crawlfrontier Documentation

Release 0.2.0

ScrapingHub

April 15, 2015

Contents

This documentation contains everything you need to know about Crawl Frontier.

First steps

1.1 Crawl Frontier at a glance

Crawl Frontier is an application framework that is meant to be used as part of a Crawling System, allowing you to easily manage and define tasks related to a *Crawling Frontier*.

Even though it was originally designed for Scrapy, it can also be used with any other Crawling Framework/System as the framework offers a generic frontier functionality.

The purpose of this document is to introduce you to the concepts behind Crawl Frontier so that you can get an idea of how it works and to decide if it is suited to your needs.

1.1.1 1. Create your crawler

Create your Scrapy project as you usually do. Enter a directory where you'd like to store your code and then run:

```
scrapy startproject tutorial
```

This will create a tutorial directory with the following contents:

```
tutorial/
    scrapy.cfg
    tutorial/
        __init__.py
        items.py
        pipelines.py
        settings.py
        spiders/
            __init__.py
            ...
```

These are basically:

- **scrapy.cfg**: the project configuration file
- **tutorial/**: the project's python module, you'll later import your code from here.
- **tutorial/items.py**: the project's items file.
- **tutorial/pipelines.py**: the project's pipelines file.
- **tutorial/settings.py**: the project's settings file.
- **tutorial/spiders/**: a directory where you'll later put your spiders.

1.1.2 2. Integrate your crawler with the frontier

Add the Scrapy Crawl Frontier middlewares to your settings:

```
SPIDER_MIDDLEWARES.update({
    'crawlfrontier.contrib.scrapy.middlewares.frontier.CrawlFrontierSpiderMiddleware': 1000,
})

DOWNLOADER_MIDDLEWARES.update({
    'crawlfrontier.contrib.scrapy.middlewares.frontier.CrawlFrontierDownloaderMiddleware': 1000,
})
```

Create a Crawl Frontier settings.py file and add it to your Scrapy settings:

```
FRONTIER_SETTINGS = 'tutorial/frontier/settings.py'
```

1.1.3 3. Choose your backend

Configure frontier settings to use a built-in backend like in-memory BFS:

```
BACKEND = 'crawlfrontier.contrib.backends.memory.heapq.BFS'
```

1.1.4 4. Run the spider

Run your Scrapy spider as usual from the command line:

```
scrapy crawl myspider
```

And that's it! You got your spider running integrated with Crawl Frontier.

1.1.5 What else?

You've seen a simple example of how to use Crawl Frontier with Scrapy, but this is just the surface. Crawl Frontier provides many powerful features for making Frontier management easy and efficient, such as:

- Easy *built-in integration with Scrapy* and *any other crawler* through its API.
- Creating different crawling logic/policies *defining your own backend*.
- Built-in support for *database storage* for crawled pages.
- Support for extending Crawl Frontier by plugging your own functionality using *middlewares*.
- Built-in middlewares for:
    - Extracting *domain info* from page URLs.
    - Create *unique fingerprints for page URLs* and *domain names*.
- Create fake sitemaps and reproduce crawling without crawler with the *graph Manager*.
- Tools for *easy frontier testing*.
- *Record your Scrapy crawls* and use it later for frontier testing.
- Logging facility that you can hook on to for catching errors and debug your frontiers.

1.1.6 What's next?

The next obvious steps are for you to *install Crawl Frontier*, and read the *architecture overview* and *API docs*. Thanks for your interest!

1.2 Installation Guide

The installation steps assume that you have the following things installed:

- Python 2.7

- pip and setuptools Python packages. Nowadays pip requires and installs setuptools if not installed.

You can install Crawl Frontier using pip (which is the canonical way to install Python packages).

To install using pip:

```
pip install crawl-frontier
```

Crawl Frontier at a glance Understand what Crawl Frontier is and how it can help you.

*Installation Guide* Get Crawl Frontier installed on your computer.

Basic concepts

2.1 What is a Crawl Frontier?

A crawl frontier is the part of a crawling system that decides the logic and policies to follow when a crawler is visiting websites (what pages should be crawled next, priorities and ordering, how often pages are revisited, etc).

A usual crawler-frontier scheme is:

The frontier is initialized with a list of start URLs, that we call the seeds. Once the frontier is initialized the crawler asks it what pages should be visited next. As the crawler starts to visit the pages and obtains results, it will inform the frontier of each page response and also of the extracted hyperlinks contained within the page. These links are added by the frontier as new requests to visit according to the frontier policies.

This process (ask for new requests/notify results) is repeated until the end condition for the crawl is reached. Some crawlers may never stop, that's what we call continuous crawls.

Frontier policies can be based in almost any logic. Common use cases are usually based in score/priority systems, computed from one or many page attributes (freshness, update times, content relevance for certain terms, etc). They can also be based in really simple logics as FIFO/LIFO or DFS/BFS page visit ordering.

Depending on frontier logic, a persistent storage system may be needed to store, update or query information about the pages. Other systems can be 100% volatile and not share any information at all between different crawls.

2.2 Architecture overview

This document describes the architecture of Crawl Frontier and how its components interact.

2.2.1 Overview

The following diagram shows an overview of the Crawl Frontier architecture with its components (referenced by numbers) and an outline of the data flow that takes place inside the system. A brief description of the components is included below with links for more detailed information about them. The data flow is also described below.
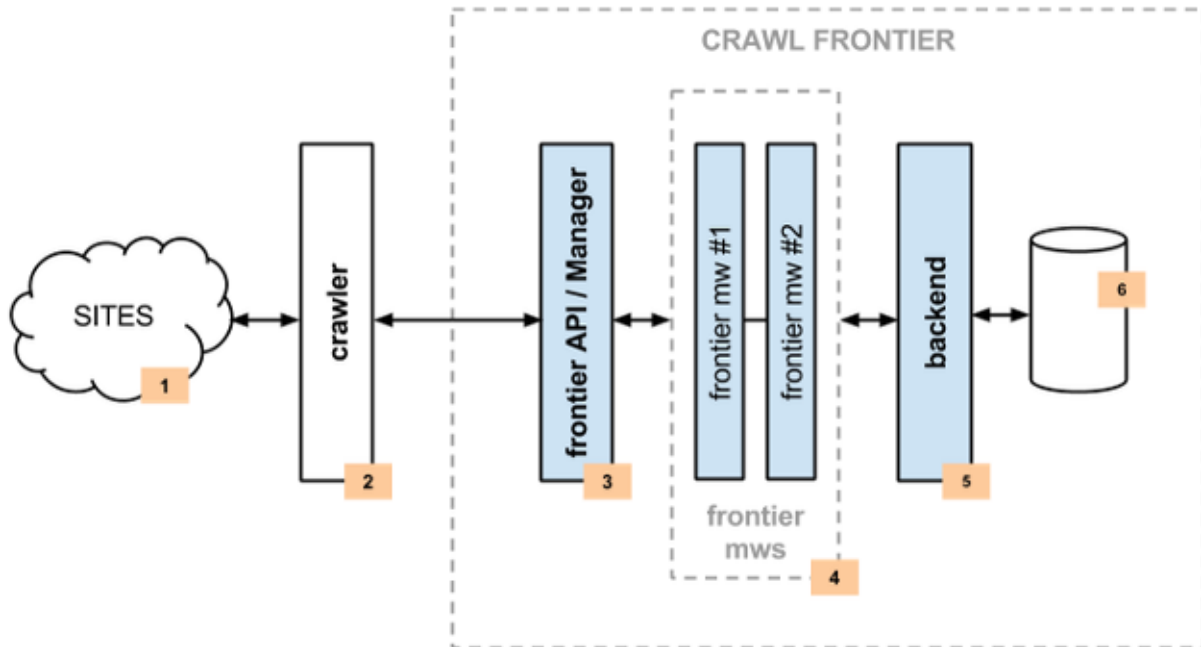
2.2.2 Components

