
django-revproxy Documentation

Release 0.11.1

Sergio Oliveira

Feb 25, 2023

Contents

1	Features	3
2	Dependencies	5
3	Install	7
4	Contents:	9
4.1	Introduction	9
4.1.1	How does it work?	9
4.2	Quickstart	10
4.2.1	Installation	10
4.2.2	Configuration	10
4.3	Usage	11
4.3.1	Proxy Views	11
4.4	API	13
4.4.1	revproxy.views	13
4.4.2	revproxy.response	13
4.4.3	revproxy.transformer	13
4.4.4	revproxy.utils	13
4.5	Changelog	14
4.5.1	0.11.1 (2023-02-26)	14
4.5.2	0.11.0 (2023-02-25)	14
4.5.3	0.10.0 (2020-02-05)	15
4.5.4	0.9.15 (2018-05-30)	15
4.5.5	0.9.14 (2018-01-11)	15
4.5.6	0.9.13 (2016-10-31)	15
4.5.7	0.9.12 (2016-05-23)	15
4.5.8	0.9.11 (2016-03-29)	15
4.5.9	0.9.10 (2016-02-03)	15
4.5.10	0.9.9 (2015-12-15)	15
4.5.11	0.9.8 (2015-12-10)	16
4.5.12	0.9.7 (2015-09-17)	16
4.5.13	0.9.6 (2015-09-09)	16
4.5.14	0.9.5 (2015-09-02)	16
4.5.15	0.9.4 (2015-08-27)	16
4.5.16	0.9.3 (2015-06-12)	16
4.5.17	0.9.2 (2015-06-09)	16

4.5.18	0.9.1 (2015-05-18)	16
4.5.19	0.9.0 (2015-03-04)	17
5	Indices and tables	19
	Python Module Index	21
	Index	23

A simple reverse proxy using Django. It allows to use Django as a reverse Proxy to HTTP requests. It also allows to use Django as an authentication Proxy.

Documentation available at <http://django-revproxy.readthedocs.org/>

CHAPTER 1

Features

- Proxies all HTTP methods: HEAD, GET, POST, PUT, DELETE, OPTIONS, TRACE, CONNECT and PATCH
- Copy all http headers sent from the client to the proxied server
- Copy all http headers sent from the proxied server to the client (except [hop-by-hop](#))
- Basic URL rewrite
- Sets the http header REQUEST_USER if the user is logged in Django
- Sets the http headers X-Forwarded-For and X-Forwarded-Proto
- Handles redirects
- Few external dependencies
- Apply XSLT transformation in the response (requires [Diazo](#))

CHAPTER 2

Dependencies

- django >= 3.0
- urllib3 >= 1.12
- diazo >= 1.0.5 (optional)
- lxml >= 3.4, < 3.5 (optional, but diazo dependency)

