HTTPie: human-friendly CLI HTTP client for the API era

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. HTTPie is designed for testing, debugging, and generally interacting with APIs & HTTP servers. The `http & https` commands allow for creating and sending arbitrary HTTP requests. They use simple and natural syntax and provide formatted and colorized output.

Contents

1 About this document

This documentation is best viewed at httpie.org/docs.

You can select your corresponding HTTPie version as well as run examples directly from the browser using a termible.io embedded terminal.

If you are reading this on GitHub, then this text covers the current *development* version. You are invited to submit fixes and improvements to the the docs by editing README.rst.

2 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
- Linux, macOS and Windows support
- Plugins
- Documentation
- Test coverage

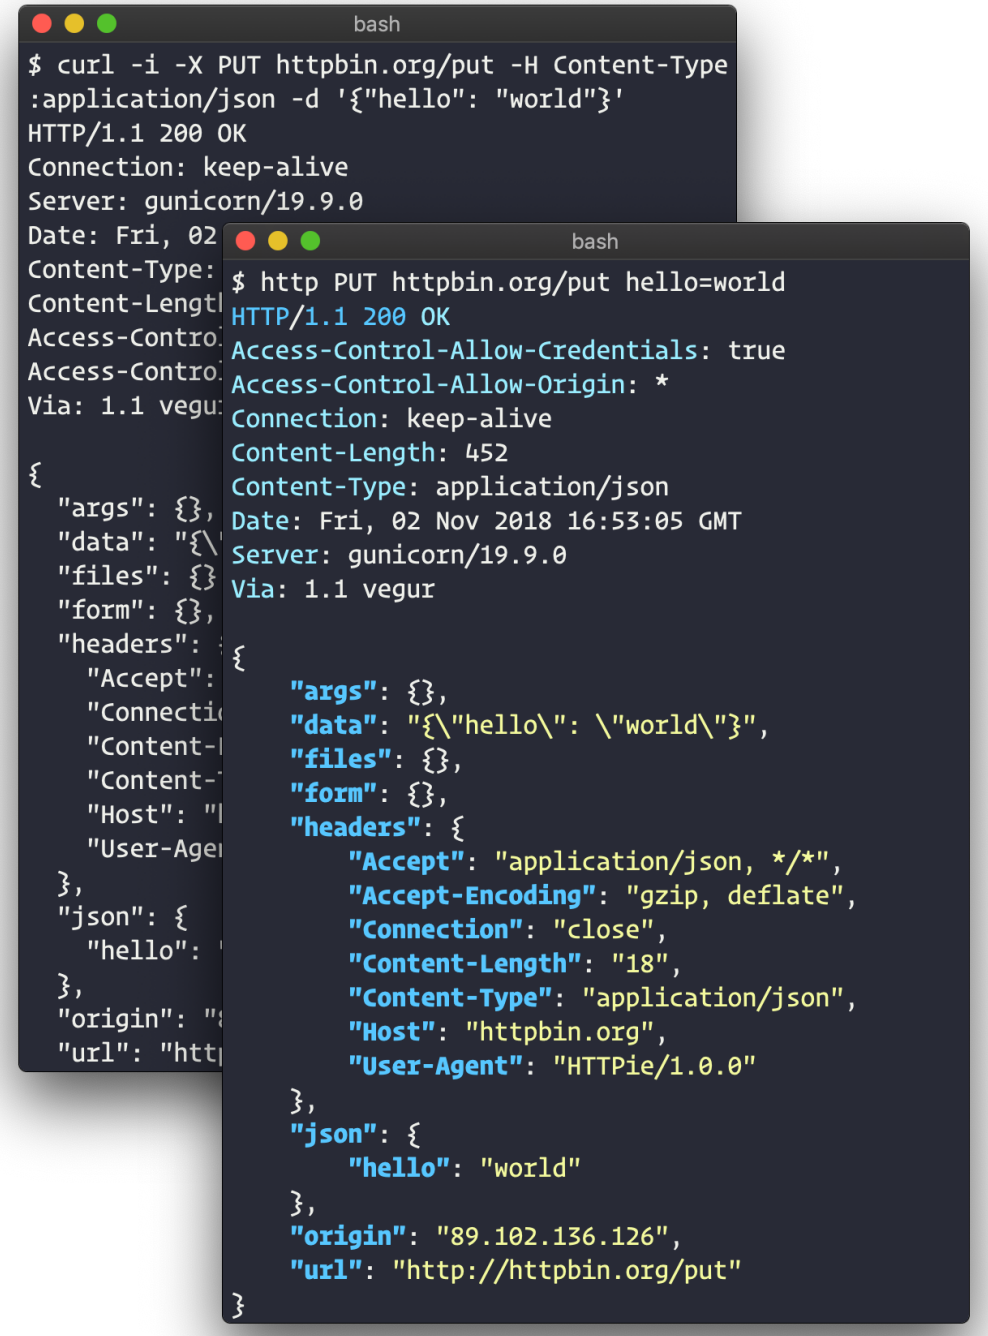
bash
$ curl -i -X PUT httpbin.org/put -H Content-Type
:application/json -d '{"hello": "world"}'
HTTP/1.1 200 OK
Connection: keep-alive
Server: gunicorn/19.9.0
Date: Fri, 02
Content-Type:
Content-Lengt
Access-Contro
Access-Contro
Via: 1.1 vegu
{
"args": {},
"data": "{\
"files": {}
"form": {},
"headers":
"Accept":
"Connecti
"Content-
"Content-
"Host": "
"User-Age
},
"json": {
"hello":
},
"origin": "
"url": "htt
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"
}

