

HTTPIe: a CLI, cURL-like tool for humans

HTTPIe (pronounced *aitch-tee-tee-pie*) is a command line HTTP client. Its goal is to make CLI interaction with web services as human-friendly as possible. It provides a simple `http` command that allows for sending arbitrary HTTP requests using a simple and natural syntax, and displays colorized output. HTTPIe can be used for testing, debugging, and generally interacting with HTTP servers.

```
bash
$ curl -i -X PUT httpbin.org/put -H Content-Type:application/json -d '{"hello": "world"}'
HTTP/1.1 200 OK
Connection: keep-alive
Server: gunicorn/19.9.0
Date: Fri, 02 Nov 2018 16:53:05 GMT
Content-Type: application/json
Content-Length: 452
Access-Control-Allow-Credentials: true
Access-Control-Allow-Origin: *
```

```
bash
$ http PUT httpbin.org/put hello=world
HTTP/1.1 200 OK
Access-Control-Allow-Credentials: true
Access-Control-Allow-Origin: *
Connection: keep-alive
Content-Length: 452
Content-Type: application/json
Date: Fri, 02 Nov 2018 16:53:05 GMT
Server: gunicorn/19.9.0
Via: 1.1 vegur

{
  "args": {},
  "data": "{\\\"hello\\\": \\\"world\\\"}",
  "files": {},
  "form": {},
  "headers": {
    "Accept": "application/json, */*",
    "Accept-Encoding": "gzip, deflate",
    "Connection": "close",
    "Content-Length": "18",
    "Content-Type": "application/json",
    "Host": "httpbin.org",
    "User-Agent": "HTTPie/1.0.0"
  },
  "json": {
    "hello": "world"
  },
  "origin": "89.102.136.126",
  "url": "http://httpbin.org/put"
}
```

Contents

1 Main features	6
2 Installation	6
2.1 macOS	6
2.2 Linux	6
2.3 Windows, etc.	6
2.4 Development version	7
2.5 Python version	7
3 Usage	7
3.1 Examples	7
4 HTTP method	8
5 Request URL	8
5.1 Querystring parameters	8
5.2 URL shortcuts for <code>localhost</code>	9
5.3 Custom default scheme	9
6 Request items	9
6.1 Escaping rules	10
7 JSON	10
7.1 Default behaviour	10
7.2 Explicit JSON	11
7.3 Non-string JSON fields	11
8 Forms	11
8.1 Regular forms	11
8.2 File upload forms	12
9 HTTP headers	12
9.1 Default request headers	12
9.2 Empty headers and header un-setting	12
10 Authentication	13
10.1 Basic auth	13
10.2 Digest auth	13
10.3 Password prompt	13
10.4 <code>.netrc</code>	13
10.5 Auth plugins	13
11 HTTP redirects	14
11.1 Follow <code>Location</code>	14
11.2 Showing intermediary redirect responses	14
11.3 Limiting maximum redirects followed	14

12	Proxies	14
12.1	Environment variables	14
12.2	SOCKS	15
13	HTTPS	15
13.1	Server SSL certificate verification	15
13.2	Custom CA bundle	15
13.3	Client side SSL certificate	15
13.4	SSL version	15
13.5	SNI (Server Name Indication)	16
14	Output options	16
14.1	What parts of the HTTP exchange should be printed	16
14.2	Viewing intermediary requests/responses	17
14.3	Conditional body download	17
15	Redirected Input	17
15.1	Request data from a filename	18
16	Terminal output	18
16.1	Colors and formatting	18
16.2	Binary data	19
17	Redirected output	19
18	Download mode	20
18.1	Downloaded file name	20
18.2	Piping while downloading	20
18.3	Resuming downloads	20
18.4	Other notes	20
19	Streamed responses	21
19.1	Disabling buffering	21
19.2	Examples use cases	21
20	Sessions	21
20.1	Named sessions	22
20.2	Anonymous sessions	22
20.3	Readonly session	22
21	Config	22
21.1	Config file location	22
21.2	Configurable options	22
21.2.1	default_options	22
21.2.2	__meta__	23
21.3	Un-setting previously specified options	23
22	Scripting	23

22.1	Best practices	23
23	Meta	23
23.1	Interface design	23
23.2	User support	24
23.3	Related projects	24
23.3.1	Dependencies	24
23.3.2	HTTPie friends	24
23.4	Contributing	24
23.5	Change log	25
23.6	Artwork	25
23.7	Licence	25
23.8	Authors	25

1 Main features

- Expressive and intuitive syntax
- Formatted and colorized terminal output
- Built-in JSON support
- Forms and file uploads
- HTTPS, proxies, and authentication
- Arbitrary request data
- Custom headers
- Persistent sessions
- Wget-like downloads
- Python 2.6, 2.7 and 3.x support
- Linux, Mac OS X and Windows support
- Plugins
- Documentation
- Test coverage

