
Cliquet Documentation

Release 2.10.0

Mozilla Services — Da French Team

October 30, 2015

1	Table of content	3
1.1	Rationale	3
1.2	Getting started	7
1.3	HTTP Protocol	10
1.4	Internals	27
1.5	Ecosystem	59
1.6	Contributing	63
1.7	Changelog	66
1.8	Contributors	79
2	Indices and tables	81
	Python Module Index	83

Cliquet is a toolkit to ease the implementation of HTTP microservices, such as data-driven REST APIs.

- [Online documentation](#)
- [Issue tracker](#)
- [Contributing](#)
- Talks [at PyBCN](#) and [at PyConFR](#)

Fig. 1: A cliquet, or ratchet, is a mechanical device that allows continuous linear or rotary motion in only one direction while preventing motion in the opposite direction.

A cliquet provides some basic but essential functionality — efficient in a variety of contexts, from bikes rear wheels to most advanced clockmaking!

Table of content

1.1 Rationale

Cliquet is a toolkit to ease the implementation of HTTP [microservices](#). It is mainly focused on data-driven REST APIs (aka *CRUD*).

1.1.1 Philosophy

- *KISS*;
- No magic;
- Works with defaults;
- Easy customization;
- Straightforward component substitution.

Cliquet doesn't try to be a framework: any project built with *Cliquet* will expose a well defined HTTP protocol for:

- Collection and records manipulation;
- HTTP status and headers handling;
- API versioning and deprecation;
- Errors formatting.

This protocol is an implementation of a series of good practices (followed at [Mozilla Services](#) and [elsewhere](#)).

The goal is to produce standardized APIs, which follow some well known patterns, encouraging genericity in clients code ¹.

Of course, *Cliquet* can be *extended* and customized in many ways. It can also be used in any kind of project, for its tooling, utilities and helpers.

1.1.2 Features

It is built around the notion of resources: resources are defined by sub-classing, and *Cliquet* brings up the HTTP endpoints automatically.

¹ Switch from custom protocol to JSON-API spec is being discussed.

Records and synchronization

- Collection of records by user;
- Optional validation from schema;
- Sorting and filtering;
- Pagination using continuation tokens;
- Polling for collection changes;
- Record race conditions handling using preconditions headers.

Generic endpoints

- Hello view at root url;
- Heartbeat for monitoring;
- Batch operations;
- API versioning and deprecation;
- Errors formatting;
- Backoff and Retry-After headers.

Toolkit

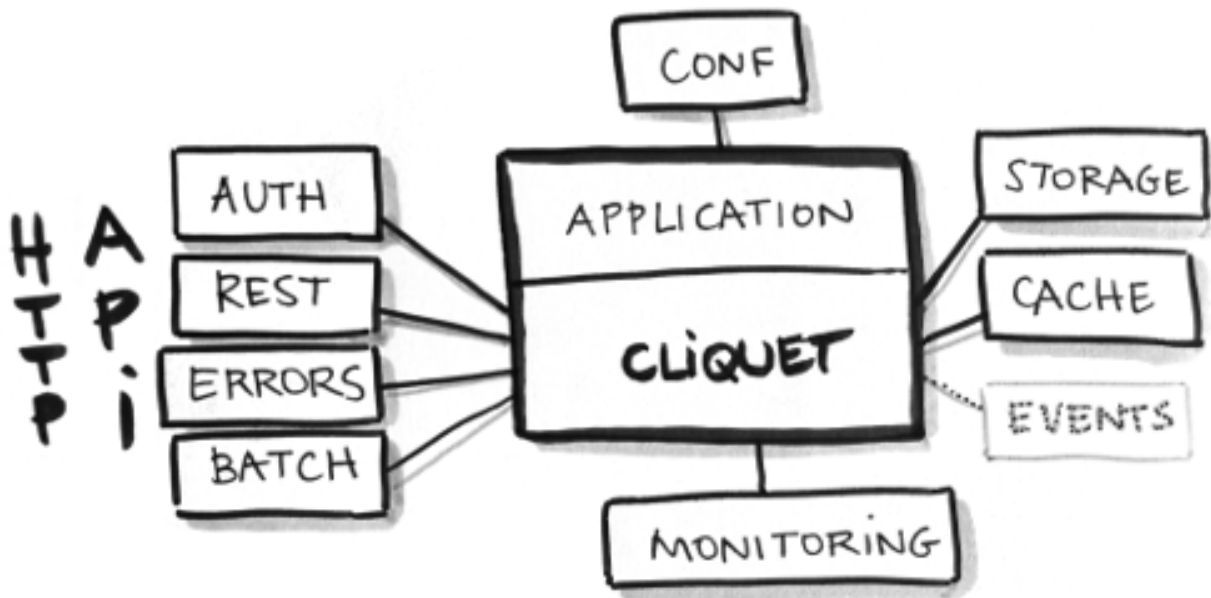


Fig. 1.1: *Cliquet* brings a set of simple but essential features to build APIs.

- Configuration through INI files or environment variables;
- Pluggable storage and cache backends;
- Pluggable authentication and user groups management;

- Pluggable authorization and permissions management;
- Structured logging;
- Monitoring tools;
- Profiling tools.

Pluggable components can be replaced by another one via configuration.

1.1.3 Dependencies

Cliquet is built on the shoulders of giants:

- [Cornice](#) for the REST helpers;
- [Pyramid](#) for the heavy HTTP stuff;

Everything else is meant to be **pluggable and optional**.

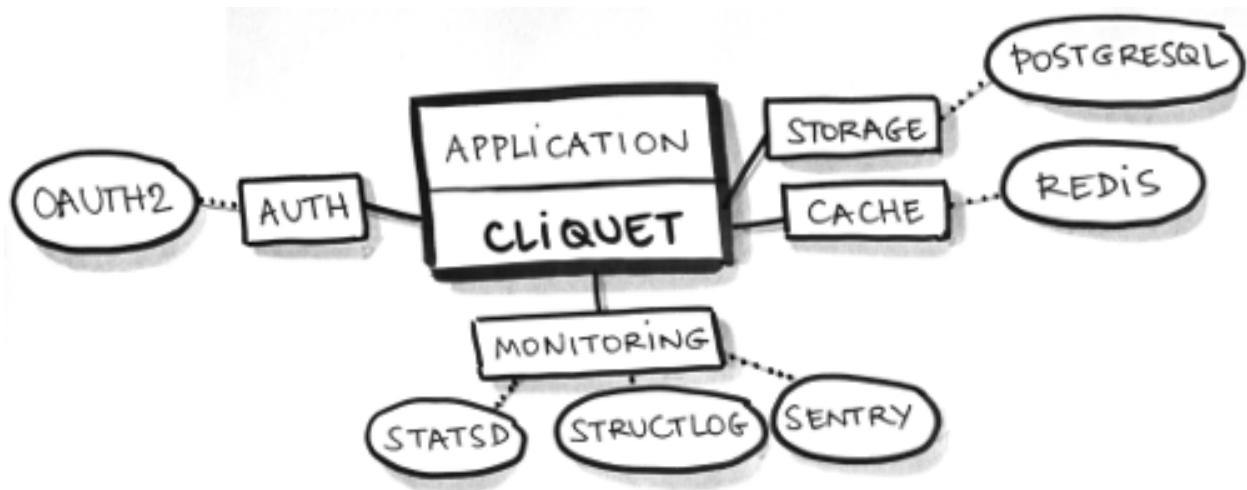


Fig. 1.2: Examples of configuration for a *Cliquet* application in production.

- *Basic Auth*, *FxA OAuth2* or any other source of authentication;
- *Default* or custom class for authorization logics;
- *PostgreSQL* for storage;
- *Redis* for key-value cache with expiration;
- *StatsD* metrics;
- *Sentry* reporting via logging;
- *NewRelic* database profiling (*for development*);
- *Werkzeug* Python code profiling (*for development*).

A *Cliquet* application can change or force default values for any setting.

1.1.4 Built with Cliquet

Some applications in the wild built with *Cliquet*:

- [Reading List](#), a service to synchronize articles between devices;
- [Kinto](#), a service to store and synchronize schema-less data.
- [Syncto](#), a service to access *Firefox Sync* using *kinto.js*.
- *Please contact us to add yours.*

Context

(to be done)

- Cloud Services team at Mozilla
- [ReadingList](#) project story
- Firefox Sync
- Cloud storage
- Firefox OS User Data synchronization and backup

1.1.5 Vision

General

Any application built with *Cliquet*:

- follows the same conventions regarding the HTTP API;
- takes advantage of its component *pluggability*;
- can be *extended* using custom code or Pyramid external packages;

Let's build a *sane ecosystem* for microservices in Python!

Roadmap

The future features we plan to implement in *Cliquet* are currently driven by the use-cases we meet internally at Mozilla. Most notable are:

- Notifications channel (e.g. run asynchronous tasks on events or listen for changes);
- Attachments on records (e.g. *Remote Storage* compatibility);
- Records generic indexing (e.g. streaming records to *ElasticSearch*).
- ... come and discuss [enhancements in the issue tracker](#)!

1.1.6 Similar projects

- [Python Eve](#), built on Flask and MongoDB;
- *Please contact us to add more if any.*

Since the protocol is language independant and follows good HTTP/REST principles, in the long term *Cliquet* should become only one among several server implementations.

Note: We encourage you to implement a clone of this project — using Node.js, Asyncio, Go, Twisted, Django or anything else — following *the same protocol*!

1.2 Getting started

1.2.1 Installation

```
$ pip install cliquet
```

More details about installation and storage backend is provided in *a dedicated section*.

1.2.2 Start a Pyramid project

As detailed in [Pyramid documentation](#), create a minimal application, or use its scaffolding tool:

```
$ pcreate -s starter MyProject
```

Include Cliquet

In the application main file (e.g. `MyProject/myproject/__init__.py`), just add some extra initialization code:

```
import pkg_resources

import cliquet
from pyramid.config import Configurator

# Module version, as defined in PEP-0396.
__version__ = pkg_resources.get_distribution(__package__).version

def main(global_config, **settings):
    config = Configurator(settings=settings)

    cliquet.initialize(config, __version__)
    return config.make_wsgi_app()
```

By doing that, basic features like authentication, monitoring, error formatting, deprecation indicators are now available, and rely on configuration present in `myproject.ini`.

Note: Shortcut!

In order to bypass the installation and configuration of *Redis* required by the default storage, permission manager and cache, use the «in-memory» backend in `development.ini`:

```
# development.ini
cliquet.cache_backend = cliquet.cache.memory
cliquet.storage_backend = cliquet.storage.memory
cliquet.permission_backend = cliquet.permission.memory
```

Now is a good time to install the *Cliquet* project locally:

```
$ pip install -e .
```

Run!

With some backends, like *PostgreSQL*, some tables and indices have to be created. A generic command is provided to accomplish this:

```
$ cliquet --ini development.ini migrate
```

Like any *Pyramid* application, it can be served locally with:

```
$ pserve development.ini --reload
```

A *hello* view is now available at <http://localhost:6543/v0/> (As well as basic endpoints like the *utilities*).

The next steps will consist in building a custom application using *Cornice* or **the Pyramid ecosystem**.

But most likely, it will consist in **defining REST resources** using *Cliquet* python API !

Authentication

Currently, if no *authentication is set in settings*, *Cliquet* relies on *Basic Auth*. It will associate a unique *user id* for every user/password combination.

Using *HTTPie*, it is as easy as:

```
$ http -v http://localhost:6543/v0/ --auth user:pass
```

Note: In the case of *Basic Auth*, there is no need of registering a user/password. Pick any combination, and include them in each request.

1.2.3 Define resources

In order to define a resource, inherit from `cliquet.resource.BaseResource`, in a subclass, in `myproject/views.py` for example:

```
from cliquet import resource

@resource.register()
class Mushroom(resource.BaseResource):
    # No schema yet.
    pass
```

In application initialization, make *Pyramid* aware of it:

```
def main(global_config, **settings):
    config = Configurator(settings=settings)

    cliquet.initialize(config, __version__)
    config.scan("myproject.views")
    return config.make_wsgi_app()
```

By doing that, a Mushroom resource API is now available at the `/mushrooms/` endpoint.

It will accept a bunch of REST operations, as defined in the [API section](#).

Warning: Without schema, a resource will not store any field at all!

The next step consists in attaching a schema to the resource, to control what fields are accepted and stored.

Schema validation

It is possible to validate records against a predefined schema, associated to the resource.

Currently, only `Colander` is supported, and it looks like this:

```
import colander
from cliquet import resource

class MushroomSchema(resource.ResourceSchema):
    name = colander.SchemaNode(colander.String())

@resource.register()
class Mushroom(resource.BaseResource):
    mapping = MushroomSchema()
```

1.2.4 What's next ?

Configuration

See [Configuration](#) to customize the application settings, such as authentication, storage or cache backends.

Resource customization

See [the resource documentation](#) to specify custom URLs, schemaless resources, read-only fields, unicity constraints, record pre-processing...

Advanced initialization

```
cliquet.initialize(config, version=None, project_name='', default_settings=None)
```

Initialize Cliquet with the given configuration, version and project name.

This will basically include cliquet in Pyramid and set route prefix based on the specified version.

Parameters

- **config** (*Configurator*) – Pyramid configuration
- **version** (*str*) – Current project version (e.g. '0.0.1') if not defined in application settings.
- **project_name** (*str*) – Project name if not defined in application settings.
- **default_settings** (*dict*) – Override cliquet default settings values.

Beyond Cliquet

Cliquet is just a component! The application can still be built and extended using the full *Pyramid* ecosystem. See [the dedicated section](#) for examples of *Cliquet* extensions.

1.3 HTTP Protocol

1.3.1 API versioning

The API versioning is based on the application version deployed. It follows the [semver](#) specifications.

During development the server will be 0.X.X, the server endpoint will be prefixed by `/v0/`.

Each non retro-compatible API change will imply the major version number to be incremented. Everything will be made to avoid retro incompatible changes.

The `/` endpoint will redirect to the last API version.

Warning: The version prefix will be **implied** throughout the rest of the API reference, to improve readability. For example, the `/` endpoint should be understood as `/v0/`.

1.3.2 Authentication

Depending on the authentication policies initialized in the application, the HTTP method to authenticate requests may differ.

A policy based on *OAuth2 bearer tokens* is recommended, but not mandatory. See [configuration](#) for further information.

In the current implementation, when multiple policies are configured, *user identifiers* are isolated by policy. In other words, there is no way to access the same set of records using different authentication methods.

By default, a relatively secure *Basic Auth* is enabled.

Basic Auth

If enabled in configuration, using a *Basic Auth* token will associate a unique *user identifier* to an username/password combination.

```
Authorization: Basic <basic_token>
```

The token shall be built using this formula `base64 ("username:password")`.

Empty passwords are accepted, and usernames can be anything (custom, UUID, etc.)

If the token has an invalid format, or if *Basic Auth* is not enabled, this will result in a 401 error response.

Warning: Since *user id* is derived from username and password, there is no way to change the password without loosing access to existing records.

OAuth Bearer token

If the configured authentication policy uses *OAuth2 bearer tokens*, authentication shall be done using this header:

```
Authorization: Bearer <oauth_token>
```

The policy will verify the provided *OAuth2 bearer token* on a remote server.

notes If the token is not valid, this will result in a 401 error response.

Firefox Accounts

In order to enable authentication with *Firefox Accounts*, install and configure [mozilla-services/cliquet-fxa](https://mozilla-services.com/cliquet-fxa).

1.3.3 Resource endpoints

GET /{collection}

Requires authentication

Returns all records of the current user for this collection.

The returned value is a JSON mapping containing:

- **data**: the list of records, with exhaustive fields;
- **permissions**: *optional* a json dict containing the permissions for the collection of records.

A **Total-Records** response header indicates the total number of records of the collection.

A **Last-Modified** response header provides a human-readable (rounded to second) of the current collection timestamp.

For cache and concurrency control, an **ETag** response header gives the value that consumers can provide in subsequent requests using **If-Match** and **If-None-Match** headers (see [section about timestamps](#)).

