GWCelery Documentation

Release 0.2.6

Leo Singer

Feb 12, 2019

Contents:

1 Quick start **3**
- 1.1 To install . . . 3
- 1.2 To test . . . 3
- 1.3 To start . . . 4

2 Design and anatomy of GWCelery **7**
- 2.1 Conceptual overview . . . 7
- 2.2 Block diagram . . . 8
- 2.3 Processes . . . 10
- 2.4 Eternal tasks . . . 10
- 2.5 Handlers . . . 11

3 Configuration **13**
- 3.1 Authentication . . . 13

4 Running under HTCondor **15**
- 4.1 Shortcuts . . . 15
- 4.2 Managing multiple deployments . . . 16

5 Monitoring and Management **17**
- 5.1 Flower . . . 17
- 5.2 Sentry . . . 18
- 5.3 Flask . . . 19
- 5.4 Nagios . . . 20
- 5.5 Command-Line Tools . . . 21

6 API Reference **23**
- 6.1 gwcelery.conf module . . . 23
- 6.2 gwcelery.sentry module . . . 26
- 6.3 gwcelery.tasks module . . . 26
- 6.4 gwcelery.util module . . . 47

7 Contributing **49**
- 7.1 Development model . . . 49
- 7.2 Where new code should go . . . 49
- 7.3 Guidelines for tasks . . . 49
- 7.4 Unit tests . . . 50

7.5 Code style . . . 50
7.6 Documentation . . . 50

8 Deployment **51**
8.1 Continuous deployment . . . 51
8.2 Making a new release . . . 52

9 Changelog **53**
9.1 0.2.6 (2019-02-12) . . . 53
9.2 0.2.5 (2019-02-01) . . . 53
9.3 0.2.4 (2018-12-17) . . . 54
9.4 0.2.3 (2018-12-16) . . . 54
9.5 0.2.2 (2018-12-14) . . . 54
9.6 0.2.1 (2018-12-14) . . . 54
9.7 0.2.0 (2018-12-14) . . . 54
9.8 0.1.7 (2018-11-27) . . . 55
9.9 0.1.6 (2018-11-14) . . . 55
9.10 0.1.5 (2018-11-13) . . . 55
9.11 0.1.4 (2018-10-29) . . . 55
9.12 0.1.3 (2018-10-26) . . . 56
9.13 0.1.2 (2018-10-11) . . . 56
9.14 0.1.1 (2018-10-04) . . . 56
9.15 0.1.0 (2018-09-26) . . . 57
9.16 0.0.31 (2018-09-04) . . . 57
9.17 0.0.30 (2018-08-02) . . . 57
9.18 0.0.29 (2018-07-31) . . . 58
9.19 0.0.28 (2018-07-25) . . . 58
9.20 0.0.27 (2018-07-22) . . . 58
9.21 0.0.26 (2018-07-20) . . . 58
9.22 0.0.25 (2018-07-19) . . . 59
9.23 0.0.24 (2018-07-18) . . . 59
9.24 0.0.23 (2018-07-18) . . . 59
9.25 0.0.22 (2018-07-11) . . . 59
9.26 0.0.21 (2018-07-10) . . . 60
9.27 0.0.20 (2018-07-09) . . . 60
9.28 0.0.19 (2018-07-09) . . . 60
9.29 0.0.18 (2018-07-06) . . . 60
9.30 0.0.17 (2018-07-05) . . . 60
9.31 0.0.16 (2018-07-05) . . . 60
9.32 0.0.15 (2018-06-29) . . . 61
9.33 0.0.14 (2018-06-28) . . . 61
9.34 0.0.13 (2018-06-28) . . . 61
9.35 0.0.12 (2018-06-28) . . . 61
9.36 0.0.11 (2018-06-27) . . . 61
9.37 0.0.10 (2018-06-13) . . . 62
9.38 0.0.9 (2018-06-06) . . . 63
9.39 0.0.8 (2018-06-06) . . . 63
9.40 0.0.7 (2018-05-31) . . . 63
9.41 0.0.6 (2018-05-26) . . . 63
9.42 0.0.5 (2018-05-08) . . . 64
9.43 0.0.4 (2018-04-28) . . . 64
9.44 0.0.3 (2018-04-27) . . . 64
9.45 0.0.2 (2018-04-27) . . . 64
9.46 0.0.1 (2018-04-27) . . . 64

10 License 65

11 Indices and tables 73

Bibliography 75

Python Module Index 77

GWCelery is a simple and reliable package for annotating and orchestrating LIGO/Virgo alerts, built from widely used open source components. It is built on the Celery distributed task queue (hence the name). This is the design and reference manual for GWCelery.

GWCelery's responsibilities include:

1. Merging related candidates from multiple online LIGO/Virgo transient searches into "superevents"
2. Correlating LIGO/Virgo events with gamma-ray bursts, neutrinos, and supernovae
3. Launching automated follow-up analyses including data quality checks, rapid sky localization, automated parameter estimation, and source classification
4. Generating and sending preliminary machine-readable GCN notices
5. Sending updated GCN notices after awaiting human input
6. Automatically composing GCN Circulars

Note: If you are a scientist, student, educator, or astronomy enthusiast looking for information about LIGO/Virgo alerts and low-latency data products, then please see our LIGO/Virgo Public Alerts User Guide.

CHAPTER 1

Quick start

These instructions are suitable for installing GWCelery for development and testing on any machine.

1.1 To install

GWCelery requires Python >= 3.6.

The easiest way to install it is with `venv` and `pip`:

```
$ python -m venv --system-site-packages ~/gwcelery
$ source ~/gwcelery/bin/activate
$ pip install gwcelery
```

Hint: **Note:** GWCelery requires a fairly new version of *setuptools*. If you get an error message that looks like this:

```
pkg_resources.VersionConflict: (setuptools 0.9.8
(/usr/lib/python2.7/site-packages),
Requirement.parse('setuptools>=30.3.0'))
```

then run `pip install --upgrade setuptools` and try again.

1.2 To test

With `setup.py`:

```
$ python setup.py test
```

1.3 To start

Before starting GWCelery, you need to authenticate for access to GraceDb and LVAlert and make sure that you have a Redis server running. Once you have completed those steps, you can start each of the GWCelery manually.

1.3.1 Authentication

To authenticate for GraceDb, obtain grid credentials from the LSC DataGrid Client by running `ligo-proxy-init`:

```
$ ligo-proxy-init albert.einstein
```

To authenticate for LVAlert, first complete the LVAlert Account Activation form once for each LVAlert server that you intend to use (generally you only need "Playground" for development purposes). Make a note of the passwords and store them in your ~/.netrc file with appropriate file permissions:

```
$ echo > ~/.netrc
$ chmod 0600 ~/.netrc
$ echo machine lvalert.cgca.uwm.edu login albert.einstein password password-for-
↪production >> ~/.netrc
$ echo machine lvalert-playground.cgca.uwm.edu login albert.einstein password␣
↪password-for-playground >> ~/.netrc
$ echo machine lvalert-test.cgca.uwm.edu login albert.einstein password password-for-
↪test >> ~/.netrc
```

1.3.2 Redis

GWCelery requires a Redis database server for task bookkeeping. Your operating system's package manager may be able to install, configure, and automatically launch a suitable Redis server for you.

Debian, Ubuntu, `apt`

Debian or Ubuntu users can install and start Redis using `apt-get`:

```
$ sudo apt-get install redis
```

macOS, MacPorts

Mac users with MacPorts can install Redis using `port install`:

```
$ sudo port install redis
```

Use `port load` to start the server:

```
$ sudo port load redis
```

From source

If none of the above options are available, then you can follow the Redis Quick Start instructions to build redis from source and start a server:

```
$ wget http://download.redis.io/redis-stable.tar.gz
$ tar xvzf redis-stable.tar.gz
$ cd redis-stable
$ make -j
$ src/redis-server
```

1.3.3 Start GWCelery components manually

GWCelery itself consists of four Celery workers and one Flask web application. Start them all by running each of the following commands:

```
$ gwcelery worker -l info -n gwcelery-worker -Q celery -B
$ gwcelery worker -l info -n gwcelery-openmp-worker -Q openmp -c 1
$ gwcelery worker -l info -n gwcelery-superevent-worker -Q superevent -c 1
$ gwcelery worker -l info -n gwcelery-exttrig-worker -Q exttrig -c 1
$ gwcelery flask run
```

Hint: With these arguments, each of the commands above will run until you type Control-C. You may want to run each of them in a separate terminal, or in the background using screen or nohup.

CHAPTER 2

Design and anatomy of GWCelery

2.1 Conceptual overview

Several online gravitational-wave transient search pipelines (currently Gstlal, PyCBC, cWB, and oLIB) upload candidates in real time to GraceDb, the central database and web portal for low-latency LIGO/Virgo analyses. Whenever an event is uploaded or altered, GraceDb pushes machine-readable notifications through LVAlert, a pubsub system based on XMPP.

The business logic for selecting and sending alerts to astronomers resides not in GraceDb itself but in GWCelery. The role of GWCelery in the LIGO/Virgo alert infrastructure is to drive the workflow of aggregating and annotating gravitational-wave candidates and sending GCN Notices to astronomers.

GWCelery interacts with GraceDb by listening for LVAlert messages and making REST API requests through the GraceDb client. GWCelery interacts with GCN by listening for GCN Notices using the PyGCN VOEvent client and sending GCN notices using the Comet VOEvent broker.

The major subsystems of GWCelery are:

- the LVAlert listener
- the GraceDb client
- the GCN listener
- the GCN broker
- the Superevent Manager, which clusters and merges related candidates into "superevents"
- the External Trigger Manager, which correlates gravitational-wave events with GRB, neutrino, and supernova events
- the Orchestrator, which executes the per-event annotation workflow

2.2 Block diagram

Below is a diagram illustrating the conceptual relationships of these subsystems. Nodes in the graph are hyperlinks to the relevant API documentation.

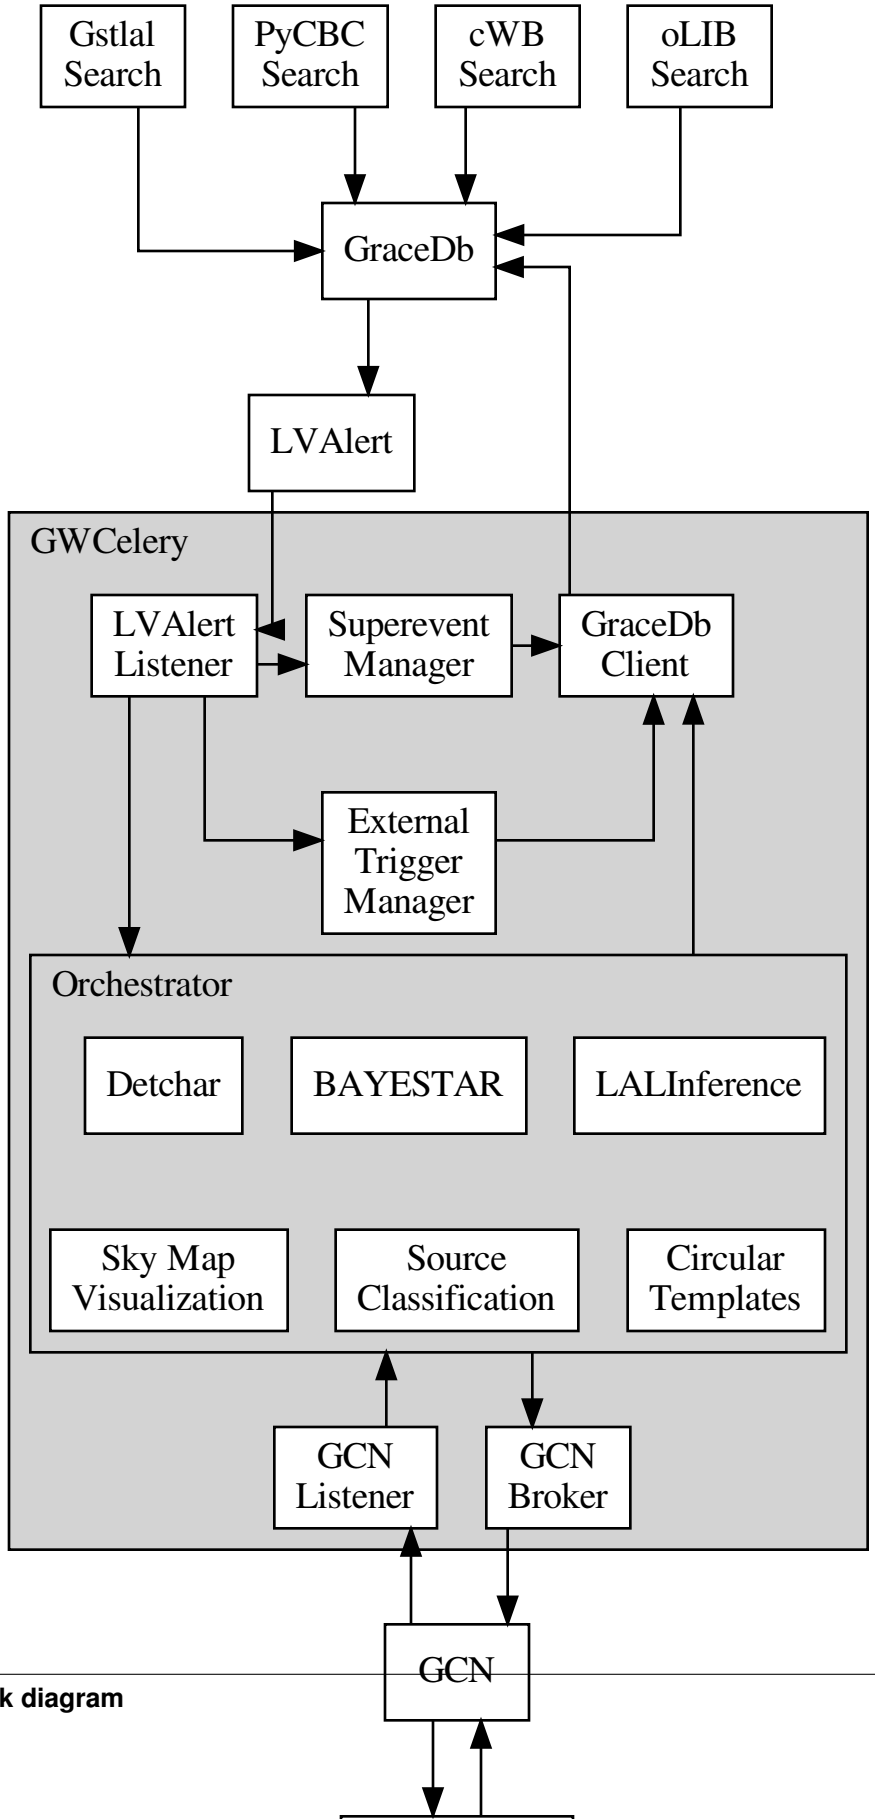
Gstlal Search
PyCBC Search
cWB Search
oLIB Search
GraceDb
LVAlert
GWCelery
LVAlert Listener
Superevent Manager
GraceDb Client
External Trigger Manager
Orchestrator
Detchar
BAYESTAR
LALInference
Sky Map Visualization
Source Classification
Circular Templates
GCN Listener
GCN Broker
GCN

2.3 Processes

A complete deployment of GWCelery (whether launched from the *shell* or from *HTCondor*) consists of several processes:

1. **Message Broker**

   Routes and distributes Celery task messages and stores results of tasks for later retrieval. See Choosing a Broker in the Celery manual for more details. For technical reasons, we use a Redis broker.