3 Installation

3.1 macOS

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

```
$ brew install httpie
```

A MacPorts *port* is also available:

```
$ port install httpie
```

3.2 Linux

Most Linux distributions provide a package that can be installed using the system package manager, for example:

```
# Debian, Ubuntu, etc.
$ apt install httpie
```

```
# Fedora
$ dnf install httpie
```

```
# CentOS, RHEL, ...
$ yum install httpie
```

```
# Gentoo
$ emerge httpie
```

```
# Arch Linux
$ pacman -S httpie
```

```
# Solus
$ eopkg install httpie
```

3.3 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:
$ python -m pip install --upgrade pip setuptools

$ python -m pip install --upgrade httpie
```

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

3.4 Python version

Python version 3.6 or greater is required.

3.5 Unstable version

You can also install the latest unreleased development version directly from the `master` branch on GitHub. It is a work-in-progress of a future stable release so the experience might be not as smooth.

On macOS you can install it with Homebrew:

```
$ brew uninstall --force httpie
$ brew install --HEAD httpie
```

Otherwise with `pip`:

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

Verify that now we have the current development version identifier with the `-dev` suffix, for example:

```
$ http --version
# 2.0.0-dev
```

4 Usage

Hello World:

```
$ https httpie.io/hello
```

Synopsis:

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

See also `http --help`.

4.1 Examples

Custom HTTP method, HTTP headers and JSON data:

```
$ http PUT pie.dev/put X-API-Token:123 name=John
```

Submitting forms:

```
$ http -f POST pie.dev/post hello=World
```

See the request that is being sent using one of the output options:

```
$ http -v pie.dev/get
```

Build and print a request without sending it using offline mode:

```
$ http --offline pie.dev/post hello=offline
```

Use GitHub API to post a comment on an issue with authentication:

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

Upload a file using redirected input:

```
$ http pie.dev/post < files/data.json
```

Download a file and save it via redirected output:

```
$ http pie.dev/image/png > image.png
```

Download a file `wget` style:

```
$ http --download pie.dev/image/png
```

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

```
$ http --session=logged-in -a username:password pie.dev/get API-Key:123
```

```
$ http --session=logged-in pie.dev/headers
```

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

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

5 HTTP method

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

```
$ http DELETE pie.dev/delete
```

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

```
DELETE /delete 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).