Request:

```
GET /articles HTTP/1.1
Accept: application/json
Authorization: Basic bWF0Og==
Host: localhost:8000
```

Response:

```
HTTP/1.1 200 OK
Access-Control-Allow-Origin: *
Access-Control-Expose-Headers: Backoff, Retry-After, Alert, Content-Length, ETag, Next-Page, Total-Records
Content-Length: 436
Content-Type: application/json; charset=UTF-8
Date: Tue, 28 Apr 2015 12:08:11 GMT
Last-Modified: Mon, 12 Apr 2015 11:12:07 GMT
ETag: "1430222877724"
Total-Records: 2

{
  "data": [
    {
      "id": "dc86afa9-a839-4ce1-ae02-3d538b75496f",
      "last_modified": 1430222877724,
```

```
    "title": "MoCo",
    "url": "https://mozilla.com",
  },
  {
    "id": "23160c47-27a5-41f6-9164-21d46141804d",
    "last_modified": 1430140411480,
    "title": "MoFo",
    "url": "https://mozilla.org",
  }
]
```

Filtering

Single value

- /collection?field=value

Minimum and maximum

Prefix attribute name with min_ or max_:

- /collection?min_field=4000

Note: The lower and upper bounds are inclusive (*i.e equivalent to greater or equal*).

Note: lt_ and gt_ can also be used to exclude the bound.

Multiple values

Prefix attribute with in_ and provide comma-separated values.

- /collection?in_status=1,2,3

Exclude

Prefix attribute name with not_:

- /collection?not_field=0

Exclude multiple values

Prefix attribute name with exclude_:

- /collection?exclude_field=0,1

Note: Will return an error if a field is unknown.

Note: The ETag and Last-Modified response headers will always be the same as the unfiltered collection.

Sorting

- /collection?_sort=-last_modified,field

Note: Ordering on a boolean field gives true values first.

Note: Will return an error if a field is unknown.

Counting

In order to count the number of records, for a specific field value for example, without fetching the actual collection, a HEAD request can be used. The Total-Records response header will then provide the total number of records.

See *batch endpoint* to count several collections in one request.

Polling for changes

The `_since` parameter is provided as an alias for `gt_last_modified`.

- `/collection?_since=1437035923844`

When filtering on `last_modified` every deleted records will appear in the list with a `deleted` flag and a `last_modified` value that corresponds to the deletion event.

If the request header `If-None-Match` is provided as described in the *section about timestamps* and if the collection was not changed, a 304 Not Modified response is returned.

Note: The `_before` parameter is also available, and is an alias for `lt_last_modified` (*strictly inferior*).

Changed in version 2.4::It will be supported until the next major version of Cliquet.

Request:

```
to was renamed before and is now deprecated HTTP/1.1
Accept: application/json
Authorization: Basic bWF0Og==
Host: localhost:8000
```

Response:

```
HTTP/1.1 200 OK
Access-Control-Allow-Origin: *
Access-Control-Expose-Headers: Backoff, Retry-After, Alert, Content-Length, ETag, Next-Page, Total-Records
Content-Length: 436
Content-Type: application/json; charset=UTF-8
Date: Tue, 28 Apr 2015 12:08:11 GMT
Last-Modified: Mon, 12 Apr 2015 11:12:07 GMT
ETag: "1430222877724"
Total-Records: 2

{
  "data": [
    {
      "id": "dc86afa9-a839-4ce1-ae02-3d538b75496f",
      "last_modified": 1430222877724,
      "title": "MoCo",
      "url": "https://mozilla.com",
    },
    {
      "id": "23160c47-27a5-41f6-9164-21d46141804d",
      "last_modified": 1430140411480,
      "title": "MoFo",
      "url": "https://mozilla.org",
    }
  ]
}
```

```
    },
    {
      "id": "11130c47-37a5-41f6-9112-32d46141804f",
      "deleted": true,
      "last_modified": 1430140411480
    }
  ]
}
```

Paginate

If the `_limit` parameter is provided, the number of records returned is limited.

If there are more records for this collection than the limit, the response will provide a `Next-Page` header with the URL for the Next-Page.

When there is no more `Next-Page` response header, there is nothing more to fetch.

Pagination works with sorting, filtering and polling.

Note: The `Next-Page` URL will contain a continuation token (`_token`).

It is recommended to add precondition headers (`If-Match` or `If-None-Match`), in order to detect changes on collection while iterating through the pages.

List of available URL parameters

- `<prefix?><attribute name>`: filter by value(s)
- `_since`, `_before`: polling changes
- `_sort`: order list
- `_limit`: pagination max size
- `_token`: pagination token

Filtering, sorting and paginating can all be combined together.

- `/collection?_sort=-last_modified&_limit=100`

HTTP Status Codes

- 200 OK: The request was processed
- 304 Not Modified: Collection did not change since value in `If-None-Match` header
- 400 Bad Request: The request querystring is invalid
- 412 Precondition Failed: Collection changed since value in `If-Match` header

POST `/collection`

Requires authentication

Used to create a record in the collection. The POST body is a JSON mapping containing:

- `data`: the values of the resource schema fields;
- `permissions`: *optional* a json dict containing the permissions for the record to be created.

The POST response body is a JSON mapping containing:

- `data`: the newly created record, if all posted values are valid;
- `permissions`: *optional* a json dict containing the permissions for the requested resource.

If the request header `If-Match` is provided, and if the record has changed meanwhile, a 412 `Precondition failed` error is returned.

Request:

```
POST /articles HTTP/1.1
Accept: application/json
Authorization: Basic bWF0Og==
Content-Type: application/json; charset=utf-8
Host: localhost:8000

{
  "data": {
    "title": "Wikipedia FR",
    "url": "http://fr.wikipedia.org"
  }
}
```

Response:

```
HTTP/1.1 201 Created
Access-Control-Allow-Origin: *
Access-Control-Expose-Headers: Backoff, Retry-After, Alert, Content-Length
Content-Length: 422
Content-Type: application/json; charset=UTF-8
Date: Tue, 28 Apr 2015 12:35:02 GMT

{
  "data": {
    "id": "cd30c031-c208-4fb9-ad65-1582d2a7ad5e",
    "last_modified": 1430224502529,
    "title": "Wikipedia FR",
    "url": "http://fr.wikipedia.org"
  }
}
```

Validation

If the posted values are invalid (e.g. *field value is not an integer*) an error response is returned with status 400.

See [details on error responses](#).

Conflicts

Since some fields can be *defined as unique* per collection (per user), some conflicts may appear when creating records.

Note: Empty values are not taken into account for field unicity.

Note: Deleted records are not taken into account for field unicity.

If a conflict occurs, an error response is returned with status 409. A `details` attribute in the response provides the offending record and field name. See [:ref:'dedicated section about errors <error-responses>'](#).

HTTP Status Codes

- 201 Created: The record was created
- 400 Bad Request: The request body is invalid
- 409 Conflict: Unicity constraint on fields is violated
- 412 Precondition Failed: Collection changed since value in If-Match header

DELETE /{collection}

Requires authentication

Delete multiple records. **Disabled by default**, see *Configuration*.

The DELETE response is a JSON mapping containing:

- `data`: list of records that were deleted, without schema fields.

It supports the same filtering capabilities as GET.

If the request header `If-Match` is provided, and if the collection has changed meanwhile, a 412 `Precondition failed` error is returned.

Request:

```
DELETE /articles HTTP/1.1
Accept: application/json
Authorization: Basic bWF0Og==
Host: localhost:8000
```

Response:

```
HTTP/1.1 200 OK
Access-Control-Allow-Origin: *
Access-Control-Expose-Headers: Backoff, Retry-After, Alert, Content-Length
Content-Length: 193
Content-Type: application/json; charset=UTF-8
Date: Tue, 28 Apr 2015 12:38:36 GMT

{
  "data": [
    {
      "deleted": true,
      "id": "cd30c031-c208-4fb9-ad65-1582d2a7ad5e",
      "last_modified": 1430224716097
    },
    {
      "deleted": true,
      "id": "dc86afa9-a839-4ce1-ae02-3d538b75496f",
      "last_modified": 1430224716098
    }
  ]
}
```

HTTP Status Codes

- 200 OK: The records were deleted;
- 405 Method Not Allowed: This endpoint is not available;
- 412 Precondition Failed: Collection changed since value in If-Match header

GET /{collection}/{<id>}

Requires authentication

Returns a specific record by its id. The GET response body is a JSON mapping containing:

- data: the record with exhaustive schema fields;
- permissions: *optional* a json dict containing the permissions for the requested record.

If the request header If-None-Match is provided, and if the record has not changed meanwhile, a 304 Not Modified is returned.

Request:

```
GET /articles/d10405bf-8161-46a1-ac93-a1893d160e62 HTTP/1.1
Accept: application/json
Authorization: Basic bWF0Og==
Host: localhost:8000
```

Response:

```
HTTP/1.1 200 OK
Access-Control-Allow-Origin: *
Access-Control-Expose-Headers: Backoff, Retry-After, Alert, Content-Length, ETag, Last-Modified
Content-Length: 438
Content-Type: application/json; charset=UTF-8
Date: Tue, 28 Apr 2015 12:42:42 GMT
ETag: "1430224945242"

{
  "data": {
    "id": "d10405bf-8161-46a1-ac93-a1893d160e62",
    "last_modified": 1430224945242,
    "title": "No backend",
    "url": "http://nobackend.org"
  }
}
```

HTTP Status Code

- 200 OK: The request was processed
- 304 Not Modified: Record did not change since value in If-None-Match header
- 412 Precondition Failed: Record changed since value in If-Match header

DELETE /{collection}/{<id>}

Requires authentication

Delete a specific record by its id.

The DELETE response is the record that was deleted. The DELETE response is a JSON mapping containing:

- `data`: the record that was deleted, without schema fields.

If the record is missing (or already deleted), a `404 Not Found` is returned. The consumer might decide to ignore it.

If the request header `If-Match` is provided, and if the record has changed meanwhile, a `412 Precondition failed` error is returned.

Note: Once deleted, a record will appear in the collection when polling for changes, with a `deleted` status (`delete=true`) and will have most of its fields empty.

HTTP Status Code

- `200 OK`: The record was deleted
- `412 Precondition Failed`: Record changed since value in `If-Match` header

PUT `/collection/<id>`

Requires authentication

Create or replace a record with its id. The PUT body is a JSON mapping containing:

- `data`: the values of the resource schema fields;
- `permissions`: *optional* a json dict containing the permissions for the record to be created.

The PUT response body is a JSON mapping containing:

- `data`: the newly created/updated record, if all posted values are valid;
- `permissions`: *optional* the newly created permissions dict, containing the permissions for the created record.

Validation and conflicts behaviour is similar to creating records (POST).

If the request header `If-Match` is provided, and if the record has changed meanwhile, a `412 Precondition failed` error is returned.

Request:

```
PUT /articles/d10405bf-8161-46a1-ac93-a1893d160e62 HTTP/1.1
Accept: application/json
Authorization: Basic bWF0Og==
Content-Type: application/json; charset=utf-8
Host: localhost:8000

{
  "data": {
    "title": "Static apps",
    "url": "http://www.staticapps.org"
  }
}
```

Response:

```

HTTP/1.1 200 OK
Access-Control-Allow-Origin: *
Access-Control-Expose-Headers: Backoff, Retry-After, Alert, Content-Length
Content-Length: 439
Content-Type: application/json; charset=UTF-8
Date: Tue, 28 Apr 2015 12:46:36 GMT
ETag: "1430225196396"

{
  "data": {
    "id": "d10405bf-8161-46a1-ac93-a1893d160e62",
    "last_modified": 1430225196396,
    "title": "Static apps",
    "url": "http://www.staticapps.org"
  }
}

```

HTTP Status Code

- 201 Created: The record was created
- 200 OK: The record was replaced
- 400 Bad Request: The record is invalid
- 409 Conflict: If replacing this record violates a field unicity constraint
- 412 Precondition Failed: Record was changed or deleted since value in If-Match header.

Note: A If-None-Match: * request header can be used to make sure the PUT won't overwrite any record.

PATCH /{collection}/{<id>}

Requires authentication

Modify a specific record by its id. The PATCH body is a JSON mapping containing:

- data: a subset of the resource schema fields (*key-value replace*);
- permissions: *optional* a json dict containing the permissions for the record to be modified.

The PATCH response body is a JSON mapping containing:

- data: the modified record (*full by default*);
- permissions: *optional* the newly created permissions dict, containing the permissions for the modified record.

If a request header Response-Behavior is set to light, only the fields whose value was changed are returned. If set to diff, only the fields whose value became different than the one provided are returned.

Request:

```

PATCH /articles/d10405bf-8161-46a1-ac93-a1893d160e62 HTTP/1.1
Accept: application/json
Authorization: Basic bWF0Og==
Content-Type: application/json; charset=utf-8
Host: localhost:8000

```

```
{
  "data": {
    "title": "No Backend"
  }
}
```

Response:

```
HTTP/1.1 200 OK
Access-Control-Allow-Origin: *
Access-Control-Expose-Headers: Backoff, Retry-After, Alert, Content-Length
Content-Length: 439
Content-Type: application/json; charset=UTF-8
Date: Tue, 28 Apr 2015 12:46:36 GMT
ETag: "1430225196396"

{
  "data": {
    "id": "d10405bf-8161-46a1-ac93-a1893d160e62",
    "last_modified": 1430225196396,
    "title": "No Backend",
    "url": "http://nobackend.org"
  }
}
```

If the record is missing (or already deleted), a 404 Not Found error is returned. The consumer might decide to ignore it.

If the request header If-Match is provided, and if the record has changed meanwhile, a 412 Precondition failed error is returned.

Note: `last_modified` is updated to the current server timestamp, only if a field value was changed.

Note: [JSON-Patch](#) is currently not supported. Any help is welcomed though!

Read-only fields

If a read-only field is modified, a 400 Bad request error is returned.

Conflicts

If changing a record field violates a field unicity constraint, a 409 Conflict error response is returned (see [error channel](#)).

HTTP Status Code

- 200 OK: The record was modified
- 400 Bad Request: The request body is invalid, or a read-only field was modified
- 409 Conflict: If modifying this record violates a field unicity constraint
- 412 Precondition Failed: Record changed since value in If-Match header

Protected resources

All of the described endpoints can be either *protected* or not. Protecting an endpoint means that only *principals* which have been granted access will be able to issue requests successfully.

In the case of a *protected* resource, body is a JSON mapping containing a `permissions` key in addition to the `data` key. Permissions can also be replaced and modified independantly from data.

On a request, `permissions` is a json dict containing the permissions for the record to be modified. It has the following signature:

```
'permissions': {'{permission}': [{list_of_principals}]}
```

{permission} is a placeholder for the permission name (e.g. *read*, *write*, *create*) and *{list_of_principals}* should be replaced by an actual list of principals.

`permissions` is also added to JSON mapping response bodies, and contains the *modified* version of the permissions in case of a modification, or the list of permissions in case of a read operation.

Changed in version 2.6:: With a `PATCH` request, the list of principals for the specified permissions is now replaced by the one provided.

1.3.4 Batch operations

POST /batch

Requires authentication

The POST body is a mapping, with the following attributes:

- `requests`: the list of requests
- `defaults`: (*optional*) default requests values in common for all requests

Each request is a JSON mapping, with the following attribute:

- `method`: HTTP verb
- `path`: URI
- `body`: a mapping
- `headers`: (*optional*), otherwise take those of batch request

```
{
  "defaults": {
    "method" : "POST",
    "path" : "/v0/articles",
    "headers" : {
      ...
    }
  },
  "requests": [
    {
      "body" : {
        "title": "MoFo",
        "url" : "http://mozilla.org",
        "added_by": "FxOS",
      }
    },
    {
```

```
    "body" : {
      "title": "MoCo",
      "url" : "http://mozilla.com"
      "added_by": "FxOS",
    }
  },
  {
    "method" : "PATCH",
    "path" : "/articles/409",
    "body" : {
      "read_position" : 3477
    }
  }
]
}
```

The response body is a list of all responses:

```
{
  "responses": [
    {
      "path" : "/articles/409",
      "status": 200,
      "body" : {
        "id": 409,
        "url": "...",
        ...
        "read_position" : 3477
      },
      "headers": {
        ...
      }
    },
    {
      "status": 201,
      "path" : "/articles",
      "body" : {
        "id": 411,
        "title": "MoFo",
        "url" : "http://mozilla.org",
        ...
      },
    },
    {
      "status": 201,
      "path" : "/articles",
      "body" : {
        "id": 412,
        "title": "MoCo",
        "url" : "http://mozilla.com",
        ...
      },
    },
  ],
}
```

HTTP Status Codes

- 200 OK: The request has been processed
- 400 Bad Request: The request body is invalid

Warning: Since the requests bodies are necessarily mappings, posting arbitrary data (*like raw text or binary*) is not supported.