2 Installation

2.1 macOS

On macOS, HTTPie can be installed via [Homebrew](#) (recommended):

```
$ brew install httpie
```

A MacPorts *port* is also available:

```
$ port install httpie
```

2.2 Linux

Most Linux distributions provide a package that can be installed using the system package manager, e.g.:

```
# Debian-based distributions such as Ubuntu:
$ apt-get install httpie

# RPM-based distributions:
$ yum install httpie

# Arch Linux
$ pacman -S httpie
```

2.3 Windows, etc.

A universal installation method (that works on Windows, Mac OS X, Linux, ..., and always provides the latest version) is to use [pip](#):

```
# Make sure we have an up-to-date version of pip and setuptools:
$ pip install --upgrade pip setuptools

$ pip install --upgrade httpie
```

(If pip installation fails for some reason, you can try `easy_install httpie` as a fallback.)

2.4 Development version

The latest development version can be installed directly from GitHub:

```
# Mac OS X via Homebrew
$ brew install httpie --HEAD

# Universal
$ pip install --upgrade https://github.com/jkbrzt/httpie/archive/master.tar.gz
```

2.5 Python version

Although Python 2.6 and 2.7 are supported as well, it is recommended to install HTTPie against the latest Python 3.x whenever possible. That will ensure that some of the newer HTTP features, such as [SNI \(Server Name Indication\)](#), work out of the box. Python 3 is the default for Homebrew installations starting with version 0.9.4. To see which version HTTPie uses, run `http --debug`.

3 Usage

Hello World:

```
$ http httpie.org
```

Synopsis:

```
$ http [flags] [METHOD] URL [ITEM [ITEM]]
```

See also `http --help`.

3.1 Examples

Custom [HTTP method](#), [HTTP headers](#) and [JSON data](#):

```
$ http PUT example.org X-API-Token:123 name=John
```

Submitting [forms](#):

```
$ http -f POST example.org hello=World
```

See the request that is being sent using one of the [output options](#):

```
$ http -v example.org
```

Use [Github API](#) to post a comment on an [issue](#) with [authentication](#):

```
$ http -a USERNAME POST https://api.github.com/repos/jkbrzt/httpie/issues/83/comments body='HTTPie is awesome! :heart:'
```

Upload a file using **redirected input**:

```
$ http example.org < file.json
```

Download a file and save it via **redirected output**:

```
$ http example.org/file > file
```

Download a file **wget** style:

```
$ http --download example.org/file
```

Use named **sessions** to make certain aspects of the communication persistent between requests to the same host:

```
$ http --session=logged-in -a username:password httpbin.org/get API-Key:123  
$ http --session=logged-in httpbin.org/headers
```

Set a custom **Host** header to work around missing DNS records:

```
$ http localhost:8000 Host:example.com
```

4 HTTP method

The name of the HTTP method comes right before the URL argument:

```
$ http DELETE example.org/todos/7
```

Which looks similar to the actual **Request-Line** that is sent:

```
DELETE /todos/7 HTTP/1.1
```

When the **METHOD** argument is omitted from the command, HTTPie defaults to either **GET** (with no request data) or **POST** (with request data).

5 Request URL

The only information HTTPie needs to perform a request is a URL. The default scheme is, somewhat unsurprisingly, `http://`, and can be omitted from the argument – `http example.org` works just fine.

5.1 Querystring parameters

If you find yourself manually constructing URLs with on the terminal, you may appreciate the `param==value` syntax for appending URL parameters. With that, you don't have to worry about escaping the `&` separators for your shell. Also, special characters in parameter values, will also automatically escaped (HTTPie otherwise expects the URL to be already escaped). To search for HTTPie logo on Google Images you could use this command:

```
$ http www.google.com search=='HTTPie logo' tbm==isch
```



```
GET /?search=HTTPIe+logo&tbm=isch HTTP/1.1
```

5.2 URL shortcuts for localhost

Additionally, curl-like shorthand for localhost is supported. This means that, for example `:3000` would expand to `http://localhost:3000`. If the port is omitted, then port 80 is assumed.

```
$ http :/foo
```

```
GET /foo HTTP/1.1
Host: localhost
```

```
$ http :3000/bar
```

```
GET /bar HTTP/1.1
Host: localhost:3000
```

```
$ http :
```

```
GET / HTTP/1.1
Host: localhost
```

5.3 Custom default scheme

You can use the `--default-scheme <URL_SCHEME>` option to create shortcuts for other protocols than HTTP:

```
$ alias https='http --default-scheme=https'
```

6 Request items

There are a few different *request item* types that provide a convenient mechanism for specifying HTTP headers, simple JSON and form data, files, and URL parameters.

