
Django Q2 Documentation

Release 1.6.0

Ilan Steemers

Oct 13, 2023

CONTENTS

1	Features	3
1.1	Installation	3
1.2	Configuration	6
1.3	Brokers	15
1.4	Tasks	18
1.5	Groups	28
1.6	Iterable	30
1.7	Chains	32
1.8	Schedules	34
1.9	Cluster	39
1.10	Monitor	42
1.11	Admin pages	46
1.12	Errors	48
1.13	Signals	48
1.14	Architecture	50
1.15	Examples	52
	Python Module Index	61
	Index	63

Django Q2 is a native Django task queue, scheduler and worker application using Python multiprocessing.

Note: Django Q2 is a fork of Django Q. Big thanks to Ilan Steemers for starting this project. Unfortunately, development of Django Q has stalled since June 2021. Django Q2 is the new updated version of Django Q, with dependencies updates, docs updates and several bug fixes. Original repository: <https://github.com/Koed00/django-q>

FEATURES

- Multiprocessing worker pools
- Asynchronous tasks
- Scheduled, cron and repeated tasks
- Signed and compressed packages
- Failure and success database or cache
- Result hooks, groups and chains
- Django Admin integration
- PaaS compatible with multiple instances
- Multi cluster monitor
- Redis, IronMQ, SQS, MongoDB or ORM
- Rollbar and Sentry support

Django Q2 is tested with: Python 3.8, 3.9, 3.10, 3.11 and 3.12. Works with Django 3.2.x, 4.1.x, 4.2.x and 5.0.x

Currently available in English, German and French.

Contents:

1.1 Installation

- Install the latest version with pip:

```
$ pip install django-q2
```

- Add *django_q* to `INSTALLED_APPS` in your projects `settings.py`:

```
INSTALLED_APPS = (  
    # other apps  
    'django_q',  
)
```

- Run Django migrations to create the database tables:

```
$ python manage.py migrate
```

- Choose a message *broker*, configure it and install the appropriate client library.

- Run Django Q2 cluster in order to handle tasks async:

```
$ python manage.py qcluster
```

1.1.1 Migrate from Django-Q to Django-Q2

If you have an application with django-q running right now, you can simply swap the libraries and you should be good to go.:

```
$ pip uninstall django-q # you might have to uninstall django-q add-ons as well
$ pip install django-q2
```

Then migrate the database to get the latest tables/fields:

```
$ python manage.py migrate
```

1.1.2 Requirements

Django Q2 is tested for Python 3.8, 3.9, 3.10, 3.11 and 3.12

- [Django](#)
Django Q2 aims to use as much of Django's standard offerings as possible. The code is tested against Django versions 3.2.x, 4.1.x, 4.2.x and 5.0.x.
- [Django-picklefield](#)
Used to store args, kwargs and result objects in the database.

Optional

- [Blessed](#) is used to display the statistics in the terminal:

```
$ pip install blessed
```

- [Redis-py](#) client by Andy McCurdy is used to interface with both the Redis:

```
$ pip install redis
```

- [Psutil](#) python system and process utilities module by Giampaolo Rodola', is an optional requirement and adds cpu affinity settings to the cluster:

```
$ pip install psutil
```

- [setproctitle](#) python module to customize the process title by Daniele Varrazzo', is an optional requirement used to set informative process titles:

```
$ pip install setproctitle
```

- [Hiredis](#) parser. This C library maintained by the core Redis team is faster than the standard PythonParser during high loads:

```
$ pip install hiredis
```


- [Boto3](#) is used for the Amazon SQS broker in favor of the now deprecating boto library:

```
$ pip install boto3
```

- [Iron-mq](#) is the official python binding for the IronMQ broker:

```
$ pip install iron-mq
```

- [Pymongo](#) is needed if you want to use MongoDB as a message broker:

```
$ pip install pymongo
```

- [Redis](#) server is the default broker for Django Q2. It provides the best performance and does not require Django's cache framework for monitoring.
- [MongoDB](#) is a highly scalable NoSQL database which makes for a very fast and reliably persistent at-least-once message broker. Usually available on most PaaS providers.
- [Pyrollbar](#) is an error notifier for [Rollbar](#) which lets you manage your worker errors in one place. Needs a [Rollbar](#) account and access key:

```
$ pip install rollbar
```

- [Croniter](#) is an optional package that is used to parse cron expressions for the scheduler:

```
$ pip install croniter
```

1.1.3 Add-ons

- [django-q-rollbar](#) is a Rollbar error reporter:

```
$ pip install django-q2[rollbar]
```

- [django-q-sentry](#) is a Sentry error reporter:

```
$ pip install django-q2[sentry]
```

- [django-q-email](#) is a compatible Django email backend that will automatically async queue your emails.

OS X

Running Django Q2 on OS X should work fine, except for the following known issues:

- `multiprocessing.Queue.qsize()` is not supported. This leads to the monitor not reporting the internal queue size of clusters running under OS X.
- CPU count through `multiprocessing.cpu_count()` does not work. Installing [psutil](#) provides Django Q2 with an alternative way of determining the number of CPU's on your system
- CPU affinity is provided by [psutil](#) which at this time does not support this feature on OSX. The code however is aware of this and will fake the CPU affinity assignment in the logs without actually assigning it. This way you can still develop with this setting.

Windows

The cluster and worker multiprocessing code depend on the OS's ability to fork, unfortunately forking is not supported under windows. You should however be able to develop and test without the cluster by setting the `sync` option to `True` in the configuration. This will run all `async` calls inline through a single cluster worker without the need for forking. Other known issues are:

- `os.getppid()` is only supported under windows since Python 3.2. If you use an older version you need to install [psutil](#) as an alternative.
- CPU count through `multiprocessing.cpu_count()` occasionally fails on servers. Installing [psutil](#) provides Django Q2 with an alternative way of determining the number of CPU's on your system
- The monitor and info commands rely on the Curses package which is not officially supported on windows. There are however some ports available like [this one](#) by Christoph Gohlke.

Python

Current tests are performed with 3.8, 3.9, 3.10, 3.11 and 3.12 If you do encounter any regressions with earlier versions, please submit an issue on [github](#)

Open-source packages

Django Q2 is always tested with the latest versions of the required and optional Python packages. We try to keep the dependencies as up to date as possible. You can reference the [requirements](#) file to determine which versions are currently being used for tests and development.

Django

We strive to be compatible with the last two major version of Django. At the moment this means we support the 3.2.x, 4.1.x, 4.2.x and 5.0.x releases.

Since we are now no longer supporting Python 2, we can also not support older versions of Django that do not support Python `>= 3.8` For this you can always use older releases, but they are no longer maintained.

1.2 Configuration

Configuration is handled via the `Q_CLUSTER` dictionary in your `settings.py`

```
# settings.py example
Q_CLUSTER = {
    'name': 'myproject',
    'workers': 8,
    'recycle': 500,
    'timeout': 60,
    'compress': True,
    'save_limit': 250,
    'queue_limit': 500,
    'cpu_affinity': 1,
    'label': 'Django Q2',
    'redis': {
        'host': '127.0.0.1',
```

(continues on next page)

(continued from previous page)

```
'port': 6379,
'db': 0, },
'ALT_CLUSTERS': {
    'long': {
        'timeout': 3000,
        'retry': 3600,
        'max_attempts': 2,
    },
    'short': {
        'timeout': 10,
        'max_attempts': 1,
    },
}
```

All configuration settings are optional:

1.2.1 name

Used to differentiate between projects using the same broker. On most broker types this will be used as the queue name. Defaults to 'default'.

Note: Tasks are signed. When a worker encounters a task with an invalid signature, it will be discarded or failed.

1.2.2 workers

The number of workers to use in the cluster. Defaults to CPU count of the current host, but can be set to a custom number.¹

1.2.3 daemonize_workers

Set the daemon flag when spawning workers. You may need to disable this flag if your worker needs to spawn child process but be careful with orphaned child processes in case of sudden termination of the main process. Defaults to True.

1.2.4 recycle

The number of tasks a worker will process before recycling . Useful to release memory resources on a regular basis. Defaults to 500.

¹ Uses `multiprocessing.cpu_count()` which can fail on some platforms. If so , please set the worker count in the configuration manually or install *psutil* to provide an alternative cpu count method.

1.2.5 max_rss

The maximum resident set size in kilobytes before a worker will recycle and release resources. Useful for limiting memory usage. Only supported on platforms that implement the python resource module or install the *psutil* module. Defaults to None.

1.2.6 timeout

The number of seconds a worker is allowed to spend on a task before it's terminated. Defaults to None, meaning it will never time out. Set this to something that makes sense for your project. Can be overridden for individual tasks.

See *retry* for details how to set values for timeout and retry.

1.2.7 time_zone

The timezone that is used for task scheduling. Use this if you are having issue with DST. The scheduler uses UTC to calculate the next date and will therefore ignore any DST changes. This will cause 1 hour or 0.5 hour changes in the schedule when time is moved one hour ahead or back. Defaults to *settings.TIME_ZONE* if *USE_TZ* is enabled.

1.2.8 ack_failures

When set to True, also acknowledge unsuccessful tasks. This causes failed tasks to be considered as successful deliveries, thereby removing them from the task queue. Can also be set per-task by passing the *ack_failure* option to *async_task()*. Defaults to False.

1.2.9 max_attempts

Limit the number of retry attempts for failed tasks. Set to 0 for infinite retries. Defaults to 0

1.2.10 retry

The number of seconds a broker will wait for a cluster to finish a task, before it's presented again. Only works with brokers that support delivery receipts. Defaults to 60.

The value must be bigger than the time it takes to complete longest task, i.e. *timeout* must be less than retry value and all tasks must complete in less time than the selected retry time. If this does not hold, i.e. the retry value is less than timeout or less than it takes to finish a task, Django-Q2 will start the task again if the used broker supports receipts.

For example, with the following code

```
# settings.py
Q_CLUSTER = {
    'retry': 5,
    'workers': 4,
    'orm': 'default',
}

# example.py

from django_q.tasks import async_task
```

(continues on next page)

(continued from previous page)

```
async_task('time.sleep', 22)
```

First, `time.sleep` is called by the first worker. After 5 seconds second worker will also call `time.sleep` because retry time has exceeded and the broker return the task again for the cluster. After 21 seconds from the call to `async_task` all four workers are running the `time.sleep(22)` call and there is one retry in queue; tasks are started after 0, 5, 10, 15 and 20 seconds after the `async_task` was called. After 22 seconds the first worker completes and the task is acknowledged in the broker and the task is not added to task queue anymore but the task that was already in the run queue will run also. So in this example, `time.sleep` was called 5 times.

Note also that the above issue might cause all workers to run the same long running task preventing new tasks from starting shortly after the task has been started by `async_task`. In this case the retry time handling could cause the task that has not been started by any worker to be put on work queue again (even multiple times).

1.2.11 compress

Compresses task packages to the broker. Useful for large payloads, but can add overhead when used with many small packages. Defaults to `False`

1.2.12 save_limit

Limits the amount of successful tasks saved to Django.

- Set to `0` for unlimited.
- Set to `-1` for no success storage at all.
- Defaults to `250`
- Failures are always saved.

1.2.13 save_limit_per

The above `save_limit` for successful tasks can be fine tuned per task type using

- Set to `"group"` to store the tasks per group
- Other possible values are `"func"`, `"name"`, `None`
- Defaults to `None`

1.2.14 guard_cycle

Guard loop sleep in seconds, must be greater than 0 and less than 60.

1.2.15 sync

When set to `True` this configuration option forces all `async_task()` calls to be run with `sync=True`. Effectively making everything synchronous. Useful for testing. Defaults to `False`.

1.2.16 queue_limit

This does not limit the amount of tasks that can be queued on the broker, but rather how many tasks are kept in memory by a single cluster. Setting this to a reasonable number, can help balance the workload and the memory overhead of each individual cluster. Defaults to `workers**2`.