Note: Responses are provided in the same order than requests.

Pros & Cons

- This respects REST principles
- This is easy for the client to handle, since it just has to pile up HTTP requests while offline
- It looks to be a convention for several REST APIs ([Neo4J](#), [Facebook](#), [Parse](#))
- Payload of response can be heavy, especially while importing huge collections
- Payload of response must all be iterated to look-up errors

Note: A form of payload optimization for massive operations is planned.

1.3.5 Utility endpoints for OPS and Devs

GET /

The returned value is a JSON mapping containing:

- `hello`: the name of the service (e.g. "reading list")
- `version`: complete version ("X.Y.Z")
- `url`: absolute URI (without a trailing slash) of the API (*can be used by client to build URIs*)
- `eos`: date of end of support in ISO 8601 format ("yyyy-mm-dd", undefined if unknown)
- `documentation`: The URL to the service documentation. (this document!)
- **settings**: a mapping with the values of relevant public settings for clients
 - `batch_max_requests`: Number of requests that can be made in a batch request.
 - `readonly`: Only requests with read operations are allowed.
- **user**: A mapping with an `id` field for the currently connected user id. The field is not present when no Authorization header is provided.

GET /__heartbeat__

Return the status of each service the application depends on. The returned value is a JSON mapping containing:

- `storage` true if storage backend is operational
- `cache` true if cache backend operational

- `permission` true if permission backend operational

If `cliquet-fxa` is installed, an additional key is present:

- `oauth` true if authentication is operational

Return 200 if the connection with each service is working properly and 503 if something doesn't work.

1.3.6 Server timestamps

In order to avoid race conditions, each change is guaranteed to increment the timestamp of the related collection. If two changes happen at the same millisecond, they will still have two different timestamps.

The ETag header with the current timestamp of the collection for the current user will be given on collection endpoints.

```
ETag: "1432208041618"
```

On record endpoints, the ETag header value will contain the timestamp of the record.

In order to bypass costly and error-prone HTTP date parsing, ETag headers are not HTTP date values.

A human readable version of the timestamp (rounded to second) is provided though in the `Last-Modified` response headers:

```
Last-Modified: Wed May 20 17:22:38 2015 +0200
```

Changed in version 2.0: In previous versions, cache and concurrency control was handled using `If-Modified-Since` and `If-Unmodified-Since`. But since the HTTP date does not include milliseconds, they contained the milliseconds timestamp as integer. The current version using ETag is HTTP compliant (see [original discussion](#).)

Note: The client may send `If-Unmodified-Since` or `If-Modified-Since` requests headers, but in the current implementation, they will be ignored.

Cache control

In order to check that the client version has not changed, a `If-None-Match` request header can be used. If the response is 304 `Not Modified` then the cached version is still good.

Concurrency control

In order to prevent race conditions, like overwriting changes occurred in the interim for example, a `If-Match` request header can be used. If the response is 412 `Precondition failed` then the resource has changed meanwhile.

The client can then choose to:

- overwrite by repeating the request without `If-Match`;
- reconcile the resource by fetching, merging and repeating the request.

1.3.7 Backoff indicators

Backoff header on heavy load

A `Backoff` header will be added to the success responses (≥ 200 and < 400) when the server is under heavy load. It provides the client with a number of seconds during which it should avoid doing unnecessary requests.

```
Backoff: 30
```

Note: The back-off time is configurable on the server.

Note: In other implementations at Mozilla, there was `X-Weave-Backoff` and `X-Backoff` but the `X-` prefix for header [has been deprecated since](#).

Retry-After indicators

A `Retry-After` header will be added to error responses (≥ 500), telling the client how many seconds it should wait before trying again.

```
Retry-After: 30
```

1.3.8 Error responses

Protocol description

Every response is JSON.

If the HTTP status is not OK (< 200 or ≥ 400), the response contains a JSON mapping, with the following attributes:

- `code`: matches the HTTP status code (e.g. 400)
- `errno`: stable application-level error number (e.g. 109)
- `error`: string description of error type (e.g. "Bad request")
- `message`: context information (e.g. "Invalid request parameters")
- `info`: online resource (e.g. URL to error details)
- `details`: additional details (e.g. list of validation errors)

Example response

```
{
  "code": 412,
  "errno": 114,
  "error": "Precondition Failed",
  "message": "Resource was modified meanwhile",
  "info": "https://server/docs/api.html#errors",
}
```

Refer yourself to the `ref:set of errors codes <errors>`.

Precondition errors

As detailed in the [timestamps](#) section, it is possible to add concurrency control using `ETag` request headers.

When a concurrency error occurs, a 412 `Precondition Failed` error response is returned.

Additional information about the record currently stored on the server will be provided in the `details` field:

```
{
  "code": 412,
  "errno": 114,
  "error": "Precondition Failed"
  "message": "Resource was modified meanwhile",
  "details": {
    "existing": {
      "last_modified": 1436434441550,
      "id": "00dd028f-16f7-4755-ab0d-e0dc0cb5da92",
      "title": "Original title"
    }
  },
}
```

Conflict errors

When a record violates unicity constraints, a 409 Conflict error response is returned.

Additional information about conflicting record and field name will be provided in the `details` field.

```
{
  "code": 409,
  "errno": 122,
  "error": "Conflict",
  "message": "Conflict of field url on record eyjafjallajokull"
  "info": "https://server/docs/api.html#errors",
  "details": {
    "field": "url",
    "record": {
      "id": "eyjafjallajokull",
      "last_modified": 1430140411480,
      "url": "http://mozilla.org"
    }
  }
}
```

Validation errors

When multiple validation errors occur on a request, the first one is presented in the message.

The full list of validation errors is provided in the `details` field.

```
{
  "code": 400,
  "errno": 109,
  "error": "Bad Request",
  "message": "Invalid posted data",
  "info": "https://server/docs/api.html#errors",
  "details": [
    {
      "description": "42 is not a string: {'name': ''}",
      "location": "body",
      "name": "name"
    }
  ]
}
```

1.3.9 Deprecation

A track of the client version will be kept to know after which date each old version can be shutdown.

The date of the end of support is provided in the API root URL (e.g. /v0)

Using the `Alert` response header, the server can communicate any potential warning messages, information, or other alerts.

The value is JSON mapping with the following attributes:

- `code`: one of the strings `"soft-eol"` or `"hard-eol"`;
- `message`: a human-readable message (optional);
- `url`: a URL at which more information is available (optional).

A 410 `Gone` error response can be returned if the client version is too old, or the service had been replaced with a new and better service using a new protocol version.

See details in [Configuration](#) to activate deprecation.

1.4 Internals

1.4.1 Installation

By default, a *Cliquet* application persists the records and cache in a local [Redis](#).

Using the [application configuration](#), other backends like « in-memory » or [PostgreSQL](#) can be enabled afterwards.

Supported Python versions

Cliquet supports Python 2.7, Python 3.4 and PyPy.

Distribute & Pip

Installing Cliquet with pip:

```
pip install cliquet
```

For *PostgreSQL* and *monitoring* support:

```
pip install cliquet[postgresql,monitoring]
```

Note: When installing cliquet with postgresql support in a virtualenv using the [PyPy](#) interpreter, the `psycpg2cffi` PostgreSQL database adapter will be installed, instead of the traditional `psycpg2`, as it provides significant [performance improvements](#).

If everything is under control *python*-wise, jump to the next chapter. Otherwise please find more details below.

Python 3.4

Linux

```
sudo apt-get install python3.4-dev
```

OS X

```
brew install python3.4
```

Cryptography libraries

Linux

On Debian / Ubuntu based systems:

```
apt-get install libffi-dev libssl-dev
```

On RHEL-derivatives:

```
apt-get install libffi-devel openssl-devel
```

OS X

Assuming `brew` is installed:

```
brew install libffi openssl pkg-config
```

Install Redis

Linux

On debian / ubuntu based systems:

```
apt-get install redis-server
```

or:

```
yum install redis
```

OS X

Assuming `brew` is installed, Redis installation becomes:

```
brew install redis
```

To restart it (Bug after configuration update):

```
brew services restart redis
```


Install PostgreSQL

Client libraries only

Install PostgreSQL client headers:

```
sudo apt-get install libpq-dev
```

Install Cliquet with related dependencies:

```
pip install cliquet[postgresql]
```

Full server

PostgreSQL version 9.4 (or higher) is required.

To install PostgreSQL on Ubuntu/Debian use:

```
sudo apt-get install postgresql-9.4
```

If your Ubuntu/Debian distribution doesn't include version 9.4 of PostgreSQL look at the [PostgreSQL Ubuntu](#) and [PostgreSQL Debian](#) pages. The PostgreSQL project provides an Apt Repository that one can use to install recent PostgreSQL versions.

By default, the `postgres` user has no password and can hence only connect if ran by the `postgres` system user. The following command will assign it:

```
sudo -u postgres psql -c "ALTER USER postgres PASSWORD 'postgres';"
```

Cliquet requires UTC to be used as the database timezone, and UTF-8 as the database encoding. You can for example use the following commands to create a database named `testdb` with the appropriate timezone and encoding:

```
sudo -u postgres psql -c "ALTER ROLE postgres SET TIMEZONE TO 'UTC';"
sudo -u postgres psql -c "CREATE DATABASE testdb ENCODING 'UTF-8';"
```

Server using Docker

Install docker, for example on Ubuntu:

```
sudo apt-get install docker.io
```

Run the official PostgreSQL container locally:

```
postgres=$(sudo docker run -d -p 5432:5432 postgres)
```

(optional) Create the test database:

```
psql -h localhost -U postgres -W
#> CREATE DATABASE "testdb";
```

Tag and save the current state with:

```
sudo docker commit $postgres cliquet-empty
```

In the future, run the tagged version of the container

```
cliquet=$(sudo docker run -d -p 5432:5432 cliquet-empty)

...

sudo docker stop $cliquet
```

1.4.2 Configuration

See [Pyramid settings documentation](#).

Environment variables

In order to ease deployment or testing strategies, *Cliquet* reads settings from environment variables, in addition to `.ini` files.

For example, `cliquet.storage_backend` is read from environment variable `CLIQUET_STORAGE_BACKEND` if defined, else from application `.ini`, else from internal defaults.

Project info

```
cliquet.project_name = project
cliquet.project_docs = https://project.rtfld.org/
# cliquet.project_version = 1.0
```

Feature settings

```
# Limit number of batch operations per request
# cliquet.batch_max_requests = 25

# Force pagination *(recommended)*
# cliquet.paginate_by = 200

# Custom record id generator class
# cliquet.id_generator = cliquet.storage.generators.UUID4
```

Disabling endpoints

It is possible to deactivate specific resources operations, directly in the settings.

To do so, a setting key must be defined for the disabled resources endpoints:

```
'cliquet.{endpoint_type}_{resource_name}_{method}_enabled'
```

Where: - **endpoint_type** is either collection or record; - **resource_name** is the name of the resource (by default, *Cliquet* uses

the name of the class);

- **method** is the http method (in lower case): For instance `put`.

For instance, to disable the PUT on records for the *Mushrooms* resource, the following setting should be declared in the `.ini` file:

```
# Disable article collection DELETE endpoint
cliquet.collection_article_delete_enabled = false

# Disable mushroom record PATCH endpoint
cliquet.record_mushroom_patch_enabled = false
```

Setting the service in readonly mode

It is also possible to deploy a *Cliquet* service in readonly mode.

Instead of having settings to disable every resource endpoint, the `readonly` setting can be set:

```
cliquet.readonly = true
```

This will disable every resources endpoints that are not accessed with one of GET, OPTIONS, or HEAD methods. Requests will receive a 405 Method not allowed error response.

This setting will also activate readonly heartbeat checks for the permission and the storage backend.

Warning: The cache backend will still needs read-write privileges, in order to cache OAuth authentication states and tokens for example.
If you do not need cache at all, set the `kinto.cache_backend` setting to an empty string to disable it.

Deployment

```
# cliquet.backoff = 10
cliquet.retry_after_seconds = 30
```

Scheme, host and port

By default *Cliquet* does not enforce requests scheme, host and port. It relies on WSGI specification and the related stack configuration. Tuning this becomes necessary when the application runs behind proxies or load balancers.

Most implementations, like *uwsgi*, provide configuration variables to adjust it properly.

However if, for some reasons, this had to be enforced at the application level, the following settings can be set:

```
# cliquet.http_scheme = https
# cliquet.http_host = production.server:7777
```

Check the `url` value returned in the hello view.

Deprecation

Activate the *service deprecation*. If the date specified in `eos` is in the future, an alert will be sent to clients. If it's in the past, the service will be declared as decommissioned.

```
# cliquet.eos = 2015-01-22
# cliquet.eos_message = "Client is too old"
# cliquet.eos_url = http://website/info-shutdown.html
```

Logging with Heka

Mozilla Services standard logging format can be enabled using:

```
cliquet.logging_renderer = cliquet.logs.MozillaHekaRenderer
```

With the following configuration, all logs are redirected to standard output (See [12factor app](#)):

```
[loggers]
keys = root

[handlers]
keys = console

[formatters]
keys = heka

[logger_root]
level = INFO
handlers = console
formatter = heka

[handler_console]
class = StreamHandler
args = (sys.stdout,)
level = NOTSET

[formatter_heka]
format = %(message)s
```

Handling exceptions with Sentry

Requires the raven package, or *Cliquet* installed with `pip install cliquet[monitoring]`.

Sentry logging can be enabled, as explained in [official documentation](#).

Note: The application sends an *INFO* message on startup, mainly for setup check.

Monitoring with StatsD

Requires the statsd package, or *Cliquet* installed with `pip install cliquet[monitoring]`.

StatsD metrics can be enabled (disabled by default):

```
cliquet.statsd_url = udp://localhost:8125
# cliquet.statsd_prefix = cliquet.project_name
```

Monitoring with New Relic

Requires the newrelic package, or *Cliquet* installed with `pip install cliquet[monitoring]`.

Enable middlewares as described [here](#).

New-Relic can be enabled (disabled by default):

```
cliquet.newrelic_config = /location/of/newrelic.ini
cliquet.newrelic_env = prod
```

Storage

```
cliquet.storage_backend = cliquet.storage.redis
cliquet.storage_url = redis://localhost:6379/1

# Safety limit while fetching from storage
# cliquet.storage_max_fetch_size = 10000

# Control number of pooled connections
# cliquet.storage_pool_size = 50
```

See *storage backend documentation* for more details.

Notifications

To activate an event listener, you can use the *event_handlers* option, which takes a list of listeners.

Each listener can get its own options, under the *event_handlers.<name>* namespace.

In the example below, the Redis listener is activated and will send data in the **cliquet** Redis list.

```
cliquet.event_handlers = redis
cliquet.event_handlers.redis.url = redis://localhost:6379/0
cliquet.event_handlers.redis.listname = cliquet
cliquet.event_handlers.redis.pool_size = 5
```

Cache

```
cliquet.cache_backend = cliquet.cache.redis
cliquet.cache_url = redis://localhost:6379/0

# Control number of pooled connections
# cliquet.storage_pool_size = 50
```

See *cache backend documentation* for more details.

It is possible to add cache control headers on a particular resource for anonymous requests. The client (or proxy) will use them to cache the resource responses for a certain amount of time.

By default, *Cliquet* indicates the clients to invalidate their cache (Cache-Control: no-cache).

```
cliquet.mushroom_cache_expires_seconds = 3600
```

Basically, this will add both Cache-Control: max-age=3600 and Expire: <server datetime + 1H> response headers to the GET responses.

If setting is set to 0, then the resource follows the default behaviour.

By default, CORS headers are cached by clients during 1H (Access-Control-Max-Age).

The duration can be set from settings. If set to empty or to 0, the header is not sent to clients.

```
cliquet.cors_max_age_seconds = 7200
```

Authentication

Since user identification is hashed in storage, a secret key is required in configuration:

```
# cliquet.userid_hmac_secret = b4c96a8692291d88fe5a97dd91846eb4
```

Authentication setup

Cliquet relies on [:github:‘pyramid_multiauth <mozilla-service/pyramid_multiauth>’_](#) to initialize authentication.