6 Request URL

The only information HTTPie needs to perform a request is a URL.

The default scheme is `http://` and can be omitted from the argument:

```
$ http example.org
# => http://example.org
```

HTTPie also installs an `https` executable, where the default scheme is `https://`:

```
$ https example.org
# => https://example.org
```

6.1 Querystring parameters

If you find yourself manually constructing URLs with querystring parameters 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. Additionally, any special characters in the parameter name or value get automatically URL-escaped (as opposed to parameters specified in the full URL, which HTTPie doesn't modify).

```
$ http https://api.github.com/search/repositories q==httpie per_page==1
```

```
GET /search/repositories?q=httpie&per_page=1 HTTP/1.1
```

6.2 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
```

6.3 Other default schemes

When HTTPie is invoked as `https` then the default scheme is `https://` (`$ https example.org` will make a request to `https://example.org`).

You can also use the `--default-scheme <URL_SCHEME>` option to create shortcuts for other protocols than HTTP (possibly supported via plugins). Example for the httpie-unixsocket plugin:

```
# Before
$ http http+unix://%2Fvar%2Frun%2Fdocker.sock/info
```

```
# Create an alias
$ alias http-unix='http --default-scheme="http+unix"'
```

```
# Now the scheme can be omitted
$ http-unix %2Fvar%2Frun%2Fdocker.sock/info
```

6.4 `--path-as-is`

The standard behaviour of HTTP clients is to normalize the path portion of URLs by squashing dot segments as a typically filesystem would:

```
$ http -v example.org/./../../etc/password
```

```
GET /etc/password HTTP/1.1
```

The `--path-as-is` option allows you to disable this behavior:

```
$ http --path-as-is -v example.org/./../../etc/password
```

```
GET /../../etc/password HTTP/1.1
```

7 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 `Name:Value`	Arbitrary HTTP header, e.g. `X-API-Token:123`.
URL parameters `name==value`	Appends the given name/value pair as a query string parameter to the URL. The == separator is used.
Data Fields `field=value`, `field=@file.txt`	Request data fields to be serialized as a JSON object (default), to be form-encoded (with `--form, -f`), or to be serialized as `multipart/form-data` (with `--multipart`).
Raw JSON fields `field:=json`, `field:=@file.json`	Useful when sending JSON and one or more fields need to be a `Boolean`, `Number`, nested `Object`, or an `Array`, e.g., `meals:='["ham","spam"]'` or `pies:=[1,2,3]` (note the quotes).
Fields upload fields `field@/dir/file` `field@file;type=mime`	Only available with `--form, -f` and `--multipart`. For example `screenshot@~/Pictures/img.png`, or `'cv@cv.txt;type=text/markdown'`. With `--form`, the presence of a file field results in a `--multipart` request.

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

7.1 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 pie.dev/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": "foo"
}
```

8 JSON

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

Simple example:

```
$ http PUT pie.dev/put name=John email=john@example.org
```

```
PUT / HTTP/1.1
Accept: application/json, */*;q=0.5
Accept-Encoding: gzip, deflate
Content-Type: application/json
Host: pie.dev

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

8.1 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, */*;q=0.5`

8.2 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, */*;q=0.5'`). Additionally, HTTPie will try to detect JSON responses even when the `Content-Type` is incorrectly `text/plain` or unknown.

8.3 Non-string JSON fields

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

```
$ http PUT pie.dev/put \
    name=John \                          # String (default)
    age:=29 \                            # Raw JSON — Number
    married:=false \                     # Raw JSON — Boolean
    hobbies:='["http", "pies"]' \        # Raw JSON — Array
    favorite:='{"tool": "HTTPie"}' \     # Raw JSON — Object
    bookmarks:=@files/data.json \        # Embed JSON file
    description=@files/text.txt          # Embed text file
```

```
PUT /person/1 HTTP/1.1
Accept: application/json, */*;q=0.5
Content-Type: application/json
Host: pie.dev