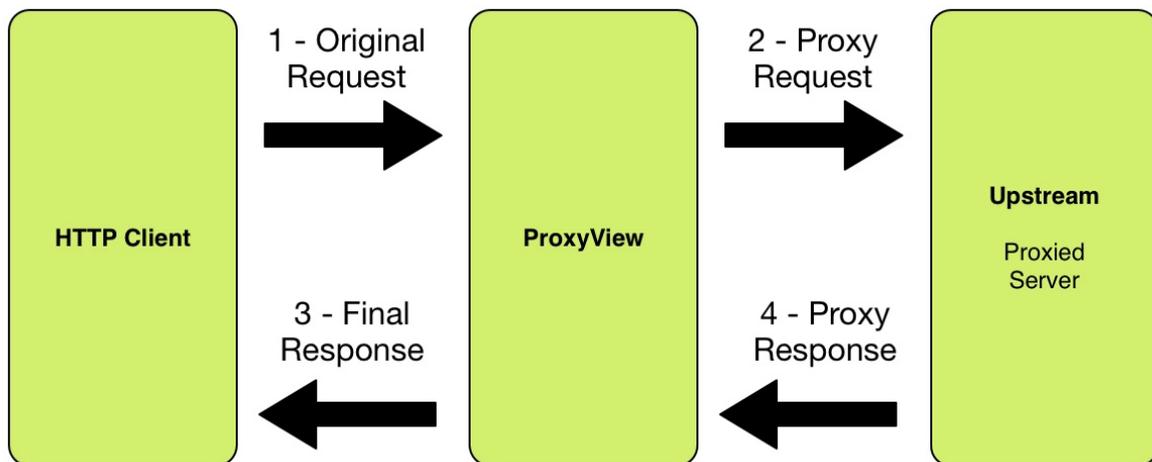
CHAPTER 3

Install

```
pip install django-revproxy
```


4.1 Introduction

4.1.1 How does it work?



At a high level, this is what happens behind the scenes in a request proxied by django-revproxy:

1. Django receives a request from the client and process it using a view that extends `revproxy.proxy.ProxyView`.
2. Revproxy will clone the client request.
3. If the user is authenticated in Django and `add_remote_user` attribute is set to `True` the HTTP header `REMOTE_USER` will be set with `request.user.username`.
4. If the `add_x_forwarded` attribute is set to `True` the HTTP headers `X-Forwarded-For` and `X-Forwarded-Proto` will be set to the IP address of the requestor and the protocol (http or https), respectively.
5. The cloned request is sent to the upstream server (set in the view).

6. After receiving the response from upstream, the view will process it to make sure all headers are set properly. Some headers like *Location* are treated as special cases.
7. The response received from the upstream server is transformed into a *django.http.HttpResponse*. For binary files *StreamingHttpResponse* is used instead to reduce memory usage.
8. If the user has setted a set of diazo rules and a theme template, a diazo/XSLT transformation will be applied on the response body.
9. Finally, the response will then be returned to the user

4.2 Quickstart

4.2.1 Installation

```
$ pip install django-revproxy
```

If you want to use `DiazoProxyView` you will also need to install `Diazo`. In that case you can use the following handy shortcut:

```
$ pip install django-revproxy[diazo]
```

4.2.2 Configuration

After installation, you'll need to configure your application to use `django-revproxy`. Start by adding `revproxy` to your `settings.py` file as follows:

```
#Add 'revproxy' to INSTALLED_APPS.
INSTALLED_APPS = (
    # ...
    'django.contrib.auth',
    'revproxy',
    # ...
)
```

Next, you'll need to create a View that extends `revproxy.views.ProxyView` and set the `upstream` attribute:

```
from revproxy.views import ProxyView

class TestProxyView(ProxyView):
    upstream = 'http://example.com'
```

And now add your view in the `urls.py`:

Alternatively you could just use the default `ProxyView` as follow:

After starting your test server you should see the content of `http://example.com/` on `http://localhost:8000/`.

See also:

An example of a project can be found here: <https://github.com/seocam/revproxy-test>

The provided test project is a simple Django project that makes uses of `revproxy`. It basically possess a `view.py` that extends from `ProxyView` and sets the `upstream` address to `'httpbin.org'`.

4.3 Usage

4.3.1 Proxy Views

This document covers the views provided by `reqproxy.views` and all its public attributes

class `reqproxy.views.ProxyView`

Proxies requests to a given upstream server and returns a Django Response.

Example urls.py:

```
from django.urls import re_path

from reqproxy.views import ProxyView

urlpatterns = [
    re_path(r'(?P<path>.*)', ProxyView.as_view(upstream='http://example.com/')),
]
```

Attributes

upstream

The URL of the proxied server. Requests will be made to this URL with `path` (extracted from `urls.py`) appended to it. This attribute is mandatory.

add_remote_user

Whether to add the `REMOTE_USER` to the request in case of an authenticated user. Defaults to `False`.

add_x_forwarded

Whether to add the `X-Forwarded-For` and `X-Forwarded-Proto` headers to the request. Defaults to `False`.

default_content_type

The *Content-Type* that will be added to the response in case the upstream server doesn't send it and if `mimetypes.guess_type` is not able to guess. Defaults to `'application/octet-stream'`.

retries

The max number of attempts for a request. This can also be an instance of `urllib3.Retry`. If set to `None` it will fail if the first attempt fails. The default value is `None`.

rewrite

A list of tuples in the style `(from, to)` where `from` must be a valid regex expression and `to` a valid URL. If `request.get_full_path` matches the `from` expression the request will be redirected to `to` with an status code 302. Matches groups can be used to pass parts from the `from` URL to the `to` URL using numbered groups. By default no rewrite is set.