They are key/value pairs specified after the URL. All have in common that they become part of the actual request that is sent and that their type is distinguished only by the separator used: `:`, `=`, `:=`, `==`, `@`, `=@`, and `:=@`. The ones with an `@` expect a file path as value.

Item Type	Description
HTTP Headers <code>Name:Value</code>	Arbitrary HTTP header, e.g. <code>X-API-Token:123</code> .
URL parameters <code>name==value</code>	Appends the given name/value pair as a query string parameter to the URL. The <code>==</code> separator is used.
Data Fields <code>field=value</code> , <code>field=@file.txt</code>	Request data fields to be serialized as a JSON object (default), or to be form-encoded (<code>--form</code> , <code>-f</code>).
Raw JSON fields <code>field:=json</code> , <code>field:=@file.json</code>	Useful when sending JSON and one or more fields need to be a Boolean, Number, nested Object, or an Array, e.g., <code>meals:='["ham","spam"]'</code> or <code>pies:=[1,2,3]</code> (note the quotes).

Form File Fields field@/dir/file	Only available with --form, -f. For example screenshot@~/Pictures/img.png. The presence of a file field results in a multipart/form-data request.
-------------------------------------	---

Note that data fields aren't the only way to specify request data: [Redirected input](#) is a mechanism for passing arbitrary data request request.

6.1 Escaping rules

You can use \ to escape characters that shouldn't be used as separators (or parts thereof). For instance, foo\==bar will become a data key/value pair (foo= and bar) instead of a URL parameter.

Often it is necessary to quote the values, e.g. foo='bar baz'.

If any of the field names or headers starts with a minus (e.g., -fieldname), you need to place all such items after the special token -- to prevent confusion with --arguments:

```
$ http httpbin.org/post -- -name-starting-with-dash=foo -Unusual-Header:bar
```

```
POST /post HTTP/1.1
-Unusual-Header: bar
Content-Type: application/json

{
  "-name-starting-with-dash": "value"
}
```

7 JSON

JSON is the *lingua franca* of modern web services and it is also the **implicit content type** HTTPie by default uses.

Simple example:

```
$ http PUT example.org name=John email=john@example.org
```

```
PUT / HTTP/1.1
Accept: application/json, */*
Accept-Encoding: gzip, deflate
Content-Type: application/json
Host: example.org

{
  "name": "John",
  "email": "john@example.org"
}
```

7.1 Default behaviour

If your command includes some data [request items](#), they are serialized as a JSON object by default. HTTPie also automatically sets the following headers, both of which can be overwritten:

Content-Type	application/json
Accept	application/json, */*

7.2 Explicit JSON

You can use `--json`, `-j` to explicitly set `Accept` to `application/json` regardless of whether you are sending data (it's a shortcut for setting the header via the usual header notation: `http url Accept:'application/json, */*'').` Additionally, HTTPie will try to detect JSON responses even when the `Content-Type` is incorrectly `text/plain` or `unknown`.

7.3 Non-string JSON fields

Non-string fields use the `:=` separator, which allows you to embed raw JSON into the resulting object. Text and raw JSON files can also be embedded into fields using `@` and `:=@`:

```
$ http PUT api.example.com/person/1 \  
  name=John \  
  age:=29 married:=false hobbies:='["http", "pies"]' \ # Raw JSON  
  description=@about-john.txt \ # Embed text file  
  bookmarks:=@bookmarks.json # Embed JSON file
```

```
PUT /person/1 HTTP/1.1  
Accept: application/json, */*  
Content-Type: application/json  
Host: api.example.com  
  
{  
  "age": 29,  
  "hobbies": [  
    "http",  
    "pies"  
  ],  
  "description": "John is a nice guy who likes pies.",  
  "married": false,  
  "name": "John",  
  "bookmarks": {  
    "HTTPie": "http://httpie.org",  
  }  
}
```

Please note that with this syntax the command gets unwieldy when sending complex data. In that case it's always better to use [redirected input](#):

```
$ http POST api.example.com/person/1 < person.json
```

8 Forms

Submitting forms is very similar to sending [JSON](#) requests. Often the only difference is in adding the `--form`, `-f` option, which ensures that data fields are serialized as, and `Content-Type` is set to, `application/x-www-form-urlencoded; charset=utf-8`. It is possible to make form data the implicit content type instead of JSON via the [config](#) file.