{
    "age": 29,
    "hobbies": [
        "http",
        "pies"
    ],
    "description": "John is a nice guy who likes pies.",
    "married": false,
    "name": "John",
    "favorite": {
        "tool": "HTTPie"
    },
    "bookmarks": {
        "HTTPie": "https://httpie.org",
    }
}
```

8.4 Raw and complex JSON

Please note that with the request items data field syntax, commands can quickly become unwieldy when sending complex structures. In such cases, it's better to pass the full raw JSON data via redirected input, for example:

```
$ echo '{"hello": "world"}' | http POST pie.dev/post
```

```
$ http POST pie.dev/post < files/data.json
```

Furthermore, this syntax only allows you to send an object as the JSON document, but not an array, etc. Here, again, the solution is to use redirected input.

9 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.

9.1 Regular forms

```
$ http --form POST pie.dev/post name='John Smith'
```

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

name=John+Smith
```

9.2 File upload forms

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

```
$ http -f POST pie.dev/post name='John Smith' cv@~/files/data.xml
```

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>
```

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

When uploading files, their content type is inferred from the file name. You can manually override the inferred content type:

```
$ http -f POST pie.dev/post name='John Smith' cv@'~/files/data.bin;type=application/pdf'
```

To perform a `multipart/form-data` request even without any files, use `--multipart` instead of `--form`:

```
$ http --multipart --offline example.org hello=world
```

```
POST / HTTP/1.1
Content-Length: 129
Content-Type: multipart/form-data; boundary=c31279ab254f40aeb06df32b433cbccb
Host: example.org

--c31279ab254f40aeb06df32b433cbccb
Content-Disposition: form-data; name="hello"

world
--c31279ab254f40aeb06df32b433cbccb--
```

File uploads are always streamed to avoid memory issues with large files.

By default, HTTPie uses a random unique string as the multipart boundary but you can use `--boundary` to specify a custom string instead:

```
$ http --form --multipart --boundary=xoxo --offline example.org hello=world
```

```
POST / HTTP/1.1
Content-Length: 129
Content-Type: multipart/form-data; boundary=xoxo
Host: example.org

--xoxo
Content-Disposition: form-data; name="hello"

world
--xoxo--
```

If you specify a custom `Content-Type` header without including the boundary bit, HTTPie will add the boundary value (explicitly specified or auto-generated) to the header automatically:

```
http --form --multipart --offline example.org hello=world Content-Type:multipart/letter
```

```
POST / HTTP/1.1
Content-Length: 129
Content-Type: multipart/letter; boundary=c31279ab254f40aeb06df32b433cbccb
Host: example.org

--c31279ab254f40aeb06df32b433cbccb
Content-Disposition: form-data; name="hello"

world
--c31279ab254f40aeb06df32b433cbccb--
```

10 HTTP headers

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

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

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

10.1 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 these can be overwritten and some of them unset (see below).

10.2 Empty headers and header un-setting

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

```
$ http pie.dev/headers Accept: User-Agent:
```

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

```
$ http pie.dev/headers 'Header;'
```

10.3 Limiting response headers

The `--max-headers=n` options allows you to control the number of headers HTTPie reads before giving up (the default `0`, i.e., there's no limit).

```
$ http --max-headers=100 pie.dev/get
```

11 Offline mode

Use `--offline` to construct HTTP requests without sending them anywhere. With `--offline`, HTTPie builds a request based on the specified options and arguments, prints it to `stdout`, and then exits. It works completely offline; no network connection is ever made. This has a number of use cases, including:

Generating API documentation examples that you can copy & paste without sending a request:

```
$ http --offline POST server.chess/api/games API-Key:ZZZ w=magnus b=hikaru t=180 i=2
```

```
$ http --offline MOVE server.chess/api/games/123 API-Key:ZZZ p=b a=R1a3 t=77
```

Generating raw requests that can be sent with any other client:

```
# 1. save a raw request to a file:
$ http --offline POST pie.dev/post hello=world > request.http
```