Crawler

The Crawler (2) is responsible for fetching web pages from the sites (1) and feeding them to the frontier which manages what pages should be crawled next.

Crawler can be implemented using Scrapy or any other crawling framework/system as the framework offers a generic frontier functionality.

Frontier API / Manager

The main entry point to Crawl Frontier API (3) is the FrontierManager object. Frontier users, in our case the Crawler (2), will communicate with the frontier through it.

Communication with the frontier can also be done through other mechanisms such as an HTTP API or a queue system. These functionalities are not available for the time being, but hopefully will be in future versions.

For more information see *Frontier API*.

Middlewares

Frontier middlewares (4) are specific hooks that sit between the Manager (3) and the Backend (5). These middlewares process `Request` and `Response` objects when they pass to and from the Frontier and the Backend. They provide a convenient mechanism for extending functionality by plugging custom code.

For more information see *Middlewares*.

Backend

The frontier backend (5) is where the crawling logic/policies lies. It's responsible for receiving all the crawl info and selecting the next pages to be crawled.

May require, depending on the logic implemented, a persistent storage (6) to manage `Request` and `Response` objects info.

For more information see *Backends*.

2.2.3 Data Flow

The data flow in Crawl Frontier is controlled by the Frontier Manager, all data passes through the manager-middlewares-backend scheme and goes like this:

1. The frontier is initialized with a list of seed requests (seed URLs) as entry point for the crawl.

2. The crawler asks for a list of requests to crawl.

3. Each url is crawled and the frontier is notified back of the crawl result as well of the extracted links the page contains. If anything went wrong during the crawl, the frontier is also informed of it.

Once all urls have been crawled, steps 2-3 are repeated until crawl of frontier end condition is reached. Each loop (steps 2-3) repetition is called a *frontier iteration*.

2.3 Frontier objects

Frontier uses 2 object types: `Request` and `Response`. They are used to represent crawling HTTP requests and responses respectively.

These classes are used by most Frontier API methods either as a parameter or as a return value depending on the method used.

Frontier also uses these objects to internally communicate between different components (middlewares and backend).

2.3.1 Request objects

2.3.2 Response objects

Fields `domain` and `fingerprint` are added by *built-in middlewares*

2.3.3 Identifying unique objects

As frontier objects are shared between the crawler and the frontier, some mechanism to uniquely identify objects is needed. This method may vary depending on the frontier logic (in most cases due to the backend used).

By default, Crawl Frontier activates the *fingerprint middleware* to generate a unique fingerprint calculated from the `Request.url` and `Response.url` fields, which is added to the `Request.meta` and `Response.meta` fields respectively. You can use this middleware or implement your own method to manage frontier objects identification.

An example of a generated fingerprint for a `Request` object:

```
>>> request.url
'http://thehackernews.com'

>>> request.meta['fingerprint']
'198d99a8b2284701d6c147174cd69a37a7dea90f'
```

2.3.4 Adding additional data to objects

In most cases frontier objects can be used to represent the information needed to manage the frontier logic/policy.

Also, additional data can be stored by components using the `Request.meta` and `Response.meta` fields.

For instance the frontier *domain middleware* adds a `domain` info field for every `Request.meta` and `Response.meta` if is activated:

```
>>> request.url
'http://www.scrapinghub.com'

>>> request.meta['domain']
{
    "name": "scrapinghub.com",
    "netloc": "www.scrapinghub.com",
    "scheme": "http",
    "sld": "scrapinghub",
    "subdomain": "www",
    "tld": "com"
}
```

2.4 Frontier API

This section documents the Crawl Frontier core API, and is intended for developers of middlewares and backends.

2.4.1 Crawl Frontier API / Manager

The main entry point to Crawl Frontier API is the `FrontierManager` object, passed to middlewares and backend through the from_manager class method. This object provides access to all Crawl Frontier core components, and is the only way for middlewares and backend to access them and hook their functionality into Crawl Frontier.

The `FrontierManager` is responsible for loading the installed middlewares and backend, as well as for managing the data flow around the whole frontier.

2.4.2 Loading from settings

Although `FrontierManager` can be initialized using parameters the most common way of doing this is using *Frontier Settings*.

This can be done through the `from_settings` class method, using either a string path:

```
>>> from crawlfrontier import FrontierManager
>>> frontier = FrontierManager.from_settings('my_project.frontier.settings')
```

or a `Settings` object instance:

```
>>> from crawlfrontier import FrontierManager, Settings
>>> settings = Settings()
>>> settings.MAX_PAGES = 0
>>> frontier = FrontierManager.from_settings(settings)
```

It can also be initialized without parameters, in this case the frontier will use the *default settings*:

```
>>> from crawlfrontier import FrontierManager, Settings
>>> frontier = FrontierManager.from_settings()
```

2.4.3 Frontier Manager

2.4.4 Starting/Stopping the frontier

Sometimes, frontier components need to perform initialization and finalization operations. The frontier mechanism to notify the different components of the frontier start and stop is done by the `start()` and `stop()` methods respectively.

By default `auto_start` frontier value is activated, this means that components will be notified once the `FrontierManager` object is created. If you need to have more fine control of when different components are initialized, deactivate `auto_start` and manually call frontier API `start()` and `stop()` methods.

> **Note:** Frontier `stop()` method is not automatically called when `auto_start` is active (because frontier is not aware of the crawling state). If you need to notify components of frontier end you should call the method manually.