8.1 Regular forms

```
$ http --form POST api.example.org/person/1 name='John Smith'
```

```
POST /person/1 HTTP/1.1
Content-Type: application/x-www-form-urlencoded; charset=utf-8

name=John+Smith
```

8.2 File upload forms

If one or more file fields is present, the serialization and content type is `multipart/form-data`:

```
$ http -f POST example.com/jobs name='John Smith' cv@~/Documents/cv.pdf
```

The request above is the same as if the following HTML form were submitted:

```
<form enctype="multipart/form-data" method="post" action="http://example.com/jobs">
  <input type="text" name="name" />
  <input type="file" name="cv" />
</form>
```

Note that `@` is used to simulate a file upload form field, whereas `=@` just embeds the file content as a regular text field value.

9 HTTP headers

To set custom headers you can use the `Header:Value` notation:

```
$ http example.org User-Agent:Bacon/1.0 'Cookie:valued-visitor=yes;foo=bar' \
  X-Foo:Bar Referer:http://httpie.org/
```

```
GET / HTTP/1.1
Accept: */*
Accept-Encoding: gzip, deflate
Cookie: valued-visitor=yes;foo=bar
Host: example.org
Referer: http://httpie.org/
User-Agent: Bacon/1.0
X-Foo: Bar
```

9.1 Default request headers

There are a couple of default headers that HTTPie sets:

```
GET / HTTP/1.1
Accept: */*
Accept-Encoding: gzip, deflate
User-Agent: HTTPie/<version>
Host: <taken-from-URL>
```

Any of those—except for `Host`—can be overwritten and some of them unset.

9.2 Empty headers and header un-setting

To unset a previously specified header (such a one of the default headers), use `Header::`

```
$ http httpbin.org/headers Accept: User-Agent:
```

To send a header with an empty value, use `Header:`:

```
$ http httpbin.org/headers 'Header:'
```

10 Authentication

The currently supported authentication schemes are Basic and Digest (see [auth plugins](#) for more). There are two flags that control authentication:

<code>--auth, -a</code>	Pass a <code>username:password</code> pair as the argument. Or, if you only specify a username (<code>-a username</code>), you'll be prompted for the password before the request is sent. To send an empty password, pass <code>username:</code> . The <code>username:password@hostname</code> URL syntax is supported as well (but credentials passed via <code>-a</code> have higher priority).
<code>--auth-type, -A</code>	Specify the auth mechanism. Possible values are <code>basic</code> and <code>digest</code> . The default value is <code>basic</code> so it can often be omitted.

10.1 Basic auth

```
$ http -a username:password example.org
```

10.2 Digest auth

```
$ http -A digest -a username:password example.org
```

10.3 Password prompt

```
$ http -a username example.org
```

10.4 .netrc

Authorization information from your `~/.netrc` file is honored as well:

```
$ cat ~/.netrc
machine httpbin.org
login httpie
password test

$ http httpbin.org/basic-auth/httpie/test
HTTP/1.1 200 OK
[...]
```

10.5 Auth plugins

Additional authentication mechanism can be installed as plugins. They can be found on the [Python Package Index](#). Here's a few picks:

- [httpie-oauth](#): OAuth

- [httpie-hmac-auth](#): HMAC
- [httpie-ntlm](#): NTLM (NT LAN Manager)
- [httpie-negotiate](#): SPNEGO (GSS Negotiate)
- [requests-hawk](#): Hawk
- [httpie-api-auth](#): ApiAuth
- [httpie-edgegrid](#): EdgeGrid
- [httpie-jwt-auth](#): JWTAuth (JSON Web Tokens)

11 HTTP redirects

By default, HTTP redirects are not followed and only the first response is shown:

```
$ http httpbin.org/redirect/3
```

11.1 Follow Location

To instruct HTTPie to follow the `Location` header of 30x responses and show the final response instead, use the `--follow`, `-F` option:

```
$ http --follow httpbin.org/redirect/3
```

11.2 Showing intermediary redirect responses

If you additionally wish to see the intermediary requests/responses, then use the `--all` option as well:

```
$ http --follow --all httpbin.org/redirect/3
```

11.3 Limiting maximum redirects followed

To change the default limit of maximum 30 redirects, use the `--max-redirects=<limit>` option:

```
$ http --follow --all --max-redirects=5 httpbin.org/redirect/3
```

12 Proxies

You can specify proxies to be used through the `--proxy` argument for each protocol (which is included in the value in case of redirects across protocols):

```
$ http --proxy=http:http://10.10.1.10:3128 --proxy=https:https://10.10.1.10:1080 example.org
```

With Basic authentication:

```
$ http --proxy=http:http://user:pass@10.10.1.10:3128 example.org
```

12.1 Environment variables

You can also configure proxies by environment variables `HTTP_PROXY` and `HTTPS_PROXY`, and the underlying Requests library will pick them up as well. If you want to disable proxies configured through the environment variables for certain hosts, you can specify them in `NO_PROXY`.