```
# 2. send it over the wire with, for example, the fantastic netcat tool:
$ nc pie.dev 80 < request.http
```

You can also use the `--offline` mode for debugging and exploring HTTP and HTTPie, and for "dry runs".

`--offline` has the side-effect of automatically activating `--print=HB`, i.e., both the request headers and the body are printed. You can customize the output with the usual output options, with the exception that there is not response to be printed. You can use `--offline` in combination with all the other options (e.g., `--session`).

12 Cookies

HTTP clients send cookies to the server as regular HTTP headers. That means, HTTPie does not offer any special syntax for specifying cookies — the usual `Header:Value` notation is used:

Send a single cookie:

```
$ http pie.dev/cookies Cookie:sessionid=foo
```

```
GET / HTTP/1.1
Accept: */*
Accept-Encoding: gzip, deflate
Connection: keep-alive
Cookie: sessionid=foo
Host: pie.dev
User-Agent: HTTPie/0.9.9
```

Send multiple cookies (note the header is quoted to prevent the shell from interpreting the `;`):

```
$ http pie.dev/cookies 'Cookie:sessionid=foo;another-cookie=bar'
```

```
GET / HTTP/1.1
Accept: */*
Accept-Encoding: gzip, deflate
Connection: keep-alive
Cookie: sessionid=foo;another-cookie=bar
Host: pie.dev
User-Agent: HTTPie/0.9.9
```

If you often deal with cookies in your requests, then chances are you'd appreciate the sessions feature.

13 Authentication

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

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

13.1 Basic auth

```
$ http -a username:password pie.dev/basic-auth/username/password
```

13.2 Digest auth

```
$ http -A digest -a username:password pie.dev/digest-auth/httpie/username/password
```

13.3 Password prompt

```
$ http -a username pie.dev/basic-auth/username/password
```

13.4 Empty password

```
$ http -a username: pie.dev/headers
```

13.5 `.netrc`

Authentication information from your `~/.netrc` file is by default honored as well.

For example:

```
$ cat ~/.netrc
machine pie.dev
login httpie
password test
```

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

This can be disabled with the `--ignore-netrc` option:

```
$ http --ignore-netrc pie.dev/basic-auth/httpie/test
HTTP/1.1 401 UNAUTHORIZED
[...]
```

13.6 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-api-auth: ApiAuth
- httpie-aws-auth: AWS / Amazon S3
- httpie-edgegrid: EdgeGrid
- httpie-hmac-auth: HMAC
- httpie-jwt-auth: JWTAuth (JSON Web Tokens)
- httpie-negotiate: SPNEGO (GSS Negotiate)
- httpie-ntlm: NTLM (NT LAN Manager)
- httpie-oauth: OAuth
- requests-hawk: Hawk

14 HTTP redirects

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

```
$ http pie.dev/redirect/3
```

14.1 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 pie.dev/redirect/3
```

14.2 Showing intermediary redirect responses

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

```
$ http --follow --all pie.dev/redirect/3
```

14.3 Limiting maximum redirects followed

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

```
$ http --follow --all --max-redirects=2 pie.dev/redirect/3
```

15 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
```

15.1 Environment variables

You can also configure proxies by environment variables `ALL_PROXY`, `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
```

15.2 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
```

16 HTTPS

16.1 Server SSL certificate verification

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

```
$ http --verify=no https://pie.dev/get
```

16.2 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
```

16.3 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
```

16.4 SSL version

Use the `--ssl=<PROTOCOL>` option 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`, `tls1.3`. (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
```

16.5 SSL ciphers

You can specify the available ciphers with `--ciphers`. It should be a string in the OpenSSL cipher list format.

```
$ http --ciphers=ECDHE-RSA-AES128-GCM-SHA256  https://pie.dev/get
```

Note: these cipher strings do not change the negotiated version of SSL or TLS, they only affect the list of available cipher suites.

To see the default cipher string, run `http --help` and see the `--ciphers` section under SSL.