2.4.5 Frontier iterations

Once frontier is running, the usual process is the one described in the *data flow* section.

Crawler asks the frontier for next pages using the `get_next_requests()` method. Each time the frontier returns a non empty list of pages (data available), is what we call a frontier iteration.

Current frontier iteration can be accessed using the `iteration` attribute.

2.4.6 Finishing the frontier

Crawl can be finished either by the Crawler or by the Crawl Frontier. Crawl frontier will finish when a maximum number of pages are returned. This limit is controlled by the `max_requests` attribute (`MAX_REQUESTS` setting).

If `max_requests` has a value of 0 (default value) the frontier will continue indefinitely.

Once the frontier is finished, no more pages will be returned by the `get_next_requests` method and `finished` attribute will be True.

2.4.7 Component objects

2.4.8 Test mode

In some cases while testing, frontier components need to act in a different way than they usually do (for instance *domain middleware* accepts non valid URLs like `'A1'` or `'B1'` when parsing domain urls in test mode).

Components can know if the frontier is in test mode via the boolean `test_mode` attribute.

2.4.9 Another ways of using the frontier

Communication with the frontier can also be done through other mechanisms such as an HTTP API or a queue system. These functionalities are not available for the time being, but hopefully will be included in future versions.

2.5 Settings

The Crawl Frontier settings allows you to customize the behaviour of all components, including the `FrontierManager`, `Middleware` and `Backend` themselves.

The infrastructure of the settings provides a global namespace of key-value mappings that can be used to pull configuration values from. The settings can be populated through different mechanisms, which are described below.

For a list of available built-in settings see: *Built-in settings reference*.

2.5.1 Designating the settings

When you use Crawl Frontier, you have to tell it which settings you're using. As `FrontierManager` is the main entry point to Frontier usage, you can do this by using the method described in the *Loading from settings* section.

When using a string path pointing to a settings file for the frontier we propose the following directory structure:

```
my_project/
    frontier/
        __init__.py
        settings.py
        middlewares.py
        backends.py
    ...
```

These are basically:

- `frontier/settings.py`: the frontier settings file.
- `frontier/middlewares.py`: the middlewares used by the frontier.
- `frontier/backends.py`: the backend(s) used by the frontier.

2.5.2 How to access settings

`Settings` can be accessed through the `FrontierManager.settings` attribute, that is passed to `Middleware.from_manager` and `Backend.from_manager` class methods:

```
class MyMiddleware(Component):

    @classmethod
    def from_manager(cls, manager):
        manager = crawler.settings
        if settings.TEST_MODE:
            print "test mode is enabled!"
```

In other words, settings can be accessed as attributes of the `Settings` object.

2.5.3 Settings class

2.5.4 Built-in frontier settings

Here's a list of all available Crawl Frontier settings, in alphabetical order, along with their default values and the scope where they apply.

AUTO_START

Default: `True`

Whether to enable frontier automatic start. See *Starting/Stopping the frontier*

BACKEND

Default: `'crawlfrontier.contrib.backends.memory.FIFO'`

The `Backend` to be used by the frontier. For more info see *Activating a backend*.

EVENT_LOGGER

Default: `'crawlfrontier.logger.events.EventLogManager'`

The EventLoggerManager class to be used by the Frontier.

LOGGER

Default: `'crawlfrontier.logger.FrontierLogger'`

The Logger class to be used by the Frontier.

MAX_NEXT_REQUESTS

Default: `0`

The maximum number of requests returned by `get_next_requests` API method. If value is 0 (default), no maximum value will be used.

MAX_REQUESTS

Default: `0`

Maximum number of returned requests after which Crawl frontier is finished. If value is 0 (default), the frontier will continue indefinitely. See *Finishing the frontier*.

MIDDLEWARES

A list containing the middlewares enabled in the frontier. For more info see *Activating a middleware*.

Default:

```
[
    'crawlfrontier.contrib.middlewares.domain.DomainMiddleware',
    'crawlfrontier.contrib.middlewares.fingerprint.UrlFingerprintMiddleware',
    'crawlfrontier.contrib.middlewares.fingerprint.DomainFingerprintMiddleware',
]
```

REQUEST_MODEL

Default: `'crawlfrontier.core.models.Request'`

The `Request` model to be used by the frontier.

RESPONSE_MODEL

Default: `'crawlfrontier.core.models.Response'`

The `Response` model to be used by the frontier.

TEST_MODE

Default: `False`

Whether to enable frontier test mode. See *Frontier test mode*

2.5.5 Built-in fingerprint middleware settings

Settings used by the *UrlFingerprintMiddleware* and *DomainFingerprintMiddleware*.

URL_FINGERPRINT_FUNCTION

Default: `crawlfrontier.utils.fingerprint.sha1`

The function used to calculate the `url` fingerprint.

DOMAIN_FINGERPRINT_FUNCTION

Default: `crawlfrontier.utils.fingerprint.sha1`

The function used to calculate the `domain` fingerprint.

2.5.6 Default settings

If no settings are specified, frontier will use the built-in default ones. For a complete list of default values see: *Built-in settings reference*. All default settings can be overridden.

Frontier default settings

Values:

```
PAGE_MODEL = 'crawlfrontier.core.models.Page'
LINK_MODEL = 'crawlfrontier.core.models.Link'
FRONTIER = 'crawlfrontier.core.frontier.Frontier'
MIDDLEWARES = [
    'crawlfrontier.contrib.middlewares.domain.DomainMiddleware',
    'crawlfrontier.contrib.middlewares.fingerprint.UrlFingerprintMiddleware',
    'crawlfrontier.contrib.middlewares.fingerprint.DomainFingerprintMiddleware',
]
BACKEND = 'crawlfrontier.contrib.backends.memory.FIFO'
TEST_MODE = False
MAX_PAGES = 0
MAX_NEXT_PAGES = 0
AUTO_START = True
```

Fingerprints middleware default settings

Values:

```
URL_FINGERPRINT_FUNCTION = 'crawlfrontier.utils.fingerprint.sha1'
DOMAIN_FINGERPRINT_FUNCTION = 'crawlfrontier.utils.fingerprint.sha1'
```

