getting your next contribution accepted.

4.1.2 Code Contributions

Steps

When contributing code, you'll want to follow this checklist:

1. Fork the repository on GitHub.
2. Run the tests to confirm they all pass on your system. If they don't, you'll need to investigate why they fail. If you're unable to diagnose this yourself, raise it as a bug report by following the guidelines in this document: *Bug Reports*.
3. Write tests that demonstrate your bug or feature. Ensure that they fail.
4. Make your change.
5. Run the entire test suite again, confirming that all tests pass *including the ones you just added*.
6. Send a GitHub Pull Request to the main repository's `development` branch. GitHub Pull Requests are the expected method of code collaboration on this project. If you object to the GitHub workflow, you may mail a patch to the maintainer.

The following sub-sections go into more detail on some of the points above.

Tests & Code Coverage

`hyper` has a substantial suite of tests, both unit tests and integration tests, and has 100% code coverage. Whenever you contribute, you must write tests that exercise your contributed code, and you must not regress the code coverage.

To run the tests, you need to install tox. Once you have, you can run the tests against all supported platforms by simply executing `tox`.

If you're having trouble running the tests, please consider raising a bug report using the guidelines in the *Bug Reports* section.

If you've done this but want to get contributing right away, you can take advantage of the fact that `hyper` uses a continuous integration system. This will automatically run the tests against any pull request raised against the main `hyper` repository. The continuous integration system treats a regression in code coverage as a failure of the test suite.

Before a contribution is merged it must have a green run through the CI system.

Code Review

Contributions will not be merged until they've been code reviewed. You should implement any code review feedback unless you strongly object to it. In the event that you object to the code review feedback, you should make your case clearly and calmly. If, after doing so, the feedback is judged to still apply, you must either apply the feedback or withdraw your contribution.

New Contributors

If you are new or relatively new to Open Source, welcome! `hyper` aims to be a gentle introduction to the world of Open Source. If you're concerned about how best to contribute, please consider mailing the maintainer and asking for help.

Please also check the *Get Early Feedback* section.

4.1.3 Documentation Contributions

Documentation improvements are always welcome! The documentation files live in the `docs/` directory of the codebase. They're written in reStructuredText, and use Sphinx to generate the full suite of documentation.

When contributing documentation, please attempt to follow the style of the documentation files. This means a soft-limit of 79 characters wide in your text files and a semi-formal prose style.

4.1.4 Bug Reports

Bug reports are hugely important! Before you raise one, though, please check through the GitHub issues, **both open and closed**, to confirm that the bug hasn't been reported before. Duplicate bug reports are a huge drain on the time of other contributors, and should be avoided as much as possible.

4.1.5 Feature Requests

Feature requests are always welcome, but please note that all the general guidelines for contribution apply. Also note that the importance of a feature request *without* an associated Pull Request is always lower than the importance of one *with* an associated Pull Request: code is more valuable than ideas.

CHAPTER 5

Frequently Asked Questions

Got a question? I might have answered it already! Take a look.

5.1 Frequently Asked Questions

`hyper` is a project that's under active development, and is in early alpha. As a result, there are plenty of rough edges and bugs. This section of the documentation attempts to address some of your likely questions.

If you find there is no answer to your question in this list, please send me an email. My email address can be found on my GitHub profile page.

5.1.1 What version of the HTTP/2 draft specification does `hyper` support?

Currently, `hyper` supports versions 14, 15, and 16 of the HTTP/2 draft specification, and version 9 of the HPACK draft specification. `hyper` will be updated to keep up with the HTTP/2 draft specifications as they progress.

5.1.2 Does `hyper` support HTTP/2 flow control?

It should! If you find it doesn't, that's a bug: please report it on GitHub.

5.1.3 Does `hyper` support Server Push?

Yes! See *Server Push*.

5.1.4 I hit a bug! What should I do?

Please tell me about it using the GitHub page for the project, here, by filing an issue. There will definitely be bugs as `hyper` is very new, and reporting them is the fastest way to get them fixed.

When you report them, please follow the contribution guidelines in the README. It'll make it a lot easier for me to fix the problem.

5.1.5 Updates

Further questions will be added here over time. Please check back regularly.

CHAPTER 6

API Documentation

The `hyper` API is documented in these pages.