Therefore, any authentication policy can be specified through configuration.

For example, using the following example, *Basic Auth*, *Persona* and *IP Auth* are enabled:

```
multiauth.policies = basicauth pyramid_persona ipauth

multiauth.policy.ipauth.use = pyramid_ipauth.IPAuthenticationPolicy
multiauth.policy.ipauth.ipaddrs = 192.168.0.*
multiauth.policy.ipauth.userid = LAN-user
multiauth.policy.ipauth.principals = trusted
```

Similarly, any authorization policies and group finder function can be specified through configuration in order to deeply customize permissions handling and authorizations.

Basic Auth

`basicauth` is mentioned among `multiauth.policies` by default.

```
multiauth.policies = basicauth
```

By default, it uses an internal *Basic Auth* policy bundled with *Cliquet*.

In order to replace it by another one:

```
multiauth.policies = basicauth
multiauth.policy.basicauth.use = myproject.authn.BasicAuthPolicy
```

Custom Authentication

Using the various [Pyramid authentication packages](#), it is possible to plug any kind of authentication.

(*Github/Twitter example to be done*)

Firefox Accounts

Enabling *Firefox Accounts* consists in including `cliquet_fxa` in configuration, mentioning `fxa` among policies and providing appropriate values for OAuth2 client settings.

See [mozilla-services/cliquet-fxa](#).

Permission configuration

ACE are usually set on objects using the permission backend.

It is also possible to configure them from settings, and it will **bypass** the permission backend.

For example, for a resource named “bucket”, the following setting will enable authenticated people to create bucket records:

```
cliquet.bucket_create_principals = system.Authenticated
```

The format of these permission settings is `<resource_name>_<permission>_principals = comma, separated, principals`.

Application profiling

It is possible to profile the application while its running. This is especially useful when trying to find slowness in the application.

Enable middlewares as described [here](#).

Update the configuration file with the following values:

```
cliquet.profiler_enabled = true
cliquet.profiler_dir = /tmp/profiling
```

Run a load test (*for example*):

```
SERVER_URL=http://localhost:8000 make bench -e
```

Render execution graphs using GraphViz:

```
sudo apt-get install graphviz
```

```
pip install gprof2dot
gprof2dot -f pstats POST.v1.batch.000176ms.1427458675.prof | dot -Tpng -o output.png
```

Enable middleware

In order to enable Cliquet middleware, wrap the application in the project main function:

```
def main(global_config, **settings):
    config = Configurator(settings=settings)
    cliquet.initialize(config, __version__)
    app = config.make_wsgi_app()
    return cliquet.install_middlewares(app, settings)
```

Initialization sequence

In order to control what part of *Cliquet* should be run during application startup, or add custom initialization steps from configuration, it is possible to change the `initialization_sequence` setting.

Warning: This is considered as a dangerous zone and should be used with caution.

Later, a better formalism should be introduced to easily allow addition or removal of steps, without repeating the whole list and without relying on internal functions location.

```
cliquet.initialization_sequence = cliquet.initialization.setup_json_serializer
                                cliquet.initialization.setup_logging
                                cliquet.initialization.setup_storage
                                cliquet.initialization.setup_cache
                                cliquet.initialization.setup_requests_scheme
                                cliquet.initialization.setup_version_redirection
                                cliquet.initialization.setup_deprecation
                                cliquet.initialization.setup_authentication
                                cliquet.initialization.setup_backoff
                                cliquet.initialization.setup_stats
```

1.4.3 Resource

Cliquet provides a basic component to build resource oriented APIs. In most cases, the main customization consists in defining the schema of the records for this resource.

Full example

```
import colander

from cliquet import resource
from cliquet import utils

class BookmarkSchema(resource.ResourceSchema):
    url = colander.SchemaNode(colander.String(), validator=colander.url)
    title = colander.SchemaNode(colander.String())
    favorite = colander.SchemaNode(colander.Boolean(), missing=False)
    device = colander.SchemaNode(colander.String(), missing='')

    class Options:
        readonly_fields = ('device',)
        unique_fields = ('url',)

@resource.register()
class Bookmark(resource.BaseResource):
    mapping = BookmarkSchema()

    def process_record(self, new, old=None):
        if new['device'] != old['device']:
            new['device'] = self.request.headers.get('User-Agent')

    return new
```

See the [ReadingList](#) and [Kinto](#) projects source code for real use cases.

Resource Schema

Override the base schema to add extra fields using the [Colander API](#).

```
class Movie(ResourceSchema):
    director = colander.SchemaNode(colander.String())
    year = colander.SchemaNode(colander.Int(),
```



```

                                validator=colander.Range(min=1850))
genre = colander.SchemaNode(colander.String(),
                                validator=colander.OneOf(['Sci-Fi', 'Comedy']))

```

Deprecated module. Moved to `cliquet.resource.schema`.

Resource class

In order to customize the resource URLs or behaviour on record processing, the resource class can be extended:

class `cliquet.resource.BaseResource` (*request*, *context=None*)

Base resource class providing every endpoint.

default_viewset

Default `cliquet.viewset.ViewSet` class to use when the resource is registered.

alias of `ViewSet`

default_model

Default `cliquet.resource.model.Model` class to use for interacting the **:module:‘cliquet.storage’** and **:module:‘cliquet.permission’** backends.

alias of `Model`

mapping = <`cliquet.resource.schema.ResourceSchema` object at 140146195526864 (named)>

Schema to validate records.

collection

The collection property.

get_parent_id (*request*)

Return the `parent_id` of the resource with regards to the current request.

Parameters **request** – The request used to create the resource.

Return type `str`

is_known_field (*field*)

Return `True` if *field* is defined in the resource mapping.

Parameters **field** (*str*) – Field name

Return type `bool`

collection_get ()

Model GET endpoint: retrieve multiple records.

Raises `HTTPNotModified` if `If-None-Match` header is provided and collection not modified in the interim.

Raises `HTTPPreconditionFailed` if `If-Match` header is provided and collection modified in the interim.

Raises `HTTPBadRequest` if filters or sorting are invalid.

collection_post ()

Model POST endpoint: create a record.

If the new record conflicts against a unique field constraint, the posted record is ignored, and the existing record is returned, with a 200 status.

Raises `HTTPPreconditionFailed` if `If-Match` header is provided and collection modified in the interim.

See also:

Add custom behaviour by overriding `cliquet.resource.BaseResource.process_record()`

collection_delete()

Model DELETE endpoint: delete multiple records.

Raises `HTTPPreconditionFailed` if If-Match header is provided and collection modified in the interim.

Raises `HTTPBadRequest` if filters are invalid.

get()

Record GET endpoint: retrieve a record.

Raises `HTTPNotFound` if the record is not found.

Raises `HTTPNotModified` if If-None-Match header is provided and record not modified in the interim.

Raises `HTTPPreconditionFailed` if If-Match header is provided and record modified in the interim.

put()

Record PUT endpoint: create or replace the provided record and return it.

Raises `HTTPPreconditionFailed` if If-Match header is provided and record modified in the interim.

Note: If If-None-Match: * request header is provided, the PUT will succeed only if no record exists with this id.

See also:

Add custom behaviour by overriding `cliquet.resource.BaseResource.process_record()`.

patch()

Record PATCH endpoint: modify a record and return its new version.

If a request header `Response-Behavior` is set to `light`, only the fields whose value was changed are returned. If set to `diff`, only the fields whose value became different than the one provided are returned.

Raises `HTTPNotFound` if the record is not found.

Raises `HTTPPreconditionFailed` if If-Match header is provided and record modified in the interim.

See also:

Add custom behaviour by overriding `cliquet.resource.BaseResource.apply_changes()` or `cliquet.resource.BaseResource.process_record()`.

delete()

Record DELETE endpoint: delete a record and return it.

Raises `HTTPNotFound` if the record is not found.

Raises `HTTPPreconditionFailed` if If-Match header is provided and record modified in the interim.

process_record(new, old=None)

Hook for processing records before they reach storage, to introduce specific logics on fields for example.

```
def process_record(self, new, old=None):
    version = old['version'] if old else 0
    new['version'] = version + 1
    return new
```

Or add extra validation based on request:

```
from cliquet.errors import raise_invalid

def process_record(self, new, old=None):
    if new['browser'] not in request.headers['User-Agent']:
        raise_invalid(self.request, name='browser', error='Wrong')
    return new
```

Parameters

- **new** (*dict*) – the validated record to be created or updated.
- **old** (*dict*) – the old record to be updated, None for creation endpoints.

Returns the processed record.

Return type dict

apply_changes (*record, changes*)
Merge *changes* into *record* fields.

Note: This is used in the context of PATCH only.

Override this to control field changes at record level, for example:

```
def apply_changes(self, record, changes):
    # Ignore value change if inferior
    if record['position'] > changes.get('position', -1):
        changes.pop('position', None)
    return super(MyResource, self).apply_changes(record, changes)
```

Raises `HTTPBadRequest` if result does not comply with resource schema.

Returns the new record with *changes* applied.

Return type dict

Interaction with storage

In order to customize the interaction of a HTTP resource with its storage, a custom collection can be plugged-in:

```
from cliquet import resource

class TrackedModel(resource.Model):
    def create_record(self, record, parent_id=None, unique_fields=None):
        record = super(TrackedModel, self).create_record(record,
                                                            parent_id,
                                                            unique_fields)

        trackid = index.track(record)
        record['trackid'] = trackid
        return record
```

```
class Payment(resource.BaseResource):
    default_model = TrackedModel
```

```
class cliquet.resource.Model(storage, id_generator=None, collection_id='', parent_id='',
                             auth=None)
```

A collection stores and manipulate records in its attached storage.

It is not aware of HTTP environment nor protocol.

Records are isolated according to the provided *name* and *parent_id*.

Those notions have no particular semantic and can represent anything. For example, the collection *name* can be the *type* of objects stored, and *parent_id* can be the current *user id* or a *group* where the collection belongs. If left empty, the collection records are not isolated.

id_field = 'id'

Name of *id* field in records

modified_field = 'last_modified'

Name of *last modified* field in records

deleted_field = 'deleted'

Name of *deleted* field in deleted records

timestamp (*parent_id=None*)

Fetch the collection current timestamp.

Parameters **parent_id** (*str*) – optional filter for parent id

Return type integer

get_records (*filters=None, sorting=None, pagination_rules=None, limit=None, include_deleted=False, parent_id=None*)

Fetch the collection records.

Override to post-process records after fetching them from storage.

Parameters

- **filters** (list of `cliquet.storage.Filter`) – Optionally filter the records by their attribute. Each filter in this list is a tuple of a field, a value and a comparison (see `cliquet.utils.COMPARISON`). All filters are combined using *AND*.
- **sorting** (list of `cliquet.storage.Sort`) – Optionnally sort the records by attribute. Each sort instruction in this list refers to a field and a direction (negative means descending). All sort instructions are cumulative.
- **pagination_rules** (list of list of `cliquet.storage.Filter`) – Optionnally paginate the list of records. This list of rules aims to reduce the set of records to the current page. A rule is a list of filters (see *filters* parameter), and all rules are combined using *OR*.
- **limit** (*int*) – Optionnally limit the number of records to be retrieved.
- **include_deleted** (*bool*) – Optionnally include the deleted records that match the filters.
- **parent_id** (*str*) – optional filter for parent id

Returns A tuple with the list of records in the current page, the total number of records in the result set.

Return type tuple

delete_records (*filters=None, parent_id=None*)

Delete multiple collection records.

Override to post-process records after their deletion from storage.

Parameters

- **filters** (list of `cliquet.storage.Filter`) – Optionally filter the records by their attribute. Each filter in this list is a tuple of a field, a value and a comparison (see `cliquet.utils.COMPARISON`). All filters are combined using *AND*.
- **parent_id** (*str*) – optional filter for parent id

Returns The list of deleted records from storage.

get_record (*record_id, parent_id=None*)

Fetch current view related record, and raise 404 if missing.

Parameters

- **record_id** (*str*) – record identifier
- **parent_id** (*str*) – optional filter for parent id

Returns the record from storage

Return type dict

create_record (*record, parent_id=None, unique_fields=None*)

Create a record in the collection.

Override to perform actions or post-process records after their creation in storage.

```
def create_record(self, record):
    record = super(MyModel, self).create_record(record)
    idx = index.store(record)
    record['index'] = idx
    return record
```

Parameters

- **record** (*dict*) – record to store
- **parent_id** (*str*) – optional filter for parent id
- **unique_fields** (*tuple*) – list of fields that should remain unique

Returns the newly created record.

Return type dict

update_record (*record, parent_id=None, unique_fields=None*)

Update a record in the collection.

Override to perform actions or post-process records after their modification in storage.

```
def update_record(self, record, parent_id=None, unique_fields=None):
    record = super(MyModel, self).update_record(record,
                                                parent_id,
                                                unique_fields)
    subject = 'Record {} was changed'.format(record[self.id_field])
    send_email(subject)
    return record
```

Parameters

- **record** (*dict*) – record to store
- **parent_id** (*str*) – optional filter for parent id
- **unique_fields** (*tuple*) – list of fields that should remain unique

Returns the updated record.

Return type dict

delete_record (*record*, *parent_id=None*)

Delete a record in the collection.

Override to perform actions or post-process records after deletion from storage for example:

```
def delete_record(self, record):
    deleted = super(MyModel, self).delete_record(record)
    erase_media(record)
    deleted['media'] = 0
    return deleted
```

Parameters

- **record** (*dict*) – the record to delete
- **record** – record to store
- **parent_id** (*str*) – optional filter for parent id

Returns the deleted record.

Return type dict

Custom record ids

By default, records ids are *UUID4* <http://en.wikipedia.org/wiki/Universally_unique_identifier>.

A custom record ID generator can be set globally in *Configuration*, or at the resource level:

```
from cliquet import resource
from cliquet import utils
from cliquet.storage import generators

class MsecId(generators.Generator):
    def __call__(self):
        return '%s' % utils.msec_time()

@resource.register()
class Mushroom(resource.BaseResource):
    def __init__(request):
        super(Mushroom, self).__init__(request)
        self.model.id_generator = MsecId()
```

Generators objects

class cliquet.storage.generators.**Generator** (*config=None*)

Base generator for records ids.

Id generators are used by storage backend during record creation, and at resource level to validate record id in requests paths.

regexp = `^[a-zA-Z0-9][a-zA-Z0-9_-]*$`

Default record id pattern. Can be changed to comply with custom ids.

match (*record_id*)

Validate that record ids match the generator. This is used mainly when a record id is picked arbitrarily (e.g with PUT requests).

Returns *True* if the specified record id matches expected format.

Return type bool

class `cliquet.storage.generators.UUID4` (*config=None*)

UUID4 record id generator.

UUID block are separated with -. (example: `'472be9ec-26fe-461b-8282-9c4e4b207ab3'`)

UUIDs are very safe in term of unicity. If 1 billion of UUIDs are generated every second for the next 100 years, the probability of creating just one duplicate would be about 50% ([source](#)).

regexp = `^[0-9a-f]{8}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{12}$`

UUID4 accurate pattern.

Custom Usage

Within views

In views, a `request` object is available and allows to use the storage configured in the application:

```
from cliquet import resource

def view(request):
    registry = request.registry

    flowers = resource.Model(storage=registry.storage,
                             collection_id='app:flowers')

    flowers.create_record({'name': 'Jonquille', 'size': 30})
    flowers.create_record({'name': 'Amapola', 'size': 18})

    min_size = resource.Filter('size', 20, resource.COMPARISON.MIN)
    records, total = flowers.get_records(filters=[min_size])

    flowers.delete_record(records[0])
```

Outside views

Outside views, an application context has to be built from scratch.

As an example, let's build a code that will copy a collection into another:

```
from cliquet import resource, DEFAULT_SETTINGS
from pyramid import Configurator

config = Configurator(settings=DEFAULT_SETTINGS)
```

```
config.add_settings({
    'cliquet.storage_backend': 'cliquet.storage.postgresql'
    'cliquet.storage_url': 'postgres://user:pass@db.server.lan:5432/dbname'
})
cliquet.initialize(config, '0.0.1')

local = resource.Model(storage=config.registry.storage,
                       parent_id='browsing',
                       collection_id='history')

remote = resource.Model(storage=config_remote.registry.storage,
                        parent_id='',
                        collection_id='history')

records, total = in remote.get_records():
for record in records:
    local.create_record(record)
```

1.4.4 Viewsets

Cliquet maps URLs and permissions to resources using *ViewSets*.