2. **Celery Beat**

   Scheduler for periodic tasks (the Celery equivalent of cron jobs). For more information, see Periodic Tasks in the Celery manual.

3. **Monitoring Console** (optional)

   You can optionally run Flower, a web monitoring console for Celery.

4. **OpenMP Worker**

   A Celery worker that has been configured to accept only computationally intensive tasks that use OpenMP parallelism. To route a task to the OpenMP worker, pass the keyword argument `queue='openmp'` to the `@app.task` decorator when you declare it.

   There are two tasks that run in the OpenMP queue:

   - *`gwcelery.tasks.bayestar.localize()`*
   - *`gwcelery.tasks.skymaps.plot_volume()`*

5. **Superevent Worker**

   A Celery worker that is dedicated to serially process triggers from low latency pipelines and create/modify superevents in GraceDb. There is only one task that runs on the Superevent queue:

   - *`gwcelery.tasks.superevents.handle()`*

6. **External Trigger Worker**

   A Celery worker that is dedicated to serially process external triggers from GRB alerts received from Fermi, Swift and neutrino alerts received from SNEWS and create/modify external trigger events in GraceDb:

   - `gwcelery.tasks.external_triggers.handle_gcn()`

7. **General-Purpose Worker**

   A Celery worker that accepts all other tasks.

8. **Flask Web Application**

   A web application that provides forms to manually initiate certain tasks, including sending an update alert or creating a mock event.

2.4 Eternal tasks

GWCelery has a few long-running tasks that do not return because they have to keep open a persistent connection with some external service. These tasks are subclasses of `celery_eternal.EternalTask` or `celery_eternal.EternalProcessTask`.

- *`gwcelery.tasks.gcn.broker()`*
- *`gwcelery.tasks.gcn.listen()`*

- *gwcelery.tasks.lvalert.listen()*

Both of these run inside the general-purpose worker process described above, and are automatically started (and restarted as necessary) by Celery Beat.

2.5 Handlers

A recurring pattern in GWCelery is that an eternal task listens continuously to a remote connection, receives packets of data over that connection, and dispatches further handling to other tasks based on packet type.

A decorator is provided to register a function as a Celery task and also plug it in as a handler for one or more packet types. This pattern is used for both GCN notices and LVAlert message handlers.

2.5.1 GCN notices

GCN notice handler tasks are declared using the *gwcelery.tasks.gcn.handler()* decorator:

```
import lxml.etree
from gwcelery.tasks import gcn

@gcn.handler(gcn.NoticeType.FERMI_GBM_GND_POS,
             gcn.NoticeType.FERMI_GBM_FIN_POS)
def handle_fermi(payload):
    root = lxml.etree.fromstring(payload)
    # do work here...
```

2.5.2 LVAlert messages

LVAlert message handler tasks are declared using the *gwcelery.tasks.lvalert.handler()* decorator:

```
from gwcelery.tasks import lvalert

@lvalert.handler('cbc_gstlal',
                 'cbc_spiir',
                 'cbc_pycbc',
                 'cbc_mbtaonline')
def handle_cbc(alert):
    # do work here...
```

CHAPTER 3

Configuration

Like any Celery application, GWCelery's configuration options are stored at run time in a global configuration object, `conf`. There are options for Celery itself such as options that affect the task and result backends; these options are documented in the Configuration and defaults section of the Celery manual.

The configuration object also holds all of the options that are specific to GWCelery and affect the behavior of individual GWCelery tasks; examples include the GraceDb and LVAlert service URLs, GCN hostnames, and frame file types and channel names. For a list of all GWCelery-specific options, see the API documentation for the *`gwcelery.conf`* module.

GWCelery provides four preset configurations, one for each GraceDb server instance (production, deployment, testing, or playground). The default configuration preset is for the playground server, `gracedb-playground.ligo.org`. The recommended way to select a different preset is to set the `CELERY_CONFIG_MODULE` environment variable before starting the workers. For example, to configure GWCelery for production:

```
$ export CELERY_CONFIG_MODULE=gwcelery.conf.production
```

3.1 Authentication

There are a few files that must be present in order to provide authentication tokens for GraceDb and LValert.

GraceDb

You must provide valid LSC DataGrid credentials in order for requests to the GraceDb REST API to work. During development and testing, you can use your personal credentials obtained from the LSC DataGrid Client by running `ligo-proxy-init`. However, credentials obtained this way expire after a few days or whenever your machine's temporary directory is wiped (e.g., at system restart).

For production deployment, you should obtain a robot certificate and store it in a location such as `~/.globus/userkey.pem` and `~/.globus/usercert.pem`.

LVAlert

You must provide a valid username and password for LVAlert. You can request an account using the LVAlert Account Activation form. The LVAlert username and password should be stored in your netrc file.

CHAPTER 4

Running under HTCondor

The recommended way to start and stop GWCelery on the LIGO Data Grid cluster is using HTCondor. See the example HTCondor submit file gwcelery.sub. This submit file will start up Redis, the worker processes, the Flask web application, and Flower. It will create some log files and a Unix domain socket, so you should first navigate to a directory where you want these files to go. For example:

```
$ mkdir -p ~/gwcelery/var && cd ~/gwcelery/var
```

Then run the submit file as follows:

```
$ gwcelery.sub
Submitting job(s)........
8 job(s) submitted to cluster 293497.
```

To stop GWCelery, run the `condor_hold` command:

```
$ condor_hold -constraint 'JobBatchName == "gwcelery"'
All jobs matching constraint (JobBatchName == "gwcelery") have been held
```

To restart GWCelery, run `condor_release`:

```
$ condor_release -constraint 'JobBatchName == "gwcelery"'
All jobs matching constraint (JobBatchName == "gwcelery") have been released
```

Note that there is normally **no need** to re-submit GWCelery if the machine is rebooted, because the jobs will persist in the HTCondor queue.

4.1 Shortcuts

The following commands are provided as shortcuts for the above operations:

```
$ gwcelery condor submit
$ gwcelery condor rm
$ gwcelery condor q
$ gwcelery condor hold
$ gwcelery condor release
```

The following command is a shortcut for `gwcelery condor rm; gwcelery condor submit`:

```
$ gwcelery condor resubmit
```

4.2 Managing multiple deployments

There should generally be at most one full deployment of GWCelery per GraceDb server running at one time. The `gwcelery condor` shortcut command is designed to protect you from accidentally starting multiple deployments of GWCelery by inspecting the HTCondor job queue before submitting new jobs. If you try to start GWCelery a second time on the same host in the same directory, you will get the following error message:

```
$ gwcelery condor submit
error: GWCelery jobs are already running in this directory.
You must first remove exist jobs with "gwcelery condor rm".
To see the status of those jobs, run "gwcelery condor q".
```

However, there are situations where you may actually want to run multiple instances of GWCelery on the same machine. For example, you may want to run one instance for the 'production' GraceDb server and one for the 'playground' server. To accomplish this, just start the two instances of gwcelery in different directories. Here is an example:

```
$ mkdir -p production
$ pushd production
$ CELERY_CONFIG_MODULE=gwcelery.conf.production gwcelery condor submit
$ popd
$ mkdir -p playground
$ pushd playground
$ CELERY_CONFIG_MODULE=gwcelery.conf.playground gwcelery condor submit
$ popd
```

CHAPTER 5

Monitoring and Management

GWCelery supports a rich selection of management and monitoring tools. Here is an introduction to a few of them. For more Celery monitoring solutions, see the Celery monitoring and management guide.

5.1 Flower

Flower is a dashboard for monitoring Celery tasks. To start Flower for monitoring during local development, run the following command and then navigate to http://localhost:5555/ in your browser:

```
$ gwcelery flower
```

To set up monitoring on a LIGO Data Grid cluster machine (e.g. `emfollow.ligo.caltech.edu`) protected by LIGO.org authentication, start Flower using the following command:

```
$ gwcelery flower --url-prefix=~${USER}/flower
```

add the following lines to the file `~/public_html/.htaccess`:

```
RewriteEngine on
RewriteRule ^flower/?(.*)$ http://emfollow.ligo.caltech.edu:5555/$1 [P]
```

Some additional firewall configuration may be required.

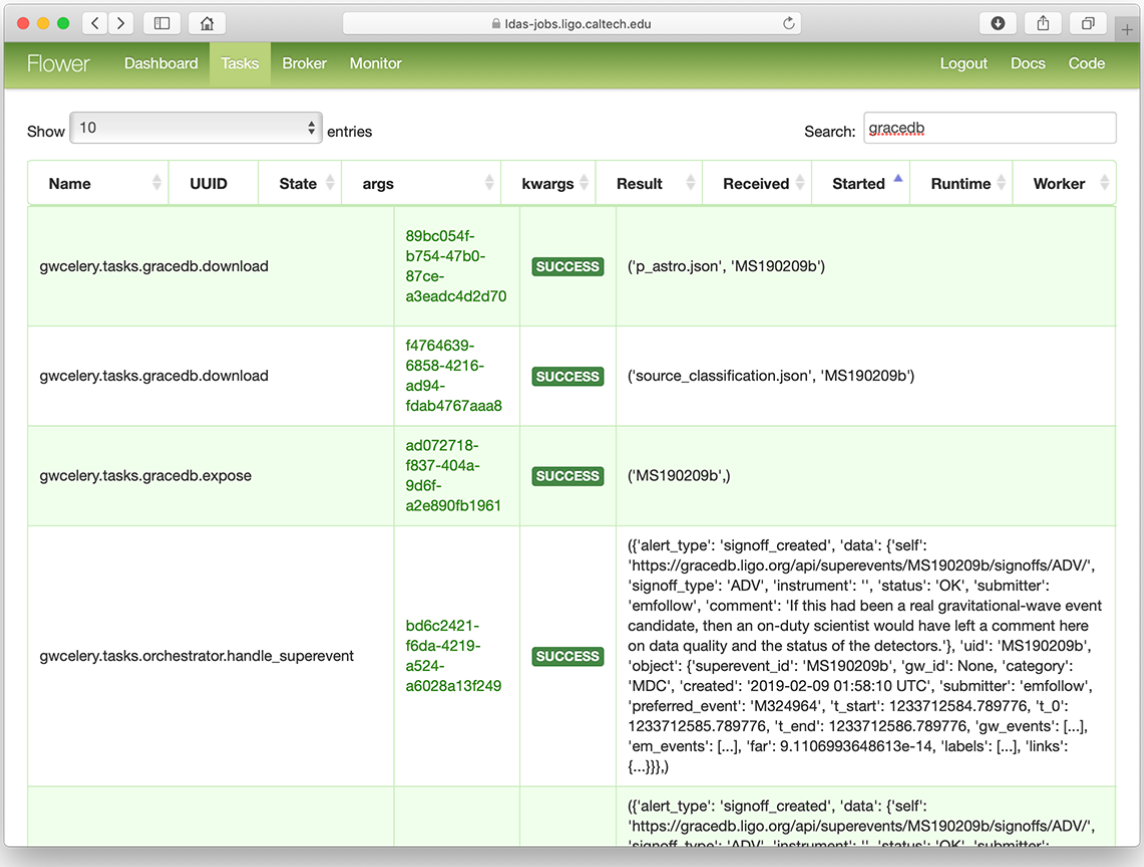

5.2 Sentry

All warnings, errors, exceptions, and tasks failures are both logged to disk and uploaded instantly to Sentry, an error monitoring and reporting platform. The Sentry instance is installed on premise at Caltech. Sentry notifies GWCelery contributors by email when a new bug occurs.

For details about the Sentry logging configuration, see the `gwcelery.sentry` module or the Celery integration module in the Sentry SDK docs.

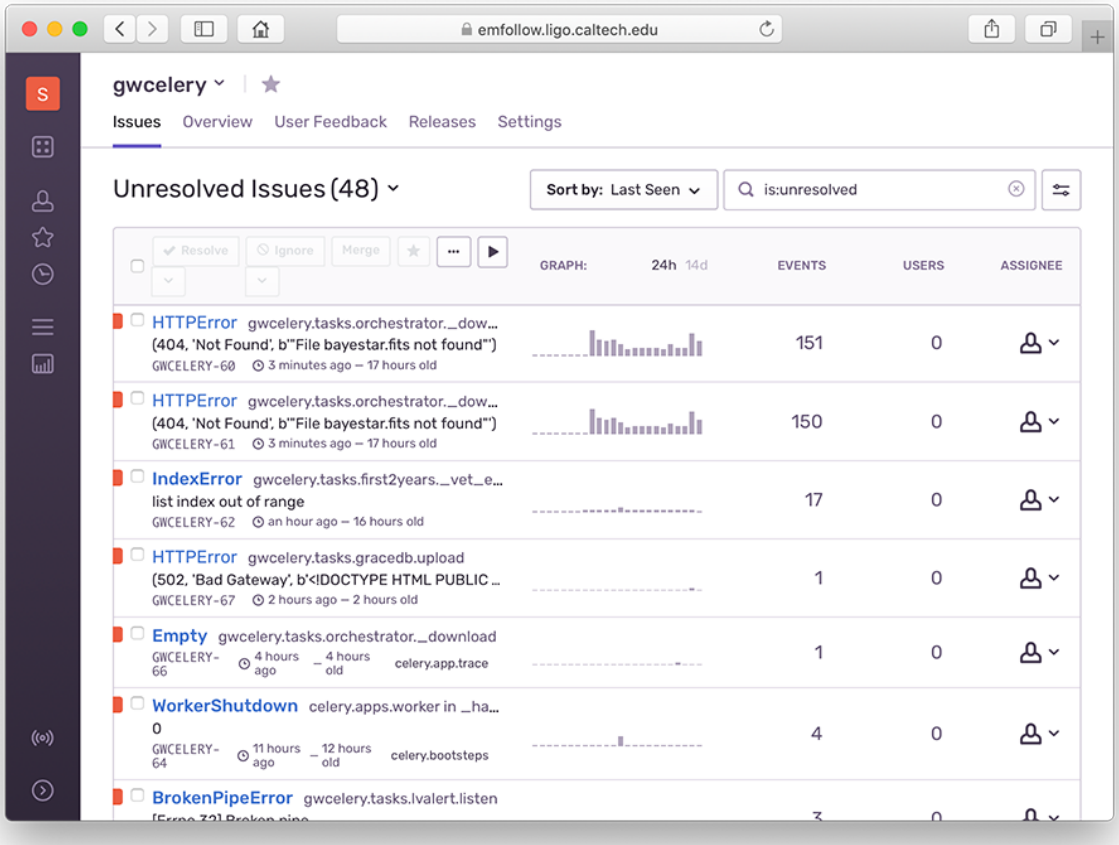

5.3 Flask

GWCelery includes a *Flask* web application that provides forms to manually initiate certain tasks.