1.2.17 label

The label used for the Django Admin page. Defaults to `'Django Q2'`

1.2.18 catch_up

The default behavior for schedules that didn't run while a cluster was down, is to play catch up and execute all the missed time slots until things are back on schedule. You can override this behavior by setting `catch_up` to `False`. This will make those schedules run only once when the cluster starts and normal scheduling resumes. Defaults to `True`.

1.2.19 redis

Connection settings for Redis. Defaults:

```
# redis defaults
Q_CLUSTER = {
    'redis': {
        'host': 'localhost',
        'port': 6379,
        'db': 0,
        'password': None,
        'socket_timeout': None,
        'charset': 'utf-8',
        'errors': 'strict',
        'unix_socket_path': None
    }
}
```

It's also possible to use a Redis connection URI:

```
Q_CLUSTER = {
    'redis': 'redis://h:asdfqwer1234asdf@ec2-111-1-1-1.compute-1.amazonaws.com:111'
}
```

For more information on these settings please refer to the [Redis-py](#) documentation

1.2.20 django_redis

If you are already using `django-redis` for your caching, you can take advantage of its excellent connection backend by supplying the name of the cache connection you want to use instead of a direct Redis connection:

```
# example django-redis connection
Q_CLUSTER = {
    'name': 'DJRedis',
    'workers': 4,
    'timeout': 90,
    'django_redis': 'default'
}
```

Tip: Django Q2 uses your `SECRET_KEY` to sign task packages and prevent task crossover. So make sure you have it set up in your Django settings.

1.2.21 iron_mq

Connection settings for IronMQ:

```
# example IronMQ connection
Q_CLUSTER = {
    'name': 'IronBroker',
    'workers': 8,
    'timeout': 30,
    'retry': 60,
    'queue_limit': 50,
    'bulk': 10,
    'iron_mq': {
        'host': 'mq-aws-us-east-1.iron.io',
        'token': 'Et1En7....0LuW39Q',
        'project_id': '500f7b....b0f302e9'
    }
}
```

All connection keywords are supported. See the `iron-mq` library for more info

1.2.22 sqs

To use Amazon SQS as a broker you need to provide the AWS region and credentials either via the config, or any other boto3 configuration method:

```
# example SQS broker connection
Q_CLUSTER = {
    'name': 'SQSExample',
    'workers': 4,
    'timeout': 60,
    'retry': 90,
```

(continues on next page)

(continued from previous page)

```
'queue_limit': 100,
'bulk': 5,
'sqs': {
    'aws_region': 'us-east-1', # optional
    'aws_access_key_id': 'ac-Idr....YwflZBaaxI', # optional
    'aws_secret_access_key': '500f7b....b0f302e9' # optional
}
```

Please make sure these credentials have proper SQS access.

Amazon SQS only supports a bulk setting between 1 and 10, with the total payload not exceeding 256kb.

1.2.23 orm

If you want to use Django's database backend as a message broker, set the `orm` keyword to the database connection you want it to use:

```
# example ORM broker connection

Q_CLUSTER = {
    'name': 'DjangORM',
    'workers': 4,
    'timeout': 90,
    'retry': 120,
    'queue_limit': 50,
    'bulk': 10,
    'orm': 'default'
}
```

Using the Django ORM backend will also enable the Queued Tasks table in the Admin.

If you need better performance, you should consider using a different database backend than the main project. Set `orm` to the name of that database connection and make sure you run migrations on it using the `--database` option.

1.2.24 mongo

To use MongoDB as a message broker you simply provide the connection information in a dictionary:

```
# example MongoDB broker connection

Q_CLUSTER = {
    'name': 'MongoDB',
    'workers': 8,
    'timeout': 60,
    'retry': 70,
    'queue_limit': 100,
    'mongo': {
        'host': '127.0.0.1',
        'port': 27017
    }
}
```


The `mongo` dictionary can contain any of the parameters exposed by `pymongo`'s `MongoClient`. If you want to use a `mongodb` uri, you can supply it as the `host` parameter.

1.2.25 mongo_db

When using the MongoDB broker you can optionally provide a database name to use for the queues. Defaults to default database if available, otherwise `django-q`

1.2.26 broker_class

You can use a custom broker class for your cluster workers:

```
# example Custom broker class connection

Q_CLUSTER = {
    'name': 'Custom',
    'workers': 8,
    'timeout': 60,
    'broker_class': 'myapp.broker.CustomBroker'
}
```

Make sure your `CustomBroker` class inherits from either the base `Broker` class or one of its children.

1.2.27 bulk

Sets the number of messages each cluster tries to get from the broker per call. Setting this on supported brokers can improve performance. Especially HTTP based or very high latency servers can benefit from bulk dequeue. Keep in mind however that settings this too high can degrade performance with multiple clusters or very large task packages.

Not supported by the default Redis broker. Defaults to 1.

1.2.28 poll

Sets the queue polling interval for database brokers that don't have a blocking call. Currently only affects the ORM and MongoDB brokers. Defaults to 0.2 (seconds).

1.2.29 cache

For some brokers, you will need to set up the Django `cache framework` to gather statistics for the monitor. You can indicate which cache to use by setting this value. Defaults to `default`.

1.2.30 cached

Switches all task and result functions from using the database backend to the cache backend. This is the same as setting the keyword `cached=True` on all task functions. Instead of a bool this can also be set to the number of seconds you want the cache to retain results. e.g. `cached=60`

1.2.31 scheduler

You can disable the scheduler by setting this option to `False`. This will reduce a little overhead if you're not using schedules, but is most useful if you want to temporarily disable all schedules. Defaults to `True`

1.2.32 error_reporter

You can redirect worker exceptions directly to various error reporters (for example [Rollbar](#) or [Sentry](#)) by installing Django Q2 with the necessary [extras](#).

To enable installed error reporters, you must provide the configuration settings required by an error reporter extension:

```
# error_reporter config--rollbar example
Q_CLUSTER = {
    'error_reporter': {
        'rollbar': {
            'access_token': '32we33a92a5224jiww8982',
            'environment': 'Django-Q2'
        }
    }
}
```

For more information on error reporters and developing error reporting plugins for Django Q2, see [errors](#).

1.2.33 cpu_affinity

Sets the number of processor each worker can use. This does not affect auxiliary processes like the sentinel or monitor and is only useful for tweaking the performance of very high traffic clusters. The affinity number has to be higher than zero and less than the total number of processors to have any effect. Defaults to using all processors:

```
# processor affinity example.

4 processors, 4 workers, cpu_affinity: 1

worker 1 cpu [0]
worker 2 cpu [1]
worker 3 cpu [2]
worker 4 cpu [3]

4 processors, 4 workers, cpu_affinity: 2

worker 1 cpu [0, 1]
worker 2 cpu [2, 3]
worker 3 cpu [0, 1]
worker 4 cpu [2, 3]
```

(continues on next page)

(continued from previous page)

```
8 processors, 8 workers, cpu_affinity: 3
```

```
worker 1 cpu [0, 1, 2]
worker 2 cpu [3, 4, 5]
worker 3 cpu [6, 7, 0]
worker 4 cpu [1, 2, 3]
worker 5 cpu [4, 5, 6]
worker 6 cpu [7, 0, 1]
worker 7 cpu [2, 3, 4]
worker 8 cpu [5, 6, 7]
```

In some cases, setting the `cpu affinity` for your workers can lead to performance improvements, especially if the load is high and consists of many repeating small tasks. Start with an affinity of 1 and work your way up. You will have to experiment with what works best for you. As a rule of thumb; `cpu_affinity 1` favors repetitive short running tasks, while no affinity benefits longer running tasks.

Note: The `cpu_affinity` setting requires the optional *psutil* module.

Psutil does not support cpu affinity on OS X at this time.

1.2.34 ALT_CLUSTERS

For multiple clusters working on multiple queues to run in one Django site. `ALT_CLUSTERS` should be a dict with `cluster_name` as its key, and the value is the configuration for the cluster with the key as its name. The configuration items are consistent with `Q_CLUSTER`, except for a few items such as `name/cluster_name/ALT_CLUSTER`, which are not available of course.

See `multiple-queues`.

Note: For a cluster, if its name is in `ALT_CLUSTERS`, the config item in `ALT_CLUSTER` will override the same config item in the `Q_CLUSTER` root. Other config items in `Q_CLUSTER` root remain in effect for this cluster.

1.3 Brokers

The broker sits between your Django instances and your Django Q2 cluster instances; accepting, saving and delivering task packages. Currently we support a variety of brokers.

The default Redis broker does not support message receipts. This means that in case of a catastrophic failure of the cluster server or worker timeouts, tasks that were being executed get lost. Keep in mind this is not the same as a failing task. If a tasks code crashes, this should only lead to a failed task status.

Even though this might be acceptable in some use cases, you might prefer brokers with message receipts support. These guarantee delivery by waiting for the cluster to send a receipt after the task has been processed. In case a receipt has not been received after a set time, the task package is put back in the queue. Django Q2 supports this behavior by setting the *retry* timer on brokers that support message receipts.

Some pointers:

- Don't set the *retry* timer to a lower or equal number than the task timeout.

- Retry time includes time the task spends waiting in the clusters internal queue.
- Don't set the `queue_limit` so high that tasks time out while waiting to be processed.
- In case a task is worked on twice, the task result will be updated with the latest results.
- In some rare cases a non-atomic broker will re-queue a task after it has been acknowledged.
- If a task runs twice and a previous run has succeeded, the new result will be discarded.
- Limiting the number of retries is handled globally in your actual broker's settings.

Support for more brokers is being worked on.

1.3.1 Redis

The default broker for Django Q2 clusters.

- Atomic
- Requires `Redis-py` client library: `pip install redis`
- Does not need cache framework for monitoring
- Does not support receipts
- Can use existing `django_redis` connections.
- Configure with `redis-py` compatible configuration

1.3.2 IronMQ

This HTTP based queue service is both available directly via [Iron.io](https://iron.io) and as an add-on on Heroku.

- Delivery receipts
- Supports bulk dequeue
- Needs Django's `Cache framework` configured for monitoring
- Requires the `iron-mq` client library: `pip install iron-mq`
- See the `iron_mq` configuration section for options.

1.3.3 Amazon SQS

Amazon's Simple Queue Service is another HTTP based message queue. Although `SQS` is not the fastest, it is stable, cheap and convenient if you already use AWS.

- Delivery receipts
- Maximum message size is 256Kb
- Supports bulk dequeue up to 10 messages with a maximum total size of 256Kb
- Needs Django's `Cache framework` configured for monitoring
- Requires the `boto3` client library: `pip install boto3`
- See the `sqs` configuration section for options.

1.3.4 MongoDB

This highly scalable NoSQL database makes for a very fast and reliably persistent at-least-once message broker. Usually available on most PaaS providers, as [open-source](#) or commercial [enterprise](#) edition.

- Delivery receipts
- Needs Django's [Cache framework](#) configured for monitoring
- Can be configured as the Django cache-backend through several open-source cache providers.
- Requires the [pymongo](#) driver: `pip install pymongo`
- See the [mongo](#) configuration section for options.

1.3.5 Django ORM

Select this to use Django's database backend as a message broker. Unless you have configured a dedicated database backend for it, this should probably not be your first choice for a high traffic setup. However for a medium message rate and scheduled tasks, this is the most convenient guaranteed delivery broker.

- Delivery receipts
- Supports bulk dequeue
- Needs Django's [Cache framework](#) configured for monitoring
- Can be [configured](#) as its own cache backend.
- Queue editable in Django Admin
- See the [orm](#) configuration on how to set it up.