17 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:

`--headers, -h`	Only the response headers are printed.
`--body, -b`	Only the response body is printed.
`--verbose, -v`	Print the whole HTTP exchange (request and response). This option also enables `--all` (see below).
`--print, -p`	Selects parts of the HTTP exchange.
`--quiet, -q`	Don't print anything to `stdout` and `stderr`.

17.1 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 pie.dev/put hello=world
```

17.2 Verbose output

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

```
$ http --verbose PUT pie.dev/put hello=world
PUT /put HTTP/1.1
Accept: application/json, */*;q=0.5
Accept-Encoding: gzip, deflate
Content-Type: application/json
Host: pie.dev
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

{
    […]
}
```

17.3 Quiet output

`--quiet` redirects all output that would otherwise go to `stdout` and `stderr` to `/dev/null` (except for errors and warnings). This doesn't affect output to a file via `--output` or `--download`.

```
# There will be no output:
$ http --quiet pie.dev/post enjoy='the silence'
```

17.4 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 pie.dev/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 pie.dev/digest-auth/auth/foo/bar
```

17.5 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 pie.dev/patch 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

18 Redirected Input

The universal method for passing request data is through redirected `stdin` (standard input)—piping.

By default, `stdin` data is buffered and then with no further processing used as the request body. If you provide `Content-Length`, then the request body is streamed without buffering. You can also use `--chunked` to enable streaming via chunked transfer encoding.

There are multiple useful ways to use piping:

Redirect from a file:

```
$ http PUT pie.dev/put X-API-Token:123 < files/data.json
```

Or the output of another program:

```
$ grep '401 Unauthorized' /var/log/httpd/error_log | http POST pie.dev/post
```

You can use `echo` for simple data:

```
$ echo '{"name": "John"}' | http PATCH pie.dev/patch X-API-Token:123
```

You can also use a Bash *here string*:

```
$ http pie.dev/post <<<'{"name": "John"}'
```

You can even pipe web services together using HTTPie:

```
$ http GET https://api.github.com/repos/httpie/httpie | http POST pie.dev/post
```

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

```
$ cat | http POST pie.dev/post
<paste>
^D
```

```
$ cat | http POST pie.dev/post 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 pie.dev/put
```

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.

18.1 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 pie.dev/put @files/data.xml
```

File uploads are always streamed to avoid memory issues with large files.

19 Chunked transfer encoding

You can use the `--chunked` flag to instruct HTTPie to use `Transfer-Encoding: chunked`:

```
$ http --chunked PUT pie.dev/put hello=world
```

```
$ http --chunked --multipart PUT pie.dev/put hello=world foo@files/data.xml
```

```
$ http --chunked pie.dev/post @files/data.xml
```

```
$ cat files/data.xml | http --chunked pie.dev/post
```

20 Terminal output

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

20.1 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. There dozens of styles available, here are just a few special or notable ones:

`auto`	Follows your terminal ANSI color styles. This is the default style used by HTTPie.
`default`	Default styles of the underlying Pygments library. Not actually used by default by HTTPie. You can enable it with `--style=default`
`monokai`	A popular color scheme. Enable with `--style=monokai`.
`fruity`	A bold, colorful scheme. Enable with `--style=fruity`.
…	See `$ http --help` for all the possible `--style` 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:

`--pretty=all`	Apply both colors and formatting. Default for terminal output.
`--pretty=colors`	Apply colors.
`--pretty=format`	Apply formatting.
`--pretty=none`	Disables output processing. Default for redirected output.

You can control the applied formatting via the `--format-options` option. The following options are available:

For example, this is how you would disable the default header and JSON key sorting, and specify a custom JSON indent size:

```
$ http --format-options headers.sort:false,json.sort_keys:false,json.indent:2 pie.dev/get
```

This is something you will typically store as one of the default options in your config file. See `http --help` for all the available formatting options.

There are also two shortcuts that allow you to quickly disable and re-enable sorting-related format options (currently it means JSON keys and headers): `--unsorted` and `--sorted`.

20.2 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 pie.dev/bytes/2000
```