Logging default settings

Values:

```
LOGGER = 'crawlfrontier.logger.FrontierLogger'
LOGGING_ENABLED = True

LOGGING_EVENTS_ENABLED = False
LOGGING_EVENTS_INCLUDE_METADATA = True
LOGGING_EVENTS_INCLUDE_DOMAIN = True
LOGGING_EVENTS_INCLUDE_DOMAIN_FIELDS = ['name', 'netloc', 'scheme', 'sld', 'tld', 'subdomain']
LOGGING_EVENTS_HANDLERS = [
    "crawlfrontier.logger.handlers.COLOR_EVENTS",
]

LOGGING_MANAGER_ENABLED = False
LOGGING_MANAGER_LOGLEVEL = logging.DEBUG
LOGGING_MANAGER_HANDLERS = [
    "crawlfrontier.logger.handlers.COLOR_CONSOLE_MANAGER",
]

LOGGING_BACKEND_ENABLED = False
LOGGING_BACKEND_LOGLEVEL = logging.DEBUG
LOGGING_BACKEND_HANDLERS = [
    "crawlfrontier.logger.handlers.COLOR_CONSOLE_BACKEND",
```

```
]

LOGGING_DEBUGGING_ENABLED = False
LOGGING_DEBUGGING_LOGLEVEL = logging.DEBUG
LOGGING_DEBUGGING_HANDLERS = [
    "crawlfrontier.logger.handlers.COLOR_CONSOLE_DEBUGGING",
]

EVENT_LOG_MANAGER = 'crawlfrontier.logger.events.EventLogManager'
```

What is a Crawl Frontier? Learn what a crawl frontier is and how to use it.

Architecture overview See how Crawl Frontier works and its different components.

Frontier objects Understand the classes used to represent links and pages.

Frontier API Learn how to use the frontier.

Settings See how to configure Crawl Frontier.

Extending Crawl Frontier

3.1 Middlewares

Frontier `Middleware` sits between `FrontierManager` and `Backend` objects, using hooks for `Request` and `Response` processing according to *frontier data flow*.

It's a light, low-level system for filtering and altering Frontier's requests and responses.

3.1.1 Activating a middleware

To activate a `Middleware` component, add it to the `MIDDLEWARES` setting, which is a list whose values can be class paths or instances of `Middleware` objects.

Here's an example:

```
MIDDLEWARES = [
    'crawlfrontier.contrib.middlewares.domain.DomainMiddleware',
]
```

Middlewares are called in the same order they've been defined in the list, to decide which order to assign to your middleware pick a value according to where you want to insert it. The order does matter because each middleware performs a different action and your middleware could depend on some previous (or subsequent) middleware being applied.

Finally, keep in mind that some middlewares may need to be enabled through a particular setting. See *each middleware documentation* for more info.

3.1.2 Writing your own middleware

Writing your own frontier backend is easy. Each `Middleware` component is a single Python class inherited from `Component`.

`FrontierManager` will communicate with all active middlewares through the methods described below.

3.1.3 Built-in middleware reference

This page describes all `Middleware` components that come with Crawl Frontier. For information on how to use them and how to write your own middleware, see the *middleware usage guide.*.

For a list of the components enabled by default (and their orders) see the `MIDDLEWARES` setting.

DomainMiddleware

UrlFingerprintMiddleware

DomainFingerprintMiddleware

3.2 Backends

Frontier `Backend` is where the crawling logic/policies lies. It's responsible for receiving all the crawl info and selecting the next pages to be crawled. It's called by the `FrontierManager` after `Middleware`, using hooks for `Request` and `Response` processing according to *frontier data flow*.

Unlike `Middleware`, that can have many different instances activated, only one `Backend` can be used per frontier.

Some backends require, depending on the logic implemented, a persistent storage to manage `Request` and `Response` objects info.

3.2.1 Activating a backend

To activate the frontier middleware component, set it through the `BACKEND` setting.

Here's an example:

```
BACKEND = 'crawlfrontier.contrib.backends.memory.FIFO'
```

Keep in mind that some backends may need to be enabled through a particular setting. See *each backend documentation* for more info.

3.2.2 Writing your own backend

Writing your own frontier backend is easy. Each `Backend` component is a single Python class inherited from `Component`.

`FrontierManager` will communicate with active `Backend` through the methods described below.

3.2.3 Built-in backend reference

This page describes all *each backend documentation* components that come with Crawl Frontier. For information on how to use them and how to write your own middleware, see the *backend usage guide.*.

To know the default activated `Backend` check the `BACKEND` setting.

Basic algorithms

Some of the built-in `Backend` objects implement basic algorithms as as FIFO/LIFO or DFS/BFS for page visit ordering.

Differences between them will be on storage engine used. For instance, `memory.FIFO` and `sqlalchemy.FIFO` will use the same logic but with different storage engines.

Memory backends

This set of `Backend` objects will use an heapq object as storage for *basic algorithms*.

class `crawlfrontier.contrib.backends.memory.`**BASE**
> Base class for in-memory heapq `Backend` objects.

class `crawlfrontier.contrib.backends.memory.`**FIFO**
> In-memory heapq `Backend` implementation of FIFO algorithm.

class `crawlfrontier.contrib.backends.memory.`**LIFO**
> In-memory heapq `Backend` implementation of LIFO algorithm.

class `crawlfrontier.contrib.backends.memory.`**BFS**
> In-memory heapq `Backend` implementation of BFS algorithm.

class `crawlfrontier.contrib.backends.memory.`**DFS**
> In-memory heapq `Backend` implementation of DFS algorithm.

class `crawlfrontier.contrib.backends.memory.`**RANDOM**
> In-memory heapq `Backend` implementation of a random selection algorithm.

SQLAlchemy backends

This set of `Backend` objects will use SQLAlchemy as storage for *basic algorithms*.

By default it uses an in-memory SQLite database as a storage engine, but any databases supported by SQLAlchemy can be used.

`Request` and `Response` are represented by a declarative sqlalchemy model:

```python
class Page(Base):
    __tablename__ = 'pages'
    __table_args__ = (
        UniqueConstraint('url'),
    )
    class State:
        NOT_CRAWLED = 'NOT CRAWLED'
        QUEUED = 'QUEUED'
        CRAWLED = 'CRAWLED'
        ERROR = 'ERROR'

    url = Column(String(1000), nullable=False)
    fingerprint = Column(String(40), primary_key=True, nullable=False, index=True, unique=True)
    depth = Column(Integer, nullable=False)
    created_at = Column(TIMESTAMP, nullable=False)
    status_code = Column(String(20))
    state = Column(String(10))
    error = Column(String(20))
```

If you need to create your own models, you can do it by using the `DEFAULT_MODELS` setting:

```python
DEFAULT_MODELS = {
    'Page': 'crawlfrontier.contrib.backends.sqlalchemy.models.Page',
}
```

This setting uses a dictionary where `key` represents the name of the model to define and `value` the model to use. If you want for instance to create a model to represent domains:

```
DEFAULT_MODELS = {
    'Page': 'crawlfrontier.contrib.backends.sqlalchemy.models.Page',
    'Domain': 'myproject.backends.sqlalchemy.models.Domain',
}
```

Models can be accessed from the Backend dictionary attribute `models`.