1.3.6 Custom Broker

You can override the [Broker](#) or any of its existing derived broker types.

```
# example Custom broker.py
from django_q.brokers import Broker

class CustomBroker(Broker):
    def info(self):
        return 'My Custom Broker'
```

Using the [broker_class](#) configuration setting you can then instruct Django Q2 to use this instead of one of the existing brokers:

```
# example Custom broker class connection

Q_CLUSTER = {
    'name': 'Custom',
    'workers': 8,
    'timeout': 60,
    'broker_class': 'myapp.broker.CustomBroker'
}
```

If you do write a custom broker for one of the many message queueing servers out there we don't support yet, please consider contributing it to the project.

1.3.7 Reference

The *Broker* class is used internally to communicate with the different types of brokers. You can override this class if you want to contribute and support your own broker.

class Broker

async_task(task)

Sends a task package to the broker queue and returns a tracking id if available.

dequeue()

Gets packages from the broker and returns a list of tuples with a tracking id and the package.

acknowledge(id)

Notifies the broker that the task has been processed. Only works with brokers that support delivery receipts.

fail(id)

Tells the broker that the message failed to be processed by the cluster. Only available on brokers that support this. Currently only occurs when a cluster fails to unpack a task package.

delete(id)

Instructs the broker to delete this message from the queue.

purge_queue()

Empties the current queue of all messages.

delete_queue()

Deletes the current queue from the broker.

queue_size()

Returns the amount of messages in the brokers queue.

lock_size()

Optional method that returns the number of messages currently awaiting acknowledgement. Only implemented on brokers that support it.

ping()

Returns True if the broker can be reached.

info()

Shows the name and version of the currently configured broker.

brokers.get_broker()

Returns a *Broker* instance based on the current configuration.

1.4 Tasks

1.4.1 async_task()

Use *async_task()* from your code to quickly offload tasks to the *Cluster*:

```
from django_q.tasks import async_task, result

# create the task
async_task('math.copysign', 2, -2)
```

(continues on next page)

(continued from previous page)

```
# or with import and storing the id
import math.copysign

task_id = async_task(copysign, 2, -2)

# get the result
task_result = result(task_id)

# result returns None if the task has not been executed yet
# you can wait for it
task_result = result(task_id, 200)

# but in most cases you will want to use a hook:

async_task('math.modf', 2.5, hook='hooks.print_result')

# hooks.py
def print_result(task):
    print(task.result)
```

`async_task()` can take the following optional keyword arguments:

hook

The function to call after the task has been executed. This function gets passed the complete [Task](#) object as its argument.

group

A group label. Check [Groups](#) for group functions.

save

Overrides the result backend's save setting for this task.

timeout

Overrides the cluster's timeout setting for this task.

See [retry](#) for details how to set values for timeout.

ack_failure

Overrides the cluster's *ack_failures* setting for this task.

sync

Simulates a task execution synchronously. Useful for testing. Can also be forced globally via the *sync* configuration option.

cached

Redirects the result to the cache backend instead of the database if set to `True` or to an integer indicating the cache timeout in seconds. e.g. `cached=60`. Especially useful with large and group operations where you don't need the all results in your database and want to take advantage of the speed of your cache backend.

broker

A broker instance, in case you want to control your own connections.

task_name

Optionally overwrites the auto-generated task name.

q_options

None of the option keywords get passed on to the task function. As an alternative you can also put them in a single keyword dict named `q_options`. This enables you to use these keywords for your function call:

```
# Async options in a dict

opts = {'hook': 'hooks.print_result',
        'group': 'math',
        'timeout': 30}

async_task('math.modf', 2.5, q_options=opts)
```

Please note that this will override any other option keywords.

Note: For tasks to be processed you will need to have a worker cluster running in the background using `python manage.py qcluster` or you need to configure Django Q2 to run in synchronous mode for testing using the *sync* option.

1.4.2 AsyncTask

Optionally you can use the `AsyncTask` class to instantiate a task and keep everything in a single object.:

```
# AsyncTask class instance example
from django_q.tasks import AsyncTask

# instantiate an async task
a = AsyncTask('math.floor', 1.5, group='math')

# you can set or change keywords afterwards
a.cached = True

# run it
a.run()

# wait indefinitely for the result and print it
# don't let the task return `None` or it will wait indefinitely
print(a.result(wait=-1))

# change the args
a.args = (2.5,)

# run it again
a.run()

# wait max 10 seconds for the result and print it

print(a.result(wait=10))
```

```
1
2
```

Once you change any of the parameters of the task after it has run, the result is invalidated and you will have to `AsyncTask.run()` it again to retrieve a new result.

1.4.3 Cached operations

You can run your tasks results against the Django cache backend instead of the database backend by either using the global `cached` setting or by supplying the `cached` keyword to individual functions. This can be useful if you are not interested in persistent results or if you run large group tasks where you only want the final result. By using a cache backend like Redis or Memcached you can speed up access to your task results significantly compared to a relational database.

When you set `cached=True`, results will be saved permanently in the cache and you will have to rely on your backend's cleanup strategies (like LRU) to manage stale results. You can also opt to set a manual timeout on the results, by setting e.g. `cached=60`. Meaning the result will be evicted from the cache after 60 seconds. This works both globally or on individual async executions.:

```
# simple cached example
from django_q.tasks import async_task, result

# cache the result for 10 seconds
```

(continues on next page)

(continued from previous page)

```
id = async_task('math.floor', 100, cached=10)

# wait max 50ms for the result to appear in the cache
result(id, wait=50, cached=True)

# or fetch the task object
task = fetch(id, cached=True)

# and then save it to the database
task.save()
```

As you can see you can easily turn a cached result into a permanent database result by calling `save()` on it.

This also works for group actions:

```
# cached group example
from django_q.tasks import async_task, result_group
from django_q.brokers import get_broker

# set up a broker instance for better performance
broker = get_broker()

# Async a hundred functions under a group label
for i in range(100):
    async_task('math.frexp',
               i,
               group='frexp',
               cached=True,
               broker=broker)

# wait max 50ms for one hundred results to return
result_group('frexp', wait=50, count=100, cached=True)
```

If you don't need hooks, that exact same result can be achieved by using the more convenient `async_iter()`.

1.4.4 Synchronous testing

`async_task()` can be instructed to execute a task immediately by setting the optional keyword `sync=True`. The task will then be injected straight into a worker and the result saved by a monitor instance:

```
from django_q.tasks import async_task, fetch

# create a synchronous task
task_id = async_task('my.buggy.code', sync=True)

# the task will then be available immediately
task = fetch(task_id)

# and can be examined
if not task.success:
    print('An error occurred: {}'.format(task.result))
```

```
An error occurred: ImportError("No module named 'my'",)
```

Note that `async_task()` will block until the task is executed and saved. This feature bypasses the broker and is intended for debugging and development. Instead of setting `sync` on each individual `async_task` you can also configure `sync` as a global override.

1.4.5 Connection pooling

Django Q2 tries to pass broker instances around its parts as much as possible to save you from running out of connections. When you are making individual calls to `async_task()` a lot though, it can help to set up a broker to reuse for `async_task()`:

```
# broker connection economy example
from django_q.tasks import async_task
from django_q.brokers import get_broker

broker = get_broker()
for i in range(50):
    async_task('math.modf', 2.5, broker=broker)
```

Tip: If you are using `django-redis` and the redis broker, you can *configure* Django Q2 to use its connection pool.

1.4.6 Reference

async_task(*func*, **args*, *hook*=None, *group*=None, *timeout*=None, *save*=None, *sync*=False, *cached*=False, *broker*=None, *q_options*=None, ***kwargs*)

Puts a task in the cluster queue

Parameters

- **func** (*object*) – The task function to execute
- **args** (*tuple*) – The arguments for the task function
- **hook** (*object*) – Optional function to call after execution
- **group** (*str*) – An optional group identifier
- **timeout** (*int*) – Overrides global cluster *timeout*.
- **save** (*bool*) – Overrides global save setting for this task.
- **ack_failure** (*bool*) – Overrides the global *ack_failures* setting for this task.
- **sync** (*bool*) – If set to True, `async_task` will simulate a task execution
- **cached** – Output the result to the cache backend. Bool or timeout in seconds
- **broker** – Optional broker connection from `brokers.get_broker()`
- **q_options** (*dict*) – Options dict, overrides option keywords
- **kwargs** (*dict*) – Keyword arguments for the task function

Returns

The uuid of the task

Return type

`str`

result(*task_id*, *wait*=0, *cached*=False)

Gets the result of a previously executed task

Parameters

- **task_id** (`str`) – the uuid or name of the task
- **wait** (`int`) – optional milliseconds to wait for a result. -1 for indefinite, but be sure the result will not be *None* otherwise it will wait indefinitely!
- **cached** (`bool`) – run this against the cache backend.

Returns

The result of the executed task

fetch(*task_id*, *wait*=0, *cached*=False)

Returns a previously executed task

Parameters

- **task_id** (`str`) – the uuid or name of the task
- **wait** (`int`) – optional milliseconds to wait for a result. -1 for indefinite
- **cached** (`bool`) – run this against the cache backend.

Returns

A task object

Return type

`Task`

Changed in version 0.2.0.

Renamed from `get_task`

queue_size()

Returns the size of the broker queue. Note that this does not count tasks currently being processed.

Returns

The amount of task packages in the broker

Return type

`int`

delete_cached(*task_id*, *broker*=None)

Deletes a task from the cache backend

Parameters

- **task_id** (`str`) – the uuid of the task
- **broker** – an optional broker instance

class Task

Database model describing an executed task

id

An `uuid.uuid4()` identifier

name

The name of the task as a humanized version of the *id*

Note: This is for convenience and can be used as a parameter for most functions that take a *task_id*. Keep in mind that it is not guaranteed to be unique if you store very large amounts of tasks in the database.

func

The function or reference that was executed

hook

The function to call after execution.

args

Positional arguments for the function.

kwargs

Keyword arguments for the function.

result

The result object. Contains the error if any occur.

started

The moment the task was created by an async command

stopped

The moment a worker finished this task

success

Was the task executed without problems?

time_taken()

Calculates the difference in seconds between started and stopped.

Note: Time taken represents the time a task spends in the cluster, this includes any time it may have waited in the queue.

group_result(*failures=False*)

Returns a list of results from this task's group. Set failures to True to include failed results.

group_count(*failures=False*)

Returns a count of the number of task results in this task's group. Returns the number of failures when failures=True

group_delete(*tasks=False*)

Resets the group label on all the tasks in this task's group. If *tasks=True* it will also delete the tasks in this group from the database, including itself.

classmethod **get_result**(*task_id*)

Gets a result directly by task uuid or name.

classmethod **get_result_group**(*group_id, failures=False*)

Returns a list of results from a task group. Set failures to *True* to include failed results.

classmethod **get_task**(*task_id*)

Fetches a single task object by uuid or name.

classmethod **get_task_group**(*group_id, failures=True*)

Gets a queryset of tasks with this group id. Set failures to *False* to exclude failed tasks.

classmethod **get_group_count**(*group_id, failures=False*)

Returns a count of the number of tasks results in a group. Returns the number of failures when *failures=True*

classmethod **delete_group**(*group_id, objects=False*)

Deletes a group label only, by default. If *objects=True* it will also delete the tasks in this group from the database.

class **Success**

A proxy model of *Task* with the queryset filtered on *Task.success* is *True*.

class **Failure**

A proxy model of *Task* with the queryset filtered on *Task.success* is *False*.

class **AsyncTask**(*func, *args, **kwargs*)

A class wrapper for the *async_task()* function.