View sets can be viewed as a set of rules which can be applied to a resource in order to define what should be inserted in the routing mechanism of pyramid.

Configuring a viewset

To use *Cliquet* in a basic fashion, there is not need to understand how viewsets work in full detail, but it might be useful to know how to extend the defaults.

Default viewset can be extended by passing viewset arguments to the *resource.register* class decorator:

```
from cliquet import resource

@resource.register(collection_methods=('GET',))
class Resource(resource.BaseResource):
    mapping = BookmarkSchema()
```

Subclassing a viewset

In case this isn't enough to update the default properties, the default *ViewSet* class can be subclassed in a more specific viewset, and then be passed during the registration phase:

```
from cliquet import resource

class MyViewSet(resource.ViewSet):

    def get_service_name(self, endpoint_type, resource):
        """Returns the name of the service, depending a given type and
        resource.
        """
        # Get the resource name from an akwards location.
        return name
```



```
@resource.register(viewset=MyViewSet())
class Resource(resource.BaseResource):
    mapping = BookmarkSchema()
```

ViewSet class

In order to customize the resource URLs or permissions, the viewset class can be extended:

class cliquet.resource.**ViewSet** (***kwargs*)

The default ViewSet object.

A viewset contains all the information needed to register any resource in the Cornice registry.

It provides the same features as `cornice.resource()`, except that it is much more flexible and extensible.

update (***kwargs*)

Update viewset attributes with provided values.

get_view_arguments (*endpoint_type, resource_cls, method*)

Return the Pyramid/Cornice view arguments for the given endpoint type and method.

Parameters

- **endpoint_type** (*str*) – either “collection” or “record”.
- **resource_cls** – the resource class.
- **method** (*str*) – the HTTP method.

get_record_schema (*resource_cls, method*)

Return the Cornice schema for the given method.

get_view (*endpoint_type, method*)

Return the view method name located on the resource object, for the given type and method.

- For collections, this will be “collection_{methodlower}”
- For records, this will be “{methodlower}”.

get_name (*resource_cls*)

Returns the name of the resource.

get_service_name (*endpoint_type, resource_cls*)

Returns the name of the service, depending a given type and resource.

is_endpoint_enabled (*endpoint_type, resource_name, method, settings*)

Returns if the given endpoint is enabled or not.

Uses the settings to tell so.

1.4.5 Storage

Backends

PostgreSQL

Redis

class cliquet.storage.redis.**Redis** (*args, **kwargs)
Storage backend implementation using Redis.

Warning: Useful for very low server load, but won't scale since records sorting and filtering are performed in memory.

Enable in configuration:

```
cliquet.storage_backend = cliquet.storage.redis
```

(Optional) Instance location URI can be customized:

```
cliquet.storage_url = redis://localhost:6379/0
```

A threaded connection pool is enabled by default:

```
cliquet.storage_pool_size = 50
```

Memory

class cliquet.storage.memory.**Memory** (*args, **kwargs)
Storage backend implementation in memory.

Useful for development or testing purposes, but records are lost after each server restart.

Enable in configuration:

```
cliquet.storage_backend = cliquet.storage.memory
```

API

Implementing a custom storage backend consists in implementing the following interface:

class cliquet.storage.**Filter** (field, value, operator)
Filtering properties.

field
Alias for field number 0

operator
Alias for field number 2

value
Alias for field number 1

class cliquet.storage.**Sort** (field, direction)
Sorting properties.

direction

Alias for field number 1

field

Alias for field number 0

class `cliquet.storage.StorageBase`

Storage abstraction used by resource views.

It is meant to be instantiated at application startup. Any operation may raise a `HTTPServiceUnavailable` error if an error occurs with the underlying service.

Configuration can be changed to choose which storage backend will persist the objects.

Raises `HTTPServiceUnavailable`

initialize_schema()

Create every necessary objects (like tables or indices) in the backend.

This is excuted when the `cliquet migrate` command is ran.

flush (*auth=None*)

Remove **every** object from this storage.

collection_timestamp (*collection_id, parent_id, auth=None*)

Get the highest timestamp of every objects in this *collection_id* for this *parent_id*.

Note: This should take deleted objects into account.

Parameters

- **collection_id** (*str*) – the collection id.
- **parent_id** (*str*) – the collection parent.

Returns the latest timestamp of the collection.

Return type `int`

create (*collection_id, parent_id, object, id_generator=None, unique_fields=None, id_field='id', modified_field='last_modified', auth=None*)

Create the specified *object* in this *collection_id* for this *parent_id*. Assign the id to the object, using the attribute `cliquet.resource.BaseResource.id_field`.

Note: This will update the collection timestamp.

Raises `cliquet.storage.exceptions.UnicityError`

Parameters

- **collection_id** (*str*) – the collection id.
- **parent_id** (*str*) – the collection parent.
- **object** (*dict*) – the object to create.

Returns the newly created object.

Return type `dict`

get (*collection_id, parent_id, object_id, id_field='id', modified_field='last_modified', auth=None*)

Retrieve the object with specified *object_id*, or raise error if not found.

Raises `cliquet.storage.exceptions.RecordNotFoundError`

Parameters

- **collection_id** (*str*) – the collection id.
- **parent_id** (*str*) – the collection parent.
- **object_id** (*str*) – unique identifier of the object

Returns the object object.

Return type dict

update (*collection_id*, *parent_id*, *object_id*, *object*, *unique_fields=None*, *id_field='id'*, *modified_field='last_modified'*, *auth=None*)
 Overwrite the *object* with the specified *object_id*.

If the specified id is not found, the object is created with the specified id.

Note: This will update the collection timestamp.

Raises `cliquet.storage.exceptions.UnicityError`

Parameters

- **collection_id** (*str*) – the collection id.
- **parent_id** (*str*) – the collection parent.
- **object_id** (*str*) – unique identifier of the object
- **object** (*dict*) – the object to update or create.

Returns the updated object.

Return type dict

delete (*collection_id*, *parent_id*, *object_id*, *with_deleted=True*, *id_field='id'*, *modified_field='last_modified'*, *deleted_field='deleted'*, *auth=None*)
 Delete the object with specified *object_id*, and raise error if not found.

Deleted objects must be removed from the database, but their ids and timestamps of deletion must be tracked for synchronization purposes. (See `cliquet.storage.StorageBase.get_all()`)

Note: This will update the collection timestamp.

Raises `cliquet.storage.exceptions.RecordNotFoundError`

Parameters

- **collection_id** (*str*) – the collection id.
- **parent_id** (*str*) – the collection parent.
- **object_id** (*str*) – unique identifier of the object
- **with_deleted** (*bool*) – track deleted record with a tombstone

Returns the deleted object, with minimal set of attributes.

Return type dict

delete_all (*collection_id*, *parent_id*, *filters=None*, *with_deleted=True*, *id_field='id'*, *modified_field='last_modified'*, *deleted_field='deleted'*, *auth=None*)

Delete all objects in this *collection_id* for this *parent_id*.

Parameters

- **collection_id** (*str*) – the collection id.
- **parent_id** (*str*) – the collection parent.
- **filters** (list of *cliquet.storage.Filter*) – Optionnally filter the objects to delete.
- **with_deleted** (*bool*) – track deleted records with a tombstone

Returns the list of deleted objects, with minimal set of attributes.

Return type list of dict

purge_deleted (*collection_id*, *parent_id*, *before=None*, *id_field='id'*, *modified_field='last_modified'*, *auth=None*)

Delete all deleted object tombstones in this *collection_id* for this *parent_id*.

Parameters

- **collection_id** (*str*) – the collection id.
- **parent_id** (*str*) – the collection parent.
- **before** (*int*) – Optionnal timestamp to limit deletion (exclusive)

Returns The number of deleted objects.

Return type int

get_all (*collection_id*, *parent_id*, *filters=None*, *sorting=None*, *pagination_rules=None*, *limit=None*, *include_deleted=False*, *id_field='id'*, *modified_field='last_modified'*, *deleted_field='deleted'*, *auth=None*)

Retrieve all objects in this *collection_id* for this *parent_id*.

Parameters

- **collection_id** (*str*) – the collection id.
- **parent_id** (*str*) – the collection parent.
- **filters** (list of *cliquet.storage.Filter*) – Optionally filter the objects by their attribute. Each filter in this list is a tuple of a field, a value and a comparison (see *cliquet.utils.COMPARISON*). All filters are combined using *AND*.
- **sorting** (list of *cliquet.storage.Sort*) – Optionnally sort the objects by attribute. Each sort instruction in this list refers to a field and a direction (negative means descending). All sort instructions are cumulative.
- **pagination_rules** (list of list of *cliquet.storage.Filter*) – Optionnally paginate the list of objects. This list of rules aims to reduce the set of objects to the current page. A rule is a list of filters (see *filters* parameter), and all rules are combined using *OR*.
- **limit** (*int*) – Optionnally limit the number of objects to be retrieved.
- **include_deleted** (*bool*) – Optionnally include the deleted objects that match the filters.

Returns the limited list of objects, and the total number of matching objects in the collection (deleted ones excluded).

Return type tuple (list, integer)

Exceptions

Exceptions raised by storage backend.

exception `cliquet.storage.exceptions.BackendError` (*original=None, message=None, *args, **kwargs*)

A generic exception raised by storage on error.

Parameters `original` (*Exception*) – the wrapped exception raised by underlying library.

exception `cliquet.storage.exceptions.RecordNotFoundError`

An exception raised when a specific record could not be found.

exception `cliquet.storage.exceptions.UnicityError` (*field, record, *args, **kwargs*)

An exception raised on unicity constraint violation.

Raised by storage backend when the creation or the modification of a record violates the unicity constraints defined by the resource.

Store custom data

Storage can be used to store arbitrary data.

```
data = {'subscribed': datetime.now()}
user_id = request.authenticated_userid

storage = request.registry.storage
storage.create(collection_id='__custom', parent_id='', record=data)
```

See the collection class to manipulate collections of records.

1.4.6 Cache

PostgreSQL

Redis

Memory

API

Implementing a custom cache backend consists on implementing the following interface:

class `cliquet.cache.CacheBase` (**args, **kwargs*)

initialize_schema()

Create every necessary objects (like tables or indices) in the backend.

This is excuted when the `cliquet migrate` command is ran.

flush()

Delete every values.

t1 (*key*)Obtain the expiration value of the specified *key*.**Parameters** **key** (*str*) – key**Returns** number of seconds or negative if no TTL.**Return type** float**expire** (*key*, *t1*)Set the expiration value *t1* for the specified *key*.**Parameters**

- **key** (*str*) – key
- **t1** (*float*) – number of seconds

set (*key*, *value*, *t1=None*)Store a value with the specified *key*. If *t1* is provided, set an expiration value.**Parameters**

- **key** (*str*) – key
- **value** (*str*) – value to store
- **t1** (*float*) – expire after number of seconds

get (*key*)Obtain the value of the specified *key*.**Parameters** **key** (*str*) – key**Returns** the stored value or None if missing.**Return type** str**delete** (*key*)Delete the value of the specified *key*.**Parameters** **key** (*str*) – key

1.4.7 Permissions

Cliquet provides a mechanism to handle authorization on the stored *objects*.

Glossary

Authorization isn't complicated, but requires the introduction of a few terms so that explanations are easier to follow:

Object: The data that is stored into *Cliquet*. *objects* usually match the resources you defined; For one resource there are two *objects*: resource's collection and resource's records.

Principal: An entity that can be authenticated. *principals* can be individual people, computers, services, or any group of such things.

Permission: An action that can be authorized or denied. *read*, *write*, *create* are *permissions*.

Access Control Entity (ACE): An association of a *principal*, an *object* and a *permission*. For instance, (Alexis, article, write).

Access Control List (ACL): A list of Access Control Entities (*ACE*).

Overview

By default, the resources defined by *Cliquet* are public, and records are isolated by user. But it is also possible to define *protected resources*, which will required the user to have access to the requested resource.

```
from cliquet import authorization
from cliquet import resource

@resource.register(factory=authorization.RouteFactory)
class Toadstool(resource.ProtectedResource):
    mapping = MushroomSchema()
```

In this example, a *route factory* is registered. Route factories are explained in more details below.

A protected resource, in addition to the *data* property of request / responses, takes a *permissions* property which contains the list of *principals* that are allowed to access or modify the current *object*.

During the creation of the *object*, the *permissions* property is stored in the permission backend, and upon access, it checks the current *principal* has access the the object, with the correct permission.

Route factory

The route factory decides which *permission* is required to access one resource or another. Here is a summary of the *permissions* that are defined by the default route factory *Cliquet* defines:

Method	<i>permission</i>
POST	create
GET / HEAD	read
PUT	create if it doesn't exist, write otherwise
PATCH	write
DELETE	write

Route factories are [best described in the pyramid documentation](#)

```
class cliquet.authorization.RouteFactory(request)
```

Authorization policy

Upon access, the authorization policy is asked if any of the current list of *principals* has access to the current resource. By default, the authorization policy *Cliquet* checks in the *permission* backend for the current *object*.

It is possible to extend this behavior, for instance if there is an inheritance tree between the defined resources (some *ACEs* should give access to its child *objects*).

In case the application should define its own inheritance tree, it should also define its own authorization policy.

To do so, subclass the default *AuthorizationPolicy* and add a specific *get_bound_permission* method.

```
from cliquet import authorization
from pyramid.security import IAuthorizationPolicy
from zope.interface import implementer

@implementer(IAuthorizationPolicy)
class AuthorizationPolicy(authorization.AuthorizationPolicy):
    def get_bound_permissions(self, *args, **kwargs):
        """Callable that takes an object ID and a permission and returns
```



```
a list of tuples (<object id>, <permission>)."""
    return build_permissions_set(*args, **kwargs)
```

```
class cliquet.authorization.AuthorizationPolicy
```

Permissions backend

The *ACLs* are stored in a *permission* backend. Currently, permission backends exists for Redis and PostgreSQL, as well as a in memory one. It is of course possible to add you own permission backend, if you wish to store your *permissions* related data in a different database.

```
class cliquet.permission.PermissionBase(*args, **kwargs)
```

```
    initialize_schema()
```

Create every necessary objects (like tables or indices) in the backend.

This is excuted with the `cliquet migrate` command.

```
    flush()
```

Delete all data stored in the permission backend.

```
    add_user_principal(user_id, principal)
```

Add an additional principal to a user.

Parameters

- **user_id** (*str*) – The user_id to add the principal to.
- **principal** (*str*) – The principal to add.

```
    remove_user_principal(user_id, principal)
```

Remove an additional principal from a user.

Parameters

- **user_id** (*str*) – The user_id to remove the principal to.
- **principal** (*str*) – The principal to remove.

```
    user_principals(user_id)
```

Return the set of additionnal principals given to a user.

Parameters **user_id** (*str*) – The user_id to get the list of groups for.

Returns The list of group principals the user is in.

Return type *set*

```
    add_principal_to_ace(object_id, permission, principal)
```

Add a principal to an Access Control Entry.

Parameters

- **object_id** (*str*) – The object to add the permission principal to.
- **permission** (*str*) – The permission to add the principal to.
- **principal** (*str*) – The principal to add to the ACE.

```
    remove_principal_from_ace(object_id, permission, principal)
```

Remove a principal to an Access Control Entry.

Parameters

- **object_id** (*str*) – The object to remove the permission principal to.
- **permission** (*str*) – The permission that should be removed.
- **principal** (*str*) – The principal to remove to the ACE.

object_permission_principals (*object_id, permission*)

Return the set of principals of a bound permission (unbound permission + object id).

Parameters

- **object_id** (*str*) – The object_id the permission is set to.
- **permission** (*str*) – The permission to query.

Returns The list of user principals

Return type *set*

principals_accessible_objects (*principals, permission, object_id_match=None, get_bound_permissions=None*)

Return the list of objects id where the specified *principals* have the specified *permission*.

Parameters

- **principal** (*list*) – List of user principals
- **permission** (*str*) – The permission to query.
- **object_id_match** (*str*) – Filter object ids based on a pattern (e.g. '**articles**').
- **get_bound_permissions** (*function*) – The methods to call in order to generate the list of permission to verify against. (ie: if you can write, you can read)

Returns The list of object ids

Return type *set*

object_permission_authorized_principals (*object_id, permission, get_bound_permissions=None*)

Return the full set of authorized principals for a given permission + object (bound permission).