In your ~/.bash_profile:

```
export HTTP_PROXY=http://10.10.1.10:3128
export HTTPS_PROXY=https://10.10.1.10:1080
export NO_PROXY=localhost,example.com
```

12.2 SOCKS

To enable SOCKS proxy support please install `requests[socks]` using `pip`:

```
$ pip install -U requests[socks]
```

Usage is the same as for other types of [proxies](#):

```
$ http --proxy=http:socks5://user:pass@host:port --proxy=https:socks5://user:pass@host:port example.org
```

13 HTTPS

13.1 Server SSL certificate verification

To skip the host's SSL certificate verification, you can pass `--verify=no` (default is `yes`):

```
$ http --verify=no https://example.org
```

13.2 Custom CA bundle

You can also use `--verify=<CA_BUNDLE_PATH>` to set a custom CA bundle path:

```
$ http --verify=/ssl/custom_ca_bundle https://example.org
```

13.3 Client side SSL certificate

To use a client side certificate for the SSL communication, you can pass the path of the cert file with `--cert`:

```
$ http --cert=client.pem https://example.org
```

If the private key is not contained in the cert file you may pass the path of the key file with `--cert-key`:

```
$ http --cert=client.crt --cert-key=client.key https://example.org
```

13.4 SSL version

Use the `--ssl=<PROTOCOL>` to specify the desired protocol version to use. This will default to SSL v2.3 which will negotiate the highest protocol that both the server and your installation of OpenSSL support. The available protocols are `ssl2.3`, `ssl3`, `tls1`, `tls1.1`, `tls1.2`. (The actually available set of protocols may vary depending on your OpenSSL installation.)

```
# Specify the vulnerable SSL v3 protocol to talk to an outdated server:
$ http --ssl=ssl3 https://vulnerable.example.org
```

13.5 SNI (Server Name Indication)

If you use HTTPie with [Python version](#) lower than 2.7.9 (can be verified with `http --debug`) and need to talk to servers that use SNI (Server Name Indication) you need to install some additional dependencies:

```
$ pip install --upgrade requests[security]
```

You can use the following command to test SNI support:

```
$ http https://sni.velox.ch
```

14 Output options

By default, HTTPie only outputs the final response and the whole response message is printed (headers as well as the body). You can control what should be printed via several options:

<code>--headers, -h</code>	Only the response headers are printed.
<code>--body, -b</code>	Only the response body is printed.
<code>--verbose, -v</code>	Print the whole HTTP exchange (request and response). This option also enables <code>--all</code> (see below).
<code>--print, -p</code>	Selects parts of the HTTP exchange.

`--verbose` can often be useful for debugging the request and generating documentation examples:

```
$ http --verbose PUT httpbin.org/put hello=world
PUT /put HTTP/1.1
Accept: application/json, */*
Accept-Encoding: gzip, deflate
Content-Type: application/json
Host: httpbin.org
User-Agent: HTTPie/0.2.7dev

{
  "hello": "world"
}

HTTP/1.1 200 OK
Connection: keep-alive
Content-Length: 477
Content-Type: application/json
Date: Sun, 05 Aug 2012 00:25:23 GMT
Server: gunicorn/0.13.4

{
  [...]
}
```

14.1 What parts of the HTTP exchange should be printed

All the other [output options](#) are under the hood just shortcuts for the more powerful `--print, -p`. It accepts a string of characters each of which represents a specific part of the HTTP exchange:

Character	Stands for
-----------	------------

H	request headers
B	request body
h	response headers
b	response body

Print request and response headers:

```
$ http --print=Hh PUT httpbin.org/put hello=world
```

14.2 Viewing intermediary requests/responses

To see all the HTTP communication, i.e. the final request/response as well as any possible intermediary requests/responses, use the `--all` option. The intermediary HTTP communication include followed redirects (with `--follow`), the first unauthorized request when HTTP digest authentication is used (`--auth=digest`), etc.

```
# Include all responses that lead to the final one:
$ http --all --follow httpbin.org/redirect/3
```

The intermediary requests/response are by default formatted according to `--print`, `-p` (and its shortcuts described above). If you'd like to change that, use the `--history-print`, `-P` option. It takes the same arguments as `--print`, `-p` but applies to the intermediary requests only.

```
# Print the intermediary requests/responses differently than the final one:
$ http -A digest -a foo:bar --all -p Hh -P H httpbin.org/digest-auth/auth/foo/bar
```

14.3 Conditional body download

As an optimization, the response body is downloaded from the server only if it's part of the output. This is similar to performing a `HEAD` request, except that it applies to any HTTP method you use.