Parameters

- **func** (*object*) – The task function to execute
- **args** (*tuple*) – The arguments for the task function
- **kwargs** (*dict*) – Keyword arguments for the task function, including *async_task* options

id

The task unique identifier. This will only be available after it has been *run()*

started

Bool indicating if the task has been run with the current parameters

func

The task function to execute

args

A tuple of arguments for the task function

kwargs

Keyword arguments for the function. Can include any of the optional `async_task` keyword attributes directly or in a `q_options` dictionary.

broker

Optional *Broker* instance to use

sync

Run this task inline instead of asynchronous.

save

Overrides the global save setting.

hook

Optional function to call after a result is available. Takes the result *Task* as the first argument.

group

Optional group identifier

cached

Run the task against the cache result backend.

run()

Send the task to a worker cluster for execution

result(*wait=0*)

The task result. Always returns None if the task hasn't been run with the current parameters.

param int wait

the number of milliseconds to wait for a result. -1 for indefinite

fetch(*wait=0*)

Returns the full *Task* result instance.

param int wait

the number of milliseconds to wait for a result. -1 for indefinite

result_group(*failures=False, wait=0, count=None*)

Returns a list of results from this task's group.

param bool failures

set this to True to include failed results

param int wait

optional milliseconds to wait for a result or count. -1 for indefinite

param int count

block until there are this many results in the group

fetch_group(*failures=True, wait=0, count=None*)

Returns a list of task results from this task's group

param bool failures

set this to False to exclude failed tasks

param int wait

optional milliseconds to wait for a task or count. -1 for indefinite

param int count

block until there are this many tasks in the group

1.5 Groups

You can group together results by passing `async_task()` the optional `group` keyword:

```
# result group example
from django_q.tasks import async_task, result_group

for i in range(4):
    async_task('math.modf', i, group='modf')

# wait until the group has 4 results
result = result_group('modf', count=4)
print(result)
```

```
[(0.0, 0.0), (0.0, 1.0), (0.0, 2.0), (0.0, 3.0)]
```

Note that this particular example can be achieved much faster with *Iterable*

Take care to not limit your results database too much and call `delete_group()` before each run, unless you want your results to keep adding up. Instead of `result_group()` you can also use `fetch_group()` to return a queryset of *Task* objects.:

```
# fetch group example
from django_q.tasks import fetch_group, count_group, result_group

# count the number of failures
failure_count = count_group('modf', failures=True)

# only use the successes
results = fetch_group('modf')
if failure_count:
    results = results.exclude(success=False)
results = [task.result for task in successes]

# this is the same as
results = fetch_group('modf', failures=False)
results = [task.result for task in successes]

# and the same as
results = result_group('modf') # filters failures by default
```

Getting results by using `result_group()` is of course much faster than using `fetch_group()`, but it doesn't offer the benefits of Django's queryset functions.

You can also access group functions from a task result instance:


```

from django_q.tasks import fetch

task = fetch('winter-speaker-alpha-ceiling')
if task.group_count() > 100:
    print(task.group_result())
    task.group_delete()
    print('Deleted group {}'.format(task.group))

```

or call them directly on *AsyncTask* object:

```

from django_q.tasks import AsyncTask

# add a task to the math group and run it cached
a = AsyncTask('math.floor', 2.5, group='math', cached=True)

# wait until this tasks group has 10 results
result = a.result_group(count=10)

```

1.5.1 Reference

result_group(*group_id*, *failures=False*, *wait=0*, *count=None*, *cached=False*)

Returns the results of a task group

Parameters

- **group_id** (*str*) – the group identifier
- **failures** (*bool*) – set this to *True* to include failed results
- **wait** (*int*) – optional milliseconds to wait for a result or count. -1 for indefinite
- **count** (*int*) – block until there are this many results in the group
- **cached** (*bool*) – run this against the cache backend

Returns

a list of results

Return type

list

fetch_group(*group_id*, *failures=True*, *wait=0*, *count=None*, *cached=False*)

Returns a list of tasks in a group

Parameters

- **group_id** (*str*) – the group identifier
- **failures** (*bool*) – set this to *False* to exclude failed tasks
- **wait** (*int*) – optional milliseconds to wait for a task or count. -1 for indefinite
- **count** (*int*) – block until there are this many tasks in the group
- **cached** (*bool*) – run this against the cache backend.

Returns

a list of *Task*

Return type`list`**count_group**(*group_id*, *failures=False*, *cached=False*)

Counts the number of task results in a group.

Parameters

- **group_id** (*str*) – the group identifier
- **failures** (*bool*) – counts the number of failures if True
- **cached** (*bool*) – run this against the cache backend.

Returns

the number of tasks or failures in a group

Return type`int`**delete_group**(*group_id*, *tasks=False*, *cached=False*)

Deletes a group label from the database.

Parameters

- **group_id** (*str*) – the group identifier
- **tasks** (*bool*) – also deletes the associated tasks if True
- **cached** (*bool*) – run this against the cache backend.

Returns

the numbers of tasks affected

Return type`int`

1.6 Iterable

If you have an iterable object with arguments for a function, you can use `async_iter()` to async them with a single command:

```
# Async Iterable example
from django_q.tasks import async_iter, result

# set up a list of arguments for math.floor
iter = [i for i in range(100)]

# async_task iter them
id=async_iter('math.floor',iter)

# wait for the collated result for 1 second
result_list = result(id, wait=1000)
```

This will individually queue 100 tasks to the worker cluster, which will save their results in the cache backend for speed. Once all the 100 results are in the cache, they are collated into a list and saved as a single result in the database. The cache results are then cleared.

You can also use an `Iter` instance which can sometimes be more convenient:

```

from django_q.tasks import Iter

i = Iter('math.copysign')

# add some arguments
i.append(1, -1)
i.append(2, -1)
i.append(3, -1)

# run it
i.run()

# get the results
print(i.result())

```

```
[-1.0, -2.0, -3.0]
```

1.6.1 Reference

async_iter(*func*, *args_iter*, ***kwargs*)

Runs iterable arguments against the cache backend and returns a single collated result. Accepts the same options as [async_task\(\)](#) except hook. See also the [Iter](#) class.

Parameters

- **func** (*object*) – The task function to execute
- **args** – An iterable containing arguments for the task function
- **kwargs** (*dict*) – Keyword arguments for the task function. Ignores hook.

Returns

The uuid of the task

Return type

`str`

class Iter(*func=None*, *args=None*, *kwargs=None*, *cached=Conf.CACHED*, *sync=Conf.SYNC*, *broker=None*)

An async task with iterable arguments. Serves as a convenient wrapper for [async_iter\(\)](#) You can pass the iterable arguments at construction or you can append individual argument tuples.

param func

the function to execute

param args

an iterable of arguments.

param kwargs

the keyword arguments

param bool cached

run this against the cache backend

param bool sync

execute this inline instead of asynchronous

param broker

optional broker instance

append(*args)

Append arguments to the iter set. Returns the current set count.

param args
the arguments for a single execution

return
the current set count

rtype
int

run()

Start queueing the tasks to the worker cluster.

return
the task result id

result(wait=0)

return the full list of results.

param int wait
how many milliseconds to wait for a result

return
an unsorted list of results

fetch(wait=0)

get the task result objects.

param int wait
how many milliseconds to wait for a result

return
an unsorted list of task objects

length()

get the length of the arguments list

return int
length of the argument list

1.7 Chains

Sometimes you want to run tasks sequentially. For that you can use the `async_chain()` function:

```
# async a chain of tasks
from django_q.tasks import async_chain, result_group

# the chain must be in the format
# [(func, (args), {kwargs}), (func, (args), {kwargs}), ...]
group_id = async_chain([('math.copysign', (1, -1)),
                       ('math.floor', (1,))])
```

(continues on next page)

(continued from previous page)

```
# get group result
result_group(group_id, count=2)
```

A slightly more convenient way is to use a [Chain](#) instance:

```
# Chain async
from django_q.tasks import Chain

# create a chain that uses the cache backend
chain = Chain(cached=True)

# add some tasks
chain.append('math.copysign', 1, -1)
chain.append('math.floor', 1)

# run it
chain.run()

print(chain.result())
```

```
[-1.0, 1]
```

1.7.1 Reference

async_chain(*chain*, *group=None*, *cached=Conf.CACHED*, *sync=Conf.SYNC*, *broker=None*)

Async a chain of tasks. See also the [Chain](#) class.

Parameters

- **chain** (*list*) – a list of tasks in the format [(func,(args),{kwargs}), (func,(args),{kwargs})]
- **group** (*str*) – an optional group name.
- **cached** (*bool*) – run this against the cache backend
- **sync** (*bool*) – execute this inline instead of asynchronous

class Chain(*chain=None*, *group=None*, *cached=Conf.CACHED*, *sync=Conf.SYNC*)

A sequential chain of tasks. Acts as a convenient wrapper for [async_chain\(\)](#) You can pass the task chain at construction or you can append individual tasks before running them.

param list chain

a list of task in the format [(func,(args),{kwargs}), (func,(args),{kwargs})]

param str group

an optional group name.

param bool cached

run this against the cache backend

param bool sync

execute this inline instead of asynchronous

append(*func*, **args*, ***kwargs*)

Append a task to the chain. Takes the same arguments as [async_task\(\)](#)

return
the current number of tasks in the chain

rtype
int

run()

Start queueing the chain to the worker cluster.

return
the chains group id

result(*wait=0*)

return the full list of results from the chain when it finishes. Blocks until timeout or result.

param int wait
how many milliseconds to wait for a result

return
an unsorted list of results

fetch(*failures=True, wait=0*)

get the task result objects from the chain when it finishes. Blocks until timeout or result.

param failures
include failed tasks

param int wait
how many milliseconds to wait for a result

return
an unsorted list of task objects

current()

get the index of the currently executing chain element

return int
current chain index

length()

get the length of the chain

return int
length of the chain

1.8 Schedules

1.8.1 Schedule

Schedules are regular Django models. You can manage them through the [Admin pages](#) or directly from your code with the [schedule\(\)](#) function or the [Schedule](#) model:

```

# Use the schedule wrapper
from django_q.tasks import schedule

schedule('math.copysign',
         2, -2,
         hook='hooks.print_result',
         schedule_type='D')

# Or create the object directly
from django_q.models import Schedule

Schedule.objects.create(func='math.copysign',
                        hook='hooks.print_result',
                        args='2,-2',
                        schedule_type=Schedule.DAILY
                        )

# In case you want to use q_options
# Specify the broker by using the property broker_name in q_options
schedule('math.sqrt',
         9,
         hook='hooks.print_result',
         q_options={'timeout': 30, 'broker_name': 'broker_1'},
         schedule_type=Schedule.HOURLY)

# Run a schedule every 5 minutes, starting at 6 today
# for 2 hours
from datetime import datetime

schedule('math.hypot',
         3, 4,
         schedule_type=Schedule.MINUTES,
         minutes=5,
         repeats=24,
         next_run=datetime.utcnow().replace(hour=18, minute=0))

# Use a cron expression
schedule('math.hypot',
         3, 4,
         schedule_type=Schedule.CRON,
         cron = '0 22 * * 1-5')

# Restrain a schedule to a specific cluster
schedule('math.hypot',
         3, 4,
         schedule_type=Schedule.DAILY,
         cluster='my_cluster')

```

1.8.2 Missed schedules

If your cluster has not run for a while, the default behavior for the scheduler is to play catch up with the schedules and keep executing them until they are up to date. In practical terms this means the scheduler will execute tasks in the past, reschedule them in the past and immediately execute them again until the schedule is set in the future. This default behavior is intended to facilitate schedules that poll or gather statistics, but might not be suitable to your particular situation. You can change this by setting the `catch_up` configuration setting to `False`. The scheduler will then skip execution of scheduled events in the past. Instead those tasks will run once when the cluster starts again and the scheduler will find the next available slot in the future according to original schedule parameters.

When `catch_up` is to `True` it may be useful for the task to know what was the date and time it was originally intended to run at. To achieve this, pass an identifier name to parameter `intended_date_kwarg` when creating the schedule. The intended datetime will then be passed - in isoformat string - as a kwarg with that identifier name to the task that has been created.