For a complete list of all settings used for sqlalchemy backends check the *settings* section.

class `crawlfrontier.contrib.backends.sqlalchemy.`**`BASE`**
    Base class for SQLAlchemy `Backend` objects.

class `crawlfrontier.contrib.backends.sqlalchemy.`**`FIFO`**
    SQLAlchemy `Backend` implementation of FIFO algorithm.

class `crawlfrontier.contrib.backends.sqlalchemy.`**`LIFO`**
    SQLAlchemy `Backend` implementation of LIFO algorithm.

class `crawlfrontier.contrib.backends.sqlalchemy.`**`BFS`**
    SQLAlchemy `Backend` implementation of BFS algorithm.

class `crawlfrontier.contrib.backends.sqlalchemy.`**`DFS`**
    SQLAlchemy `Backend` implementation of DFS algorithm.

class `crawlfrontier.contrib.backends.sqlalchemy.`**`RANDOM`**
    SQLAlchemy `Backend` implementation of a random selection algorithm.

Middlewares Filter or alter information for links and pages.

Backends Define your own crawling logic.

Built-in services and tools

4.1 Using the Frontier with Scrapy

Using Crawl Frontier is quite easy, it includes a set of Scrapy middlewares that encapsulates frontier usage and can be easily configured using Scrapy settings.

4.1.1 Activating the frontier

The frontier uses 2 different middlewares: `CrawlFrontierSpiderMiddleware` and `CrawlFrontierDownloaderMiddleware`.

To activate the frontier in your Scrapy project, just add them to the SPIDER_MIDDLEWARES and DOWN-LOADER_MIDDLEWARES settings:

```
SPIDER_MIDDLEWARES.update({
    'crawlfrontier.contrib.scrapy.middlewares.frontier.CrawlFrontierSpiderMiddleware': 1000,
})

DOWNLOADER_MIDDLEWARES.update({
    'crawlfrontier.contrib.scrapy.middlewares.frontier.CrawlFrontierDownloaderMiddleware': 1000,
})
```

Create a Crawl Frontier `settings.py` file and add it to your Scrapy settings:

```
FRONTIER_SETTINGS = 'tutorial/frontier/settings.py'
```

4.1.2 Organizing files

When using frontier with a Scrapy project, we propose the following directory structure:

```
my_scrapy_project/
    my_scrapy_project/
        frontier/
            __init__.py
            settings.py
            middlewares.py
            backends.py
        spiders/
            ...
        __init__.py
```

```
    settings.py
  scrapy.cfg
```

These are basically:

- `my_scrapy_project/frontier/settings.py`: the frontier settings file.

- `my_scrapy_project/frontier/middlewares.py`: the middlewares used by the frontier.

- `my_scrapy_project/frontier/backends.py`: the backend(s) used by the frontier.

- `my_scrapy_project/spiders`: the Scrapy spiders folder

- `my_scrapy_project/settings.py`: the Scrapy settings file

- `scrapy.cfg`: the Scrapy config file

4.1.3 Running the Crawl

Just run your Scrapy spider as usual from the command line:

```
scrapy crawl myspider
```

In case you need to disable frontier, you can do it by overriding the `FRONTIER_ENABLED` setting:

```
scrapy crawl myspider -s FRONTIER_ENABLED=False
```

4.1.4 Frontier Scrapy settings

Here's a list of all available Crawl Frontier Scrapy settings, in alphabetical order, along with their default values and the scope where they apply:

FRONTIER_ENABLED

Default: `True`

Whether to enable frontier in your Scrapy project.

FRONTIER_SCHEDULER_CONCURRENT_REQUESTS

Default: `256`

Number of concurrent requests that the middleware will maintain while asking for next pages.

FRONTIER_SCHEDULER_INTERVAL

Default: `0.01`

Interval of number of requests check in seconds. Indicates how often the frontier will be asked for new pages if there is gap for new requests.

FRONTIER_SETTINGS

Default: `None`

A file path pointing to Crawl Frontier settings.

4.2 Using the Frontier with Requests

To integrate frontier with Requests library, there is a `RequestsFrontierManager` class available.

This class is just a simple `FrontierManager` wrapper that uses Requests objects (`Request/Response`) and converts them from and to frontier ones for you.

Use it in the same way that `FrontierManager`, initialize it with your settings and use Requests `Request` and `Response` objects. `get_next_requests` method will return a Requests `Request` object.

An example:

```python
import re

import requests

from urlparse import urljoin

from crawlfrontier.contrib.requests.manager import RequestsFrontierManager
from crawlfrontier import Settings

SETTINGS = Settings()
SETTINGS.BACKEND = 'crawlfrontier.contrib.backends.memory.FIFO'
SETTINGS.LOGGING_MANAGER_ENABLED = True
SETTINGS.LOGGING_BACKEND_ENABLED = True
SETTINGS.MAX_REQUESTS = 100
SETTINGS.MAX_NEXT_REQUESTS = 10

SEEDS = [
    'http://www.imdb.com',
]

LINK_RE = re.compile(r'href="(.*?)"')


def extract_page_links(response):
    return [urljoin(response.url, link) for link in LINK_RE.findall(response.text)]

if __name__ == '__main__':

    frontier = RequestsFrontierManager(SETTINGS)
    frontier.add_seeds([requests.Request(url=url) for url in SEEDS])
    while True:
        next_requests = frontier.get_next_requests()
        if not next_requests:
            break
        for request in next_requests:
                try:
                    response = requests.get(request.url)
                    links = [requests.Request(url=url) for url in extract_page_links(response)]
                    frontier.page_crawled(response=response, links=links)
                except requests.RequestException, e:
```

```
                    error_code = type(e).__name__
                    frontier.request_error(request, error_code)
```

4.3 Graph Manager

The Graph Manager is a tool to represent web sitemaps as a graph.

It can easily be used to test frontiers. We can "fake" crawler request/responses by querying pages to the graph manager, and also know the links extracted for each one without using a crawler at all. You can make your own fake tests or use the *Frontier Tester tool*.

You can use it by defining your own sites for testing or use the *Scrapy Recorder* to record crawlings that can be reproduced later.

4.3.1 Defining a Site Graph

Pages from a web site and its links can be easily defined as a directed graph, where each node represents a page and the edges the links between them.