Parameters

- **object_id** (*str*) – The object_id the permission is set to.
- **permission** (*str*) – The permission to query.
- **get_bound_permissions** (*function*) – The methods to call in order to generate the list of permission to verify against. (ie: if you can write, you can read)

Returns The list of user principals

Return type *set*

check_permission (*object_id, permission, principals, get_bound_permissions=None*)

Test if a principal set have got a permission on an object.

Parameters

- **object_id** (*str*) – The identifier of the object concerned by the permission.
- **permission** (*str*) – The permission to test.
- **principals** (*set*) – A set of user principals to test the permission against.
- **get_bound_permissions** (*function*) – The method to call in order to generate the set of permission to verify against. (ie: if you can write, you can read)

object_permissions (*object_id*, *permissions=None*)

Return the set of principals for each object permission.

Parameters

- **object_id** (*str*) – The object_id the permission is set to.
- **permissions** (*list*) – List of permissions to retrieve. If not define will try to find them all.

Returns The dictionary with the list of user principals for each object permissions

Return type dict

replace_object_permissions (*object_id*, *permissions*)

Replace given object permissions.

Parameters

- **object_id** (*str*) – The object to replace permissions to.
- **permissions** (*str*) – The permissions dict to replace.

delete_object_permissions (**object_id_list*)

Delete all listed object permissions.

Parameters **object_id** (*str*) – Remove given objects permissions.

1.4.8 Errors

`cliquet.errors.ERRORS`

Predefined errors as specified by the protocol.

status code	errno	description
401	104	Missing Authorization Token
401	105	Invalid Authorization Token
400	106	request body was not valid JSON
400	107	invalid request parameter
400	108	missing request parameter
400	109	invalid posted data
404	110	Invalid Token / id
404	111	Missing Token / id
411	112	Content-Length header was not provided
413	113	Request body too large
412	114	Resource was modified meanwhile
405	115	Method not allowed on this end point
404	116	Requested version not available on this server
429	117	Client has sent too many requests
403	121	Resource's access forbidden for this user
409	122	Another resource violates constraint
500	999	Internal Server Error
503	201	Service Temporary unavailable due to high load
410	202	Service deprecated

alias of Enum

`cliquet.errors.http_error` (*httpexception*, *errno=None*, *code=None*, *error=None*, *message=None*, *info=None*, *details=None*)

Return a JSON formatted response matching the error protocol.

Parameters

- **httpexception** – Instance of `httpexceptions`
- **errno** – stable application-level error number (e.g. 109)
- **code** – matches the HTTP status code (e.g 400)
- **error** – string description of error type (e.g. “Bad request”)
- **message** – context information (e.g. “Invalid request parameters”)
- **info** – information about error (e.g. URL to troubleshooting)
- **details** – additional structured details (conflicting record)

Returns the formatted response object

Return type `pyramid.httpexceptions.HTTPException`

`cliquet.errors.json_error_handler(errors)`

Cornice JSON error handler, returning constant JSON formatted errors from schema validation errors.

This is meant to be used in custom services in your applications.

```
upload = Service(name="upload", path='/upload',
                 error_handler=errors.json_error_handler)
```

Warning: Only the first error of the list is formatted in the response. (c.f. protocol).

`cliquet.errors.raise_invalid(request, location='body', name=None, description=None, **kwargs)`

Helper to raise a validation error.

Parameters

- **location** – location in request (e.g. 'querystring')
- **name** – field name
- **description** – detailed description of validation error

Raises `HTTPBadRequest`

`cliquet.errors.send_alert(request, message=None, url=None, code='soft-eol')`

Helper to add an Alert header to the response.

Parameters

- **code** – The type of error 'soft-eol', 'hard-eol'
- **message** – The description message.
- **url** – The URL for more information, default to the documentation url.

1.4.9 Utils

`cliquet.utils.strip_whitespace(v)`

Remove whitespace, newlines, and tabs from the beginning/end of a string.

Parameters `v (str)` – the string to strip.

Return type `str`

`cliquet.utils.msec_time()`

Return current epoch time in milliseconds.

Return type int

`cliquet.utils.classname(obj)`

Get a classname from an object.

Return type str

`cliquet.utils.merge_dicts(a, b)`

Merge *b* into *a* recursively, without overwriting values.

Parameters *a* (*dict*) – the dict that will be altered with values of *b*.

Return type None

`cliquet.utils.random_bytes_hex(bytes_length)`

Return a hexstring of *bytes_length* cryptographic-friendly random bytes.

Parameters *bytes_length* (*integer*) – number of random bytes.

Return type str

`cliquet.utils.native_value(value)`

Convert string value to native python values.

Parameters *value* (*str*) – value to interpret.

Returns the value coerced to python type

`cliquet.utils.read_env(key, value)`

Read the setting key from environment variables.

Parameters

- **key** – the setting name
- **value** – default value if undefined in environment

Returns the value from environment, coerced to python type

`cliquet.utils.encode64(content, encoding='utf-8')`

Encode some content in base64.

Return type str

`cliquet.utils.decode64(encoded_content, encoding='utf-8')`

Decode some base64 encoded content.

Return type str

`cliquet.utils.hmac_digest(secret, message, encoding='utf-8')`

Return hex digest of a message HMAC using secret

`cliquet.utils.reapply_cors(request, response)`

Reapply cors headers to the new response with regards to the request.

We need to re-apply the CORS checks done by Cornice, in case we're recreating the response from scratch.

`cliquet.utils.current_service(request)`

Return the Cornice service matching the specified request.

Returns the service or None if unmatched.

Return type cornice.Service

`cliquet.utils.build_request (original, dict_obj)`

Transform a dict object into a `pyramid.request.Request` object.

Parameters

- **original** – the original request.
- **dict_obj** – a dict object with the sub-request specifications.

`cliquet.utils.build_response (response, request)`

Transform a `pyramid.response.Response` object into a serializable dict.

Parameters

- **response** – a response object, returned by Pyramid.
- **request** – the request that was used to get the response.

`cliquet.utils.follow_subrequest (request, subrequest)`

Run a subrequest (e.g. batch), and follow the redirection if any.

Return type tuple

Returns the response and the redirection request (or *subrequest* if no redirection happened.)

`cliquet.utils.encode_header (value, encoding='utf-8')`

Make sure the value is of type `str` in both PY2 and PY3.

`cliquet.utils.decode_header (value, encoding='utf-8')`

Make sure the header is an unicode string.

`cliquet.utils.strip_uri_prefix (path)`

Remove potential version prefix in URI.

1.4.10 Glossary

CRUD Acronym for Create, Read, Update, Delete

endpoint An endpoint handles a particular HTTP verb at a particular URL.

extensible «Extensible» means that the component behaviour can be overridden via lines of code. It differs from «*pluggable*».

Firefox Accounts Account account system run by Mozilla (<https://accounts.firefox.com>).

KISS «Keep it simple, stupid» is a design principle which states that most systems work best if they are kept simple rather than made complicated.

pluggable «Pluggable» means that the component can be replaced via configuration. It differs from «*extensible*».

resource A resource is a collection of records.

user id, user identifier, user identifiers A string that identifies a user. By default, *Cliquet* uses a HMAC on authentication credentials to generate users identifications strings.

See [Pyramid authentication](#).

object, objects The data that is stored into *Cliquet*. Objects usually match the resources you defined; For one resource there are two objects: resource's collection and resource's records.

tombstone, tombstones When a record is deleted in a resource, a tombstone is created to keep track of the deletion when polling for changes. A tombstone only contains the `id` and `last_modified` fields, everything else is really deleted.

principal, principals An entity that can be authenticated. Principals can be individual people, computers, services, or any group of such things.

permission, permissions An action that can be authorized or denied. read, write, create are permissions.

ACE, ACEs, Access Control Entity An association of a principal, an object and a permission. For instance, (Alexis, article, write).

ACL, ACLs, Access Control List A list of Access Control Entities (ACE).

1.5 Ecosystem

This section gathers information about extending *Cliquet*, and third-party packages.

1.5.1 Packages

- [mozilla-services/cliquet-fxa](#): Add support of *Firefox Accounts* OAuth2 authentication in *Cliquet*

Note: If you build a package that you would like to see listed here, just get in touch with us!

1.5.2 Extending Cliquet

Pluggable components

Pluggable components can be substituted from configuration files, as long as the replacement follows the original component API.

```
# myproject.ini
cliquet.logging_renderer = cliquet_fluent.FluentRenderer
```

This is the simplest way to extend *Cliquet*, but will be limited to its existing components (cache, storage, log renderer, ...).

In order to add extra features, including external packages is the way to go!

Include external packages

Appart from usual python «import and use», *Pyramid* can include external packages, which can bring views, event listeners etc.

```
import cliquet
from pyramid.config import Configurator

def main(global_config, **settings):
    config = Configurator(settings=settings)

    cliquet.initialize(config, '0.0.1')
    config.scan("myproject.views")

    config.include('cliquet_elasticsearch')

    return config.make_wsgi_app()
```

Alternatively, packages can also be included via configuration:

```
# myproject.ini
cliquet.includes = cliquet_elasticsearch
                  pyramid_debugtoolbar
```

There are ‘**many available packages**’, and it is straightforward to build one.

Include me

In order to be included, a package must define an `includeme(config)` function.

For example, in `cliquet_elasticsearch/init.py`:

```
def includeme(config):
    settings = config.get_settings()

    config.add_view(...)
```

Configuration

In order to ease the management of settings, *Cliquet* provides a helper that reads values from *environment variables* and uses default application values.

```
import cliquet
from pyramid.settings import asbool

DEFAULT_SETTINGS = {
    'cliquet_elasticsearch.refresh_enabled': False
}

def includeme(config):
    cliquet.load_default_settings(config, DEFAULT_SETTINGS)
    settings = config.get_settings()

    refresh_enabled = settings['cliquet_elasticsearch.refresh_enabled']
    if asbool(refresh_enabled):
        ...

    config.add_view(...)
```

In this example, if the environment variable `CLIQUET_ELASTICSEARCH_REFRESH_ENABLED` is set to `true`, the value present in configuration file is ignored.

1.5.3 Custom backend

As a simple example, let’s add another kind of cache backend to *Cliquet*.

`cliquet_riak/cache.py`:

```
from cliquet.cache import CacheBase
from riak import RiakClient
```



```

class Riak(CacheBase):
    def __init__(self, **kwargs):
        self._client = RiakClient(**kwargs)
        self._bucket = self._client.bucket('cache')

    def set(self, key, value, ttl=None):
        key = self._bucket.new(key, data=value)
        key.store()
        if ttl is not None:
            # ...

    def get(self, key):
        fetched = self._bucket.get(key)
        return fetched.data

    #
    # ...see cache documentation for a complete API description.
    #

def load_from_config(config):
    settings = config.get_settings()
    uri = settings['cliquet.cache_url']
    uri = urlparse.urlparse(uri)

    return Riak(pb_port=uri.port or 8087)

```

Once its package installed and available in Python path, this new backend type can be specified in application configuration:

```

# myproject.ini
cliquet.cache_backend = cliquet_riak.cache

```

1.5.4 Adding features

Another use-case would be to add extra-features, like indexing for example.

- Initialize an indexer on startup;
- Add a `/search/{collection}/` end-point;
- Index records manipulated by resources.

Inclusion and startup in `cliquet_indexing/__init__.py`:

```

DEFAULT_BACKEND = 'cliquet_indexing.elasticsearch'

def includeme(config):
    settings = config.get_settings()
    backend = settings.get('cliquet.indexing_backend', DEFAULT_BACKEND)
    indexer = config.maybe_dotted(backend)

    # Store indexer instance in registry.
    config.registry.indexer = indexer.load_from_config(config)

    # Activate end-points.
    config.scan('cliquet_indexing.views')

```

End-point definitions in `cliquet_indexing/views.py`:

```
from cornice import Service

search = Service(name="search",
                 path='/search/{collection_id}/',
                 description="Search")

@search.post()
def get_search(request):
    collection_id = request.matchdict['collection_id']
    query = request.body

    # Access indexer from views using registry.
    indexer = request.registry.indexer
    results = indexer.search(collection_id, query)

    return results
```

Example indexer class in cliquet_indexing/elasticsearch.py:

```
class Indexer(...):
    def __init__(self, hosts):
        self.client = elasticsearch.Elasticsearch(hosts)

    def search(self, collection_id, query, **kwargs):
        try:
            return self.client.search(index=collection_id,
                                     doc_type=collection_id,
                                     body=query,
                                     **kwargs)
        except ElasticsearchException as e:
            logger.error(e)
            raise

    def index_record(self, collection_id, record, id_field):
        record_id = record[id_field]
        try:
            index = self.client.index(index=collection_id,
                                     doc_type=collection_id,
                                     id=record_id,
                                     body=record,
                                     refresh=True)

            return index
        except ElasticsearchException as e:
            logger.error(e)
            raise
```

Indexed resource in cliquet_indexing/resource.py:

```
class IndexedModel(cliquet.resource.Model):
    def create_record(self, record):
        r = super(IndexedModel, self).create_record(self, record)

        self.indexer.index_record(self, record)

        return r

class IndexedResource(cliquet.resource.BaseResource):
    def __init__(self, request):
        super(IndexedResource, self).__init__(request)
```

```
self.model.indexer = request.registry.indexer
```

Note: In this example, `IndexedResource` must be used explicitly as a base resource class in applications. A nicer pattern would be to trigger *Pyramid* events in *Cliquet* and let packages like this one plug listeners. If you're interested, we started to discuss it!

1.5.5 JavaScript client

One of the main goal of *Cliquet* is to ease the development of REST microservices, most likely to be used in a JavaScript environment.

A client could look like this:

```
var client = new cliquet.Client({
  server: 'https://api.server.com',
  store: localforage
});

var articles = client.resource('/articles');

articles.create({title: "Hello world"})
  .then(function (result) {
    // success!
  });

articles.get('id-1234')
  .then(function (record) {
    // Read from local if offline.
  });

articles.filter({
  title: {'$eq': 'Hello'}
})
  .then(function (results) {
    // List of records.
  });

articles.sync()
  .then(function (result) {
    // Synchronize offline store with server.
  })
  .catch(function (err) {
    // Error happened.
    console.error(err);
  });
```

1.6 Contributing

Thank you for considering to contribute to *Cliquet*!

Note: No contribution is too small; we welcome fixes about typos and grammar bloopers. Don't hesitate to send us a pull request!

Note: Open a pull-request even if your contribution is not ready yet! It can be discussed and improved collaboratively, and avoid having you doing a lot of work without getting feedback.

1.6.1 Setup your development environment

To prepare your system with Python 3.4, PostgreSQL and Redis, please refer to the [Installation](#) guide.

You might need to install [curl](#), if you don't have it already.

Prepare your project environment by running:

```
$ make install-dev
```

```
$ pip install tox
```

OS X

On OSX especially you might get the following error when running tests:

```
$ ValueError: unknown locale: UTF-8
```

If this is the case add the following to your `~/.bash_profile`:

```
export LC_ALL=en_US.UTF-8
export LANG=en_US.UTF-8
```

Then run:

```
$ source ~/.bash_profile
```

1.6.2 Run tests

Currently, running the complete test suite implies to run every type of backend.

That means:

- Run Redis on `localhost:6379`
- Run a PostgreSQL 9.4 `testdb` database on `localhost:5432` with user `postgres/postgres`. The database encoding should be UTF-8, and the database timezone should be UTC.

```
make tests
```

Run a single test

For Test-Driven Development, it is possible to run a single test case, in order to speed-up the execution:

```
nosetests -s --with-mocha-reporter cliquet.tests.test_views_hello:HelloViewTest.test_returns_info_about_hello
```

1.6.3 Definition of done

In order to have your changes incorporated, you need to respect these rules:

- **Tests pass**; Travis-CI will build the tests for you on the branch when you push it.

- **Code added comes with tests;** We try to have a 100% coverage on the codebase to avoid surprises. No code should be untested :) If you fail to see how to test your changes, feel welcome to say so in the pull request, we'll gladly help you to find out.
- **Documentation is up to date;**

1.6.4 IRC channel

If you want to discuss with the team behind *Cliquet*, please come and join us on `#storage` on `irc.mozilla.org`.