1.8.3 Management Commands

If you want to schedule regular Django management commands, you can use the `django.core.management` module to call them directly:

```
from django_q.tasks import schedule

# run `manage.py clearsession` every hour
schedule('django.core.management.call_command',
        'clearsessions',
        schedule_type='H')
```

Or you can make a wrapper function which you can then schedule in Django Q:

```
# tasks.py
from django.core import management

# wrapping `manage.py clearsessions`
def clear_sessions_command():
    return management.call_command('clearsessions')

# now you can schedule it to run every hour
from django_q.tasks import schedule

schedule('tasks.clear_sessions_command', schedule_type='H')
```

Check out the *Shell* examples if you want to schedule regular shell commands

Note: Schedules needs the optional *Croniter* package installed to parse cron expressions.

1.8.4 Reference

schedule(*func*, **args*, *name*=None, *hook*=None, *schedule_type*='O', *minutes*=None, *repeats*=-1, *next_run*=now(), *q_options*=None, ***kwargs*)

Creates a schedule

Parameters

- **func** (*str*) – the function to schedule. Dotted strings only.
- **args** – arguments for the scheduled function.
- **name** (*str*) – An optional name for your schedule.
- **hook** (*str*) – optional result hook function. Dotted strings only.
- **schedule_type** (*str*) – (O)nce, M(I)utes, (H)ourly, (D)aily, (W)eekly, (M)onthly, (Q)uarterly, (Y)early or (C)ron [Schedule.TYPE](#)
- **minutes** (*int*) – Number of minutes for the Minutes type.
- **cron** (*str*) – Cron expression for the Cron type.
- **repeats** (*int*) – Number of times to repeat schedule. -1=Always, 0=Never, n =n.
- **next_run** (*datetime*) – Next or first scheduled execution datetime.
- **cluster** (*str*) – optional cluster name. Task will be executed only on a cluster with a matching *name*.
- **intended_date_kwarg** (*str*) – optional identifier to pass intended schedule date.
- **q_options** (*dict*) – options passed to `async_task` for this schedule
- **kwargs** – optional keyword arguments for the scheduled function.

Note: `q_options` does not accept the ‘broker’ key with a broker instance but accepts a ‘broker_name’ key instead. This can be used to specify the broker connection name to assign the task. If a broker with the specified name does not exist or is not running at the moment of placing the task in queue it fallbacks to the random broker/queue that handled the schedule.

class Schedule

A database model for task schedules.

id

Primary key

name

A name for your schedule. Tasks created by this schedule will assume this or the primary key as their group id.

func

The function to be scheduled

hook

Optional hook function to be called after execution.

args

Positional arguments for the function.

kwargs

Keyword arguments for the function

schedule_type

The type of schedule. Follows *Schedule.TYPE*

TYPE

ONCE, MINUTES, HOURLY, DAILY, WEEKLY, BIWEEKLY, MONTHLY, BIMONTHLY, QUARTERLY, YEARLY, CRON

minutes

The number of minutes the *MINUTES* schedule should use. Is ignored for other schedule types.

cron

A cron string describing the schedule. You need the optional *croniter* package installed for this.

repeats

Number of times to repeat the schedule. -1=Always, 0=Never, n =n. When set to -1, this will keep counting down.

cluster

Task will be executed only on a cluster with a matching *name*.

intended_date_kwarg

Name of kwarg to pass intended schedule date.

next_run

Datetime of the next scheduled execution.

task

Id of the last task generated by this schedule.

last_run()

Admin link to the last executed task.

success()

Returns the success status of the last executed task.

ONCE

'O' the schedule will only run once. If it has a negative *repeats* it will be deleted after it has run. If you want to keep the result, set *repeats* to a positive number.

MINUTES

'T' will run every *minutes* after its first run.

HOURLY

'H' the scheduled task will run every hour after its first run.

DAILY

'D' the scheduled task will run every day at the time of its first run.

WEEKLY

‘W’ the task will run every week on they day and time of the first run.

BIWEEKLY

‘BW’ the task will run once every two weeks on they day and time of the first run.

MONTHLY

‘M’ the tasks runs every month on they day and time of the last run.

Note: Months are tricky. If you schedule something on the 31st of the month and the next month has only 30 days or less, the task will run on the last day of the next month. It will however continue to run on that day, e.g. the 28th, in subsequent months.

BIMONTHLY

‘BM’ the tasks runs once every two months on they day and time of the last run.

Note: Months are tricky. If you schedule something on the 31st of the month and the next month has only 30 days or less, the task will run on the last day of the next month. It will however continue to run on that day, e.g. the 28th, in subsequent months.

QUARTERLY

‘Q’ this task runs once every 3 months on the day and time of the last run.

YEARLY

‘Y’ only runs once a year. The same caution as with months apply; If you set this to february 29th, it will run on february 28th in the following years.

CRON

‘C’ uses the optional *croniter* package to determine a schedule based on a cron expression.

1.9 Cluster

Django Q2 uses Python’s multiprocessing module to manage a pool of workers that will handle your tasks. Start your cluster using Django’s `manage.py` command:

```
$ python manage.py qcluster
```

You should see the cluster starting

```
10:57:40 [Q] INFO Q Cluster freddie-uncle-twenty-ten starting.
10:57:40 [Q] INFO Process-ede257774c4444c980ab479f10947acc ready for work at 31784
10:57:40 [Q] INFO Process-ed580482da3f42968230baa2e4253e42 ready for work at 31785
10:57:40 [Q] INFO Process-8a370dc2bc1d49aa9864e517c9895f74 ready for work at 31786
10:57:40 [Q] INFO Process-74912f9844264d1397c6e54476b530c0 ready for work at 31787
10:57:40 [Q] INFO Process-b00edb26c6074a6189e5696c60aeb35b ready for work at 31788
10:57:40 [Q] INFO Process-b0862965db04479f9784a26639ee51e0 ready for work at 31789
10:57:40 [Q] INFO Process-7e8abbb8ca2d4d9bb20a937dd5e2872b ready for work at 31790
```

(continues on next page)

(continued from previous page)

```
10:57:40 [Q] INFO Process-b0862965db04479f9784a26639ee51e0 ready for work at 31791
10:57:40 [Q] INFO Process-67fa9461ac034736a766cd813f617e62 monitoring at 31792
10:57:40 [Q] INFO Process-eac052c646b2459797cee98bdb84c85d guarding cluster at 31783
10:57:40 [Q] INFO Process-5d98deb19b1e4b2da2ef1e5bd6824f75 pushing tasks at 31793
10:57:40 [Q] INFO Q Cluster freddie-uncle-twenty-ten running.
```

Stopping the cluster with ctrl-c or either the SIGTERM and SIGKILL signals, will initiate the *Stop procedure*:

```
16:44:12 [Q] INFO Q Cluster freddie-uncle-twenty-ten stopping.
16:44:12 [Q] INFO Process-eac052c646b2459797cee98bdb84c85d stopping cluster processes
16:44:13 [Q] INFO Process-5d98deb19b1e4b2da2ef1e5bd6824f75 stopped pushing tasks
16:44:13 [Q] INFO Process-b0862965db04479f9784a26639ee51e0 stopped doing work
16:44:13 [Q] INFO Process-7e8abbb8ca2d4d9bb20a937dd5e2872b stopped doing work
16:44:13 [Q] INFO Process-b0862965db04479f9784a26639ee51e0 stopped doing work
16:44:13 [Q] INFO Process-b00edb26c6074a6189e5696c60aeb35b stopped doing work
16:44:13 [Q] INFO Process-74912f9844264d1397c6e54476b530c0 stopped doing work
16:44:13 [Q] INFO Process-8a370dc2bc1d49aa9864e517c9895f74 stopped doing work
16:44:13 [Q] INFO Process-ed580482da3f42968230baa2e4253e42 stopped doing work
16:44:13 [Q] INFO Process-ed257774c4444c980ab479f10947acc stopped doing work
16:44:14 [Q] INFO Process-67fa9461ac034736a766cd813f617e62 stopped monitoring results
16:44:15 [Q] INFO Q Cluster freddie-uncle-twenty-ten has stopped.
```

The number of workers, optional timeouts, recycles and cpu_affinity can be controlled via the *Configuration* settings.

1.9.1 Multiple Clusters

You can have multiple clusters on multiple machines, working on the same queue as long as:

- They connect to the same *broker*.
- They use the same cluster name. See *Configuration*
- They share the same SECRET_KEY for Django.

1.9.2 Multiple Queues

You can have multiple queues in one Django site, and use multiple cluster to work on each queue. Different queues are identified by different queue names which are also cluster names. To run an alternate cluster, e.g. to work on the ‘long’ queue, start your cluster with command:

```
# On Linux
$ Q_CLUSTER_NAME=long python manage.py qcluster

# On Windows
$ python manage.py qcluster --name long
```

You can set different Q_CLUSTER options for alternative clusters, such as ‘timeout’, ‘queue_limit’ and any other options which are valid in *Configuration*. See *ALT_CLUSTERS*.

Note: To use multiple queue, use the keyword argument *cluster* in *async_task()* and *schedule()*:

- if *cluster* is not set (the default), *async_task()* and *schedule()* will be handled by the default cluster;

- if *cluster* is set, only clusters with matching cluster name will run the task or do the schedule.

1.9.3 Using a Procfile

If you host on [Heroku](#) or you are using [Honcho](#) you can start the cluster from a Procfile with an entry like this:

```
worker: python manage.py qcluster
```

1.9.4 Process managers

While you certainly can run a Django Q2 with a process manager like [Supervisor](#) or [Circus](#) it is not strictly necessary. The cluster has an internal sentinel that checks the health of all the processes and recycles or reincarnates according to your settings or in case of unexpected crashes. Because of the multiprocessing daemonic nature of the cluster, it is impossible for a process manager to determine the clusters health and resource usage.

An example `circus.ini`

```
[circus]
check_delay = 5
endpoint = tcp://127.0.0.1:5555
pubsub_endpoint = tcp://127.0.0.1:5556
stats_endpoint = tcp://127.0.0.1:5557

[watcher:django_q]
cmd = python manage.py qcluster
numprocesses = 1
copy_env = True
```

Note that we only start one process. It is not a good idea to run multiple instances of the cluster in the same environment since this does nothing to increase performance and in all likelihood will diminish it. Control your cluster using the `workers`, `recycle` and `timeout` settings in your [Configuration](#)

An example `supervisor.conf`

```
[program:django-q]
command = python manage.py qcluster
stopasgroup = true
```

Supervisor's `stopasgroup` will ensure that the single process doesn't leave orphan process on stop or restart.

1.9.5 Reference

class Cluster

`start()`

Spawns a cluster and then returns

`stop()`

Initiates *Stop procedure* and waits for it to finish.

stat()

returns a *Stat* object with the current cluster status.

pid

The cluster process id.

host

The current hostname

sentinel

returns the `multiprocessing.Process` containing the *Sentinel*.

timeout

The clusters timeout setting in seconds

start_event

A `multiprocessing.Event` indicating if the *Sentinel* has finished starting the cluster

stop_event

A `multiprocessing.Event` used to instruct the *Sentinel* to initiate the *Stop procedure*

is_starting

Bool. Indicating that the cluster is busy starting up

is_running

Bool. Tells you if the cluster is up and running.

is_stopping

Bool. Shows that the stop procedure has been started.

has_stopped

Bool. Tells you if the cluster has finished the stop procedure

1.10 Monitor

Warning: Blessed needs to be installed to get this to work! See: <https://pypi.org/project/blessed/>

The cluster monitor shows live information about all the Q clusters connected to your project.

Start the monitor with Django's *manage.py* command:

```
$ python manage.py qmonitor
```

Host	Id	State	Pool	TQ	RQ	RC	Up
Orion	17004	Working	4	23	0	0	0:01:15
Phoenix	14710	Working	8	63	14	0	0:10:18

For all broker types except the Redis broker, the monitor utilizes Django's cache framework to store statistics of running clusters. This can be any type of cache backend as long as it can be shared among Django instances. For this reason, the local memory backend will not work.

1.10.1 Legend

Host

Shows the hostname of the server this cluster is running on.

Id

The cluster Id. Same as the cluster process ID or pid.