Let's use a really simple site representation with a starting page *A* that have links inside to tree pages *B, C, D*. We can represent the site with this graph:

We use a list to represent the different site pages and one tuple to define the page and its links, for the previous example:

```
site = [
    ('A', ['B', 'C', 'D']),
]
```

Note that we don't need to define pages without links, but we can also use it as a valid representation:

```
site = [
    ('A', ['B', 'C', 'D']),
    ('B', []),
    ('C', []),
    ('D', []),
]
```

A more complex site:

Can be represented as:

```
site = [
    ('A', ['B', 'C', 'D']),
    ('D', ['A', 'D', 'E', 'F']),
]
```

Note that *D* is linking to itself and to his parent *A*.

In the same way, a page can have several parents:

```
site = [
    ('A', ['B', 'C', 'D']),
    ('B', ['C']),
    ('D', ['C']),
]
```

In order to simplify examples we're not using urls for page representation, but of course urls are the intended use for site graphs:


```
site = [
    ('http://example.com', ['http://example.com/anotherpage', 'http://othersite.com']),
]
```

4.3.2 Using the Graph Manager

Once we have defined our site represented as a graph, we can start using it with the Graph Manager.

We must first create our graph manager:

```
>>> from crawlfrontier import graphs
>>> g = graphs.Manager()
```

And add the site using the *add_site* method:

```
>>> site = [('A', ['B', 'C', 'D'])]
>>> g.add_site(site)
```

The manager is now initialized and ready to be used.

We can get all the pages in the graph:

```
>>> g.pages
[<1:A*>, <2:B>, <3:C>, <4:D>]
```

Asterisk represents that the page is a seed, if we want to get just the seeds of the site graph:

```
>>> g.seeds
[<1:A*>]
```

We can get individual pages using *get_page*, if a page does not exists None is returned

```
>>> g.get_page('A')
<1:A*>

>>> g.get_page('F')
None
```

4.3.3 CrawlPage objects

Pages are represented as a `CrawlPage` object:

class **CrawlPage**

> A `CrawlPage` object represents an Graph Manager page, which is usually generated in the Graph Manager.
>
> **id**
> > Autonumeric page id.
>
> **url**
> > The url of the page.
>
> **status**
> > Represents the HTTP code status of the page.
>
> **is_seed**
> > Boolean value indicating if the page is seed or not.
>
> **links**
> > List of pages the current page links to.
>
> **referers**
> > List of pages that link to the current page.

In our example:

```
>>> p = g.get_page('A')
>>> p.id
1

>>> p.url
```

```
u'A'
```

```
>>> p.status   # defaults to 200
u'200'
```

```
>>> p.is_seed
True
```

```
>>> p.links
[<2:B>, <3:C>, <4:D>]
```

```
>>> p.referers   # No referers for A
[]
```

```
>>> g.get_page('B').referers   # referers for B
[<1:A*>]
```

4.3.4 Adding pages and Links

Site graphs can be also defined adding pages and links individually, the same graph from our example can be defined this way:

```
>>> g = graphs.Manager()
>>> a = g.add_page(url='A', is_seed=True)
>>> b = g.add_link(page=a, url='B')
>>> c = g.add_link(page=a, url='C')
>>> d = g.add_link(page=a, url='D')
```

add_page and *add_link* can be combined with *add_site* and used anytime:

```
>>> site = [('A', ['B', 'C', 'D'])]
>>> g = graphs.Manager()
>>> g.add_site(site)
>>> d = g.get_page('D')
>>> g.add_link(d, 'E')
```

4.3.5 Adding multiple sites

Multiple sites can be added to the manager:

```
>>> site1 = [('A1', ['B1', 'C1', 'D1'])]
>>> site2 = [('A2', ['B2', 'C2', 'D2'])]
```

```
>>> g = graphs.Manager()
>>> g.add_site(site1)
>>> g.add_site(site2)
```

```
>>> g.pages
[<1:A1*>, <2:B1>, <3:C1>, <4:D1>, <5:A2*>, <6:B2>, <7:C2>, <8:D2>]
```

```
>>> g.seeds
[<1:A1*>, <5:A2*>]
```

Or as a list of sites with *add_site_list* method:

```
>>> site_list = [
    [('A1', ['B1', 'C1', 'D1'])],
    [('A2', ['B2', 'C2', 'D2'])],
]
>>> g = graphs.Manager()
>>> g.add_site_list(site_list)
```

4.3.6 Graphs Database

Graph Manager uses SQLAlchemy to store and represent graphs.

By default it uses an in-memory SQLite database as a storage engine, but any databases supported by SQLAlchemy can be used.

An example using SQLite:

```
>>> g = graphs.Manager(engine='sqlite:///graph.db')
```

Changes are committed with every new add by default, graphs can be loaded later:

```
>>> graph = graphs.Manager(engine='sqlite:///graph.db')
>>> graph.add_site(('A', []))

>>> another_graph = graphs.Manager(engine='sqlite:///graph.db')
>>> another_graph.pages
[<1:A1*>]
```

A database content reset can be done using *clear_content* parameter:

```
>>> g = graphs.Manager(engine='sqlite:///graph.db', clear_content=True)
```

4.3.7 Using graphs with status codes

In order to recreate/simulate crawling using graphs, HTTP response codes can be defined for each page.

Example for a 404 error:

```
>>> g = graphs.Manager()
>>> g.add_page(url='A', status=404)
```

Status codes can be defined for sites in the following way using a list of tuples:

```
>>> site_with_status_codes = [
    ((200, "A"), ["B", "C"]),
    ((404, "B"), ["D", "E"]),
    ((500, "C"), ["F", "G"]),
]
>>> g = graphs.Manager()
>>> g.add_site(site_with_status_codes)
```

Default status code value is 200 for new pages.

4.3.8 A simple crawl faking example

Frontier tests can better be done using the *Frontier Tester tool*, but here's an example of how fake a crawl with a frontier:

```
from crawlfrontier import FrontierManager, graphs, Request, Response

if __name__ == '__main__':
    # Load graph from existing database
    graph = graphs.Manager('sqlite:///graph.db')

    # Create frontier from default settings
    frontier = FrontierManager.from_settings()

    # Create and add seeds
    seeds = [Request(seed.url) for seed in graph.seeds]
    frontier.add_seeds(seeds)

    # Get next requests
    next_requets = frontier.get_next_requests()

    # Crawl pages
    while (next_requests):
        for request in next_requests:

            # Fake page crawling
            crawled_page = graph.get_page(request.url)

            # Create response
            response = Response(url=crawled_page.url, status_code=crawled_page.status)

            # Update Page
            page = frontier.page_crawled(response=response
                                         links=[link.url for link in crawled_page.links])
            # Get next requests
            next_requets = frontier.get_next_requests()
```

4.3.9 Rendering graphs

Graphs can be rendered to png files:

```
>>> g.render(filename='graph.png', label='A simple Graph')
```

Rendering graphs uses pydot, a Python interface to Graphviz's Dot language.