6.1 Interface

This section of the documentation covers the interface portions of `hyper`.

6.1.1 Primary HTTP/2 Interface

class `hyper.`**`HTTP20Connection`**(*host*, *port=None*, *window_manager=None*, *enable_push=False*, ***kwargs*)

An object representing a single HTTP/2 connection to a server.

This object behaves similarly to the Python standard library's `HTTPConnection` object, with a few critical differences.

Most of the standard library's arguments to the constructor are irrelevant for HTTP/2 or not supported by hyper.

Parameters

- **host** – The host to connect to. This may be an IP address or a hostname, and optionally may include a port: for example, `'twitter.com'`, `'twitter.com:443'` or `'127.0.0.1'`.
- **port** – (optional) The port to connect to. If not provided and one also isn't provided in the `host` parameter, defaults to 443.
- **window_manager** – (optional) The class to use to manage flow control windows. This needs to be a subclass of the `BaseFlowControlManager`. If not provided, `FlowControlManager` will be used.
- **enable_push** – (optional) Whether the server is allowed to push resources to the client (see `getpushes()`).

`close`()

Close the connection to the server.

Returns Nothing.

`connect`()

Connect to the server specified when the object was created. This is a no-op if we're already connected.

Returns Nothing.

endheaders(*message_body=None*, *final=False*, *stream_id=None*)
Sends the prepared headers to the server. If the `message_body` argument is provided it will also be sent to the server as the body of the request, and the stream will immediately be closed. If the `final` argument is set to True, the stream will also immediately be closed: otherwise, the stream will be left open and subsequent calls to `send()` will be required.

Parameters

- **message_body** – (optional) The body to send. May not be provided assuming that `send()` will be called.
- **final** – (optional) If the `message_body` parameter is provided, should be set to `True` if no further data will be provided via calls to `send()`.
- **stream_id** – (optional) The stream ID of the request to finish sending the headers on.

Returns Nothing.

getpushes(*stream_id=None*, *capture_all=False*)
Returns a generator that yields push promises from the server. **Note that this method is not idempotent**: promises returned in one call will not be returned in subsequent calls. Iterating through generators returned by multiple calls to this method simultaneously results in undefined behavior.

Parameters

- **stream_id** – (optional) The stream ID of the request for which to get push promises.
- **capture_all** – (optional) If `False`, the generator will yield all buffered push promises without blocking. If `True`, the generator will first yield all buffered push promises, then yield additional ones as they arrive, and terminate when the original stream closes.

Returns A generator of `HTTP20Push` objects corresponding to the streams pushed by the server.

getresponse(*stream_id=None*)
Should be called after a request is sent to get a response from the server. If sending multiple parallel requests, pass the stream ID of the request whose response you want. Returns a `HTTP20Response` instance. If you pass no `stream_id`, you will receive the oldest `HTTPResponse` still outstanding.

Parameters **stream_id** – (optional) The stream ID of the request for which to get a response.

Returns A `HTTP20Response` object.

network_buffer_size = **None**
The size of the in-memory buffer used to store data from the network. This is used as a performance optimisation. Increase buffer size to improve performance: decrease it to conserve memory. Defaults to 64kB.

putheader(*header*, *argument*, *stream_id=None*)
Sends an HTTP header to the server, with name `header` and value `argument`.

Unlike the `httplib` version of this function, this version does not actually send anything when called. Instead, it queues the headers up to be sent when you call `endheaders()`.

This method ensures that headers conform to the HTTP/2 specification. In particular, it strips out the `Connection` header, as that header is no longer valid in HTTP/2. This is to make it easy to write code that runs correctly in both HTTP/1.1 and HTTP/2.

Parameters

- **header** – The name of the header.
- **argument** – The value of the header.

- **stream_id** – (optional) The stream ID of the request to add the header to.

Returns Nothing.

putrequest(*method*, *selector*, ***kwargs*)
This should be the first call for sending a given HTTP request to a server. It returns a stream ID for the given connection that should be passed to all subsequent request building calls.

Parameters

- **method** – The request method, e.g. `'GET'`.
- **selector** – The path selector.