You will nearly instantly see something like this:

```
HTTP/1.1 200 OK
Content-Type: application/octet-stream

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

21 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 pie.dev/image/png > image.png
```

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 pie.dev/get | 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;
}
```

22 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/httpie/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)
```

22.1 Downloaded filename

There are three mutually exclusive ways through which HTTPie determines the output filename (with decreasing priority):

1. You can explicitly provide it via `--output, -o`. The file gets overwritten if it already exists (or appended to with `--continue, -c`).
2. The server may specify the filename in the optional `Content-Disposition` response header. Any leading dots are stripped from a server-provided filename.
3. The last resort HTTPie uses is to generate the filename from a combination of the request URL and the response `Content-Type`. The initial URL is always used as the basis for the generated filename — even if there has been one or more redirects.

To prevent data loss by overwriting, HTTPie adds a unique numerical suffix to the filename when necessary (unless specified with `--output, -o`).

22.2 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/httpie/httpie/archive/master.tar.gz |  tar zxf -
```

22.3 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
```

22.4 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).
- `--download` also implies `--check-status` (error HTTP status will result in a non-zero exist static code).
- HTTPie exits with status code `1` (error) if the body hasn't been fully downloaded.
- `Accept-Encoding` cannot be set with `--download`.

23 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.

23.1 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.

23.2 Examples use cases

Prettified streamed response:

```
$ http --stream pie.dev/stream/3
```

Streamed output by small chunks à la `tail -f`:

```
# Send each new line (JSON object) to another URL as soon as it arrives from a streaming API:
$ http --stream pie.dev/stream/3 | while read line; do echo "$line" | http pie.dev/post ; done
```

24 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 HTTP headers (except for the ones starting with `Content-` or `If-`), authentication, and cookies (manually specified or sent by the server) persist between requests to the same host.

```
# Create a new session:
$ http --session=./session.json pie.dev/headers API-Token:123
```

```
# Inspect / edit the generated session file:
$ cat session.json
```

```
# Re-use the existing session — the API-Token header will be set:
$ http --session=./session.json pie.dev/headers
```

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. It also means that they can be read by anyone who has access to the session file.

24.1 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 `pie.dev`:

```
$ http --session=user1 -a user1:password pie.dev/get X-Foo:Bar
```

From now on, you can refer to the session by its name (`user1`). When you choose to use the session again, any previously specified authentication or HTTP headers will automatically be set:

```
$ http --session=user1 pie.dev/get
```

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

```
$ http --session=user2 -a user2:password pie.dev/get X-Bar:Foo
```

Named sessions's data is stored in JSON files inside the `sessions` subdirectory of the config directory, typically: `~/.config/httpie/sessions/<host>/<name>.json` (`%APPDATA%\httpie\sessions\<host>\<name>.json` on Windows).

If you have executed the above commands on a unix machine, you should be able list the generated sessions files using:

```
$ ls -l ~/.config/httpie/sessions/pie.dev
```

24.2 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:

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

```
# Use the session to make a request to another host:
$ http --session=/tmp/session.json admin.example.org
```

```
# You can also refer to a previously created named session:
$ http --session=~/.config/httpie/sessions/another.example.org/test.json example.org
```

When creating anonymous sessions, please remember to always include at least one `/`, even if the session files is located in the current directory (i.e., `--session=./session.json` instead of just `--session=session.json`), otherwise HTTPie assumes a named session instead.