4.3.10 How to use it

Graph Manager can be used to test frontiers in conjunction with *Frontier Tester* and also with *Scrapy Recordings*.

4.4 Testing a Frontier

Frontier Tester is a helper class for easy frontier testing.

Basically it runs a fake crawl against a Frontier, crawl info is faked using a *Graph Manager* instance.

4.4.1 Creating a Frontier Tester

FrontierTester needs a *Graph Manager* and a `FrontierManager` instances:

```
>>> from crawlfrontier import FrontierManager, FrontierTester, graphs
>>> graph = graphs.Manager('sqlite:///graph.db')  # Crawl fake data loading
>>> frontier = FrontierManager.from_settings()  # Create frontier from default settings
>>> tester = FrontierTester(frontier, graph)
```

4.4.2 Running a Test

The tester is now initialized, to run the test just call the method *run*:

```
>>> tester.run()
```

When run method is called the tester will:

1. Add all the seeds from the graph.

2. Ask the frontier about next pages.

3. Fake page response and inform the frontier about page crawl and its links.

Steps 1 and 2 are repeated until crawl or frontier ends.

Once the test is finished, the crawling page `sequence` is available as a list of frontier `Request` objects:

4.4.3 Test Parameters

In some test cases you may want to add all graph pages as seeds, this can be done with the parameter `add_all_pages`:

```
>>> tester.run(add_all_pages=True)
```

Maximum number of returned pages per `get_next_requests` call can be set using frontier settings, but also can be modified when creating the FrontierTester with the `max_next_pages` argument:

```
>>> tester = FrontierTester(frontier, graph, max_next_pages=10)
```

4.4.4 An example of use

A working example using test data from graphs and *basic backends*:

```
from crawlfrontier import FrontierManager, Settings, FrontierTester, graphs


def test_backend(backend):
    # Graph
    graph = graphs.Manager()
    graph.add_site_list(graphs.data.SITE_LIST_02)

    # Frontier
    settings = Settings()
    settings.BACKEND = backend
    settings.TEST_MODE = True
    frontier = FrontierManager.from_settings(settings)

    # Tester
    tester = FrontierTester(frontier, graph)
    tester.run(add_all_pages=True)
```

```
    # Show crawling sequence
    print '-'*40
    print frontier.backend.name
    print '-'*40
    for page in tester.sequence:
        print page.url

if __name__ == '__main__':
    test_backend('crawlfrontier.contrib.backends.memory.heapq.FIFO')
    test_backend('crawlfrontier.contrib.backends.memory.heapq.LIFO')
    test_backend('crawlfrontier.contrib.backends.memory.heapq.BFS')
    test_backend('crawlfrontier.contrib.backends.memory.heapq.DFS')
```

4.5 Recording a Scrapy crawl

Scrapy Recorder is a set of Scrapy middlewares that will allow you to record a scrapy crawl and store it into a *Graph Manager*.

This can be useful to perform frontier tests without having to crawl the entire site again or even using Scrapy.

4.5.1 Activating the recorder

The recorder uses 2 different middlewares: `CrawlRecorderSpiderMiddleware` and `CrawlRecorderDownloaderMiddleware`.

To activate the recording in your Scrapy project, just add them to the SPIDER_MIDDLEWARES and DOWN-LOADER_MIDDLEWARES settings:

```
SPIDER_MIDDLEWARES.update({
    'crawlfrontier.contrib.scrapy.middlewares.recording.CrawlRecorderSpiderMiddleware': 1000,
})

DOWNLOADER_MIDDLEWARES.update({
    'crawlfrontier.contrib.scrapy.middlewares.recording.CrawlRecorderDownloaderMiddleware': 1000,
})
```

4.5.2 Choosing your storage engine

As *Graph Manager* is internally used by the recorder to store crawled pages, you can choose between *different storage engines*.

We can set the storage engine with the `RECORDER_STORAGE_ENGINE` setting:

```
RECORDER_STORAGE_ENGINE = 'sqlite:///my_record.db'
```

You can also choose to reset database tables or just reset data with this settings:

```
RECORDER_STORAGE_DROP_ALL_TABLES = True
RECORDER_STORAGE_CLEAR_CONTENT = True
```

4.5.3 Running the Crawl

Just run your Scrapy spider as usual from the command line:

```
scrapy crawl myspider
```

Once it's finished you should have the recording available and ready for use.

In case you need to disable recording, you can do it by overriding the RECORDER_ENABLED setting:

```
scrapy crawl myspider -s RECORDER_ENABLED=False
```

4.5.4 Recorder settings

Here's a list of all available Scrapy Recorder settings, in alphabetical order, along with their default values and the scope where they apply.

RECORDER_ENABLED

Default: `True`

Activate or deactivate recording middlewares.

RECORDER_STORAGE_CLEAR_CONTENT

Default: `True`

Deletes table content from *storage database* in Graph Manager.

RECORDER_STORAGE_DROP_ALL_TABLES

Default: `True`

Drop *storage database* tables in Graph Manager.

RECORDER_STORAGE_ENGINE

Default: `None`

Sets *Graph Manager storage engine* used to store the recording.

4.6 Scrapy Seed Loaders

Crawl Frontier has some built-in Scrapy middlewares for seed loading.

Seed loaders use the `process_start_requests` method to generate requests from a source that are added later to the `FrontierManager`.

4.6.1 Activating a Seed loader

Just add the Seed Loader middleware to the `SPIDER_MIDDLEWARES` scrapy settings:

```
SPIDER_MIDDLEWARES.update({
    'crawl_frontier.contrib.scrapy.middlewares.seeds.FileSeedLoader': 650
})
```

4.6.2 FileSeedLoader

Load seed URLs from a file. The file must be formatted contain one URL per line:

```
http://www.asite.com
http://www.anothersite.com
...
```

Yo can disable URLs using the # character:

```
...
#http://www.acommentedsite.com
...
```

Settings:

- SEEDS_SOURCE: Path to the seeds file

4.6.3 S3SeedLoader

Load seeds from a file stored in an Amazon S3 bucket

File format should the same one used in *FileSeedLoader*.

Settings:

- SEEDS_SOURCE: Path to S3 bucket file. eg: `s3://some-project/seed-urls/`
- SEEDS_AWS_ACCESS_KEY: S3 credentials Access Key
- SEEDS_AWS_SECRET_ACCESS_KEY: S3 credentials Secret Access Key

Using the Frontier with Scrapy Learn how to use Crawl Frontier with Scrapy.

Using the Frontier with Requests Learn how to use Crawl Frontier with Requests.