To start Flask for monitoring during local development, run the following command and then navigate to http://localhost:5000/ in your browser:

```
$ gwcelery flask run
```

Setting up the Flask web application on a LIGO Data Grid cluster machine is similar to setting up Flower. Run the Flask application with the following additional command line arguments:

```
$ FLASK_URL_PREFIX=/~${USER}/gwcelery gwcelery flask run --port=5556 --with-threads --
↪host 0.0.0.0
```

add the following lines to the file `~/public_html/.htaccess`:

```
RewriteEngine on
RewriteRule ^gwcelery/?(.*)$ http://emfollow.ligo.caltech.edu:5556/$1 [P]
```

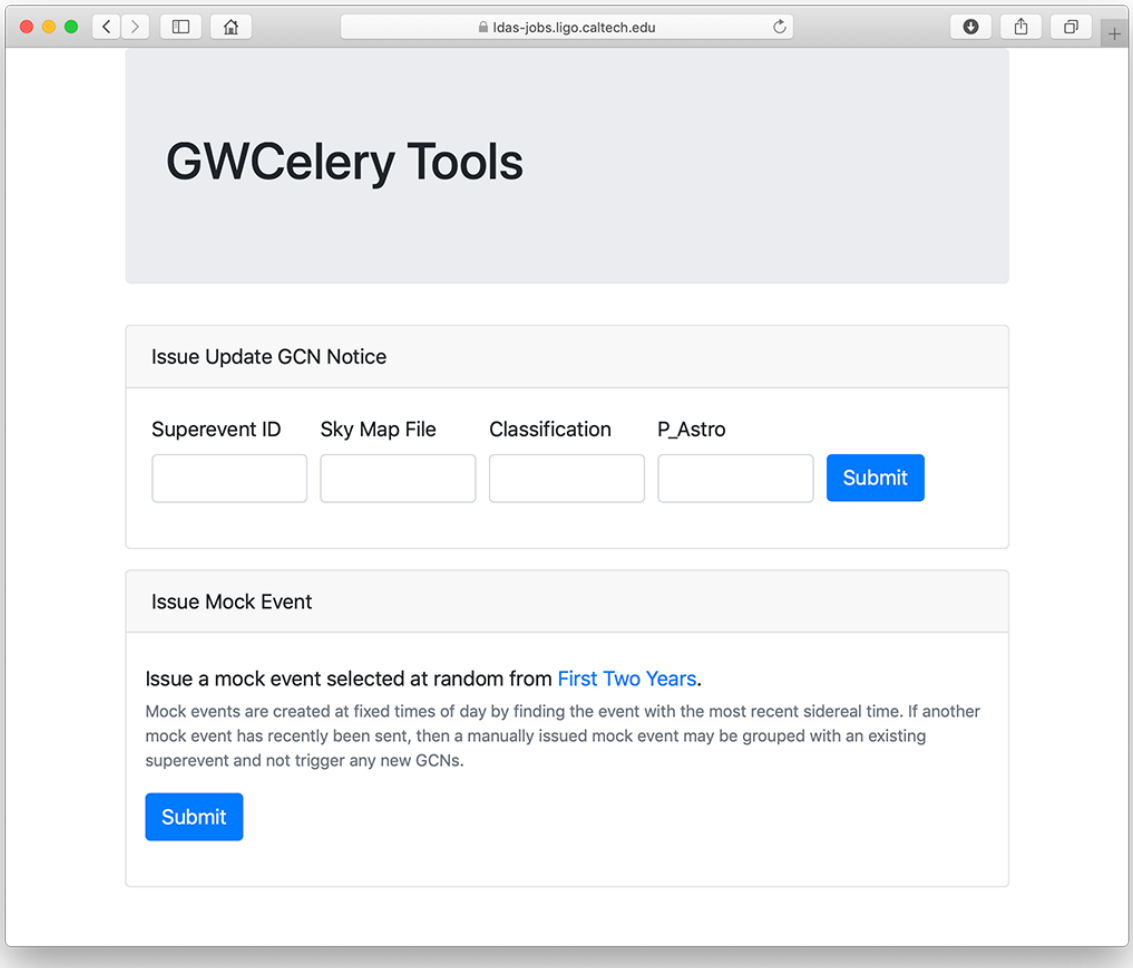

5.4 Nagios

Note: The GWCelery Nagios plugin is tailored to GWCelery and is not sufficiently general to use with other Celery applications.

The dashboard.ligo.org and monitor.ligo.org services use Nagios to monitor and report on the health of all of the components of the low-latency analysis infrastructure.

GWCelery provides the command `gwcelery nagios` to check the status of the application and provide a report in the format that Nagios expects.

You can run it manually from the command line:

```
$ gwcelery nagios
OK: GWCelery is running normally
```

To configure Nagios itself, see the Nagios configuration overview, or if GWCelery and Nagios are running on different hosts, the Nagios Remote Plugin Executor (NRPE) documentation.

5.5 Command-Line Tools

All Celery application provide command-line monitoring and management utilities, including the following:

- `gwcelery shell`: Start an interactive Python or IPython interpreter for interacting with Celery. All tasks as well as the *app* application instance are automatically imported and available as globals. Example:

```
$ gwcelery shell
Python 3.6.6 (default, Jun 28 2018, 05:43:53)
Type 'copyright', 'credits' or 'license' for more information
IPython 6.5.0 -- An enhanced Interactive Python. Type '?' for help.

In [1]: download.s('coinc.xml', 'M6757').delay().get()
```

- `gwcelery call`: Call a task from the command line by passing it arguments in JSON format. The output is the unique identifier of the result. Example:

```
$ gwcelery call gwcelery.tasks.gracedb.download --args='["coinc.xml", "M6757"]'
d11099e7-75e5-4aa3-800b-b122b667757c
```

- `gwcelery result`: Get the result of a previously called task. Example:

```
$ gwcelery result ab4aa6d7-9f21-420c-8401-cbe6863cf7dc
(b'<?xml version=\'1.0\' encoding=\'utf-8\'?>\n<!DOCTYPE LIGO_LW SYSTEM "htt'
 b'p://ldas-sw.ligo.caltech.edu/doc/ligolwAPI/html/ligolw_dtd.txt">\n<LIGO_L'
 ...
 b'\t</Stream>\n\t</Table>\n</LIGO_LW>\n')
```

- `gwcelery events`: A text UI monitoring tool that offers functionality similar to Flower. Example:

```
$ gwcelery events
```

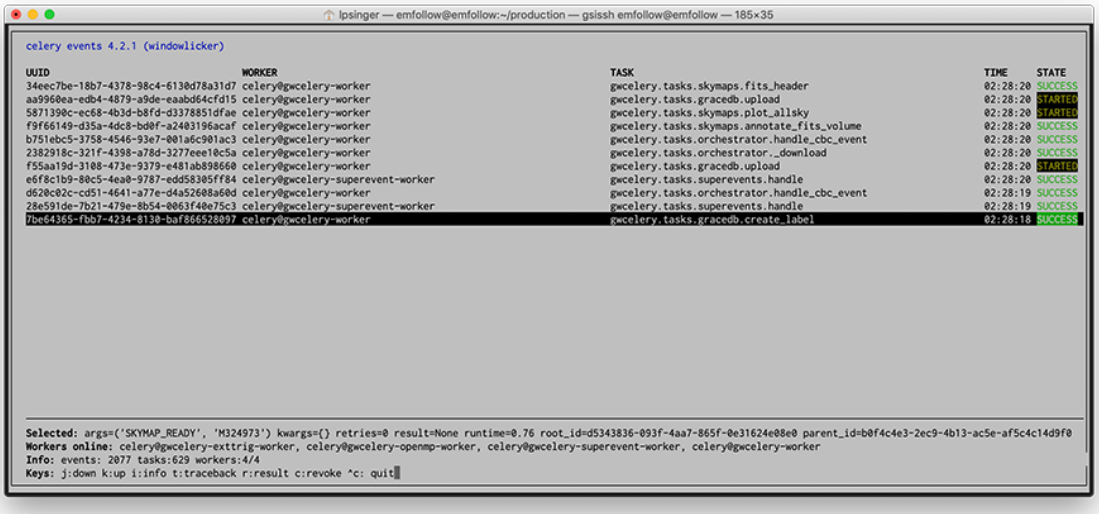

CHAPTER 6

API Reference

Celery application initialization.

`gwcelery.`**`app = <Celery gwcelery>`**
Celery application object.

6.1 gwcelery.conf module

GWCelery application configuration.

This module defines configuration variables and default values, including both generic options for Celery as well as options that control the behavior of specific GWCelery *tasks*.

To override the configuration, define the `CELERY_CONFIG_MODULE` environment variable to the fully qualified name of any Python module that can be located in `sys.path`, including any of the following presets:

- *gwcelery.conf.development*
- *gwcelery.conf.playground* (the default)
- *gwcelery.conf.production*
- *gwcelery.conf.test*

`gwcelery.conf.`**`lvalert_host = 'lvalert-playground.cgca.uwm.edu'`**
LVAlert host.

`gwcelery.conf.`**`gracedb_host = 'gracedb-playground.ligo.org'`**
GraceDb host.

`gwcelery.conf.`**`gcn_broker_address = ':5341'`**
The VOEvent broker will bind to this address to send GCNs. This should be a string of the form *host:port*. If *host* is empty, then listen on all available interfaces.

`gwcelery.conf.`**`gcn_broker_accept_addresses = []`**
List of hosts from which the broker will accept connections. If empty, then completely disable the broker's broadcast capability.

gwcelery.conf.**gcn_client_address = ''**
The VOEvent listener will connect to this address to receive GCNs. If empty, then completely disable the GCN listener.

gwcelery.conf.**superevent_d_t_start = {'gstlal': 1.0, 'mbtaonline': 1.0, 'pycbc': 1.0, 's**
Pipeline based lower extent of superevent segments. For cwb and lib this is decided from extra attributes.

gwcelery.conf.**superevent_d_t_end = {'gstlal': 1.0, 'mbtaonline': 1.0, 'pycbc': 1.0, 'spi**
Pipeline based upper extent of superevent segments For cwb and lib this is decided from extra attributes.

gwcelery.conf.**superevent_query_d_t_start = 100.0**
Lower extent of superevents query

gwcelery.conf.**superevent_query_d_t_end = 100.0**
Upper extent of superevents query

gwcelery.conf.**superevent_default_d_t_start = 1.0**
Default lower extent of superevent segments

gwcelery.conf.**superevent_default_d_t_end = 1.0**
Default upper extent for superevent segments

gwcelery.conf.**superevent_far_threshold = 0.0002777777777777778**
Maximum false alarm rate to consider events superevents.

gwcelery.conf.**preliminary_alert_far_threshold = {'burst': 3.1709791983764586e-08, 'cbc':**
Group specific maximum false alarm rate to consider sending preliminary alerts.

gwcelery.conf.**preliminary_alert_trials_factor = {'burst': 5.0, 'cbc': 5.0}**
Trials factor corresponding to trigger categories. For CBC, trials factor is the number of pipelines plus the external coincidence search. For Burst, this is the total number of searches plus the external coincidence search. CBC pipelines are gstlal, pycbc, mbtaonline, spiir. Burst searches are cwb.allsky, cwb.bbh, cwb.imbh and olib.allsky.

gwcelery.conf.**orchestrator_timeout = 15.0**
The orchestrator will wait this many seconds from the time of the creation of a new superevent to the time that annotations begin, in order to let the superevent manager's decision on the preferred event stabilize.

gwcelery.conf.**check_vector_prepost = {'CWB': [0.5, 0.5], 'Fermi': [2, 2], 'HardwareInject**
Seconds before and after the superevent start and end times which the DQ vector check will include in its check. Pipeline dependent.

gwcelery.conf.**uses_gatedhoft = {'CWB': True, 'Fermi': False, 'HardwareInjection': False,**
Whether or not a pipeline uses gated h(t). Determines whether or not the DMT-DQ_VECTOR will be analyzed for data quality.

gwcelery.conf.**llhoft_glob = '/dev/shm/llhoft/{detector}/*.gwf'**
File glob for low-latency h(t) frames.

gwcelery.conf.**llhoft_channels = {'H1:DMT-DQ_VECTOR': 'dmt_dq_vector_bits', 'H1:GDS-CALIB_S**
Low-latency h(t) state vector configuration. This is a dictionary consisting of a channel and its bitmask, as defined in *gwcelery.tasks.detchar*.

gwcelery.conf.**idq_channels = ['H1:IDQ-PGLITCH_RANDOM_FOREST_16_4096', 'L1:IDQ-PGLITCH_RANDO**
Low-latency iDQ p(glitch) channel names

gwcelery.conf.**idq_pglitch_thresh = 0.95**
Minimum p(glitch) reported by iDQ required before notice is posted to GraceDb

gwcelery.conf.**p_astro_gstlal_ln_likelihood_threshold = 6**
log likelihood threshold

gwcelery.conf.**low_latency_frame_types = {'H1': 'H1_O2_llhoft', 'L1': 'L1_O2_llhoft', 'V1**
Types of frames used in Parameter Estimation with LALInference (see *gwcelery.tasks.lalinference*)

gwcelery.conf.**high_latency_frame_types = {'H1': 'None', 'L1': 'None', 'V1': 'None'}**
Types of nonllhoft-frames used in Parameter Estimation with LALInference. They do not exist for O2Replay data. (see *gwcelery.tasks.lalinference*)

gwcelery.conf.**strain_channel_names = {'H1': 'H1:GDS-CALIB_STRAIN_O2Replay', 'L1': 'L1:GDS**
Names of h(t) channels used in Parameter Estimation with LALInference (see *gwcelery.tasks.lalinference*)

gwcelery.conf.**state_vector_channel_names = {'H1': 'H1:GDS-CALIB_STATE_VECTOR', 'L1': 'L1**
Names of state vector channels used in Parameter Estimation with LALInference (see *gwcelery.tasks.lalinference*)

gwcelery.conf.**pe_threshold = 8.267195767195768e-07**
FAR threshold in Hz for Parameter Estimation. PE group now applies 1/(2 weeks) as a threshold. 86400 seconds = 1 day and 14 days = 2 weeks.