Example:

```
class CustomProxyView(ProxyView):
    upstream = 'http://www.example.com'
    rewrite = (
        (r'^/yellow/star/$', r'/black/hole/'),
        (r'^/red/?$', r'http://www.mozilla.org/'),

        # Example with numbered match groups
        (r'^/foo/(.*)$', r'/bar\1'),
    )
```

strict_cookies

Whether to only accept RFC-compliant cookies. If set to `True`, any cookies received from the upstream server that do not conform to the RFC will be dropped.

streaming_amount

The buffering amount for streaming HTTP response(in bytes), response will be buffered until it's length exceeds this value. `None` means using default value, override this variable to change.

Methods

class `revproxy.views.DiazoProxyView`

In addition to `ProxyView` behavior this view also performs Diazo transformations on the response before sending it back to the original client. Furthermore, it's possible to pass context data to the view thanks to `ContextMixin` behavior through `get_context_data()` method.

See also:

Diazo is an awesome tool developed by Plone Community to perform XSLT transformations in a simpler way. In order to use all Diazo power please refer to: <http://docs.diazo.org/en/latest/>

Example urls.py:

```
from django.urls import re_path

from revproxy.views import DiazoProxyView

proxy_view = DiazoProxyView.as_view(
    upstream='http://example.com/',
    html5=True,
    diazo_theme_template='base.html',
)

urlpatterns = [
    re_path(r'(?P<path>.*)', proxy_view),
]
```

Example base.html

```
<html>
  <head>...</head>
  <body>
    ...
    <div id="content"></div>
    ...Fix all links in the docs (and README file etc) from old to new repo
```

Example diazo.xml

```
<rules
  xmlns="http://namespaces.plone.org/diazo"
  xmlns:css="http://namespaces.plone.org/diazo/css"
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform">

  <!-- Adds 'body' content from example.com into theme #content -->
  <before css:theme-children="#content" css:content-children="body" />
</rules>
```

Attributes

diazo_theme_template

The Django template to be used as Diazo theme. If set to `None` Diazo will be disabled. By default

`diazo.html` will be used.

diazo_rules

The absolute path for the diazo rules file. By default it will look for the file `diazo.xml` on the Django application directory. If set to `None` Diazo will be disabled.

html5

By default Diazo changes the doctype for html5 to html4. If this attribute is set to `True` the doctype will be kept. This attribute only works if Diazo transformations are enabled.

Methods

4.4 API

4.4.1 `revproxy.views`

4.4.2 `revproxy.response`

4.4.3 `revproxy.transformer`

4.4.4 `revproxy.utils`

`revproxy.utils.DEFAULT_CHARSET = 'latin-1'`

Variable that represent the default charset used

`revproxy.utils.HTML_CONTENT_TYPES = ('text/html', 'application/xhtml+xml')`

List containing string constants that represents possible html content type

`revproxy.utils.IGNORE_HEADERS = ('HTTP_ACCEPT_ENCODING', 'HTTP_HOST', 'HTTP_REMOTE_USER')`

List containing string constant that are used to represent headers that can be ignored in the `required_header` function

`revproxy.utils.MIN_STREAMING_LENGTH = 4096`

Variable used to represent a minimal content size required for response to be turned into stream

`revproxy.utils.cookie_from_string(cookie_string, strict_cookies=False)`

Parser for HTTP header set-cookie The return from this function will be used as parameters for django's `response.set_cookie` method. Because `set_cookie` doesn't have parameter comment, this cookie attribute will be ignored.