- Because of differing time zones, you may not get an immediate response to your question, but please be patient and stay logged into IRC — someone will almost always respond if you wait long enough (it may take a few hours).
- If you don't have an IRC client handy, use [the webchat](#) for quick feedback.
- You can direct your IRC client to the channel using this IRC link or you can manually join the `#storage` IRC channel on the mozilla IRC network.

1.6.5 How to release

In order to prepare a new release, we are following the following steps.

The *prerelease* and *postrelease* commands are coming from [zest.releaser](#).

Install *zest.releaser* with the *recommended* dependencies. They contain *wheel* and *twine*, which are required to release a new version.

```
$ pip install "zest.releaser[recommended]"
```

Step 1

- Merge remaining pull requests
- Update `CHANGELOG.rst`
- Update version in `cliquet_docs/conf.py`
- Known good versions of dependencies in `requirements.txt`
- Update `CONTRIBUTORS.rst` using: `git shortlog -sne | awk '{ $1=""; sub(" ", ""); print }' | awk -F'<' '!x[$1]++' | awk -F'<' '!x[$2]++' | sort`

```
$ git checkout -b prepare-X.Y.Z
$ prerelease
$ vim cliquet_docs/conf.py
$ make build-requirements
$ git commit -a --amend
$ git push origin prepare-X.Y.Z
```

- Open a pull-request with to release the version.

Step 2

Once the pull-request is validated, merge it and do a release. Use the `release` command to invoke the `setup.py`, which builds and uploads to PyPI

```
$ git checkout master
$ git merge --no-ff prepare-X.Y.Z
$ release
$ postrelease
```

Step 3

As a final step:

- Close the milestone in Github
- Add entry in Github release page
- Create next milestone in Github in the case of a major release
- Configure the version in ReadTheDocs
- Send mail to ML (If major release)

That's all folks!

1.6.6 Cleaning your environment

There are three levels of cleaning your environment:

- `make clean` will remove `*.pyc` files and `__pycache__` directory.
- `make distclean` will also remove `*.egg-info` files and `*.egg`, `build` and `dist` directories.
- `make maintainer-clean` will also remove the `.tox` and the `.venv` directories.

1.7 Changelog

This document describes changes between each past release.

1.7.1 2.10.0 (2015-10-30)

Protocol

- Moved `userid` attribute to a dedicated `user` mapping in the hello view.
- Fixed 503 error message to mention backend errors in addition to unavailability.
- Set cache headers only when anonymous (fixes #449)
- Follow redirections in batch subrequests (fixes #511)
- When recreating a record that was previously deleted, status code is now 201 (ref #530).

New features

- Follow redirections in batch subrequests (fixes #511)
- Add a `readonly` setting to run the service in read-only mode. (#525)
- If no client cache is set, add `Cache-Control: no-cache` by default, so that clients are forced to revalidate their cache against the server (#522, ref Kinto/kinto#231)

Bug fixes

- Fix PostgreSQL error when deleting an empty collection in a protected resource (fixes #528)
- Fix PUT not using `create()` method in storage backend when tombstone exists (fixes #530)
- Delete tombstone when record is re-created (fixes #518)
- Fix crash with empty body for PATCH (fixes #477, fixes #516)
- Fix english typo in 404 error message (fixes #527)

Internal changes

- Better `__pycache__` cleaning

1.7.2 2.9.0 (2015-10-27)

New features

- Added Pyramid events, triggered when the content of a resource has changed. (#488)
- Added `cliquet.includes` setting allowing loading of plugins once Cliquet is initialized (unlike `pyramid.includes`). (#504)

Protocol

- Remove the broken git revision `commit` field in the hello page. (#495).

Breaking changes

- Renamed internal backend classes for better consistency. Settings remain unchanged, but if you imported the backend classes in your Cliquet application, it will break (#491).
- `cliquet.schema` is now deprecated, and was moved to a `cliquet.resource` module. (#505)
- Resource collection attribute is now deprecated. Use `model` attribute instead. (#506)

Internal changes

- Rework PostgreSQL backends to use composition instead of inheritance for the client code. (#491)
- Replace DROP INDEX by a conditional creation in PostgreSQL schemas (#487, #496 thanks @rodo)
- Documentation and minor refactors in viewset code (#490, #498, #502)
- Add the `build-requirements`, `distclean` and `maintainer-clean` Makefile rules.
- Documentation JSON patch format. (#484)
- Fix for permission among record fields in 412 errors. (#499)

1.7.3 2.8.2 (2015-10-22)

Bug fixes

- Fix crash on settings with list values (#481)
- Fix crash in Redis permission backend (ref Kinto/kinto#215)

Internal changes

- Use tox installed in virtualenv (#486)
- Skip python versions unavailable in tox (#486)

1.7.4 2.8.1 (2015-10-14)

- Expose public settings without prefix, except if we explicitly configure `public_settings` to expose them (with `cliquet.` or `project_name.`) (ref #476)

1.7.5 2.8.0 (2015-10-06)

Breaking changes

- Deprecated settings `cliquet.cache_pool_maxconn`, `cliquet.storage_pool_maxconn` and `cliquet.basic_auth_enabled` were removed (ref #448)
- Prefixed settings will not work if `project_name` is not defined. (either with `cliquet.initialize()` or with the `cliquet.project_name` configuration variable).
- Settings should now be read without their prefix in the code: `request.registry.settings['max_duration']` rather than `request.registry.settings['cliquet.max_duration']`

New features

- Add cache CORS headers. (ref #466)
- Use the project name as setting prefix (ref #472)

Internal changes

- Expose statsd client so that projects using cliquet can send statsd metrics. (ref #465)
- Refactor `BaseWebTest`. (ref #468)
- Remove hard coded CORS origins in order to be able to override it with config. (ref #467)
- Allow overriding 405 response error to give context (ref #471)
- Allow overriding 503 response error to give context (ref #473)

1.7.6 2.7.0 (2015-09-23)

Breaking changes

- Backends are not instantiated by default anymore (used to be with *Redis*) (#461)

New features

- Redirect to remove trailing slash in URLs (fixes Kinto/kinto#112)
- Add resource cache control headers via settings (fixes #401)
- Add request `bound_data` attribute, shared with subrequests. Useful to share context or cache values between BATCH requests for example (#459)

Bug fixes

- Fix Werkzeug profiling setup docs and code (#451)
- Fix logger encoding error with UTF-8 output (#455)
- Do not instantiate backends if not configured (fixes #386)

Internal changes

- Huge refactoring the interaction between `Resource` and `Permission` backend (#454)
- Fetch record only once from storage with PUT requests on resources (#452)

- Index permissions columns, bringing huge performance gain for shared collections (#458, ref #354)
- Add instructions to mention contributors list in documentation (#408)
- Explicitly call to collection create_record on PUT (#460)

1.7.7 2.6.2 (2015-09-09)

Bug fixes

- Expose CORS headers on subrequest error response and for non service errors (#435).
- Make sure a tuple is passed for Postgresql list comparisons even for ids (#443).

Internal changes

- Use the `get_bound_permissions` callback to select shared records in collection list (#444).

1.7.8 2.6.1 (2015-09-08)

Bug fixes

- Make sure a tuple is passed for Postgresql in conditions (#441).

1.7.9 2.6.0 (2015-09-08)

Protocol

- Fix consistency in API to modify permissions with PATCH (#437, ref Kinto/kinto#155). The list of principals for each specified permission is now replaced by the one provided.

New features

- Partial collection of records for `ProtectedResource` when user has no `read` permission (fixes #354). Alice can now obtain a list of Bob records on which she has read/write permission.

Internal changes

- Fix Wheel packaging for Pypy (fixes Kinto/kinto#177)
- Add additional test to make sure 400 errors returns CORS Allowed Headers

1.7.10 2.5.0 (2015-09-04)

Protocol

- Collection records can now be filtered using multiple values (`?in_status=1,2,3`) (fixes #39)
- Collection records can now be filtered excluding multiple values (`?exclude_status=1,2,3`) (fixes mozilla-services/readinglist#68)

Internal changes

- We can obtains accessible `objects_id` in a collection from user principals (fixes #423)

1.7.11 2.4.3 (2015-08-26)

Bug fixes

- Fix the packaging for cliquet (#430)

1.7.12 2.4.2 (2015-08-26)

Internal changes

- Remove the symlink to cliquet_docs and put the documentation inside *cliquet_docs* directly (#426)

1.7.13 2.4.1 (2015-08-25)

Internal changes

- Make documentation available from outside by using *cliquet_docs* (#413)

1.7.14 2.4.0 (2015-08-14)

Protocol

- Userid is now provided when requesting the hello endpoint with an `Authorization` header (#319)
- UUID validation now accepts any kind of UUID, not just v4 (fixes #387)
- Querystring parameter `_to` was renamed to `_before` (*the former is now deprecated*) (#391)

New features

- `Cliquet Service` class now has the default error handler attached (#388)
- Allow to configure info link in error responses with `cliquet.error_info_link` setting (#395)
- Storage backend now has a `purge_deleted()` to get rid of *tombstones* (#400)

Bug fixes

- Fix missing `Backoff` header for 304 responses (fixes #416)
- Fix Python3 encoding errors (#328)
- `data` is not mandatory in request body if the resource does not define any schema or if no field is mandatory (fixes mozilla-services/kinto#63)
- Fix no validation error on PATCH with unknown attribute (fixes #374)
- Fix permissions not validated on PATCH (fixes #375)
- Fix CORS header missing in 404 responses for unknown URLs (fixes #414)

Internal changes

- Renamed main documentation sections to *HTTP Protocol* and *Internals* (#394)
- Remove mentions of storage in documentation to avoid confusions with the *Kinto* project.
- Add details in timestamp documentation.
- Mention talk at Python Meetup Barcelona in README
- Fix documentation about postgres-contrib dependancy (#409)

- Add `cliquet.utils` to *Internals* documentation (#407)
- Default id generator now accepts dashes and underscores (#411)

1.7.15 2.3.1 (2015-07-15)

Bug fixes

- Fix crash on hello view when application is not deployed from Git repository (fixes #382)
- Expose Content-Length header to Kinto.js (#390)

1.7.16 2.3 (2015-07-13)

New features

- Provide details about existing record in 412 error responses (fixes mozilla-services/kinto#122)
- Add ETag on record PUT/PATCH responses (fixes #352)
- Add StatsD counters for the permission backend

Bug fixes

- Fix crashes in permission backends when permission set is empty (fixes #368, #371)
- Fix value of ETag on record: provide collection timestamp on collection endpoints only (fixes #356)
- Default resources do accept `permissions` attribute in payload anymore
- Default resources do not require a root factory (fixes #348)
- Default resources do not hit the permission backend anymore
- Default viewset was split and does not handle permissions anymore (fixes #322)
- Permissions on views is now set only on resources
- Fix missing `last_modified` field in PATCH response when no field was changed (fixes #371)
- Fix lost querystring during version redirection (fixes #364)

Internal changes

- Document the list of public settings in hello view (mozilla-services/kinto#133)

1.7.17 2.2.1 (2015-07-06)

Bug fixes

- Fix permissions handling on PATCH /resource (#358)

1.7.18 2.2.0 (2015-07-02)

New features

- Add public settings in hello view (#318)

Bug fixes

- Fix version redirection behaviour for unsupported versions (#341)

- PostgreSQL dependencies are now fully optional in code (#340)
- Prevent overriding final settings from `default_settings` parameter in `cliquet.initialize()` (#343)

Internal changes

- Fix installation documentation regarding PostgreSQL 9.4 (#338, thanks @elemoine!)
- Add detail about UTC and UTF-8 for PostgreSQL (#347, thanks @elemoine!)
- Remove `UserWarning` exception when running tests (#339, thanks @elemoine!)
- Move `build_request` and `build_response` to `cliquet.utils` (#344)
- Pypy is now tested on Travis CI (#337)

1.7.19 2.1.0 (2015-06-26)

New features

- Cliquet does not require authentication policies to prefix user ids anymore (fixes #299).
- Pypy support (thanks Balthazar Rouberol #325)
- Allow to override parent id of resources (#333)

Bug fixes

- Fix crash in authorization on `OPTIONS` requests (#331)
- Fix crash when `If-Match` is provided without `If-None-Match` (#335)

Internal changes

- Fix docstrings and documentation (#329)

1.7.20 2.0.0 (2015-06-16)

New features

- Authentication and authorization policies, as well as group finder function can now be specified via configuration (fixes #40, #265)
- Resources can now be protected by fine-grained permissions (#288 via #291, #302)

Minor

- Preserve provided `id` field of records using `POST` on collection (#293 via #294)
- Logging value for authentication type is now available for any kind of authentication policy.
- Any resource endpoint can now be disabled from settings (#46 via #268)

Bug fixes

- Do not limit cache values to string (#279)
- When `PUT` creates the record, the HTTP status code is now 201 (#298, #300)
- Add safety check in `utils.current_service()` (#316)

Breaking changes

- `cliquet.storage.postgresql` now requires PostgreSQL version 9.4, since it now relies on *JSONB*. Data will be migrated automatically using the `migrate` command.

- the `@crud` decorator was replaced by `@register()` (fixes #12, #268)
- Firefox Accounts code was removed and published as external package *cliquet-fxa*
- The *Cloud storage* storage backend was removed out of *Cliquet* and should be revamped in *Kinto* repository (mozilla-services/kinto#45)

API

- Resource endpoints now expect payloads to have a `data` attribute (#254, #287)
- Resource endpoints switched from `If-Modified-Since` and `If-Unmodified-Since` to `Etags` (fixes #251 via #275), thanks @michielbdejong!

Minor

- `existing` attribute of conflict errors responses was moved inside a generic `details` attribute that is also used to list validation errors.
- Setting `cliquet.basic_auth_enabled` is now deprecated. Use `pyramid_multiauth` configuration instead to specify authentication policies.
- Logging value for authentication type is now `authn_type` (with `FxOAuth` or `BasicAuth` as default values).

Internal changes

- Cliquet resource code was split into `Collection` and `Resource` (fixes #243, #282)
- Cleaner separation of concern between `Resource` and the new notion of `ViewSet` (#268)
- Quickstart documentation improvement (#271, #312) thanks @N1k0 and @brouberol!
- API versioning documentation improvements (#313)
- Contribution documentation improvement (#306)

1.7.21 1.8.0 (2015-05-13)

Breaking changes

- Switch PostgreSQL storage to JSONB: requires 9.4+ (#104)
- Resource name is not a Python property anymore (ref #243)
- Return existing record instead of raising 409 on POST (fixes #75)
- `cliquet.storage.postgresql` now requires version PostgreSQL 9.4, since it now relies on *JSONB*. Data will be migrated automatically using the `migrate` command.
- Conflict errors responses `existing` attribute was moved inside a generic `details` attribute that is also used to list validation errors.
- In heartbeat end-point response, `database` attribute was renamed to `storage`

New features

- Storage records ids are now managed in python (fixes #71, #208)
- Add setting to disable version redirection (#107, thanks @hiromipaw)
- Add response behaviour headers for PATCH on record (#234)
- Provide details in error responses (#233)
- Expose new function `cliquet.load_default_settings()` to ease reading of settings from defaults and environment (#264)

- Heartbeat callback functions can now be registered during startup (#261)

Bug fixes

- Fix migration behaviour when metadata table is flushed (#221)
- Fix backoff header presence if disabled in settings (#238)

Internal changes

- Require 100% of coverage for tests to pass
- Add original error message to storage backend error
- A lots of improvements in documentation (#212, #225, #228, #229, #237, #246, #247, #248, #256, #266, thanks Michiel De Jong)
- Migrate *Kinto* storage schema on startup (#218)
- Fields `id` and `last_modified` are not part of resource schema anymore (#217, mozilla-services/readinlist#170)
- Got rid of redundant indices in storage schema (#208, ref #138)
- Disable Cornice schema request binding (#172)
- Do not hide FxA errors (fixes mozilla-services/readinglist#70)
- Move initialization functions to dedicated module (ref #137)
- Got rid of request custom attributes for storage and cache (#245)

1.7.22 1.7.0 (2015-04-10)

Breaking changes

- A **command must be ran during deployment** for database schema migration:

```
$ cliquet -ini production.ini migrate
```
- Sentry custom code was removed. Sentry logging is now managed through the logging configuration, as explained [in docs](#).