Graph Manager Define fake crawlings for websites to test your frontier.

Testing a Frontier Test your frontier in an easy way.

Recording a Scrapy crawl Create Scrapy crawl recordings and reproduce them later.

Scrapy Seed Loaders Scrapy middlewares for seed loading

All the rest

5.1 Examples

The project repo includes an `examples` folder with some scripts and projects using CrawlFrontier:

```
examples/
    requests/
    scrapy_frontier/
    scrapy_recording/
    scripts/
```

- **requests**: Example script with Requests library.
- **scrapy_frontier**: Scrapy Frontier example project.
- **scrapy_recording**: Scrapy Recording example project.
- **scripts**: Some simple scripts.

Note: **This examples may need to install additional libraries in order to work**.

You can install them using pip:

```
pip install -r requirements/examples.txt
```

5.1.1 requests

A simple script that follow all the links from a site using Requests library.

How to run it:

```
python links_follower.py
```

5.1.2 scrapy_frontier

A simple script with a spider that follows all the links for the sites defined in a `seeds.txt` file.

How to run it:

```
scrapy crawl example
```

5.1.3 scrapy_recording

A simple script with a spider that follows all the links for a site, recording crawling results.

How to run it:

```
scrapy crawl recorder
```

5.1.4 scripts

Some sample scripts on how to use different frontier components.

5.2 Tests

Crawl Frontier tests are implemented using the pytest tool.

You can install pytest and the additional required libraries used in the tests using pip:

```
pip install -r requirements/tests.txt
```

5.2.1 Running tests

To run all tests go to the root directory of source code and run:

```
py.test
```

5.2.2 Writing tests

All functionality (including new features and bug fixes) must include a test case to check that it works as expected, so please include tests for your patches if you want them to get accepted sooner.

5.2.3 Backend testing

A base pytest class for `Backend` testing is provided: `BackendTest`

Let's say for instance that you want to test to your backend `MyBackend` and create a new frontier instance for each test method call, you can define a test class like this:

```python
class TestMyBackend(backends.BackendTest):

    backend_class = 'crawlfrontier.contrib.backend.abackend.MyBackend'

    def test_one(self):
        frontier = self.get_frontier()
        ...

    def test_two(self):
        frontier = self.get_frontier()
        ...

    ...
```

And let's say too that it uses a database file and you need to clean it before and after each test:

```python
class TestMyBackend(backends.BackendTest):

    backend_class = 'crawlfrontier.contrib.backend.abackend.MyBackend'

    def setup_backend(self, method):
        self._delete_test_db()

    def teardown_backend(self, method):
        self._delete_test_db()

    def _delete_test_db(self):
        try:
            os.remove('mytestdb.db')
        except OSError:
            pass

    def test_one(self):
        frontier = self.get_frontier()
        ...

    def test_two(self):
        frontier = self.get_frontier()
        ...

    ...
```

5.2.4 Testing backend sequences

To test `Backend` crawling sequences you can use the `BackendTest` class.

`BackendTest` class will run a complete crawl of the passed site graphs and return the sequence used by the backend for visiting the different pages.

Let's say you want to test to a backend that sort pages using alphabetic order. You can define the following test:

```python
class TestAlphabeticSortBackend(backends.BackendSequenceTest):

    backend_class = 'crawlfrontier.contrib.backend.abackend.AlphabeticSortBackend'

    SITE_LIST = [
        [
            ('C', []),
            ('B', []),
            ('A', []),
        ],
    ]

    def test_one(self):
        # Check sequence is the expected one
        self.assert_sequence(site_list=self.SITE_LIST,
                             expected_sequence=['A', 'B', 'C'],
                             max_next_requests=0)

    def test_two(self):
        # Get sequence and work with it
        sequence = self.get_sequence(site_list=SITE_LIST,
```

```
                                    max_next_requests=0)
        assert len(sequence) > 2


    ...
```

5.2.5 Testing basic algorithms

If your backend uses any of the *basic algorithms logics*, you can just inherit the correponding test base class for each logic and sequences will be automatically tested for it:

```python
from crawlfrontier.tests import backends


class TestMyBackendFIFO(backends.FIFOBackendTest):
    backend_class = 'crawlfrontier.contrib.backends.abackend.MyBackendFIFO'


class TestMyBackendLIFO(backends.LIFOBackendTest):
    backend_class = 'crawlfrontier.contrib.backends.abackend.MyBackendLIFO'


class TestMyBackendDFS(backends.DFSBackendTest):
    backend_class = 'crawlfrontier.contrib.backends.abackend.MyBackendDFS'


class TestMyBackendBFS(backends.BFSBackendTest):
    backend_class = 'crawlfrontier.contrib.backends.abackend.MyBackendBFS'


class TestMyBackendRANDOM(backends.RANDOMBackendTest):
    backend_class = 'crawlfrontier.contrib.backends.abackend.MyBackendRANDOM'
```

5.3 Release Notes

5.3.1 0.2.0 (released 2015-01-12)

- Added documentation (Scrapy Seed Loaders+Tests+Examples) (8e5f60d)
- Refactored backend tests (00910bf, 5702bef, 9567566)
- Added requests library example (8796011)
- Added requests library manager and object converters (d6590b6)
- Added FrontierManagerWrapper (4f04a48)
- Added frontier object converters (7da51a4)
- Fixed script examples for new changes (101ea27)
- Optional Color logging (only if available) (c0ba0ba)
- Changed Scrapy frontier and recorder integration to scheduler+middlewares (cbe5f4f / 2fcdc06 / f7bf02b / 0d15dc1)
- Changed default frontier backend (03cd307)
- Added comment support to seeds (7d48973)

- Added doc requirements for RTD build (27daea4)

- Removed optional dependencies for setup.py and requirements (c6099f3 / 79a4e4d / e6910e3)

- Changed tests to pytest (848d2bf / edc9c01 / c318d14)

- Updated docstrings and documentation (fdccd92 / 9dec38c / 71d626f / 0977bbf)

- Changed frontier componets (Backend and Middleware) to abc (1e74467)

- Modified Scrapy frontier example to use seed loaders (0ad905d)

- Refactored Scrapy Seed loaders (a0eac84)

- Added new fields to `Request` and `Response` frontier objects (bb64afb)

- Added `ScrapyFrontierManager` (Scrapy wrapper for Frontier Manager) (8e50dc0)

- Changed frontier core objects (`Page`/`Link` to `Request`/`Response`) (74b54c8)

5.3.2 0.1

First release of Crawl Frontier.

Examples Some example projects and scripts using Crawl Frontier.

Tests How to run and write Crawl Frontier tests.

Release Notes See what has changed in recent Crawl Frontier versions.