6.1.1 gwcelery.conf.development module

Application configuration for `gracedb-dev1.ligo.org`. Inherits all settings from *gwcelery.conf.test*, with the exceptions below.

gwcelery.conf.development.**gracedb_host = 'gracedb-dev1.ligo.org'**
GraceDb host.

gwcelery.conf.development.**sentry_environment = 'development'**
Record this *environment tag <https://docs.sentry.io/enriching-error-data/environments/>*)‘ in Sentry log messages.

6.1.2 gwcelery.conf.playground module

Application configuration for `gracedb-playground.ligo.org`.

gwcelery.conf.playground.**sentry_environment = 'playground'**
Record this *environment tag <https://docs.sentry.io/enriching-error-data/environments/>*)‘ in Sentry log messages.

6.1.3 gwcelery.conf.production module

Application configuration for `gracedb.ligo.org`. Inherits all settings from *gwcelery.conf.playground*, with the exceptions below.

gwcelery.conf.production.**lvalert_host = 'lvalert.cgca.uwm.edu'**
LVAlert host.

gwcelery.conf.production.**gracedb_host = 'gracedb.ligo.org'**
GraceDb host.

gwcelery.conf.production.**gcn_broker_accept_addresses = ['capella2.gsfc.nasa.gov']**
List of hosts from which the broker will accept connections. If empty, then completely disable the broker's broadcast capability.

gwcelery.conf.production.**gcn_client_address = '68.169.57.253:8096'**
The VOEvent listener will connect to this address to receive GCNs. If empty, then completely disable the GCN listener.

We are temporarily using the pre-registered port 8096 for receiving proprietary LIGO/Virgo alerts on emfollow.ligo.caltech.edu. This means that the capability to receive GCNs requires setting up a site configuration in advance with Scott Barthelmey.

Once we switch to sending public alerts exclusively, then we can switch back to using port 8099 for anonymous access, requiring no prior site configuration.

`gwcelery.conf.production.`**`low_latency_frame_types = {'H1': 'H1_llhoft', 'L1': 'L1_llhoft',`**

Types of frames used in Parameter Estimation with LALInference (see *gwcelery.tasks.lalinference*)

`gwcelery.conf.production.`**`high_latency_frame_types = {'H1': 'H1_HOFT_C00', 'L1': 'L1_HOFT_`**

Types of nonllhoft-frames used in Parameter Estimation with LALInference (see *gwcelery.tasks.lalinference*)

`gwcelery.conf.production.`**`strain_channel_names = {'H1': 'H1:GDS-CALIB_STRAIN', 'L1': 'L1:C`**

Names of h(t) channels used in Parameter Estimation with LALInference (see *gwcelery.tasks.lalinference*)

`gwcelery.conf.production.`**`sentry_environment = 'production'`**

Record this *environment tag <https://docs.sentry.io/enriching-error-data/environments/>*)‘ in Sentry log messages.

6.1.4 gwcelery.conf.test module

Application configuration for `gracedb-test.ligo.org`. Inherits all settings from *gwcelery.conf.playground*, with the exceptions below.

`gwcelery.conf.test.`**`lvalert_host = 'lvalert-test.cgca.uwm.edu'`**

LVAlert host.

`gwcelery.conf.test.`**`gracedb_host = 'gracedb-test.ligo.org'`**

GraceDb host.

`gwcelery.conf.test.`**`sentry_environment = 'test'`**

Record this *environment tag <https://docs.sentry.io/enriching-error-data/environments/>*)‘ in Sentry log messages.

6.2 gwcelery.sentry module

Integration of the Celery logging system with Sentry.

`gwcelery.sentry.`**`DSN = 'http://localhost:9000/2'`**

Sentry data source name (DSN).

`gwcelery.sentry.`**`configure`**`()`

Configure Sentry logging integration for Celery according to the official instructions.

Add the API key username/pasword pair to your netrc file.

6.3 gwcelery.tasks module

All Celery tasks are declared in submodules of this module.

6.3.1 gwcelery.tasks.bayestar module

Rapid sky localization with BAYESTAR.

(task) gwcelery.tasks.bayestar.**localize**(*coinc_psd*, *graceid*, *filename='bayestar.fits.gz'*, *disabled_detectors=None*)

Generate a rapid sky localization using BAYESTAR.

Parameters

- **coinc_psd** (*tuple*) – Tuple consisting of the byte contents of the input event's `coinc.xml` and `psd.xml.gz` files.
- **graceid** (*str*) – The GraceDB ID, used for FITS metadata and recording log messages to GraceDb.
- **filename** (*str, optional*) – The name of the FITS file.
- **disabled_detectors** (*list, optional*) – List of detectors to disable.

Returns The byte contents of the finished FITS file.

Return type bytes

Notes

This task is adapted from the command-line tool bayestar-localize-lvalert.

It should execute in a special queue for computationally intensive, multithreaded, OpenMP tasks.

6.3.2 gwcelery.tasks.circulars module

Generate and upload automated circulars.

(task) gwcelery.tasks.circulars.**create_circular**(*graceid*)

Create and return circular txt.

6.3.3 gwcelery.tasks.condor module

Submit and monitor HTCondor jobs[1].

Notes

Internally, we use the XML condor log format[2] for easier parsing.

References

exception gwcelery.tasks.condor.**JobAborted**

Bases: `Exception`

Raised if an HTCondor job was aborted (e.g. by `condor_rm`).

[1] http://research.cs.wisc.edu/htcondor/manual/latest/condor_submit.html
[2] http://research.cs.wisc.edu/htcondor/classad/refman/node3.html

exception gwcelery.tasks.condor.**JobRunning**

Bases: Exception

Raised if an HTCondor job is still running.

exception gwcelery.tasks.condor.**JobFailed**(*returncode*, *cmd*, *output=None*, *stderr=None*)

Bases: subprocess.CalledProcessError

Raised if an HTCondor job fails.

(task) gwcelery.tasks.condor.**submit**(*submit_file*, *log=None*)

Submit a job using HTCondor.

Parameters

- **submit_file** (*str*) – Path of the submit file.
- **log** (*str*) – Used internally to track job state. Caller should not set.

Raises

- *JobAborted* – If the job was aborted (e.g. by running condor_rm).
- *JobFailed* – If the job terminates and returns a nonzero exit code.
- *JobRunning* – If the job is still running. Causes the task to be re-queued until the job is complete.

Example

```
>>> submit.s('example.sub',
...          accounting_group='ligo.dev.o3.cbc.explore.test')
```

(task) gwcelery.tasks.condor.**check_output**(*args*, *log=None*, *error=None*, *output=None*, ***kwargs*)

Call a process using HTCondor.

Call an external process using HTCondor, in a manner patterned after subprocess.check_output(). If successful, returns its output on stdout. On failure, raise an exception.

Parameters

- **args** (*list*) – Command line arguments, as if passed to subprocess.check_call().
- **error, output** (*log,*) – Used internally to track job state. Caller should not set.
- ****kwargs** – Extra submit description file commands. See the documentation for condor_submit for possible values.

Returns Captured output from command.

Return type str

Raises

- *JobAborted* – If the job was aborted (e.g. by running condor_rm).
- *JobFailed* – If the job terminates and returns a nonzero exit code.
- *JobRunning* – If the job is still running. Causes the task to be re-queued until the job is complete.

Example

```
>>> check_output.s(['sleep', '10'],
...                accounting_group='ligo.dev.o3.cbc.explore.test')
```

6.3.4 gwcelery.tasks.detchar module

Data quality and detector characterization tasks.

These tasks are mostly focused on checking interferometer state vectors. By design, the [LIGO] and [Virgo] state vectors share the same definitions for the first 8 fields.

LIGO also has a [DMT] DQ vector that provides some additional instrumental checks.

References

`gwcelery.tasks.detchar.`**`dmt_dq_vector_bits`**
DMT DQ vector bits (LIGO only).

`gwcelery.tasks.detchar.`**`state_vector_bits`**
State vector bitfield definitions for LIGO and Virgo.

`gwcelery.tasks.detchar.`**`no_dq_veto_pycbc_bits`**
No DQ veto stream bitfield definitions for Virgo. NOTE: Since the results for these bits will be NOT()ed, the bit definitions are the NO_* versions of what the bit * actually is. This is an inelegant but the simplest solution since the logic used in these channels are opposite to those in all the other checked channels.

`gwcelery.tasks.detchar.`**`create_cache`**(*ifo*, *start*, *end*)
Find .gwf files and create cache. Will first look in the llhoft, and if the frames have expired from llhoft, will call gwdatafind.

Parameters

- **`ifo`** (*`str`*) – Interferometer name (e.g. `H1`).
- **`end`** (*`start,`*) – GPS start and end times desired.

Returns

Return type `glue.lal.Cache`

Example

```
>>> create_cache('H1', 1198800018, 1198800618)
[<glue.lal.CacheEntry at 0x7fbae6b71278>,
  <glue.lal.CacheEntry at 0x7fbae6ae5b38>,
  <glue.lal.CacheEntry at 0x7fbae6ae5c50>,
 ...
  <glue.lal.CacheEntry at 0x7fbae6b15080>,
  <glue.lal.CacheEntry at 0x7fbae6b15828>]
```

`gwcelery.tasks.detchar.`**`dqr_json`**(*state*, *summary*)
Generate DQR-compatible json-ready dictionary from process results, as described in `data-quality-report.design`.

Parameters

- **state** (*{'pass', 'fail'}*) – State of the detchar checks.
- **summary** (*str*) – Summary of results from the process.

Returns Ready to be converted into json.

Return type dict

`gwcelery.tasks.detchar.check_idq`(*cache*, *channel*, *start*, *end*)

Looks for iDQ frame and reads them.

Parameters

- **cache** (`glue.lal.Cache`) – Cache from which to check.
- **channel** (*str*) – which idq channel (pglitch)
- **end** (*start,*) – GPS start and end times desired.

Returns Tuple mapping iDQ channel to its maximum P(glitch).

Return type tuple

Example

```
>>> check_idq(cache, 'H1:IDQ-PGLITCH-OVL-100-1000',
              1216496260, 1216496262)
('H1:IDQ-PGLITCH-OVL-100-1000', 0.87)
```

`gwcelery.tasks.detchar.check_vector`(*cache*, *channel*, *start*, *end*, *bits*, *logic_type='all'*)

Check timeseries of decimals against a bitmask. This is inclusive of the start time and exclusive of the end time, i.e. [start, . . . , end).

Parameters

- **cache** (`glue.lal.Cache`) – Cache from which to check.
- **channel** (*str*) – Channel to look at, e.g. `H1:DMT-DQ_VECTOR`.
- **end** (*start,*) – GPS start and end times desired.
- **bits** (`gwpy.TimeSeries.Bits`) – Definitions of the bits in the channel.
- **logic_type** (*str, optional*) – Type of logic to apply for vetoing. If `all`, then all samples in the window must pass the bitmask. If `any`, then one or more samples in the window must pass.

Returns Maps each bit in channel to its state.

Return type dict

Example

```
>>> check_vector(cache, 'H1:GDS-CALIB_STATE_VECTOR', 1216496260,
                 1216496262, state_vector_bits)
{'H1:HOFT_OK': True,
 'H1:OBSERVATION_INTENT': True,
 'H1:NO_STOCH_HW_INJ': True,
 'H1:NO_CBC_HW_INJ': True,
 'H1:NO_BURST_HW_INJ': True,
 'H1:NO_DETCHAR_HW_INJ': True}
```

(task) gwcelery.tasks.detchar.**check_vectors**(*event*, *graceid*, *start*, *end*)

Perform data quality checks for an event and labels/logs results to GraceDb.

Depending on the pipeline, a certain amount of time (specified in `check_vector_prepost`) is appended to either side of the superevent start and end time. This is to catch DQ issues slightly before and after the event, such as that appearing in L1 just before GW170817.

A cache is then created for H1, L1, and V1, regardless of the detectors involved in the event. Then, the bits and channels specified in the configuration file (`llhoft_channels`) are checked. If an injection is found in the active detectors, 'INJ' is labeled to GraceDb. If an injection is found in any detector, a message with the injection found is logged to GraceDb. If no injections are found across all detectors, this is logged to GraceDb.

A similar task is performed for the DQ states described in the DMT-DQ_VECTOR, LIGO GDS-CALIB_STATE_VECTOR, and Virgo DQ_ANALYSIS_STATE_VECTOR. If no DQ issues are found in active detectors, 'DQOK' is labeled to GraceDb. Otherwise, 'DQV' is labeled. In all cases, the DQ states of all the state vectors checked are logged to GraceDb.

This skips MDC events.

Parameters

- **event** (*dict*) – Details of event.
- **graceid** (*str*) – GraceID of event to which to log.
- **end** (*start,*) – GPS start and end times desired.

6.3.5 gwcelery.tasks.em_bright module

Qualitative source classification for CBC events.

gwcelery.tasks.em_bright.**source_classification**(*m1*, *m2*, *c1*, *threshold=3.0*)

This is the place-holder function for the source classfication pipeline. In the future, the actual source classification pipeline will be integrated in three steps. First step will be the simple integration of the point-estimate code that will be using the em_progenitors code from PyCBC. In the second step, rapid_pe needs to be made Python3 compatible so that the ambiguity ellipsoid feature can be brough back into action. And, finally the O3 implementation will be incorporated which is currently a work in progress. This placeholder code will only act upon the mass2 point estimate value and classify the systems as whether they have a neutron or not. It does not attempt to classify for the remnant mass, returns a NaN value for that probability.