Parameters

- **cookie_string** – A string representing a valid cookie
- **strict_cookies** – Whether to only accept RFC-compliant cookies

Returns A dictionary containing the `cookie_string` attributes

`revproxy.utils.encode_items(items)`

Function that encode all elements in the list of items passed as a parameter

Parameters **items** – A list of tuple

Returns A list of tuple with all items encoded in 'utf-8'

`revproxy.utils.get_charset(content_type)`

Function used to retrieve the charset from a content-type. If there is no charset in the content type then the charset defined on `DEFAULT_CHARSET` will be returned

Parameters `content_type` – A string containing a Content-Type header

Returns A string containing the charset

`revproxy.utils.is_html_content_type` (*content_type*)

Function used to verify if the parameter is a proper html content type

Parameters `content_type` – String variable that represent a content-type

Returns A boolean value stating if the `content_type` is a valid html content type

`revproxy.utils.normalize_request_headers` (*request*)

Function used to transform header, replacing 'HTTP_' to '' and replace '_' to '-'

Parameters `request` – A HttpRequest that will be transformed

Returns A dictionary with the normalized headers

`revproxy.utils.required_header` (*header*)

Function that verify if the header parameter is a essential header

Parameters `header` – A string represented a header

Returns A boolean value that represent if the header is required

`revproxy.utils.set_response_headers` (*response*, *response_headers*)

`revproxy.utils.should_stream` (*proxy_response*)

Function to verify if the `proxy_response` must be converted into a stream. This will be done by checking the `proxy_response` content-length and verify if its length is bigger than one stipulated by `MIN_STREAMING_LENGTH`.

Parameters `proxy_response` – An Instance of `urllib3.response.HTTPResponse`

Returns A boolean stating if the `proxy_response` should be treated as a stream

`revproxy.utils.unquote` (*value*)

Remove wrapping quotes from a string.

Parameters `value` – A string that might be wrapped in double quotes, such as a HTTP cookie value.

Returns Beginning and ending quotes removed and escaped quotes (") unescaped

4.5 Changelog

4.5.1 0.11.1 (2023-02-26)

- No changes.

4.5.2 0.11.0 (2023-02-25)

- Add X-Forwarded-For and X-Forwarded-Proto headers. Fixes #79.
- Add Django 3.2, 4.0 and 4.1 support. Fixes #126.
- Add Python 3.8, 3.9, 3.10 and 3.11 support
- Drop Python 3.4, 3.5 and 3.6 support
- Drop Django <3.0 support

- Fixed README badges

4.5.3 0.10.0 (2020-02-05)

- Fix `add_remote_user` when run without `AuthenticationMiddleware`. Fix #86
- Add `get_encoded_query_params` method
- Add support for Python 3.7 and 3.8.
- Add support for Django 2.2 and 3.0.

4.5.4 0.9.15 (2018-05-30)

- Fix issues with latest `urllib3`. Fixes #75.
- Fix issues with parsing cookies. Fixes #84.
- Drop Python 3.3, 3.4, and PyPy support.
- Add Python 3.6 support.

4.5.5 0.9.14 (2018-01-11)

- Move construction of proxied path to method [[@dimrozakis](#)]
- `User.get_username()` rather than `User.name` to support custom User models [[@acordiner](#)]

4.5.6 0.9.13 (2016-10-31)

- Added support to Django 1.10 (support to 1.7 was dropped)

4.5.7 0.9.12 (2016-05-23)

- Fixed error 500 caused by content with wrong encoding [[@lucaskanashiro](#), [@macartur](#)]

4.5.8 0.9.11 (2016-03-29)

- Updated `urllib3` to 1.12 (at least)

4.5.9 0.9.10 (2016-02-03)

- Fixed Python 3 compatibility issue (see #59 and #61). Thanks [@stefanklug](#) and [@macro1!](#)

4.5.10 0.9.9 (2015-12-15)

- Reorder header prior to `httplib` request. `Host` should be always the first request header.