Let's say that there is an API that returns the whole resource when it is updated, but you are only interested in the response headers to see the status code after an update:

```
$ http --headers PATCH example.org/Really-Huge-Resource name='New Name '
```

Since we are only printing the HTTP headers here, the connection to the server is closed as soon as all the response headers have been received. Therefore, bandwidth and time isn't wasted downloading the body which you don't care about. The response headers are downloaded always, even if they are not part of the output

15 Redirected Input

The universal method for passing request data is through redirected `stdin` (standard input)—piping. Such data is buffered and then with no further processing used as the request body. There are multiple useful ways to use piping:

Redirect from a file:

```
$ http PUT example.com/person/1 X-API-Token:123 < person.json
```

Or the output of another program:

```
$ grep '401 Unauthorized' /var/log/httpd/error_log | http POST example.org/intruders
```

You can use `echo` for simple data:

```
$ echo '{"name": "John"}' | http PATCH example.com/person/1 X-API-Token:123
```

You can even pipe web services together using HTTPie:

```
$ http GET https://api.github.com/repos/jkbrzt/httpie | http POST httpbin.org/post
```

You can use `cat` to enter multiline data on the terminal:

```
$ cat | http POST example.com
<paste>
^D
```

```
$ cat | http POST example.com/todos Content-Type:text/plain
- buy milk
- call parents
^D
```

On OS X, you can send the contents of the clipboard with `pbpaste`:

```
$ pbpaste | http PUT example.com
```

Passing data through `stdin` cannot be combined with data fields specified on the command line:

```
$ echo 'data' | http POST example.org more=data # This is invalid
```

To prevent HTTPie from reading `stdin` data you can use the `--ignore-stdin` option.

15.1 Request data from a filename

An alternative to redirected `stdin` is specifying a filename (as `@/path/to/file`) whose content is used as if it came from `stdin`.

It has the advantage that the `Content-Type` header is automatically set to the appropriate value based on the filename extension. For example, the following request sends the verbatim contents of that XML file with `Content-Type: application/xml`:

```
$ http PUT httpbin.org/put @/data/file.xml
```

16 Terminal output

HTTPie does several things by default in order to make its terminal output easy to read.

16.1 Colors and formatting

Syntax highlighting is applied to HTTP headers and bodies (where it makes sense). You can choose your preferred color scheme via the `--style` option if you don't like the default one (see `$ http --help` for the possible values).

Also, the following formatting is applied:

- HTTP headers are sorted by name.
- JSON data is indented, sorted by keys, and unicode escapes are converted to the characters they represent.

One of these options can be used to control output processing:

<code>--pretty=all</code>	Apply both colors and formatting. Default for terminal output.
<code>--pretty=colors</code>	Apply colors.
<code>--pretty=format</code>	Apply formatting.
<code>--pretty=none</code>	Disables output processing. Default for redirected output.

16.2 Binary data

Binary data is suppressed for terminal output, which makes it safe to perform requests to URLs that send back binary data. Binary data is suppressed also in redirected, but prettified output. The connection is closed as soon as we know that the response body is binary,

```
$ http example.org/Movie.mov
```

You will nearly instantly see something like this:

```
HTTP/1.1 200 OK
Accept-Ranges: bytes
Content-Encoding: gzip
Content-Type: video/quicktime
Transfer-Encoding: chunked

+-----+
| NOTE: binary data not shown in terminal |
+-----+
```

17 Redirected output

HTTPie uses a different set of defaults for redirected output than for [terminal output](#). The differences being:

- Formatting and colors aren't applied (unless `--pretty` is specified).
- Only the response body is printed (unless one of the [output options](#) is set).
- Also, binary data isn't suppressed.

The reason is to make piping HTTPie's output to another programs and downloading files work with no extra flags. Most of the time, only the raw response body is of an interest when the output is redirected.

Download a file:

```
$ http example.org/Movie.mov > Movie.mov
```

Download an image of Octocat, resize it using ImageMagick, upload it elsewhere:

```
$ http octodex.github.com/images/original.jpg | convert - -resize 25% - | http example.org/Octocats
```

Force colorizing and formatting, and show both the request and the response in `less` pager:

```
$ http --pretty=all --verbose example.org | less -R
```

The `-R` flag tells `less` to interpret color escape sequences included HTTPie's output.

You can create a shortcut for invoking HTTPie with colorized and paged output by adding the following to your `~/.bash_profile`:

```
function httpless {  
    # `httpless example.org`  
    http --pretty=all --print=hb "$@" | less -R;  
}
```

18 Download mode

HTTPie features a download mode in which it acts similarly to `wget`.