(task) gwcelery.tasks.em_bright.**classifier**(*coinc_psd*, *graceid*)

This function is currently actually calculating the simple source classification probability (m1 < 3.0 M_sun). In the future this code will call a classification code that will be put on lalinference.

6.3.6 gwcelery.tasks.external_triggers module

This module listens to the *GCNs* from SNEWS and the Fermi and Swift missions. It is also responsible for carrying out tasks related to external trigger-gravitational wave coincidences, including looking for temporal coincidences, creating combined GRB-GW sky localization probability maps, and computing their joint temporal and spatio-temporal false alarm rates.

There are two GCN and two LVAlert message handlers in the *~gwcelery.tasks.external_triggers* module:

- `handle_sn_gcn()` is called for each SNEWS GCN.
- `handle_grb_gcn()` is called for each Fermi and Swift GCN.
- `handle_sn_lvalert()` is called for each SNEWS external trigger and superevent LVAlert.
- `handle_grb_lvalert()` is called for each Fermi and Swift external trigger and superevent LVAlert.

Flow Chart

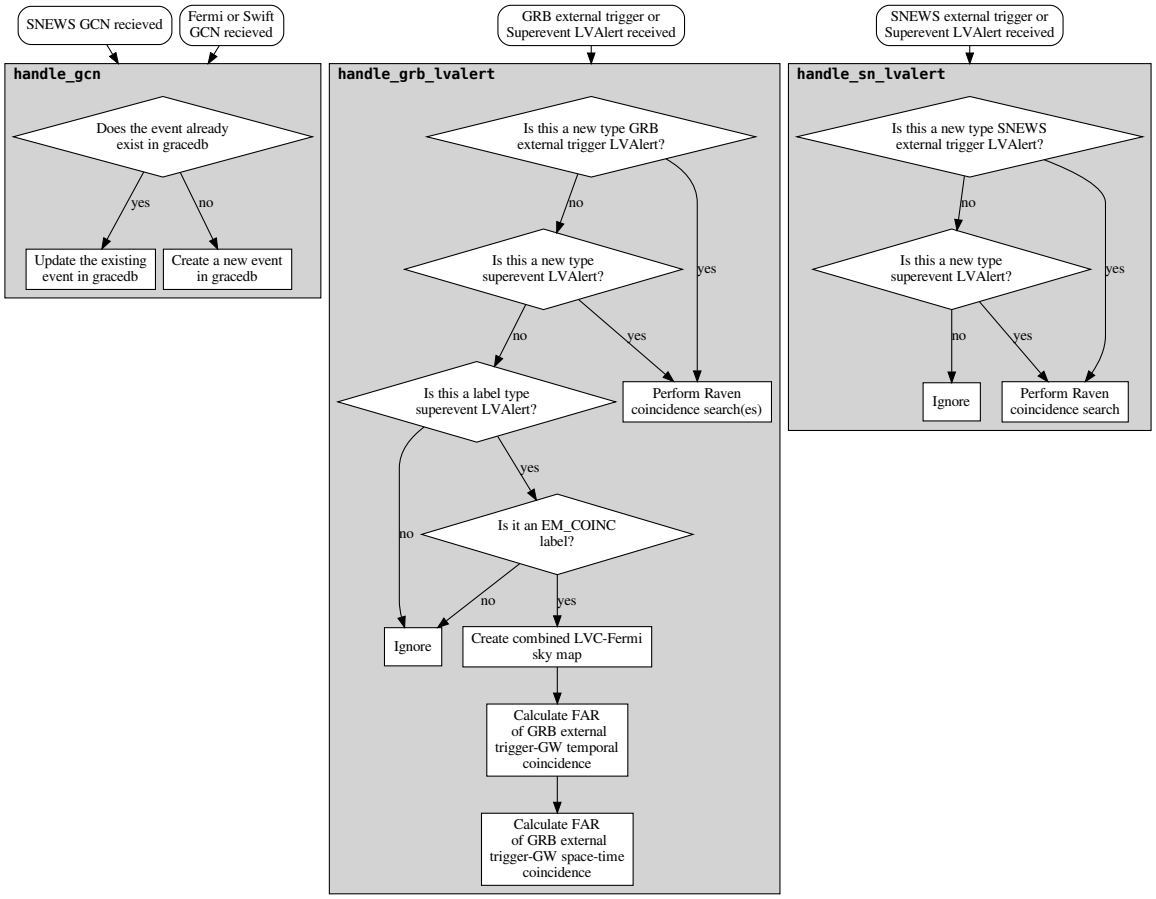

Tasks

(task) `gwcelery.tasks.external_triggers.`**`handle_sn_gcn`**(*payload*)

Handles the payload from SNEWS alerts. Prepares the alert to be sent to graceDB as 'E' events.

(task) `gwcelery.tasks.external_triggers.`**`handle_grb_gcn`**(*payload*)

Handles the payload from Fermi and Swift alerts. Prepares the alert to be sent to graceDB as 'E' events.

(task) `gwcelery.tasks.external_triggers.`**`handle_grb_lvalert`**(*alert*)

Parse an LVAlert message related to superevents/GRB external triggers and dispatch it to other tasks.

Notes

This LVAlert message handler is triggered by creating a new superevent or GRB external trigger event, or applying the `EM_COINC` label to any superevent:

- Any new event triggers a coincidence search with *`gwcelery.tasks.raven.coincidence_search()`*.

- The `EM_COINC` label triggers the creation of a combined GW-GRB sky map using *gwcelery.tasks.ligo_fermi_skymaps.create_combined_skymap()*.

(task) `gwcelery.tasks.external_triggers.`**`handle_sn_lvalert`**(*alert*)

Parse an LVAlert message related to superevents/SN external triggers and dispatch it to other tasks.

Notes

This LVAlert message handler is triggered by creating a new superevent or SN external trigger event, or applying the `EM_COINC` label to any superevent:

- Any new event triggers a coincidence search with *gwcelery.tasks.raven.coincidence_search()*.

6.3.7 gwcelery.tasks.first2years module

Create mock events from the "First Two Years" paper.

(task) `gwcelery.tasks.first2years.`**`pick_coinc`**

Pick a coincidence from the "First Two Years" paper.

(task) `gwcelery.tasks.first2years.`**`upload_event`**

Upload a random event from the "First Two Years" paper.

After 2 minutes, randomly either retract or confirm the event to send a retraction or initial notice respectively.

6.3.8 gwcelery.tasks.ligo_fermi_skymaps module

Create and upload LVC-Fermi sky maps.

`gwcelery.tasks.ligo_fermi_skymaps.`**`create_combined_skymap`**(*graceid*)

Creates and uploads the combined LVC-Fermi skymap. This also uploads the external trigger skymap to the external trigger GraceDb page.

(task) `gwcelery.tasks.ligo_fermi_skymaps.`**`get_preferred_skymap`**(*graceid*)

Get the LVC skymap fits filename. If not available, will try again 10 seconds later, then 20, then 40, etc. until up to 10 minutes after initial attempt.

(task) `gwcelery.tasks.ligo_fermi_skymaps.`**`combine_skymaps`**(*skymap1filebytes*, *skymap2filebytes*)

This task combines the two input skymaps, in this case the external trigger skymap and the LVC skymap and writes to a temporary output file. It then returns the contents of the file as a byte array.

(task) `gwcelery.tasks.ligo_fermi_skymaps.`**`external_trigger`**(*graceid*)

Returns the associated external trigger GraceDB ID.

(task) `gwcelery.tasks.ligo_fermi_skymaps.`**`external_trigger_heasarc`**(*external_id*)

Returns the HEASARC fits file link

(task) `gwcelery.tasks.ligo_fermi_skymaps.`**`get_external_skymap`**(*heasarc_link*)

Download the Fermi sky map fits file and return the contents as a byte array. If not available, will try again 10 seconds later, then 20, then 40, etc. until up to 10 minutes after initial attempt.

6.3.9 gwcelery.tasks.gcn module

Subsystem for sending, receiving, and processing Gamma-ray Coordinates Network [GCN] notices.

References

(task) gwcelery.tasks.gcn.**broker**
Run an embedded Comet VOEvent broker to send GCNs.

(task) gwcelery.tasks.gcn.**send**(*message*)
Send a VOEvent to the local Comet instance for forwarding to GCN.

Internally, this just calls comet-sendvo.

gwcelery.tasks.gcn.**handler = {<sphinx.ext.autodoc.importer._MockObject object>: [<@task:**
Function decorator to register a handler callback for specified GCN notice types. The decorated function is turned into a Celery task, which will be automatically called whenever a matching GCN notice is received.

Parameters

- ***keys** – List of GCN notice types to accept
- ****kwargs** – Additional keyword arguments for `celery.Celery.task()`.

Examples

Declare a new handler like this:

```
@gcn.handler(gcn.NoticeType.FERMI_GBM_GND_POS,
             gcn.NoticeType.FERMI_GBM_FIN_POS)
def handle_fermi(payload):
    root = lxml.etree.fromstring(payload)
    # do work here...
```

(task) gwcelery.tasks.gcn.**listen**
Listen to GCN notices forever. GCN notices are dispatched asynchronously to tasks that have been registered with *`gwcelery.tasks.gcn.handler()`*.

6.3.10 gwcelery.tasks.gcn.validate module

Validate LIGO/Virgo GCN notices to make sure that their contents match the original VOEvent notices that we sent.

(task) gwcelery.tasks.gcn.validate.**validate_voevent**(*payload*)
Check that the contents of a public LIGO/Virgo GCN matches the original VOEvent in GraceDB.

6.3.11 gwcelery.tasks.gracedb module

Communication with GraceDB.

gwcelery.tasks.gracedb.**task**(**args*, ***kwargs*)

(task) gwcelery.tasks.gracedb.**create_event**(*filecontents*, *search*, *pipeline*, *group*)
Create an event in GraceDb.

(task) gwcelery.tasks.gracedb.**create_label**(*label*, *graceid*)
Create a label in GraceDb.

(task) gwcelery.tasks.gracedb.**create_signoff**(*status*, *comment*, *signoff_type*, *graceid*)
Create a label in GraceDb.

(task) gwcelery.tasks.gracedb.**create_tag**(*filename*, *tag*, *graceid*)
Create a tag in GraceDb.

(task) gwcelery.tasks.gracedb.**create_voevent**(*graceid*, *voevent_type*, ***kwargs*)
Create a VOEvent.

Returns The filename of the new VOEvent.

Return type str

(task) gwcelery.tasks.gracedb.**download**(*filename*, *graceid*)
Download a file from GraceDB.

(task) gwcelery.tasks.gracedb.**expose**(*graceid*)
Expose an event to the public.

(task) gwcelery.tasks.gracedb.**get_events**(*query=None*, *orderby=None*, *count=None*, *columns=None*)
Get events from GraceDb.

(task) gwcelery.tasks.gracedb.**get_event**(*graceid*)
Retrieve an event from GraceDb.

(task) gwcelery.tasks.gracedb.**get_labels**(*graceid*)
Get all labels for an event in GraceDb.

(task) gwcelery.tasks.gracedb.**get_log**(*graceid*)
Get all log messages for an event in GraceDb.

(task) gwcelery.tasks.gracedb.**get_superevent**(*graceid*)
Retrieve a superevent from GraceDb.

(task) gwcelery.tasks.gracedb.**replace_event**(*graceid*, *payload*)
Get an event from GraceDb.

(task) gwcelery.tasks.gracedb.**upload**(*filecontents*, *filename*, *graceid*, *message*, *tags=None*)
Upload a file to GraceDB.

(task) gwcelery.tasks.gracedb.**get_superevents**(*query*)
List matching superevents in gracedb.

Parameters **query** (*str*) – query to be passed to `superevents()`

Returns **superevents** – The list of the superevents.

Return type list

(task) gwcelery.tasks.gracedb.**update_superevent**(*superevent_id*, *t_start=None*, *t_end=None*, *t_0=None*, *preferred_event=None*)
Update superevent information. Wrapper around `updateSuperevent()`

Parameters

- **superevent_id** (*str*) – superevent uid
- **t_start** (*float*) – start of superevent time window, unchanged if None
- **t_end** (*float*) – end of superevent time window, unchanged if None
- **t_0** (*float*) – superevent t_0, unchanged if None
- **preferred_event** (*str*) – uid of the preferred event, unchanged if None

(task) gwcelery.tasks.gracedb.**create_superevent**(*graceid*, *t0*, *d_t_start*, *d_t_end*, *category*)
Create new superevent in GraceDb with *graceid*

Parameters

- **graceid** (*str*) – graceid with which superevent is created.
- **t0** (*float*) – *t_0* parameter of superevent
- **d_t_start** (*float*) – superevent *t_start = t0 - d_t_start*
- **d_t_end** (*float*) – superevent *t_end = t0 + t_end*
- **category** (*str*) – superevent category

(task) gwcelery.tasks.gracedb.**add_event_to_superevent**(*superevent_id*, *graceid*)
Add an event to a superevent in GraceDb.

6.3.12 gwcelery.tasks.lalinference module

Source Parameter Estimation with LALInference.

(task) gwcelery.tasks.lalinference.**prepare_ini**(*event*, *superevent_id=None*)
Determine an appropriate PE settings for the target event and return ini file content

(task) gwcelery.tasks.lalinference.**dag_prepare**(*rundir*, *ini_contents*, *preferred_event_id*, *superevent_id*)
Create a Condor DAG to run LALInference on a given event.

Parameters

- **rundir** (*str*) – The path to a run directory where the DAG file exits
- **ini_contents** (*str*) – The content of online_pe.ini
- **preferred_event_id** (*str*) – The GraceDb ID of a target preferred event
- **superevent_id** (*str*) – The GraceDb ID of a target superevent

Returns **submit_file** – The path to the .sub file

Return type str

(task) gwcelery.tasks.lalinference.**job_error_notification**(*request*, *exc*, *traceback*, *superevent_id*)
Upload notification when condor.submit terminates unexpectedly.

Parameters

- **request** (*Context (placeholder)*) – Task request variables
- **exc** (*Exception*) – Exception rased by condor.submit
- **traceback** (*str (placeholder)*) – Traceback message from a task
- **superevent_id** (*str*) – The GraceDb ID of a target superevent

(task) gwcelery.tasks.lalinference.**upload_result**(*webdir*, *filename*, *graceid*, *message*, *tag*)
Upload a PE result

Parameters **graceid** (*str*) – The GraceDb ID.