New features

- Add PostgreSQL schema migration system (#139)
- Add cache and oauth in heartbeat view (#184)
- Add monitoring features using NewRelic (#189)
- Add profiling features using Werkzeug (#196)
- Add ability to override default settings in initialization (#136)
- Add more statsd counter for views and authentication (#200)
- Add in-memory cache class (#127)

Bug fixes

- Fix crash in DELETE on collection with PostgreSQL backend
- Fix Heka logging format of objects (#199)
- Fix performance of record insertion using ordered index (#138)
- Fix 405 errors not JSON formatted (#88)

- Fix basic auth prompt when disabled (#182)

Internal changes

- Improve development setup documentation (thanks @hiromipaw)
- Deprecated `cliquet.initialize_cliquet`, renamed to `cliquet.initialize`.
- Code coverage of tests is now 100%
- Skip unstable tests on TravisCI, caused by `fsync = off` in their PostgreSQL.
- Perform random creation and deletion in heartbeat view (#202)

1.7.23 1.6.0 (2015-03-30)

New features

- Split schema initialization from application startup, using a command-line tool.

```
cliquet --ini production.ini init
```

Bug fixes

- Fix connection pool no being shared between cache and storage (#176)
- Default connection pool size to 10 (instead of 50) (#176)
- Warn if PostgreSQL session has not UTC timezone (#177)

Internal changes

- Deprecated `cliquet.storage_pool_maxconn` and `cliquet.cache_pool_maxconn` settings (renamed to `cliquet.storage_pool_size` and `cliquet.cache_pool_size`)

1.7.24 1.5.0 (2015-03-27)

New features

- Measure calls on the authentication policy (#167)

Breaking changes

- Prefix statsd metrics with the value of `cliquet.statsd_prefix` or `cliquet.project_name` (#162)
- `http_scheme` setting has been replaced by `cliquet.http_scheme` and `cliquet.http_host` was introduced ((#151, #166)
- URL in the hello view now has version prefix (#165)

Bug fixes

- Fix Next-Page url if service has key in url (#158)
- Fix some PostgreSQL connection bottlenecks (#170)

Internal changes

- Update of PyFxA to get it working with gevent monkey patching (#168)
- Reload kinto on changes (#158)

1.7.25 1.4.1 (2015-03-25)

Bug fixes

- Rely on Pyramid API to build pagination Next-Url (#147)

1.7.26 1.4.0 (2015-03-24)

Breaking changes

- Make monitoring dependencies optional (#121)

Bug fixes

- Force PostgreSQL session timezone to UTC (#122)
- Fix basic auth ofuscation and prefix (#128)
- Make sure the *paginate_by* setting overrides the passed *limit* argument (#129)
- Fix limit comparison under Python3 (#143)
- Do not serialize using JSON if not necessary (#131)
- Fix crash of classic logger with unicode (#142)
- Fix crash of CloudStorage backend when remote returns 500 (#142)
- Fix behaviour of CloudStorage with backslashes in querystring (#142)
- Fix python3.4 segmentation fault (#142)
- Add missing port in Next-Page header (#147)

Internal changes

- Use ujson again, it was removed in the 1.3.2 release (#132)
- Add index for `as_epoch(last_modified)` (#130). Please add the following statements to SQL for the migration:

```
ALTER FUNCTION as_epoch(TIMESTAMP) IMMUTABLE;  
CREATE INDEX idx_records_last_modified_epoch ON records(as_epoch(last_modified));  
CREATE INDEX idx_deleted_last_modified_epoch ON deleted(as_epoch(last_modified));
```

- Prevent fetching to many records for one user collection (#130)
- Use UPSERT for the heartbeat (#141)
- Add missing OpenSSL in installation docs (#146)
- Improve tests of basic auth (#128)

1.7.27 1.3.2 (2015-03-20)

- Revert ujson usage (#132)

1.7.28 1.3.1 (2015-03-20)

Bug fixes

- Fix packaging (#118)

1.7.29 1.3.0 (2015-03-20)

New features

- Add PostgreSQL connection pooling, with new settings `cliquet.storage_pool_maxconn` and `cliquet.cache_pool_maxconn` (*Default: 50*) (#112)
- Add [StatsD](#) support, enabled with `cliquet.statsd_url = udp://server:port` (#114)
- Add [Sentry](#) support, enabled with `cliquet.sentry_url = http://user:pass@server/1` (#110)

Bug fixes

- Fix FxA verification cache not being used (#103)
- Fix heartbeat database check (#109)
- Fix PATCH endpoint crash if request has no body (#115)

Internal changes

- Switch to [ujson](#) for JSON de/serialization optimizations (#108)

1.7.30 1.2.1 (2015-03-18)

- Fix tests about unicode characters in BATCH querystring patch
- Remove CREATE CAST for the postgresql backend
- Fix environment variable override

1.7.31 1.2 (2015-03-18)

Breaking changes

- `cliquet.storage.postgresql` now uses UUID as record primary key (#70)
- Settings `cliquet.session_backend` and `cliquet.session_url` were renamed `cliquet.cache_backend` and `cliquet.cache_url` respectively.
- FxA user ids are not hashed anymore (#82)
- Setting `cliquet.retry_after` was renamed `cliquet.retry_after_seconds`
- OAuth2 redirect url now requires to be listed in `fxa-oauth.webapp.authorized_domains` (e.g. `*.mozilla.com`)
- Batch are now limited to 25 requests by default (#90)

New features

- Every setting can be specified via an environment variable (e.g. `cliquet.storage_url` with `CLIQUET_STORAGE_URL`)
- Logging now relies on [structlog](#) (#78)
- Logging output can be configured to stream JSON (#78)
- New cache backend for PostgreSQL (#44)
- Documentation was improved on various aspects (#64, #86)
- Handle every backend errors and return 503 errors (#21)
- State verification for OAuth2 dance now expires after 1 hour (#83)

Bug fixes

- FxA OAuth views errors are now JSON formatted (#67)
- Prevent error when pagination token has bad format (#72)
- List of CORS exposed headers were fixed in POST on collection (#54)

Internal changes

- Added a method in *cliquet.resource.Resource* to override known fields (*required by Kinto*)
- Every setting has a default value
- Every end-point requires authentication by default
- Session backend was renamed to cache (#96)

1.7.32 1.1.4 (2015-03-03)

- Update deleted_field support for postgres (#62)

1.7.33 1.1.3 (2015-03-03)

- Fix include_deleted code for the redis backend (#60)
- Improve the update_record API (#61)

1.7.34 1.1.2 (2015-03-03)

- Fix packaging to include .sql files.

1.7.35 1.1.1 (2015-03-03)

- Fix packaging to include .sql files.

1.7.36 1.1 (2015-03-03)

New features

- Support filter on deleted using since (#51)

Internal changes

- Remove python 2.6 support (#50)
- Renamed Resource.deleted_mark to Resource.deleted_field (#51)
- Improve native_value (#56)
- Fixed Schema options inheritance (#55)
- Re-build the virtualenv when setup.py changes
- Renamed storage.url to cliquet.storage_url (#49)
- Refactored the tests/support.py file (#38)

1.7.37 1.0 (2015-03-02)

- Initial version, extracted from Mozilla Services Reading List project (#1)

New features

- Expose CORS headers so that client behind CORS policy can access them (#5)
- Postgresql Backend (#8)
- Use RedisSession as a cache backend for PyFxA (#10)
- Delete multiple records via DELETE on the collection_path (#13)
- Batch default prefix for endpoints (#14 / #16)
- Use the app version in the / endpoint (#22)
- Promote Basic Auth as a proper authentication backend (#37)

Internal changes

- Backends documentation (#15)
- Namedtuple for filters and sort (#17)
- Multiple DELETE in Postgresql (#18)
- Improve Resource API (#29)
- Refactoring of error management (#41)
- Default Options for Schema (#47)

1.8 Contributors

- Alexis Metaireau <alexis@mozilla.com>
- Andy McKay <amckay@mozilla.com>
- Balthazar Rouberol <br@imap.cc>
- Dan Phrawzty <phrawzty+github@gmail.com>
- Éric Lemoine <eric.lemoine@gmail.com>
- Hiromipaw <silvia@nopressure.co.uk>
- Mathieu Leplatre <mathieu@mozilla.com>
- Michiel de Jong <michiel@unhosted.org>
- Nicolas Perriault <nperriault@mozilla.com>
- Rémy Hubscher <rhubscher@mozilla.com>
- Rodolphe Quiédeville <rodolphe@quiedeville.org>
- Tarek Ziade <tarek@mozilla.com>

Indices and tables

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

C

`cliquet.errors`, [55](#)
`cliquet.schema`, [37](#)
`cliquet.storage`, [46](#)
`cliquet.storage.exceptions`, [50](#)
`cliquet.storage.generators`, [42](#)
`cliquet.utils`, [56](#)

A

Access Control Entity, [59](#)
 Access Control List, [59](#)
 ACE, [59](#)
 ACEs, [59](#)
 ACL, [59](#)
 ACLs, [59](#)
 add_principal_to_ace() (cli-
 quet.permission.PermissionBase method),
 [53](#)
 add_user_principal() (cliquet.permission.PermissionBase
 method), [53](#)
 apply_changes() (cliquet.resource.BaseResource
 method), [39](#)
 AuthorizationPolicy (class in cliquet.authorization), [53](#)

B

BackendError, [50](#)
 BaseResource (class in cliquet.resource), [37](#)
 build_request() (in module cliquet.utils), [57](#)
 build_response() (in module cliquet.utils), [58](#)

C

CacheBase (class in cliquet.cache), [50](#)
 check_permission() (cliquet.permission.PermissionBase
 method), [54](#)
 classname() (in module cliquet.utils), [57](#)
 cliquet.errors (module), [55](#)
 cliquet.schema (module), [37](#)
 cliquet.storage (module), [46](#)
 cliquet.storage.exceptions (module), [50](#)
 cliquet.storage.generators (module), [42](#)
 cliquet.utils (module), [56](#)
 collection (cliquet.resource.BaseResource attribute), [37](#)
 collection_delete() (cliquet.resource.BaseResource
 method), [38](#)
 collection_get() (cliquet.resource.BaseResource method),
 [37](#)
 collection_post() (cliquet.resource.BaseResource
 method), [37](#)

collection_timestamp() (cliquet.storage.StorageBase
 method), [47](#)
 create() (cliquet.storage.StorageBase method), [47](#)
 create_record() (cliquet.resource.Model method), [41](#)
 CRUD, [58](#)
 current_service() (in module cliquet.utils), [57](#)

D

decode64() (in module cliquet.utils), [57](#)
 decode_header() (in module cliquet.utils), [58](#)
 default_model (cliquet.resource.BaseResource attribute),
 [37](#)
 default_viewset (cliquet.resource.BaseResource at-
 tribute), [37](#)
 delete() (cliquet.cache.CacheBase method), [51](#)
 delete() (cliquet.resource.BaseResource method), [38](#)
 delete() (cliquet.storage.StorageBase method), [48](#)
 delete_all() (cliquet.storage.StorageBase method), [48](#)
 delete_object_permissions() (cli-
 quet.permission.PermissionBase method),
 [55](#)
 delete_record() (cliquet.resource.Model method), [42](#)
 delete_records() (cliquet.resource.Model method), [40](#)
 deleted_field (cliquet.resource.Model attribute), [40](#)
 direction (cliquet.storage.Sort attribute), [46](#)

E

encode64() (in module cliquet.utils), [57](#)
 encode_header() (in module cliquet.utils), [58](#)
 endpoint, [58](#)
 ERRORS (in module cliquet.errors), [55](#)
 expire() (cliquet.cache.CacheBase method), [51](#)
 extensible, [58](#)

F

field (cliquet.storage.Filter attribute), [46](#)
 field (cliquet.storage.Sort attribute), [47](#)
 Filter (class in cliquet.storage), [46](#)
 Firefox Accounts, [58](#)
 flush() (cliquet.cache.CacheBase method), [50](#)

flush() (cliquet.permission.PermissionBase method), 53
 flush() (cliquet.storage.StorageBase method), 47
 follow_subrequest() (in module cliquet.utils), 58

G

Generator (class in cliquet.storage.generators), 42
 get() (cliquet.cache.CacheBase method), 51
 get() (cliquet.resource.BaseResource method), 38
 get() (cliquet.storage.StorageBase method), 47
 get_all() (cliquet.storage.StorageBase method), 49
 get_name() (cliquet.resource.ViewSet method), 45
 get_parent_id() (cliquet.resource.BaseResource method), 37
 get_record() (cliquet.resource.Model method), 41
 get_record_schema() (cliquet.resource.ViewSet method), 45
 get_records() (cliquet.resource.Model method), 40
 get_service_name() (cliquet.resource.ViewSet method), 45
 get_view() (cliquet.resource.ViewSet method), 45
 get_view_arguments() (cliquet.resource.ViewSet method), 45

H

hmac_digest() (in module cliquet.utils), 57
 http_error() (in module cliquet.errors), 55

I

id_field (cliquet.resource.Model attribute), 40
 initialize() (in module cliquet), 9
 initialize_schema() (cliquet.cache.CacheBase method), 50
 initialize_schema() (cliquet.permission.PermissionBase method), 53
 initialize_schema() (cliquet.storage.StorageBase method), 47
 is_endpoint_enabled() (cliquet.resource.ViewSet method), 45
 is_known_field() (cliquet.resource.BaseResource method), 37

J

json_error_handler() (in module cliquet.errors), 56

K

KISS, 58

M

mapping (cliquet.resource.BaseResource attribute), 37
 match() (cliquet.storage.generators.Generator method), 43
 Memory (class in cliquet.storage.memory), 46
 merge_dicts() (in module cliquet.utils), 57

Model (class in cliquet.resource), 40
 modified_field (cliquet.resource.Model attribute), 40
 msec_time() (in module cliquet.utils), 56

N

native_value() (in module cliquet.utils), 57

O

object, 58
 object_permission_authorized_principals() (cli-
 quet.permission.PermissionBase method), 54
 object_permission_principals() (cli-
 quet.permission.PermissionBase method), 54
 object_permissions() (cliquet.permission.PermissionBase method), 54
 objects, 58
 operator (cliquet.storage.Filter attribute), 46

P

patch() (cliquet.resource.BaseResource method), 38
 permission, 59
 PermissionBase (class in cliquet.permission), 53
 permissions, 59
 pluggable, 58
 principal, 59
 principals, 59
 principals_accessible_objects() (cli-
 quet.permission.PermissionBase method), 54
 process_record() (cliquet.resource.BaseResource method), 38
 purge_deleted() (cliquet.storage.StorageBase method), 49
 put() (cliquet.resource.BaseResource method), 38

R

raise_invalid() (in module cliquet.errors), 56
 random_bytes_hex() (in module cliquet.utils), 57
 read_env() (in module cliquet.utils), 57
 reapply_cors() (in module cliquet.utils), 57
 RecordNotFoundError, 50
 Redis (class in cliquet.storage.redis), 46
 regexp (cliquet.storage.generators.Generator attribute), 43
 regexp (cliquet.storage.generators.UUID4 attribute), 43
 remove_principal_from_ace() (cli-
 quet.permission.PermissionBase method), 53
 remove_user_principal() (cli-
 quet.permission.PermissionBase method), 53

replace_object_permissions() (cli-
quet.permission.PermissionBase method),
55
resource, 58
RouteFactory (class in cliquet.authorization), 52

S

send_alert() (in module cliquet.errors), 56
set() (cliquet.cache.CacheBase method), 51
Sort (class in cliquet.storage), 46
StorageBase (class in cliquet.storage), 47
strip_uri_prefix() (in module cliquet.utils), 58
strip_whitespace() (in module cliquet.utils), 56

T

timestamp() (cliquet.resource.Model method), 40
tombstone, 58
tombstones, 58
ttl() (cliquet.cache.CacheBase method), 50

U

UnicityError, 50
update() (cliquet.resource.ViewSet method), 45
update() (cliquet.storage.StorageBase method), 48
update_record() (cliquet.resource.Model method), 41
user id, 58
user identifier, 58
user identifiers, 58
user_principals() (cliquet.permission.PermissionBase
method), 53
UUID4 (class in cliquet.storage.generators), 43

V

value (cliquet.storage.Filter attribute), 46
ViewSet (class in cliquet.resource), 45