State

Current state of the cluster:

- **Starting** The cluster is spawning workers and getting ready.
- **Idle** Everything is ok, but there are no tasks to process.
- **Working** Processing tasks like a good cluster should.
- **Stopping** The cluster does not take on any new tasks and is finishing.
- **Stopped** All tasks have been processed and the cluster is shutting down.

Pool

The current number of workers in the cluster pool.

TQ

Task Queue counts the number of tasks in the queue¹

If this keeps rising it means you are taking on more tasks than your cluster can handle. You can limit this by settings the `queue_limit` in your cluster configuration, after which it will turn green when that limit has been reached. If your task queue is always hitting its limit and your running out of resources, it may be time to add another cluster.

RQ

Result Queue shows the number of results in the queue.¹

Since results are only saved by a single process which has to access the database. It's normal for the result queue to take slightly longer to clear than the task queue.

¹ Uses `multiprocessing.Queue.qsize()` which is not implemented on OS X and always returns 0.

RC

Reincarnations shows the amount of processes that have been reincarnated after a recycle, sudden death or timeout. If this number is unusually high, you are either suffering from repeated task errors or severe timeouts and you should check your logs for details.

Up

Uptime the amount of time that has passed since the cluster was started.

Press *q* to quit the monitor and return to your terminal.

1.10.2 Info

If you just want to see a one-off summary of your cluster stats you can use the *qinfo* management command:

```
$ python manage.py qinfo
```

-- djq summary --					
Clusters	1	Workers	4	Restarts	0
Queued	0	Successes	1000	Failures	2
Schedules	4	Tasks/hour	7.38	Avg time	0.1612

All stats are summed over all available clusters.

Task rate is calculated over the last 24 hours and will show the number of tasks per second, minute, hour or day depending on the amount. Average execution time (*Avg time*) is calculated in seconds over the last 24 hours.

Since some of these numbers are based on what is available in your tasks database, limiting or disabling the result backend will skew them.

Like with the monitor, these statistics come from a Redis server or Django's cache framework. So make sure you have either one configured.

To print out the current configuration run:

```
$ python manage.py qinfo --config
```

1.10.3 Status

You can check the status of your clusters straight from your code with the *Stat* class:

```
from django_q.status import Stat

for stat in Stat.get_all():
    print(stat.cluster_id, stat.status)

# or if you know the cluster id
cluster_id = 1234
stat = Stat.get(cluster_id)
print(stat.status, stat.workers)
```


1.10.4 Reference

class Stat

Cluster status object.

cluster_id

Id of this cluster. Corresponds with the process id.

tob

Time Of Birth

uptime()

Shows the number of seconds passed since the time of birth

reincarnations

The number of times the sentinel had to start a new worker process.

status

String representing the current cluster status.

task_q_size

The number of tasks currently in the task queue.¹

done_q_size

The number of tasks currently in the result queue.^{Page 43, 1}

pusher

The pid of the pusher process

monitor

The pid of the monitor process

sentinel

The pid of the sentinel process

workers

A list of process ids of the workers currently in the cluster pool.

empty_queues()

Returns true or false depending on any tasks still present in the task or result queue.

classmethod get(cluster_id, broker=None)

Gets the current [Stat](#) for the cluster id. Takes an optional broker connection.

classmethod get_all(broker=None)

Returns a list of [Stat](#) objects for all active clusters. Takes an optional broker connection.

1.11 Admin pages

Django Q2 does not use custom pages, but instead leverages what is offered by Django’s model admin by default. When you open Django Q2’s admin pages you will see three models:

1.11.1 Successful tasks

Shows all successfully executed tasks. Meaning they did not encounter any errors during execution. From here you can look at details of each task or delete them. Use the group filter to filter your results by schedule name or group id. The table is searchable by *name*, *func* and *group*

Search						Filter
Action: ----- Go 0 of 100 selected						By group
<input type="checkbox"/>	Name	Func	Started	Stopped	Time taken	Group
<input type="checkbox"/>	mockingbird-west-speaker-carpet	<function parzen_estimation at 0x7ff1caac8378>	July 30, 2015, 7:37 p.m.	July 30, 2015, 7:37 p.m.	12.402957	parzen
<input type="checkbox"/>	pizza-berlin-earth-romeo	<function parzen_estimation at 0x7ff1caac8378>	July 30, 2015, 7:37 p.m.	July 30, 2015, 7:37 p.m.	12.379033	parzen
<input type="checkbox"/>	winner-nitrogen-vermont-sad	<function parzen_estimation at 0x7ff1caac8378>	July 30, 2015, 7:37 p.m.	July 30, 2015, 7:37 p.m.	12.371897	parzen
<input type="checkbox"/>	neptune-early-orange-sodium	<function parzen_estimation at 0x7ff1caac8378>	July 30, 2015, 7:37 p.m.	July 30, 2015, 7:37 p.m.	12.371839	parzen
<input type="checkbox"/>	stream-chicken-dakota-uranus	<function parzen_estimation at 0x7ff1caac8378>	July 30, 2015, 7:37 p.m.	July 30, 2015, 7:37 p.m.	12.283605	parzen
<input type="checkbox"/>	avocado-blossom-bulldog-fifteen	<function parzen_estimation at 0x7ff1caac8378>	July 30, 2015, 7:37 p.m.	July 30, 2015, 7:37 p.m.	12.258351	parzen
<input type="checkbox"/>	sierra-autumn-romeo-eight	<function parzen_estimation at 0x7ff1caac8378>	July 30, 2015, 7:37 p.m.	July 30, 2015, 7:37 p.m.	12.256816	parzen
<input type="checkbox"/>	paris-uniform-fillet-georgia	<function parzen_estimation at 0x7ff1caac8378>	July 30, 2015, 7:37 p.m.	July 30, 2015, 7:37 p.m.	12.22891	parzen

Uses the *Success* proxy model.

Tip: The maximum number of successful tasks can be set using the *save_limit* option.

1.11.2 Failed tasks

Failed tasks have encountered an error, preventing them from finishing execution. The worker will try to put the error in the *result* field of the task so you can review what happened.

You can resubmit a failed task back to the queue using the admins action menu.

Uses the *Failure* proxy model

Customize the admin UI by creating your own `admin.ModelAdmin` class and use `admin.site.unregister` and `admin.site.register` to replace the default for example:

```
from django_q import models as q_models
from django_q import admin as q_admin

admin.site.unregister([q_models.Failure])
@admin.register(q_models.Failure)
class ChildClassAdmin(q_admin.FailAdmin):
    list_display = (
        'name',
        'func',
        'result',
        'started',
        # add attempt_count to list_display
        'attempt_count'
    )
```

1.11.3 Scheduled tasks

Here you can check on the status of your scheduled tasks, create, edit or delete them.

Select Scheduled task to change Add Scheduled task +

Q Search

Action: ----- Go 0 of 3 selected

<input type="checkbox"/>	ID	Name	Func	Schedule Type	Repeats	Next Run	Last run	Success
<input type="checkbox"/>	2	Time test	djq.loadtest.time_test	Hourly	-71	Aug. 5, 2015, 5:02 a.m.	[beer-rugby-mars-delaware]	✓
<input type="checkbox"/>	1	Math	math.copysign	Hourly	-82	Aug. 5, 2015, 5:12 a.m.	[pizza-item-lemon-asparagus]	✓
<input type="checkbox"/>	3	Parzen	djq.tasks.parzen_async	Daily	-5	Aug. 5, 2015, 10:03 p.m.	[georgia-eleven-vegan-ceiling]	✓

3 Scheduled tasks

Filter

By Next Run

Any date

Today

Past 7 days

This month

This year

By Schedule Type

All

Once

Hourly

Daily

Weekly

Monthly

Quarterly

Yearly

Repeats

If you want a schedule to only run a finite amount of times, e.g. every hour for the next 24 hours, you can do that using the `Schedule.repeats` attribute. In this case you would set the schedule type to `Schedule.HOURLY` and the repeats to 24. Every time the schedule runs the repeats count down until it hits zero and schedule is no longer run.

When you set repeats to -1 the schedule will continue indefinitely and the repeats will still count down. This can be used as an indicator of how many times the schedule has been executed.

An exception to this are schedules of type `Schedule.ONCE`. Negative repeats for this schedule type will cause it to be deleted from the database. This behavior is useful if you have many delayed actions which you do not necessarily need a result for. A positive number will keep the ONCE schedule, but it will not run again.

You can pause a schedule by setting its repeats value to zero.

Note: To run a ONCE schedule again, change the repeats to something other than 0. Set a new run time before you do this or let it execute immediately.

Next run

Shows you when this task will be added to the queue next.

Last run

Links to the task result of the last scheduled run. Shows nothing if the schedule hasn't run yet or if task result has been deleted.

Success

Indicates the success status of the last scheduled task, if any.

Note: if you have set the `save_limit` configuration option to not save successful tasks to the database, you will only see the failed results of your schedules.

Uses the `Schedule` model

1.11.4 Queued tasks

This admin view is only enabled when you use the `Django ORM` broker. It shows all tasks packages currently in the broker queue. The `lock` column shows the moment at which this package was picked up by the cluster and is used to determine whether it has expired or not. For development purposes you can edit and delete queued tasks from here.

1.12 Errors

Django Q2 uses a pluggable error reporter system based upon python `extras`, allowing anyone to develop plugins for their favorite error reporting and monitoring integration. Currently implemented examples include `Rollbar` and `Sentry`.

Error reporting plugins register a class which implements a `report` method, which is invoked when a Django Q2 cluster encounters an error, passing information to the particular service. Error reporters must be *configured* via the `Q_CLUSTER` dictionary in your `settings.py`. These settings are passed as kwargs upon initiation of the Error Reporter. Therefore, in order to implement a new plugin, a package must expose a class which will be instantiated with the necessary information via the `Q_CLUSTER` settings and implements a single `report` method.

For example implementations, see `django-q-rollbar` and `django-q-sentry`

1.13 Signals

1.13.1 Available signals

Django Q2 emits the following signals during its lifecycle.

Before enqueueing a task

The `django_q.signals.pre_enqueue` signal is emitted before a task is enqueued. The task dictionary is given as the task argument.

After spawning a worker process

The `django_q.signals.post_spawn` signal is emitted after a worker process has spawned. The process name is given as the `proc_name` argument (string).

Before executing a task

The `django_q.signals.pre_execute` signal is emitted before a task is executed by a worker. This signal provides two arguments:

- `task`: the task dictionary.
- `func`: the actual function that will be executed. If the task was created with a function path, this argument will be the callable function nonetheless.

After executing a task

The `django_q.signals.post_execute` signal is emitted after a task is executed by a worker and processed by the monitor. It included the `task` dictionary with the result.