(task) gwcelery.tasks.lalinference.**clean_up**(*rundir*)
Clean up a run directory.

Parameters **rundir** (*str*) – The path to a run directory where the DAG file exits

gwcelery.tasks.lalinference.**dag_finished**(*rundir*, *preferred_event_id*, *superevent_id*)
Upload PE results and clean up run directory

Parameters

- **rundir** (*str*) – The path to a run directory where the DAG file exits
- **preferred_event_id** (*str*) – The GraceDb ID of a target preferred event
- **superevent_id** (*str*) – The GraceDb ID of a target superevent

Returns **tasks** – The work-flow for uploading PE results

Return type canvas

(task) gwcelery.tasks.lalinference.**start_pe**(*ini_contents*, *preferred_event_id*, *superevent_id*)

Run LALInference on a given event.

Parameters

- **ini_contents** (*str*) – The content of online_pe.ini
- **preferred_event_id** (*str*) – The GraceDb ID of a target preferred event
- **superevent_id** (*str*) – The GraceDb ID of a target superevent

6.3.13 gwcelery.tasks.lvalert module

LVAlert client.

gwcelery.tasks.lvalert.**handler = {'burst_cwb': [<@task: gwcelery.tasks.superevents.handle**

Function decorator to register a handler callback for specified LVAlert message types. The decorated function is turned into a Celery task, which will be automatically called whenever a matching LVAlert message is received.

Parameters

- ***keys** – List of LVAlert message types to accept
- ****kwargs** – Additional keyword arguments for `celery.Celery.task()`.

Examples

Declare a new handler like this:

```
@lvalert.handler('cbc_gstlal',
                 'cbc_spiir',
                 'cbc_pycbc',
                 'cbc_mbtaonline')
def handle_cbc(alert_content):
    # do work here...
```

(task) gwcelery.tasks.lvalert.**listen**

Listen for LVAlert messages forever. LVAlert messages are dispatched asynchronously to tasks that have been registered with *gwcelery.tasks.lvalert.handler()*.

6.3.14 gwcelery.tasks.orchestrator module

This module implements the alert orchestrator, which responsible for the vetting and annotation workflow to produce preliminary, initial, and update alerts for gravitational-wave event candidates.

The orchestrator consists of two LVAlert message handlers:

- `handle_superevent()` is called for each superevent. It waits for a short duration of `orchestrator_timeout` seconds for the selection of the superevent by the `superevent manager` to stabilize, then performs data quality checks. If the data quality checks pass, then it calls `preliminary_alert()` to copy annotations from the preferred event and send the preliminary GCN notice.
- `handle_cbc_event()` is called for each CBC event. It performs some CBC-specific annotations that depend closely on the CBC matched-filter parameters estimates and that might influence selection of the preferred event: rapid sky localization with BAYESTAR and rapid source classification.

  Note that there is no equivalent of this task for burst events because both burst searches (cWB, LIB) have integrated source localization and have no other annotations.

Preliminary Alerts

The flow chart below illustrates the operation of these two tasks.

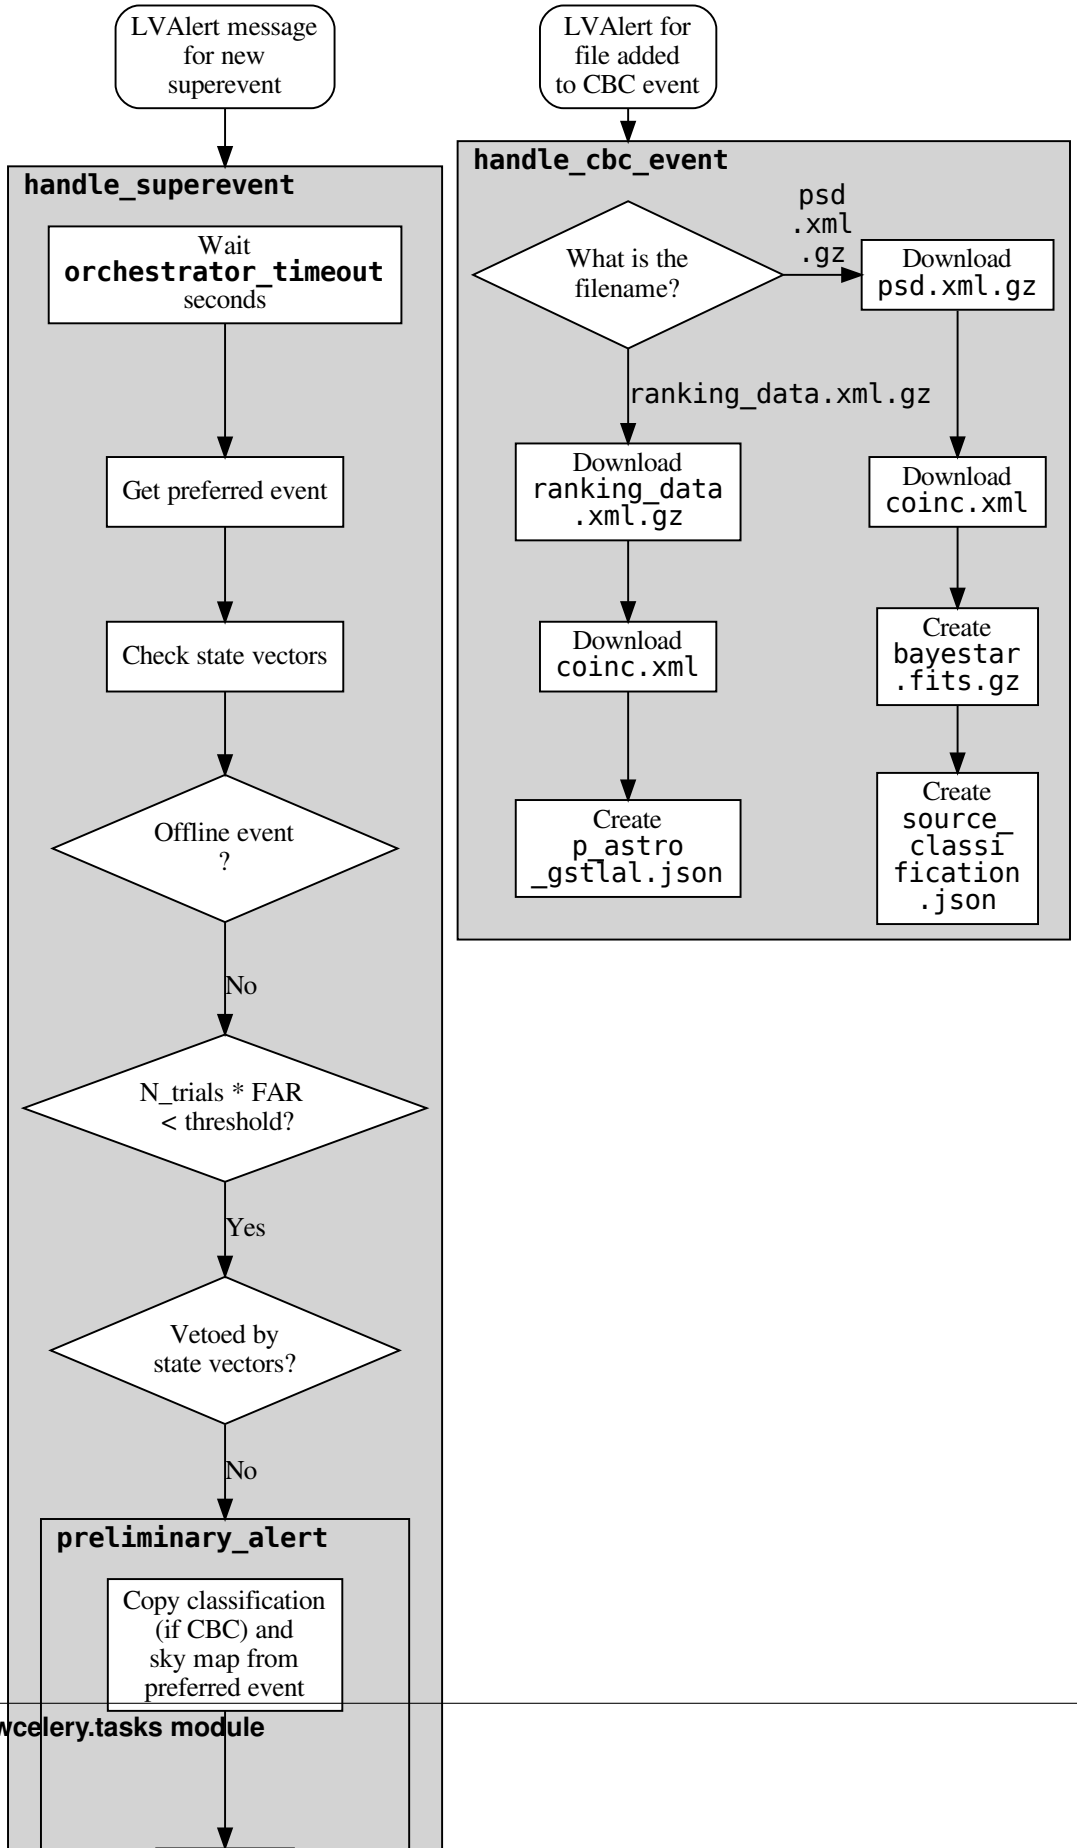
LVAlert message for new superevent
handle_superevent
Wait orchestrator_timeout seconds
Get preferred event
Check state vectors
Offline event ?
No
N_trials * FAR < threshold?
Yes
Vetoed by state vectors?
No
preliminary_alert
Copy classification (if CBC) and sky map from preferred event
LVAlert for file added to CBC event
handle_cbc_event
What is the filename?
psd .xml .gz
Download psd.xml.gz
ranking_data.xml.gz
Download ranking_data .xml.gz
Download coinc.xml
Download coinc.xml
Create bayestar .fits.gz
Create p_astro _gstlal.json
Create source_ classi fication .json

Initial and Update Alerts

The `initial_alert()` and `update_alert()` tasks create Initial and Update alerts respectively. At the moment, there is no handler or user interface to trigger these tasks, and they must be invoked manually (see *Command-Line Tools*). A flow chart for the initial alerts is shown below; the flow chart for update alerts is the same.

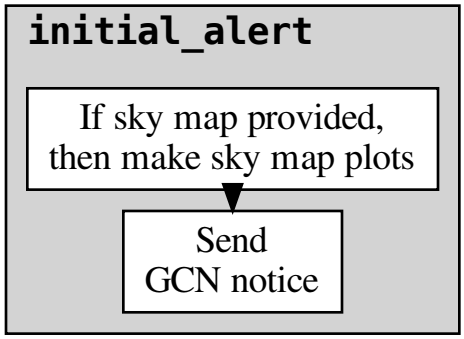

Retraction Alerts

Likewise, the `retraction_alert()` task creates Retraction alerts, and at the moment must be invoked manually. A flow chart is shown below.

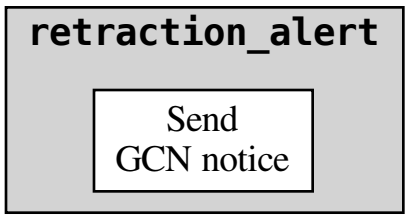

Tasks

Tasks that comprise the alert orchestrator, which responsible for the vetting and annotation workflow to produce preliminary, initial, and update alerts for gravitational-wave event candidates.

(task) `gwcelery.tasks.orchestrator.`**`handle_superevent`**(*alert*)

Schedule annotations for new superevents.

After waiting for a time specified by the `orchestrator_timeout` configuration variable for the choice of preferred event to settle down, this task peforms data quality checks with `gwcelery.tasks.detchar.check_vectors()` and calls `preliminary_alert()` to send a preliminary GCN notice.

(task) gwcelery.tasks.orchestrator.**handle_cbc_event**(*alert*)

Peform annotations for CBC events that depend on pipeline-specific matched-filter parameter estimates.

Notes

This LVAlert message handler is triggered by updates that include the files `psd.xml.gz` and `ranking_data.xml.gz`. The table below lists which files are created as a result, and which tasks generate them.

File	Task
`bayestar.fits`	`gwcelery.tasks.bayestar.localize()`
`source_classification.json`	`gwcelery.tasks.em_bright.classifier()`
`p_astro.json`	`gwcelery.tasks.p_astro_gstlal.compute_p_astro()`

(task) gwcelery.tasks.orchestrator.**preliminary_alert**(*event*, *superevent_id*)

Produce a preliminary alert by copying any sky maps.

This consists of the following steps:

1. Copy any sky maps and source classification from the preferred event to the superevent.
2. Create standard annotations for sky maps including all-sky plots by calling `gwcelery.tasks.skymaps.annotate_fits()`.
3. Create a preliminary VOEvent.
4. Send the VOEvent to GCN.
5. Apply the GCN_PRELIM_SENT label to the superevent.
6. Create and upload a GCN Circular draft.

(task) gwcelery.tasks.orchestrator.**parameter_estimation**(*event*, *superevent_id*)

Tasks for Parameter Estimation Followup with LALInference

This consists of the following steps:

1. Upload an ini file which is suitable for the target event.
2. Start Parameter Estimation if FAR is smaller than the PE threshold.

(task) gwcelery.tasks.orchestrator.**initial_or_update_alert**(*superevent_id*, *alert_type*, *skymap_filename=None*, *source_classification_filename=None*, *p_astro_filename=None*)

Create and send initial or update GCN notice.

Parameters

- **superevent_id** (*str*) – The superevent ID.
- **alert_type** (*{'initial', 'update'}*) – The alert type.

- **skymap_filename** (*str, optional*) – The sky map to send. If None, then most recent public sky map is used.
- **source_classification_filename** (*str, optional*) – The source classification file to use. If None, then most recent one is used.
- **p_astro_filename** (*str, optional*) – The p_astro file to use. If None, then most recent one is used.

(task) gwcelery.tasks.orchestrator.**initial_alert**(*superevent_id, skymap_filename=None, source_classification_filename=None, p_astro_filename=None*)

Produce an initial alert.

This does nothing more than call *initial_or_update_alert()* with `alert_type='initial'`.

Parameters

- **superevent_id** (*str*) – The superevent ID.
- **skymap_filename** (*str, optional*) – The sky map to send. If None, then most recent public sky map is used.
- **source_classification_filename** (*str, optional*) – The source classification file to use. If None, then most recent one is used.
- **p_astro_filename** (*str, optional*) – The p_astro file to use. If None, then most recent one is used.