When enabled using the `--download`, `-d` flag, response headers are printed to the terminal (`stderr`), and a progress bar is shown while the response body is being saved to a file.

```
$ http --download https://github.com/jkbrzt/httpie/archive/master.tar.gz
```

```
HTTP/1.1 200 OK  
Content-Disposition: attachment; filename=httpie-master.tar.gz  
Content-Length: 257336  
Content-Type: application/x-gzip  
  
Downloading 251.30 kB to "httpie-master.tar.gz"  
Done. 251.30 kB in 2.73862s (91.76 kB/s)
```

18.1 Downloaded file name

If not provided via `--output`, `-o`, the output filename will be determined from `Content-Disposition` (if available), or from the URL and `Content-Type`. If the guessed filename already exists, HTTPie adds a unique suffix to it.

18.2 Piping while downloading

You can also redirect the response body to another program while the response headers and progress are still shown in the terminal:

```
$ http -d https://github.com/jkbrzt/httpie/archive/master.tar.gz | tar xzf -
```

18.3 Resuming downloads

If `--output`, `-o` is specified, you can resume a partial download using the `--continue`, `-c` option. This only works with servers that support Range requests and 206 Partial Content responses. If the server doesn't support that, the whole file will simply be downloaded:

```
$ http -dco file.zip example.org/file
```

18.4 Other notes

- The `--download` option only changes how the response body is treated.

- You can still set custom headers, use sessions, `--verbose`, `-v`, etc.
- `--download` always implies `--follow` (redirects are followed).
- HTTPie exits with status code 1 (error) if the body hasn't been fully downloaded.
- `Accept-Encoding` cannot be set with `--download`.

19 Streamed responses

Responses are downloaded and printed in chunks which allows for streaming and large file downloads without using too much memory. However, when [colors and formatting](#) is applied, the whole response is buffered and only then processed at once.

19.1 Disabling buffering

You can use the `--stream`, `-S` flag to make two things happen:

1. The output is flushed in much smaller chunks without any buffering, which makes HTTPie behave kind of like `tail -f` for URLs.
2. Streaming becomes enabled even when the output is prettified: It will be applied to each line of the response and flushed immediately. This makes it possible to have a nice output for long-lived requests, such as one to the Twitter streaming API.

19.2 Examples use cases

Prettified streamed response:

```
$ http --stream -f -a YOUR-TWITTER-NAME https://stream.twitter.com/1/statuses/filter.json track='Justin Bieber'
```

Streamed output by small chunks alá `tail -f`:

```
# Send each new tweet (JSON object) mentioning "Apple" to another
# server as soon as it arrives from the Twitter streaming API:
$ http --stream -f -a YOUR-TWITTER-NAME https://stream.twitter.com/1/statuses/filter.json track=Apple \
| while read tweet; do echo "$tweet" | http POST example.org/tweets ; done
```

20 Sessions

By default, every request HTTPie makes is completely independent of any previous ones to the same host.

However, HTTPie also supports persistent sessions via the `--session=SESSION_NAME_OR_PATH` option. In a session, custom headers—except for the ones starting with `Content-` or `If-`—, authorization, and cookies (manually specified or sent by the server) persist between requests to the same host.

```
# Create a new session
$ http --session=/tmp/session.json example.org API-Token:123

# Re-use an existing session — API-Token will be set:
$ http --session=/tmp/session.json example.org
```

All session data, including credentials, cookie data, and custom headers are stored in plain text. That means session files can also be created and edited manually in a text editor—they are regular JSON.

20.1 Named sessions

You can create one or more named session per host. For example, this is how you can create a new session named `user1` for `example.org`:

```
$ http --session=user1 -a user1:password example.org X-Foo:Bar
```

From now onw, you can refer to the session by its name. When you choose to use the session again, any the previously used authorization and HTTP headers will automatically be set:

```
$ http --session=user1 example.org
```

To create or reuse a different session, simple specify a different name:

```
$ http --session=user2 -a user2:password example.org X-Bar:Foo
```

Named sessions' data is stored in JSON files in the directory `~/.httpie/sessions/<host>/<name>.json` (`%APPDATA%\httpie\sessions\<host>\<name>.json` on Windows).

20.2 Anonymous sessions

Instead of a name, you can also directly specify a path to a session file. This allows for sessions to be re-used across multiple hosts:

```
$ http --session=/tmp/session.json example.org
$ http --session=/tmp/session.json admin.example.org
$ http --session=~/.httpie/sessions/another.example.org/test.json example.org
$ http --session-read-only=/tmp/session.json example.org
```

20.3 Readonly session

To use an existing session file without updating it from the request/response exchange once it is created, specify the session name via `--session-read-only=SESSION_NAME_OR_PATH` instead.

21 Config

HTTPie uses a simple JSON config file.

21.1 Config file location

The default location of the configuration file is `~/.httpie/config.json` (or `%APPDATA%\httpie\config.json` on Windows). The config directory location can be changed by setting the `HTTPIE_CONFIG_DIR` environment variable. To view the exact location run `http --debug`.