1.13.2 Subscribing to a signal

Connecting to a Django Q2 signal is done the same as any other Django signal:

```
from django.dispatch import receiver
from django_q.signals import pre_enqueue, pre_execute, post_execute, post_spawn

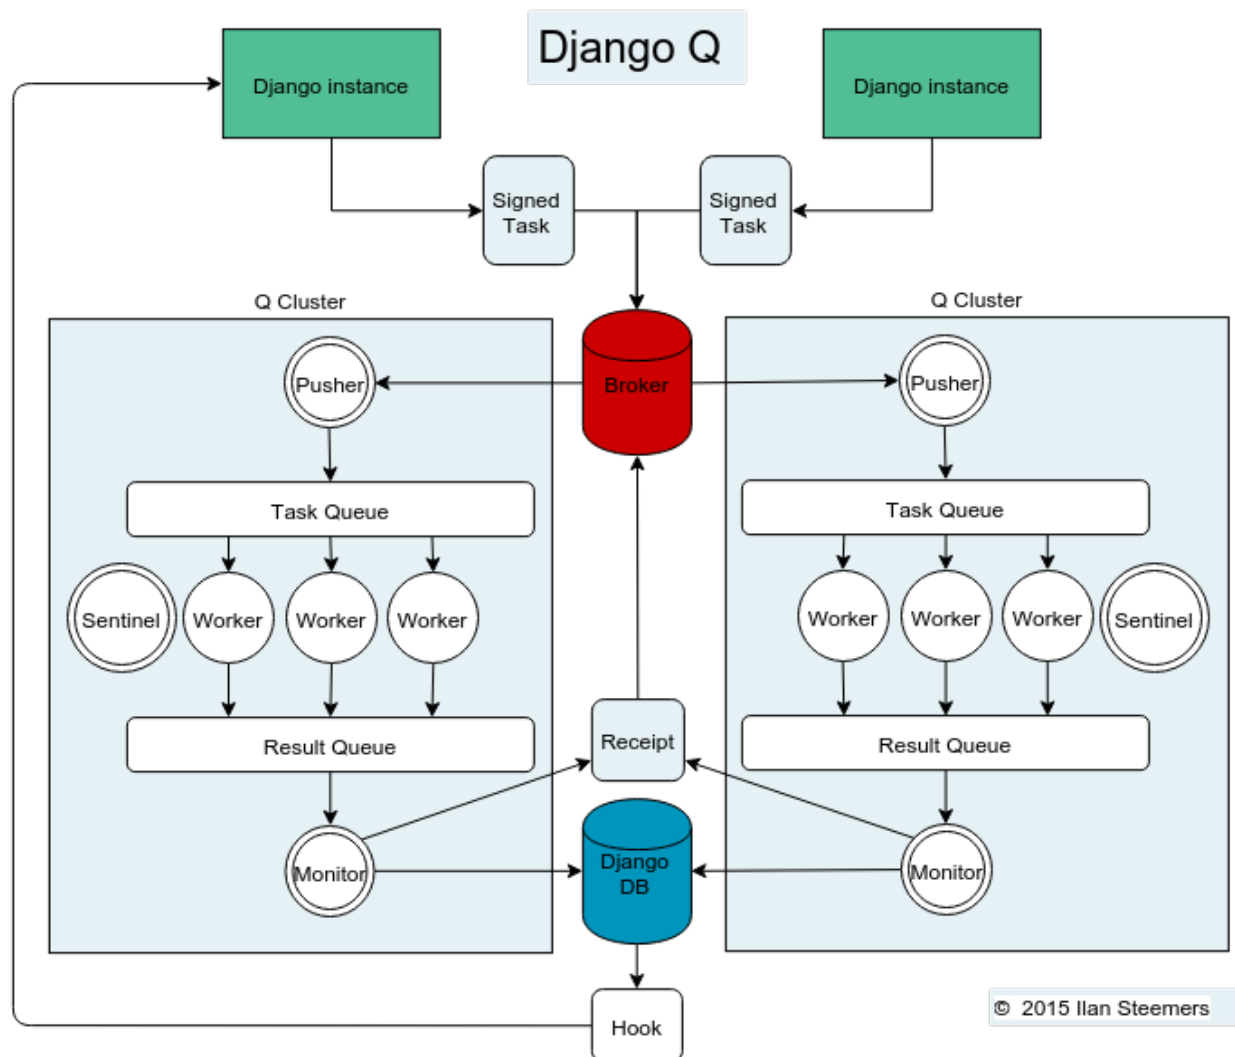
@receiver(pre_enqueue)
def my_pre_enqueue_callback(sender, task, **kwargs):
    print(f"Task {task['name']} will be queued")

@receiver(pre_execute)
def my_pre_execute_callback(sender, func, task, **kwargs):
    print(f"Task {task['name']} will be executed by calling {func}")

@receiver(post_execute)
def my_post_execute_callback(sender, task, **kwargs):
    print(f"Task {task['name']} was executed with result {task['result']}")

@receiver(post_spawn)
def my_post_spawn_callback(sender, proc_name, **kwargs):
    print(f"Process {proc_name} has spawned")
```

1.14 Architecture



1.14.1 Signed Tasks

Tasks are first pickled and then signed using Django's own `django.core.signing` module using the `SECRET_KEY` and cluster name as salt, before being sent to a message broker. This ensures that task packages on the broker can only be executed and read by clusters and django servers who share the same secret key and cluster name. If a package fails to unpack, it will be marked failed with the broker and discarded. Optionally the packages can be compressed before transport.

1.14.2 Broker

The broker collects task packages from the django instances and queues them for pick up by a cluster. If the broker supports message receipts, it will keep a copy of the tasks around until a cluster acknowledges the processing of the task. Otherwise it is put back in the queue after a timeout period. This ensure at-least-once delivery. Most failed deliveries will be the result of a worker or the cluster crashing before the task was saved.

Note: When the `ack_failures` option is set to `False` (the default), a task is considered a failed delivery when it raises an `Exception`. Set this option to `True` to acknowledge failed tasks as successful.

1.14.3 Pusher

The pusher process continuously checks the broker for new task packages. It checks the signing and unpacks the task to the internal Task Queue. The amount of tasks in the Task Queue can be configured to control memory usage and minimize data loss in case of a failure.

1.14.4 Worker

A worker process pulls a task of the Task Queue and it sets a shared countdown timer with *Sentinel* indicating it is about to start work. The worker then tries to execute the task and afterwards the timer is reset and any results (including errors) are saved to the package. Irrespective of the failure or success of any of these steps, the package is then pushed onto the Result Queue.

1.14.5 Monitor

The result monitor checks the Result Queue for processed packages and saves both failed and successful packages to the Django database or cache backend. If the broker supports it, a delivery receipt is sent. In case the task was part of a chain, the next task is queued.

1.14.6 Sentinel

The sentinel spawns all process and then checks the health of all workers, including the pusher and the monitor. This includes checking timers on each worker for timeouts. In case of a sudden death or timeout, it will reincarnate the failing processes. When a stop signal is received, the sentinel will halt the pusher and instruct the workers and monitor to finish the remaining items. See *Stop procedure*

1.14.7 Timeouts

Before each task execution the worker sets a countdown timer on the sentinel and resets it again after execution. Meanwhile the sentinel checks if the timers don't reach zero, in which case it will terminate the worker and reincarnate a new one.

1.14.8 Scheduler

Twice a minute the scheduler checks for any scheduled tasks that should be starting.

- Creates a task from the schedule
- Subtracts 1 from `django_q.Schedule.repeats`
- Sets the next run time if there are repeats left or if it has a negative value.

1.14.9 Stop procedure

When a stop signal is received, the sentinel exits the guard loop and instructs the pusher to stop pushing. Once this is confirmed, the sentinel pushes poison pills onto the task queue and will wait for all the workers to exit. This ensures that the task queue is emptied before the workers exit. Afterwards the sentinel waits for the monitor to empty the result queue and the stop procedure is complete.

- Send stop event to pusher
- Wait for pusher to exit
- Put poison pills in the Task Queue
- Wait for all the workers to clear the queue and stop
- Put a poison pill on the Result Queue
- Wait for monitor to process remaining results and exit
- Signal that we have stopped

Warning: If you force the cluster to terminate before the stop procedure has completed, you can lose tasks or results still being held in memory. You can manage the amount of tasks in a clusters memory by setting the `queue_limit`.

1.15 Examples

1.15.1 Emails

Sending an email can take a while so why not queue it:

```
# Welcome mail with follow up example
from datetime import timedelta
from django.utils import timezone
from django_q.tasks import async_task, schedule
from django_q.models import Schedule

def welcome_mail(user):
    msg = 'Welcome to our website'
    # send this message right away
    async_task('django.core.mail.send_mail',
               'Welcome',
               msg,
               'from@example.com',
```

(continues on next page)

(continued from previous page)

```

        [user.email])
    # and this follow up email in one hour
    msg = 'Here are some tips to get you started...'
    schedule('django.core.mail.send_mail',
            'Follow up',
            msg,
            'from@example.com',
            [user.email],
            schedule_type=Schedule.ONCE,
            next_run=timezone.now() + timedelta(hours=1))

    # since the `repeats` defaults to -1
    # this schedule will erase itself after having run

```

Since you're only telling Django Q2 to take care of the emails, you can quickly move on to serving web pages to your user.

1.15.2 Signals

A good place to use async tasks are Django's model signals. You don't want to delay the saving or creation of objects, but sometimes you want to trigger a lot of actions:

```

# Message on object change
from django.contrib.auth.models import User
from django.db.models.signals import pre_save
from django.dispatch import receiver
from django_q.tasks import async_task

# set up the pre_save signal for our user
@receiver(pre_save, sender=User)
def email_changed(sender, instance, **kwargs):
    try:
        user = sender.objects.get(pk=instance.pk)
    except sender.DoesNotExist:
        pass # new user
    else:
        # has his email changed?
        if not user.email == instance.email:
            # tell everyone
            async_task('tasks.inform_everyone', instance)

```

The task will send a message to everyone else informing them that the users email address has changed. Note that this adds almost no overhead to the save action:

```

# tasks.py
def inform_everyone(user):
    mails = []
    for u in User.objects.exclude(pk=user.pk):
        msg = f'Dear {u.username}, {user.username} has a new email address: {user.email}'
        mails.append(('New email', msg,
                     'from@example.com', [u.email]))
    return send_mass_mail(mails)

```

```
# or do it async again
def inform_everyone_async(user):
    for u in User.objects.exclude(pk=user.pk):
        msg = f'Dear {u.username}, {user.username} has a new email address: {user.email}'
        async_task('django.core.mail.send_mail',
                   'New email', msg, 'from@example.com', [u.email])
```

Of course you can do other things beside sending emails. These are just generic examples. You can use signals with async to update fields in other objects too. Let's say this users email address is not just on the User object, but you stored it in some other places too without a reference. By attaching an async action to the save signal, you can now update that email address in those other places without impacting the the time it takes to return your views.

1.15.3 Reports

In this example the user requests a report and we let the cluster do the generating, while handling the result with a hook.

```
# Report generation with hook example
from django_q.tasks import async_task

# views.py
# user requests a report.
def create_report(request):
    async_task('tasks.create_html_report',
              request.user,
              hook='tasks.email_report')
```

```
# tasks.py
from django_q.tasks import async_task

# report generator
def create_html_report(user):
    html_report = 'We had a great quarter!'
    return html_report

# report mailer
def email_report(task):
    if task.success:
        # Email the report
        async_task('django.core.mail.send_mail',
                  'The report you requested',
                  task.result,
                  'from@example.com',
                  task.args[0].email)
    else:
        # Tell the admins something went wrong
        async_task('django.core.mail.mail_admins',
                  'Report generation failed',
                  task.result)
```

The hook is practical here, because it allows us to detach the sending task from the report generation function and to report on possible failures.

1.15.4 Haystack

If you use [Haystack](#) as your projects search engine, here's an example of how you can have Django Q2 take care of your indexes in real time using model signals:

```
# Real time Haystack indexing
from .models import Document
from django.db.models.signals import post_save
from django.dispatch import receiver
from django_q.tasks import async_task

# hook up the post save handler
@receiver(post_save, sender=Document)
def document_changed(sender, instance, **kwargs):
    async_task('tasks.index_object', sender, instance, save=False)
    # turn off result saving to not flood your database
```

```
# tasks.py
from haystack import connection_router, connections

def index_object(sender, instance):
    # get possible backends
    backends = connection_router.for_write(instance=instance)

    for backend in backends:
        # get the index for this model
        index = connections[backend].get_unified_index()\
            .get_index(sender)
        # update it
        index.update_object(instance, using=backend)
```

Now every time a Document is saved, your indexes will be updated without causing a delay in your save action. You could expand this to dealing with deletes, by adding a `post_delete` signal and calling `index.remove_object` in the `async_task` function.