(task) gwcelery.tasks.orchestrator.**update_alert**(*superevent_id, skymap_filename=None, source_classification_filename=None, p_astro_filename=None*)

Produce an update alert.

This does nothing more than call *initial_or_update_alert()* with `alert_type='update'`.

Parameters

- **superevent_id** (*str*) – The superevent ID.
- **skymap_filename** (*str, optional*) – The sky map to send. If None, then most recent public sky map is used.
- **source_classification_filename** (*str, optional*) – The source classification file to use. If None, then most recent one is used.
- **p_astro_filename** (*str, optional*) – The p_astro file to use. If None, then most recent one is used.

(task) gwcelery.tasks.orchestrator.**retraction_alert**(*superevent_id*)

Produce a retraction alert. This is currently just a stub and does nothing more than create and send a VOEvent.

6.3.15 gwcelery.tasks.p_astro_gstlal module

Module containing the computation of p_astro by source category See https://dcc.ligo.org/LIGO-T1800072 for details.

gwcelery.tasks.p_astro_gstlal.**p_astro_update**(*category, event_bayesfac_dict, mean_values_dict*)

(task) gwcelery.tasks.p_astro_gstlal.**compute_p_astro**(*files*)

Task to compute *p_astro* by source category.

Parameters **files** (*tuple*) – Tuple of byte content from (coinc.xml, ranking_data.xml.gz)

Returns **p_astros** – JSON dump of the p_astro by source category

Return type str

Example

```
>>> p_astros = json.loads(compute_p_astro(files))
>>> p_astros
{'BNS': 0.999, 'BBH': 0.0, 'NSBH': 0.0, 'Terr': 0.001}
```

6.3.16 gwcelery.tasks.raven module

Search for GRB-GW coincidences with ligo-raven.

`gwcelery.tasks.raven.`**`calculate_spacetime_coincidence_far`**(*gracedb_id*, *group*)

Compute spatio-temporal coincidence FAR for GRB external trigger and superevent coincidence by calling ligo.raven.search.calc_signif_gracedb. Note: this will only run if skymaps from both triggers are available to download.

Parameters

- **gracedb_id** (*str*) – ID of the superevent trigger used by GraceDb
- **group** (*str*) – CBC or Burst; group of the preferred_event associated with the gracedb_id superevent

`gwcelery.tasks.raven.`**`calculate_coincidence_far`**(*gracedb_id*, *group*)

Compute temporal coincidence FAR for external trigger and superevent coincidence by calling ligo.raven.search.calc_signif_gracedb.

Parameters

- **gracedb_id** (*str*) – ID of the superevent trigger used by GraceDb
- **group** (*str*) – CBC or Burst; group of the preferred_event associated with the gracedb_id superevent

(task) `gwcelery.tasks.raven.`**`calc_signif`**(*se*, *exttrig*, *tl*, *th*, *incl_sky*)

Calculate FAR of GRB exttrig-GW coincidence

`gwcelery.tasks.raven.`**`coincidence_search`**(*gracedb_id*, *alert_object*, *group=None*, *pipelines=[]*)

Perform ligo-raven search for coincidences. The ligo.raven.search.search method applies EM_COINC label on its own.

Parameters

- **gracedb_id** (*str*) – ID of the trigger used by GraceDb
- **alert_object** (*dict*) – lvalert['object']
- **group** (*str*) – Burst or CBC
- **pipelines** (*list*) – list of external trigger pipeline names

(task) `gwcelery.tasks.raven.`**`search`**(*gracedb_id*, *alert_object*, *tl=-5*, *th=5*, *group=None*, *pipelines=[]*)

Perform ligo-raven search for coincidences. The ligo.raven.search.search method applies EM_COINC label on its own.

Parameters

- **gracedb_id** (*str*) – ID of the trigger used by GraceDb
- **alert_object** (*dict*) – lvalert['object']
- **tl** (*int*) – number of seconds to search before
- **th** (*int*) – number of seconds to search after
- **group** (*str*) – Burst or CBC
- **pipelines** (*list*) – list of external trigger pipelines for performing coincidence search against

Returns

Return type list with the dictionaries of related gracedb events

(task) gwcelery.tasks.raven.**add_exttrig_to_superevent**(*raven_search_results*, *gracedb_id*)

Add external trigger to the list of em_events after ligo.raven.search.search finds a coincidence

Parameters

- **raven_search_results** (*list*) – list of dictionaries of each related gracedb trigger
- **gracedb_id** (*str*) – ID of either a superevent or external trigger

6.3.17 gwcelery.tasks.skymaps module

Annotations for sky maps.

gwcelery.tasks.skymaps.**annotate_fits**(*versioned_filename*, *graceid*, *tags*)

Perform annotations on a sky map.

This function downloads a FITS file and then generates and uploads all derived images as well as an HTML dump of the FITS header.

gwcelery.tasks.skymaps.**is_3d_fits_file**(*filecontents*)

Determine if a FITS file has distance information.

(task) gwcelery.tasks.skymaps.**annotate_fits_volume**(*filecontents*, **args*)

Perform annotations that are specific to 3D sky maps.

(task) gwcelery.tasks.skymaps.**fits_header**(*filecontents*, *filename*)

Dump FITS header to HTML.

(task) gwcelery.tasks.skymaps.**plot_allsky**(*filecontents*)

Plot a Mollweide projection of a sky map using the command-line tool ligo-skymap-plot.

(task) gwcelery.tasks.skymaps.**plot_volume**(*filecontents*)

Plot a 3D volume rendering of a sky map using the command-line tool ligo-skymap-plot-volume.

(task) gwcelery.tasks.skymaps.**flatten**(*filecontents*, *filename*)

Convert a HEALPix FITS file from multi-resolution UNIQ indexing to the more common IMPLICIT indexing using the command-line tool ligo-skymap-flatten.

6.3.18 gwcelery.tasks.superevents module

Superevents are a new abstraction of gravitational-wave candidates introduced in the third LIGO/Virgo observing (O3). Each superevent is intended to represent a single astrophysical event. A superevent consists of one or more

event candidates, possibly from different pipelines, that are neighbors in `gpstime`. One event belonging to the superevent is identified as the preferred event.

Flow Chart

The flow chart below illustrates the decision process for selection of the preferred event.

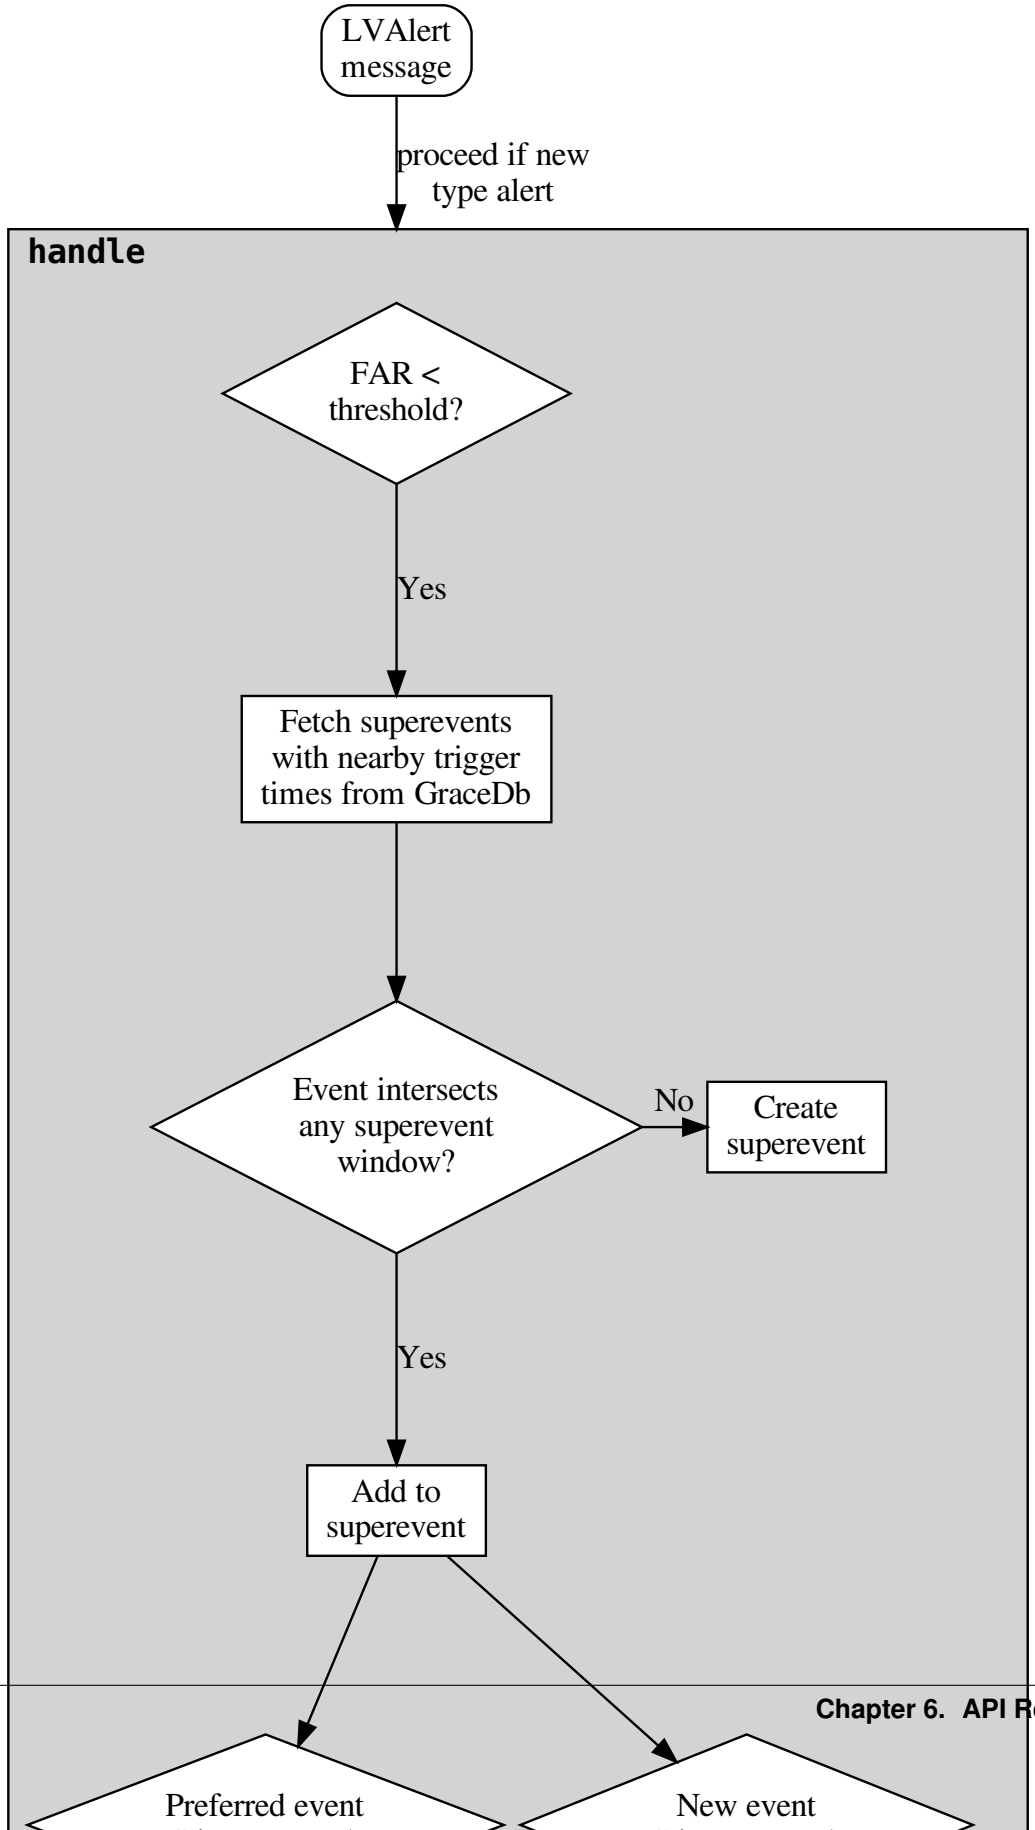
LVAlert message
proceed if new type alert
handle
FAR < threshold?
Yes
Fetch superevents with nearby trigger times from GraceDb
Event intersects any superevent window?
No
Create superevent
Yes
Add to superevent
Preferred event
New event

Tasks

Module containing the functionality for creation and management of superevents.

- There is serial processing of triggers from low latency pipelines.
- Dedicated **superevent** queue for this purpose.
- Primary logic to respond to low latency triggers contained in `handle()` function.

(task) `gwcelery.tasks.superevents.`**`handle`**(*payload*)

LVAlert handler for superevent manager. Recieves payload from test and production nodes and serially processes them to create/modify superevents

6.4 gwcelery.util module

Miscellaneous utilities that are useful inside many different tasks.

class `gwcelery.util.`**`PromiseProxy`**(**args*, ***kwargs*)

Bases: `object`

`gwcelery.util.`**`NamedTemporaryFile`**(*content=None*, ***kwargs*)

Convenience wrapper for `tempfile.NamedTemporaryFile()` that writes some data to the file before handing it to the calling code.

Parameters

- **`content`** (*`str`, `bytes`, `None`*) – Initial contents of the file.
- **`**kwargs`** – Additional keyword arguments to pass to `tempfile.NamedTemporaryFile()`.

CHAPTER 7

Contributing

Contributors may familiarize themselves with Celery itself by going through the First Steps with Celery tutorial.

7.1 Development model

GWCelery operates on a fork-and-merge development model (see GitLab basics for an introduction).

To contribute to GWCelery development, follow these steps:

1. Create a personal fork of GWCelery.
2. Make your changes on a branch.
3. Open a merge request.

Note that GWCelery uses fast-forward merges.

7.2 Where new code should go

New code will generally consist of adding Celery tasks. Tasks are organized by functionality into submodules of `gwcelery.tasks`. If your new task does not match with one of the existing submodules, please create a new submodule.

7.3 Guidelines for tasks

- **Tasks should be short.** When deciding where a new task should go, start from the following loose rules of thumb:
    1. If it's less than a screenful of code, and related to functionality in an existing module, then put the code in a new task in that module.

2. If it's up to a few screenfuls of code, or not related to functionality in an existing module, then try to break it into a few smaller functions or tasks and put it in a new module.
3. If it's more than a few screenfuls of code, or adds many additional dependencies, then it should go in a separate package.