21.2 Configurable options

The JSON file contains an object with the following keys:

21.2.1 default_options

An Array (by default empty) of default options that should be applied to every invocation of HTTPie.

For instance, you can use this option to change the default style and output options: `"default_options": ["--style=fruity", "--body"]` Another useful default option could be

"--session=default" to make HTTPie always use [sessions](#) (one named default will automatically be used). Or you could change the implicit request content type from JSON to form by adding --form to the list.

21.2.2 __meta__

HTTPie automatically stores some of its metadata here. Please do not change.

21.3 Un-setting previously specified options

Default options from the config file, or specified any other way, can be unset for a particular invocation via --no-OPTION arguments passed on the command line (e.g., --no-style or --no-session).

22 Scripting

When using HTTPie from shell scripts, it can be handy to set the --check-status flag. It instructs HTTPie to exit with an error if the HTTP status is one of 3xx, 4xx, or 5xx. The exit status will be 3 (unless --follow is set), 4, or 5, respectively.

```
#!/bin/bash

if http --check-status --ignore-stdin --timeout=2.5 HEAD example.org/health &> /dev/null; then
    echo 'OK!'
else
    case $? in
        2) echo 'Request timed out!' ;;
        3) echo 'Unexpected HTTP 3xx Redirection!' ;;
        4) echo 'HTTP 4xx Client Error!' ;;
        5) echo 'HTTP 5xx Server Error!' ;;
        6) echo 'Exceeded --max-redirects=<n> redirects!' ;;
        *) echo 'Other Error!' ;;
    esac
fi
```

22.1 Best practices

The default behaviour of automatically reading stdin is typically not desirable during non-interactive invocations. You most likely want use the --ignore-stdin option to disable it.

It is a common gotcha that without this option HTTPie seemingly hangs. What happens is that when HTTPie is invoked for example from a cron job, stdin is not connected to a terminal. Therefore, rules for [redirected input](#) apply, i.e., HTTPie starts to read it expecting that the request body will be passed through. And since there's no data nor EOF, it will be stuck. So unless you're piping some data to HTTPie, this flag should be used in scripts.

Also, it's might be good to override the default 30 second --timeout to something that suits you.

23 Meta

23.1 Interface design

The syntax of the command arguments closely corresponds to the actual HTTP requests sent over the wire. It has the advantage that it's easy to remember and read. It is often possible to translate an HTTP request to an HTTPie argument list just by inlining the request elements. For example, compare this HTTP request:

```
POST /collection HTTP/1.1
X-API-Key: 123
```

```
User-Agent: Bacon/1.0
Content-Type: application/x-www-form-urlencoded

name=value&name2=value2
```

with the HTTPie command that sends it:

```
$ http -f POST example.org/collection \
  X-API-Key:123 \
  User-Agent:Bacon/1.0 \
  name=value \
  name2=value2
```

Notice that both the order of elements and the syntax is very similar, and that only a small portion of the command is used to control HTTPie and doesn't directly correspond to any part of the request (here it's only `-f` asking HTTPie to send a form request).

The two modes, `--pretty=all` (default for terminal) and `--pretty=none` (default for redirected output), allow for both user-friendly interactive use and usage from scripts, where HTTPie serves as a generic HTTP client.

As HTTPie is still under heavy development, the existing command line syntax and some of the `--OPTIONS` may change slightly before HTTPie reaches its final version 1.0. All changes are recorded in the [change log](#).

23.2 User support

Please use the following support channels:

- [GitHub issues](#) for bug reports and feature requests.
- [Our Gitter chat room](#) to ask questions, discuss features, and for general discussion.
- [StackOverflow](#) to ask questions (please make sure to use the [httpie](#) tag).
- Tweet directly to [@clihttp](#).
- You can also tweet directly to [@jnbrzt](#).

23.3 Related projects

23.3.1 Dependencies

Under the hood, HTTPie uses these two amazing libraries:

- [Requests](#) — Python HTTP library for humans
- [Pygments](#) — Python syntax highlighter

23.3.2 HTTPie friends

HTTPie plays exceptionally well with the following tools:

- [jq](#) — CLI JSON processor that works great in conjunction with HTTPie
- [http-prompt](#) — interactive shell for HTTPie featuring autocomplete and command syntax highlighting

23.4 Contributing

See [CONTRIBUTING.rst](#).

23.5 Change log

See [CHANGELOG](#).

23.6 Artwork

See [claudiatd/httpie-artwork](#)

23.7 Licence

BSD-3-Clause: [LICENSE](#).

23.8 Authors

[Jakub Roztocil](#) (@jkrzt) created HTTPie and [these fine people](#) have contributed.