1.15.5 Shell

You can execute or schedule shell commands using Python's `subprocess` module:

```
from django_q.tasks import async_task, result

# make a backup copy of setup.py
async_task('subprocess.call', ['cp', 'setup.py', 'setup.py.bak'])

# call ls -l and dump the output
task_id=async_task('subprocess.check_output', ['ls', '-l'])

# get the result
dir_list = result(task_id)
```

In Python 3.5 the `subprocess` module has changed quite a bit and returns a `subprocess.CompletedProcess` object instead:

```
from django_q.tasks import async_task, result

# make a backup copy of setup.py
tid = async_task('subprocess.run', ['cp', 'setup.py', 'setup.py.bak'])

# get the result
r=result(tid, 500)
# we can now look at the original arguments
>>> r.args
['cp', 'setup.py', 'setup.py.bak']
# and the returncode
>>> r.returncode
0

# to capture the output we'll need a pipe
from subprocess import PIPE

# call ls -l and pipe the output
tid = async_task('subprocess.run', ['ls', '-l'], stdout=PIPE)
# get the result
res = result(tid, 500)
# print the output
print(res.stdout)
```

Instead of `async_task()` you can of course also use `schedule()` to schedule commands.

For regular Django management commands, it is easier to call them directly:

```
from django_q.tasks import async_task, schedule

async_task('django.core.management.call_command', 'clearsessions')

# or clear those sessions every hour

schedule('django.core.management.call_command',
        'clearsessions',
        schedule_type='H')
```

1.15.6 Groups

A group example with Kernel density estimation for probability density functions using the Parzen-window technique. Adapted from Sebastian Raschka's blog

```
# Group example with Parzen-window estimation
import numpy

from django_q.tasks import async_task, result_group, delete_group

# the estimation function
def parzen_estimation(x_samples, point_x, h):
    k_n = 0
    for row in x_samples:
```

(continues on next page)

(continued from previous page)

```

    x_i = (point_x - row[:, numpy.newaxis]) / h
    for row in x_i:
        if numpy.abs(row) > (1 / 2):
            break
    else:
        k_n += 1
    return h, (k_n / len(x_samples)) / (h ** point_x.shape[1])

# create 100 calculations and return the collated result
def parzen_async():
    # clear the previous results
    delete_group('parzen', cached=True)
    mu_vec = numpy.array([0, 0])
    cov_mat = numpy.array([[1, 0], [0, 1]])
    sample = numpy.random. \
        multivariate_normal(mu_vec, cov_mat, 10000)
    widths = numpy.linspace(1.0, 1.2, 100)
    x = numpy.array([[0], [0]])
    # async_task them with a group label to the cache backend
    for w in widths:
        async_task(parzen_estimation, sample, x, w,
                    group='parzen', cached=True)
    # return after 100 results
    return result_group('parzen', count=100, cached=True)

```

Django Q2 is not optimized for distributed computing, but this example will give you an idea of what you can do with task *Groups*.

Alternatively the `parzen_async()` function can also be written with `async_iter()`, which automatically utilizes the cache backend and groups to return a single result from an iterable:

```

# create 100 calculations and return the collated result
def parzen_async():
    mu_vec = numpy.array([0, 0])
    cov_mat = numpy.array([[1, 0], [0, 1]])
    sample = numpy.random. \
        multivariate_normal(mu_vec, cov_mat, 10000)
    widths = numpy.linspace(1.0, 1.2, 100)
    x = numpy.array([[0], [0]])
    # async_task them with async_task iterable
    args = [(sample, x, w) for w in widths]
    result_id = async_iter(parzen_estimation, args, cached=True)
    # return the cached result or timeout after 10 seconds
    return result(result_id, wait=10000, cached=True)

```

1.15.7 Http Health Check

An example of a python http server you can use (localhost:8080) to validate the health status of all the clusters in your environment. Example is http only.

Requires cache to be enabled. Save file in your Django project's root directory and run with command: `python worker_hc.py` in your container or other environment. Can be customized to show whatever you'd like from the Stat class or modified as needed.

```
from http.server import BaseHTTPRequestHandler, HTTPServer
from mtt_app.settings.base import EMAIL_USE_TLS

import os
import django

# Set the correct path to you settings module
os.environ.setdefault("DJANGO_SETTINGS_MODULE", "my.settings.path")

# All django stuff has to come after the setup:
django.setup()

from django_q.status import Stat
from django_q.conf import Conf

# Set host and port settings
hostName = "localhost"
serverPort = 8080

class HealthCheckServer(BaseHTTPRequestHandler):
    def do_GET(self):

        # Count the clusters and their status
        happy_clusters = 0
        total_clusters = 0

        for stat in Stat.get_all():
            total_clusters += 1
            if stat.status in [Conf.IDLE, Conf.WORKING]:
                happy_clusters += 1

        # Return 200 response if there is at least 1 cluster running,
        # and make sure all running clusters are happy
        if total_clusters and happy_clusters == total_clusters:
            response_code = 200
        else:
            response_code = 500

        self.send_response(response_code)
        self.send_header("Content-type", "text/html")
        self.end_headers()

        self.wfile.write(
            bytes("<html><head><title>Django-Q Heath Check</title></head>", "utf-8")
```

(continues on next page)

(continued from previous page)

```

    )
    self.wfile.write(
        bytes(f"<p>Health check returned {response_code} response</p>", "utf-8")
    )
    self.wfile.write(
        bytes(
            f"<p>{happy_clusters} of {total_clusters} cluster(s) are happy</p></html>
↪",
            "utf-8",
        )
    )

if __name__ == "__main__":
    webServer = HTTPServer((hostName, serverPort), HealthCheckServer)
    print("Server started at http://%s:%s" % (hostName, serverPort))

    try:
        webServer.serve_forever()
    except KeyboardInterrupt:
        pass

    webServer.server_close()
    print("Server stopped.")

```

Note: If you have an example you want to share, please submit a pull request on [github](#).

- genindex

PYTHON MODULE INDEX

d

django_q, [15](#)

A

acknowledge() (*Broker method*), 18
 append() (*Chain method*), 33
 append() (*Iter method*), 32
 args (*AsyncTask attribute*), 26
 args (*Schedule attribute*), 37
 args (*Task attribute*), 25
 async_chain() (*in module django_q*), 33
 async_iter() (*in module django_q*), 31
 async_task() (*Broker method*), 18
 async_task() (*in module django_q*), 23
 AsyncTask (*class in django_q*), 26

B

BIMONTHLY (*Schedule attribute*), 39
 BIWEEKLY (*Schedule attribute*), 39
 broker (*AsyncTask attribute*), 27
 Broker (*built-in class*), 18
 brokers.get_broker()
 built-in function, 18
 built-in function
 brokers.get_broker(), 18

C

cached (*AsyncTask attribute*), 27
 Chain (*class in django_q*), 33
 Cluster (*class in django_q*), 41
 cluster (*Schedule attribute*), 38
 cluster_id (*Stat attribute*), 45
 count_group() (*in module django_q*), 30
 CRON (*Schedule attribute*), 39
 cron (*Schedule attribute*), 38
 current() (*Chain method*), 34

D

DAILY (*Schedule attribute*), 38
 delete() (*Broker method*), 18
 delete_cached() (*in module django_q*), 24
 delete_group() (*in module django_q*), 30
 delete_group() (*Task class method*), 26
 delete_queue() (*Broker method*), 18
 dequeue() (*Broker method*), 18

django_q
 module, 15
 done_q_size (*Stat attribute*), 45

E

empty_queues() (*Stat method*), 45

F

fail() (*Broker method*), 18
 Failure (*class in django_q*), 26
 fetch() (*AsyncTask method*), 27
 fetch() (*Chain method*), 34
 fetch() (*in module django_q*), 24
 fetch() (*Iter method*), 32
 fetch_group() (*AsyncTask method*), 27
 fetch_group() (*in module django_q*), 29
 func (*AsyncTask attribute*), 26
 func (*Schedule attribute*), 37
 func (*Task attribute*), 25

G

get() (*Stat class method*), 45
 get_all() (*Stat class method*), 45
 get_group_count() (*Task class method*), 26
 get_result() (*Task class method*), 26
 get_result_group() (*Task class method*), 26
 get_task() (*Task class method*), 26
 get_task_group() (*Task class method*), 26
 group (*AsyncTask attribute*), 27
 group_count() (*Task method*), 25
 group_delete() (*Task method*), 25
 group_result() (*Task method*), 25

H

has_stopped (*Cluster attribute*), 42
 hook (*AsyncTask attribute*), 27
 hook (*Schedule attribute*), 37
 hook (*Task attribute*), 25
 host (*Cluster attribute*), 42
 HOURLY (*Schedule attribute*), 38

I

`id` (*AsyncTask attribute*), 26
`id` (*Schedule attribute*), 37
`id` (*Task attribute*), 24
`info()` (*Broker method*), 18
`intended_date_kwarg` (*Schedule attribute*), 38
`is_running` (*Cluster attribute*), 42
`is_starting` (*Cluster attribute*), 42
`is_stopping` (*Cluster attribute*), 42
`Iter` (*class in django_q*), 31

K

`kwargs` (*AsyncTask attribute*), 26
`kwargs` (*Schedule attribute*), 37
`kwargs` (*Task attribute*), 25

L

`last_run()` (*Schedule method*), 38
`length()` (*Chain method*), 34
`length()` (*Iter method*), 32
`lock_size()` (*Broker method*), 18

M

`MINUTES` (*Schedule attribute*), 38
`minutes` (*Schedule attribute*), 38
`module`
 `django_q`, 15
`monitor` (*Stat attribute*), 45
`MONTHLY` (*Schedule attribute*), 39

N

`name` (*Schedule attribute*), 37
`name` (*Task attribute*), 25
`next_run` (*Schedule attribute*), 38

O

`ONCE` (*Schedule attribute*), 38

P

`pid` (*Cluster attribute*), 42
`ping()` (*Broker method*), 18
`purge_queue()` (*Broker method*), 18
`pusher` (*Stat attribute*), 45

Q

`QUARTERLY` (*Schedule attribute*), 39
`queue_size()` (*Broker method*), 18
`queue_size()` (*in module django_q*), 24

R

`reincarnations` (*Stat attribute*), 45
`repeats` (*Schedule attribute*), 38

`result` (*Task attribute*), 25
`result()` (*AsyncTask method*), 27
`result()` (*Chain method*), 34
`result()` (*in module django_q*), 24
`result()` (*Iter method*), 32
`result_group()` (*AsyncTask method*), 27
`result_group()` (*in module django_q*), 29
`run()` (*AsyncTask method*), 27
`run()` (*Chain method*), 34
`run()` (*Iter method*), 32

S

`save` (*AsyncTask attribute*), 27
`Schedule` (*class in django_q*), 37
`schedule()` (*in module django_q*), 37
`schedule_type` (*Schedule attribute*), 38
`sentinel` (*Cluster attribute*), 42
`sentinel` (*Stat attribute*), 45
`start()` (*Cluster method*), 41
`start_event` (*Cluster attribute*), 42
`started` (*AsyncTask attribute*), 26
`started` (*Task attribute*), 25
`Stat` (*built-in class*), 45
`stat()` (*Cluster method*), 41
`status` (*Stat attribute*), 45
`stop()` (*Cluster method*), 41
`stop_event` (*Cluster attribute*), 42
`stopped` (*Task attribute*), 25
`Success` (*class in django_q*), 26
`success` (*Task attribute*), 25
`success()` (*Schedule method*), 38
`sync` (*AsyncTask attribute*), 27

T

`Task` (*class in django_q*), 24
`task` (*Schedule attribute*), 38
`task_q_size` (*Stat attribute*), 45
`time_taken()` (*Task method*), 25
`timeout` (*Cluster attribute*), 42
`tob` (*Stat attribute*), 45
`TYPE` (*Schedule attribute*), 38

U

`uptime()` (*Stat method*), 45

W

`WEEKLY` (*Schedule attribute*), 38
`workers` (*Stat attribute*), 45

Y

`YEARLY` (*Schedule attribute*), 39