4.5.11 0.9.8 (2015-12-10)

- Added support to Django 1.9 (dropped support to Django 1.6)
- Added `get_request_headers` to make easier to set and override request headers

4.5.12 0.9.7 (2015-09-17)

- Bug fixed: property preventing to set upstream and diazo_rules (#53, #54) [@vdemin]
- Security issue fixed: when colon is present at URL path urljoin ignores the upstream and the request is redirected to the path itself allowing content injection

4.5.13 0.9.6 (2015-09-09)

- Fixed connections pools
- Use wsgiref to check for hop-by-hop headers [#50]
- Refactored tests
- Fixed security issue that allowed remote-user header injection

4.5.14 0.9.5 (2015-09-02)

- Added `extras_require` to make easier diazo installation

4.5.15 0.9.4 (2015-08-27)

- Allow to send context dict to transformation template. [@chaws, @macartur]

4.5.16 0.9.3 (2015-06-12)

- Use StringIO instead of BytesIO on theme compilation (transformation)

4.5.17 0.9.2 (2015-06-09)

Thanks @rafamanzo for the reports.

- Append a backslash on upstream when needed
- Validate upstream URL to make sure it has a scheme
- Added branch test coverage

4.5.18 0.9.1 (2015-05-18)

- More permissive URL scheme (#41).
- Refactored code to allow setting custom headers by extending method (#40) [@marciomazza]

4.5.19 0.9.0 (2015-03-04)

- urllib2 replaced by urllib3 (#10)
- No Diazo transformation if header X-Diazo-Off is set to true - either request or response (#15)
- Removed double memory usage when reading response body (#16)
- Fixed bug caused by many set-cookies coming from upstream (#23) - by @thiagovsk and @macartur
- Added stream support for serving big files with an acceptable memory footprint (#17 and #24). Thanks to @lucasmoura, @macartur, @carloshfoliveira and @thiagovsk.
- Moved Diazo functionalities to DiazoProxyView.
- Logging improved (#21).
- Added options for default_content_type and retries [@gldnspud].
- Sphinx docs (#25).
- 100% test coverage.

CHAPTER 5

Indices and tables

- `genindex`
- `modindex`
- `search`

r

`revproxy.utils`, 13

A

`add_remote_user` (*revproxy.views.ProxyView* attribute), 11
`add_x_forwarded` (*revproxy.views.ProxyView* attribute), 11

C

`cookie_from_string()` (*in module revproxy.utils*), 13

D

`DEFAULT_CHARSET` (*in module revproxy.utils*), 13
`default_content_type` (*revproxy.views.ProxyView* attribute), 11
`diazo_rules` (*revproxy.views.DiazoProxyView* attribute), 13
`diazo_theme_template` (*revproxy.views.DiazoProxyView* attribute), 12

E

`encode_items()` (*in module revproxy.utils*), 13

G

`get_charset()` (*in module revproxy.utils*), 13

H

`html5` (*revproxy.views.DiazoProxyView* attribute), 13
`HTML_CONTENT_TYPES` (*in module revproxy.utils*), 13

I

`IGNORE_HEADERS` (*in module revproxy.utils*), 13
`is_html_content_type()` (*in module revproxy.utils*), 14

M

`MIN_STREAMING_LENGTH` (*in module revproxy.utils*), 13

N

`normalize_request_headers()` (*in module revproxy.utils*), 14

R

`required_header()` (*in module revproxy.utils*), 14
`retries` (*revproxy.views.ProxyView* attribute), 11
`revproxy.utils` (*module*), 13
`revproxy.views.DiazoProxyView` (*built-in class*), 12
`revproxy.views.ProxyView` (*built-in class*), 11
`rewrite` (*revproxy.views.ProxyView* attribute), 11

S

`set_response_headers()` (*in module revproxy.utils*), 14
`should_stream()` (*in module revproxy.utils*), 14
`streaming_amount` (*revproxy.views.ProxyView* attribute), 12
`strict_cookies` (*revproxy.views.ProxyView* attribute), 11

U

`unquote()` (*in module revproxy.utils*), 14
`upstream` (*revproxy.views.ProxyView* attribute), 11