24.3 Readonly session

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

```
# If the session file doesn't exist, then it is created:
$ http --session-read-only=./ro-session.json pie.dev/headers Custom-Header:orig-value
```

```
# But it is not updated:
$ http --session-read-only=./ro-session.json pie.dev/headers Custom-Header:new-value
```

24.4 Cookie Storage Behaviour

TL;DR: Cookie storage priority: Server response > Command line request > Session file

To set a cookie within a Session there are three options:

1. Get a `Set-Cookie` header in a response from a server

```
$ http --session=./session.json pie.dev/cookie/set?foo=bar
```

2. Set the cookie name and value through the command line as seen in cookies

```
$ http --session=./session.json pie.dev/headers Cookie:foo=bar
```

3. Manually set cookie parameters in the json file of the session

```
{
    "__meta__": {
    "about": "HTTPie session file",
    "help": "https://httpie.org/doc#sessions",
    "httpie": "2.2.0-dev"
    },
    "auth": {
        "password": null,
        "type": null,
        "username": null
    },
    "cookies": {
        "foo": {
            "expires": null,
            "path": "/",
            "secure": false,
            "value": "bar"
            }
    }
}
```

Cookies will be set in the session file with the priority specified above. For example, a cookie set through the command line will overwrite a cookie of the same name stored in the session file. If the server returns a `Set-Cookie` header with a cookie of the same name, the returned cookie will overwrite the preexisting cookie.

Expired cookies are never stored. If a cookie in a session file expires, it will be removed before sending a new request. If the server expires an existing cookie, it will also be removed from the session file.

25 Config

HTTPie uses a simple `config.json` file. The file doesn't exist by default but you can create it manually.

25.1 Config file directory

To see the exact location for your installation, run `http --debug` and look for `config_dir` in the output.

The default location of the configuration file on most platforms is `$XDG_CONFIG_HOME/httpie/config.json` (defaulting to `~/.config/httpie/config.json`).

For backwards compatibility, if the directory `~/.httpie` exists, the configuration file there will be used instead.

On Windows, the config file is located at `%APPDATA%\httpie\config.json`.

The config directory can be changed by setting the `$HTTPIE_CONFIG_DIR` environment variable:

```
$ export HTTPIE_CONFIG_DIR=/tmp/httpie
$ http pie.dev/get
```

25.2 Configurable options

Currently HTTPie offers a single configurable option:

25.2.1 `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 config option to change your default color theme:

```
$ cat ~/.config/httpie/config.json
```

```
{
    "default_options": [
      "--style=fruity"
    ]
}
```

Even though it is technically possible to include there any of HTTPie's options, it is not recommended to modify the default behaviour in a way that would break your compatibility with the wider world as that can generate a lot of confusion.

25.3 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`).

26 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 pie.dev/get &> /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
```

26.1 Best practices

The default behaviour of automatically reading `stdin` is typically not desirable during non-interactive invocations. You most likely want to 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 might be good to set a connection `--timeout` limit to prevent your program from hanging if the server never responds.

27 Meta

27.1 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 /post HTTP/1.1
Host: pie.dev
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 pie.dev/post \
  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.

27.2 Community and Support

HTTPie has the following community channels:

- GitHub issues for bug reports and feature requests.
- Discord server to ask questions, discuss features, and for general API development discussion.

- StackOverflow to ask questions (please make sure to use the httpie tag).
- Tweet directly to @httpie.
- You can also tweet directly to @jakubroztocil.

27.3 Related projects

27.3.1 Dependencies

Under the hood, HTTPie uses these two amazing libraries:

- Requests — Python HTTP library for humans
- Pygments — Python syntax highlighter

27.3.2 HTTPie friends

HTTPie plays exceptionally well with the following tools:

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

Helpers to convert from other client tools:

- CurliPie help convert cURL command line to HTTPie command line.

27.3.3 Alternatives

- httpcat — a lower-level sister utility of HTTPie for constructing raw HTTP requests on the command line.
- curl — a "Swiss knife" command line tool and an exceptional library for transferring data with URLs.

27.4 Contributing

See CONTRIBUTING.rst.

27.5 Change log

See CHANGELOG.

27.6 Artwork

- Logo by Cláudia Delgado.
- Animation by Allen Smith of GitHub.

27.7 Licence

BSD-3-Clause: LICENSE.

27.8 Authors

Jakub Roztocil (@jakubroztocil) created HTTPie and these fine people have contributed.