Returns A stream ID for the request.

receive_frame(*frame*)
Handles receiving frames intended for the stream.

request(*method*, *url*, *body=None*, *headers={}*)
This will send a request to the server using the HTTP request method `method` and the selector `url`. If the `body` argument is present, it should be string or bytes object of data to send after the headers are finished. Strings are encoded as UTF-8. To use other encodings, pass a bytes object. The Content-Length header is set to the length of the body field.

Parameters

- **method** – The request method, e.g. `'GET'`.
- **url** – The URL to contact, e.g. `'/path/segment'`.
- **body** – (optional) The request body to send. Must be a bytestring or a file-like object.
- **headers** – (optional) The headers to send on the request.

Returns A stream ID for the request.

send(*data*, *final=False*, *stream_id=None*)
Sends some data to the server. This data will be sent immediately (excluding the normal HTTP/2 flow control rules). If this is the last data that will be sent as part of this request, the `final` argument should be set to `True`. This will cause the stream to be closed.

Parameters

- **data** – The data to send.
- **final** – (optional) Whether this is the last bit of data to be sent on this request.
- **stream_id** – (optional) The stream ID of the request to send the data on.

Returns Nothing.

class `hyper.`**HTTP20Response**(*headers*, *stream*)
An `HTTP20Response` wraps the HTTP/2 response from the server. It provides access to the response headers and the entity body. The response is an iterable object and can be used in a with statement (though due to the persistent connections used in HTTP/2 this has no effect, and is done soley for compatibility).

close()
Close the response. In effect this closes the backing HTTP/2 stream.

Returns Nothing.

fileno()
Return the `fileno` of the underlying socket. This function is currently not implemented.

getheader(*name*, *default=None*)
Return the value of the header `name`, or `default` if there is no header matching `name`. If there is more than one header with the value `name`, return all of the values joined by ', '. If `default` is any iterable other than a single string, its elements are similarly returned joined by commas.

Parameters

- **name** – The name of the header to get the value of.
- **default** – (optional) The return value if the header wasn't sent.

Returns The value of the header.

getheaders()
Get all the headers sent on the response.

Returns A list of `(header, value)` tuples.

gettrailer(*name*, *default=None*)
Return the value of the trailer `name`, or `default` if there is no trailer matching `name`. If there is more than one trailer with the value `name`, return all of the values joined by ', '. If `default` is any iterable other than a single string, its elements are similarly returned joined by commas.

> **Warning:** Note that this method requires that the stream is totally exhausted. This means that, if you have not completely read from the stream, all stream data will be read into memory.

Parameters

- **name** – The name of the trailer to get the value of.
- **default** – (optional) The return value if the trailer wasn't sent.

Returns The value of the trailer.

gettrailers()
Get all the trailers sent on the response.

> **Warning:** Note that this method requires that the stream is totally exhausted. This means that, if you have not completely read from the stream, all stream data will be read into memory.

Returns A list of `(header, value)` tuples.

read(*amt=None*, *decode_content=True*)
Reads the response body, or up to the next `amt` bytes.

Parameters

- **amt** – (optional) The amount of data to read. If not provided, all the data will be read from the response.
- **decode_content** – (optional) If `True`, will transparently decode the response data.

Returns The read data. Note that if `decode_content` is set to `True`, the actual amount of data returned may be different to the amount requested.

reason = None
The reason phrase returned by the server. This is not used in HTTP/2, and so is always the empty string.

status = None
The status code returned by the server.

class `hyper.`**HTTP20Push**(*request_headers*, *stream*)
Represents a request-response pair sent by the server through the server push mechanism.

authority = None
The authority of the simulated request (usually host:port)

cancel()
Cancel the pushed response and close the stream.

Returns Nothing.

getrequestheader(*name*, *default=None*)
Return the value of the simulated request header `name`, or `default` if there is no header matching `name`. If there is more than one header with the value `name`, return all of the values joined by `', '`. If `default` is any iterable other than a single string, its elements are similarly returned joined by commas.

Parameters

- **name** – The name of the header to get the value of.
- **default** – (optional) The return value if the header wasn't sent.

Returns The value of the header.

getrequestheaders()
Get all the simulated request headers.

Returns A list of `(header, value)` tuples.

getresponse()
Get the pushed response provided by the server.

Returns A `HTTP20Response` object representing the pushed response.

method = None
The method of the simulated request (must be safe and cacheable, e.g. GET)

path = None
The path of the simulated request

scheme = None
The scheme of the simulated request

6.1.2 Requests Transport Adapter

class `hyper.contrib.`**HTTP20Adapter**(**args*, ***kwargs*)
A Requests Transport Adapter that uses hyper to send requests over HTTP/2. This implements some degree of connection pooling to maximise the HTTP/2 gain.

build_response(*request*, *resp*)
Builds a Requests' response object. This emulates most of the logic of the standard fuction but deals with the lack of the `.headers` property on the HTTP20Response object.

Additionally, this function builds in a number of features that are purely for HTTPie. This is to allow maximum compatibility with what urllib3 does, so that HTTPie doesn't fall over when it uses us.

connections = None
A mapping between HTTP netlocs and `HTTP20Connection` objects.

get_connection(*netloc*)
Gets an appropriate HTTP/2 connection object based on netloc.

send(*request*, *stream=False*, ***kwargs*)
Sends a HTTP message to the server.

6.1.3 Flow Control

class hyper.http20.window.**BaseFlowControlManager**(*initial_window_size*, *document_size=None*)

The abstract base class for flow control managers.

This class defines the interface for pluggable flow-control managers. A flow-control manager defines a flow-control policy, which basically boils down to deciding when to increase the flow control window.

This decision can be based on a number of factors:

- the initial window size,
- the size of the document being retrieved,
- the size of the received data frames,
- any other information the manager can obtain

A flow-control manager may be defined at the connection level or at the stream level. If no stream-level flow-control manager is defined, an instance of the connection-level flow control manager is used.

A class that inherits from this one must not adjust the member variables defined in this class. They are updated and set by methods on this class.

blocked()

Called whenever the remote endpoint reports that it is blocked behind the flow control window.

When this method is called the remote endpoint is signaling that it has more data to send and that the transport layer is capable of transmitting it, but that the HTTP/2 flow control window prevents it being sent.

This method should return the size by which the window should be incremented, which may be zero. This method should *not* adjust any of the member variables of this class.

Returns The amount to increase the receive window by. Return zero if the window should not be increased.

document_size = None

The size of the document being retrieved, in bytes. This is retrieved from the Content-Length header, if provided. Note that the total number of bytes that will be received may be larger than this value due to HTTP/2 padding. It should not be assumed that simply because the the document size is smaller than the initial window size that there will never be a need to increase the window size.

increase_window_size(*frame_size*)

Determine whether or not to emit a WINDOWUPDATE frame.

This method should be overridden to determine, based on the state of the system and the size of the received frame, whether or not a WindowUpdate frame should be sent for the stream.

This method should *not* adjust any of the member variables of this class.

Note that this method is called before the window size is decremented as a result of the frame being handled.

Parameters **frame_size** – The size of the received frame. Note that this *may* be zero. When this parameter is zero, it's possible that a WINDOWUPDATE frame may want to be emitted anyway. A zero-length frame size is usually associated with a change in the size of the receive window due to a SETTINGS frame.

Returns The amount to increase the receive window by. Return zero if the window should not be increased.

initial_window_size = None
The initial size of the connection window in bytes. This is set at creation time.

window_size = None
The current size of the connection window. Any methods overridden by the user must not adjust this value.

class hyper.http20.window.**FlowControlManager**(*initial_window_size*, *document_size=None*)
hyper's default flow control manager.

This implements hyper's flow control algorithms. This algorithm attempts to reduce the number of WINDOWUPDATE frames we send without blocking the remote endpoint behind the flow control window.

This algorithm will become more complicated over time. In the current form, the algorithm is very simple:

- When the flow control window gets less than 1/4 of the maximum size, increment back to the maximum.
- Otherwise, if the flow control window gets to less than 1kB, increment back to the maximum.

6.1.4 Exceptions

class hyper.http20.exceptions.**HTTP20Error**
The base class for all of hyper's HTTP/2-related exceptions.

class hyper.http20.exceptions.**HPACKEncodingError**
An error has been encountered while performing HPACK encoding.

class hyper.http20.exceptions.**HPACKDecodingError**
An error has been encountered while performing HPACK decoding.

class hyper.http20.exceptions.**ConnectionError**
The remote party signalled an error affecting the entire HTTP/2 connection, and the connection has been closed.

h

S

W