See also the note on Granularity in the Celery manual's Tips and Best Practices section.

- **Tasks should avoid saving files to disk.** Output should be placed directly in GraceDb. Temporary files that are written in `/tmp` are OK but should be cleaned up promptly.

  See also the Celery manual's notes on Data locality and State.

- **Dependencies should be installable by pip.** Dependencies of tasks should be listed in the requirements.txt file so that they are installed automatically when GWCelery is installed with pip.

  There are two extra steps involved in making changes to the dependencies:

  1. The Sphinx-generated documentation (that is to say, this manual) is generally built without most of the dependencies installed. Whenever you add a new package to requirements.txt, you should also add any modules that are imported from that package to the `autodoc_mock_imports` list in the Sphinx configuration file, doc/conf.py.
  2. We use pipenv to make the precise versions of packages reproducible in our deployment. If you make changes to requirements.txt, then run `pipenv update` and commit the changes to Pipfile.lock.

7.4 Unit tests

Unit tests and code coverage measurement are run automatically for every branch and for every merge request. New code contributions must have 100% test coverage. Modifications to existing code must not decrease test coverage. To run the unit tests and measure code coverage, run the following commands in the top directory of your local source checkout:

```
$ pip install pytest-cov
$ python setup.py test --addopts='--cov --cov-report html'
```

This will save a coverage report that you can view in a web browser as `htmlcov/index.html`.

7.5 Code style

Code should be written in the **PEP 8** style and must pass linting by Flake8. To check code style, run the following commands in the top of your source directory:

```
$ pip install flake8 pep8-naming
$ flake8 --show-source .
```

7.6 Documentation

Documentation strings should be written in the Numpydoc style.

CHAPTER 8

Deployment

8.1 Continuous deployment

GWCelery is automatically deployed using GitLab's continuous deployment features, configured through the project's .gitlab-ci.yml file. Deployment can be managed through the GitLab project's Environments page.

Python dependencies in the deployment environment are managed automatically using pipenv. Note that there are some runtime environment settings such as the `~/.bashrc` file that are managed in the repository https://git.ligo.org/emfollow/deployment/ligo.caltech.edu, which is a holdover from a previous continuous deployment strategy.

There are two instances of GWCelery that are running on the LIGO-Caltech computing cluster and that are managed in this manner:

- **Playground**: The playground instance is re-deployed *on every push to master that passes the unit tests*. It uses the `gwcelery.conf.playground` configuration preset.
- **Production**: The production instance is re-deployed *only when manually triggered through GitLab*. It uses the `gwcelery.conf.production` configuration preset.

When we observe that the Playground instance shows correct end-to-end behavior, we have the option of triggering a re-deployment to Production. Deployment to production should preferably occur at a release. The procedure for performing a release is described below.

Danger: It is possible to start an interactive session inside the GWCelery production environment by logging in to the LIGO-Caltech cluster, but this measure should be **reserved for emergencies only**.

Any manual changes to the environment **may disrupt the logging and monitoring subsystems**. Any files that are manually changed, added to, or removed from the deployment environment **will not be captured in version control** and may be **rolled back without warning** the next time that the continuous deployment is triggered.

8.2 Making a new release

We always prepare releases from the tip of the `master` branch. GitLab is configured through the project's .gitlab-ci.yml file to automatically build and push any tagged release to the Python Package Index (PyPI). Follow these steps when issuing a release in order to maintain a consistent and orderly change log.

1. **Check the pipeline status.** Before you begin, first make sure that the unit tests, documentation, and packaging jobs are passing. Consult the project's GitLab pipeline status to make sure that all of the continuous integration jobs are passing on `master`.

   If necessary, fix any bugs that are preventing the pipeline from passing, push the changes to master, and repeat until all jobs pass.

2. **Update the change log.** The first subsection of the change log file, CHANGES.rst, should have the title *`MAJOR.MINOR.PATCH`* `(unreleased)`, where *`MAJOR.MINOR.PATCH`* will be the version number of the new release. Review the git commit log.

   Make any necessary changes to CHANGES.rst so that this subsection of the change log accurately summarizes all of the significant changes since the last release and is free of spelling, grammatical, or reStructuredText formatting errors.

   Review the list of changes and make sure that the new version number is appropriate. We follow SemVer *very* loosely, and also generally bump at least the minor version number at the start of a new LSC/Virgo engineering or observing run.

   Commit and push any corrections to CHANGES.rst to `master`.

3. **Tag the release.** Change the title of the first section of CHANGES.rst to *`MAJOR.MINOR.PATCH`* `(`*`YYYY-MM-DD`*`)` where *`YYYY-MM-DD`* is today's date. Commit with the message `Update changelog for version` *`MAJOR.MINOR.PATCH`*.

   Create a git tag to mark the release by running the following command:

   ```
   $ git tag vMAJOR.MINOR.PATCH -m "Version MAJOR.MINOR.PATCH"
   ```

4. **Create a change log section for the next release.** Add a new section to CHANGES.rst with the title *`NEXT_MAJOR.NEXT_MINOR.NEXT_PATCH`* `(unreleased)`, where *`NEXT_MAJOR.NEXT_MINOR.NEXT_PATCH`* is a provisional version number for the next release. Add a single list item with the text `No changes yet.` Commit with the message `Back to development.`

5. **Push the new tag and updated change log.** Push the new tag and updated change log:

   ```
   git push && git push --tags
   ```

6. Wait a couple minutes, and then verify that the new release has been published on our PyPI project page, https://pypi.org/project/gwcelery/.

7. If desired, navigate to the GitLab project's Environments page and trigger a deployment to production.

CHAPTER 9

Changelog

9.1 0.2.6 (2019-02-12)

- Report an environment tag to Sentry corresponding to the GWCelery configuration module (`production`, `test`, `playground`, or `development`) in order to differentiate log messages from different deployments.
- The `gwcelery condor` command now identifies jobs that it owns by matching both the job batch name and the working directory. This makes it possible to run multiple isolated instances of GWCelery under HTCondor on the same cluster in different working directories.
- Change the conditions for starting parameter estimation. For every CBC superevent, create an `online_pe.ini` file suitable for starting LALInference. However, only start LALInference if the false alarm rate is less than once per 2 weeks.
- Determine PSD segment length for LALInference automatically based on data availability and data quality.
- Add a Flask-based web interface for manually triggering certain tasks such as sending updated GCN notices.

9.2 0.2.5 (2019-02-01)

- Pass along the GWCelery version number to Sentry.
- Upload stdout and stderr when dag creation fails and notifications when submitted job fails in Parameter Estimation
- Allow detchar module's `create_cache` to use gwdatafind when frames are no longer in llhoft.
- The Nagios monitoring plugin will now report on the status of LVAlert subscriptions.
- Change trials factor to 5 for both CBC and Burst categories. CBC includes the 4 CBC pipelines. Burst includes the 4 searches performed in total by the 2 Burst pipelines. An additional external coincidence search.
- Automatically set up PE ini file depending on source parameters reported by detection pipelines.

9.3 0.2.4 (2018-12-17)

- Fix broken links in log messages due to changes in GraceDb URL routes.
- Whenever we send a public VOEvent using GCN, also make the corresponding VOEvent file in GraceDb public.
- Don't include Mollweide projection PNG file in VOEvents. The sky map visualizations take longer to generate than the FITS files themselves, so they were unnecessarily slowing down the preliminary alerts.
- Preliminary GCN FAR threshold is modified to be group (CBC, Burst, Test) specific.

9.4 0.2.3 (2018-12-16)

- Update frame type used in LALInference Parameter Estimation.
- Handle cases where `p_astro_gstlal.compute_p_astro` returns NaNs by falling back to `p_astro_other.compute_p_astro`.
- Fix a bug that prevented annotations that are specific to 3D sky maps from being performed for multi-resolution FITS files.
- Fetch the graceid for the new event added from the gracedb logs since superevent packet does not provide information as to which event is added in case of type event_added.

9.5 0.2.2 (2018-12-14)

- Add error handling for nonexistent iDQ frames in detchar module.

9.6 0.2.1 (2018-12-14)

- Update detchar module configuration for ER13.

9.7 0.2.0 (2018-12-14)

- This is the release of GWCelery for ER13.
- Run two separate instances of Comet, one to act as a broker and one to act as a client. This breaks a cycle that would cause retransmission of GRB notices back to GCN.
- Fix a race condition that could cause preliminary alerts to be sent out for events for which data quality checks had failed.
- Unpin the `redis` package version because recent updates to Kombu and Billiard seem to have fixed the Nagios unit tests.
- Start the Comet VOEvent broker as a subprocess intead of using `multiprocessing` and go back to using PyGCN instead of Comet as the VOEvent client. This is a workaround for suspected instability due to a bad interaction between `redis-py` and `multiprocessing`.
- Reset Matplotlib's style before running `ligo-skymap-plot` and `ligo-skymap-plot-volume`. There is some other module (probably in LALSuite) that is messing with the rcparams at module scope, which was causing Mollweide plots to come out with unusual aspect ratios.

- Run `check_vectors` upon addition of an event to a superevent if the superevent already has an `DQV` label.
- Do not check the DMT-DQ_VECTOR for pipelines which use gated h(t).
- Remove static example VOEvents from the Open Alert Users Guide. We never used them because activating sample alerts got help until ER13.
- Disable running the Orchestrator for test events for ER13. After ER13 is over, we need to carefully audit the code and make sure that test events are handled appropriately.
- Enable public GraceDb entries and public GCNs for mock (MDC) events. For **real** events in ER13, disable public preliminary GCNs. Instead, advocate signoffs will trigger making events and GCN notices public: `ADVOK` for initial notices and `ADVNO` for retraction notices.
- Include source classification output (BNS/NSBH/BBH/Terrestrial) in GCN Notices.

9.8 0.1.7 (2018-11-27)

- Pin the `redis` package version at <3 because the latest version of redis breaks the Nagios unit tests.
- Ditch our own homebrew VOEvent broker and use Comet instead.
- In addition to traditional flat, fixed-nside sky maps, BAYESTAR will now also upload an experimental multiresolution format described in LIGO-G1800186-v4.

9.9 0.1.6 (2018-11-14)

- Update URL for static example event.

9.10 0.1.5 (2018-11-13)

- Add tasks for submitting HTCondor DAGs.
- Add a new module, `gwcelery.tasks.lalinference`, which provides tasks to start parameter estimation with LALInference and upload the results to GraceDB.
- Depend on lalsuite nightly build from 2018-11-04 to pick up changes to LALInference for Python 3 support.
- Send static example VOEvents from the Open Alert Users Guide. This will provide a stream of example alerts for astronomers until GraceDb is ready for public access.
- Add trials factor correction to the event FAR when comparing against FAR threshold to send out preliminary GCN.
- Require that LIGO/Virgo VOEvents that we receive from GCN match the original VOEvents from GraceDb byte-for-byte, since GCN will now pass through our VOEvents without modification.

9.11 0.1.4 (2018-10-29)

- Work around a bug in astropy.visualization.wcsaxes that affected all-sky plots when Matplotlib's `text.usetex` rcparam is set to `True` (https://github.com/astropy/astropy/issues/8004). This bug has evidently been present since at least astropy 1.3, but was not being triggered until recently: it is likely that some other package that we import (e.g. lalsuite) is now globally setting `text.usetex` to `True`.

- A try except is added around updateSuperevent to handle a bad request error from server side when updating superevent parameters which have nearby values.
- Send automatic preliminary alerts only for events with a false alarm rate below a maximum value specified by a new configuration variable, `preliminary_alert_far_threshold`.
- State vector vetoes will not suppress processing of preliminary sky maps and source classification. They will still suppress sending preliminary alerts.
- Set `open_alert` to `True` for all automated VOEvents.

9.12 0.1.3 (2018-10-26)

- Preliminary GCN is not sent for superevents created from offline gw events.
- Add `dqr_json` function to `gwcelery.tasks.detchar`, which uploads a DQR-compatible json to GraceDb with the results of the detchar checks.
- Depend on ligo.skymap >= 0.0.17.
- Fix a bug in sending initial, update, and retraction GCN notices: we were sending the VOEvent filenames instead of the file contents.

9.13 0.1.2 (2018-10-11)

- Setted `vetted` flag to true for all initial, update, and retraction alerts that are triggered by GraceDb signoffs.
- Write GraceDb signoffs, instead of just labels, to simulate initial and retraction alerts for mock events, because merely creating the `ADVNO` or `ADVOK` label does not cause GraceDb to erase the `ADVREQ` label. This change makes mock alerts more realistic.
- Change filename of cWB sky maps from `skyprobcc_cWB.fits` to `cWB.fits.gz` for consistency with other pipelines.
- Any time that we send a VOEvent, first change the GraceDb permissions on the corresponding superevent so that it is visible to the public. Note that this has no effect during the ongoing software engineering runs because LVEM and unauthenticated access are currently disabled in GraceDb.

9.14 0.1.1 (2018-10-04)

- Use the `public` tag instead of the `lvem` tag to mark preliminary sky maps for public access rather than LV-EM partner access. Note that GraceDb has not yet actually implemented unauthenticated access, so this should have no effect during our ongoing software engineering runs.
- Add `check_idq` function to detchar module, which reads probabilities generated by iDQ.
- Automated `DQV` labels should not trigger retraction notices because they prevent preliminary notices from being sent in the first place.
- The criterion for selecting a superevent's preferred event now prefers multiple-detector events to single-detector events, with precedence over source type (CBC versus burst). Any remaining tie is broken by using SNR for CBC and FAR for Burst triggers.
- By default, initial and update alerts will find and send the most recently added public sky map.

- The initial and update sky maps no longer perform sky map annotations, because they would only be duplicating the annotations performed as part of the preliminary alert.
- Mock events now include example initial and retraction notices. Two minutes after each mock event is uploaded, there will be either an `ADVOK` or an `ADVNO` label applied at random, triggering either an initial or a retraction notice respectively.
- Depend on ligo-gracedb >= 2.0.1 in order to pull in a bug fix for VOEvents with ProbHasNS or ProbHasRemnant set to 0.0.
- Use the `sentry-sdk` package instead of the deprecated `raven` package for Sentry integration.

9.15 0.1.0 (2018-09-26)

- Separated the external GCN listening handlers into two: one that listens to GCNs about SNEWS triggers and another that listens to Fermi and Swift.
- Fixed calls to the raven temporal coincidence search so that search results separate SNEWS triggers from Fermi and Swift.
- Add space-time FAR calculation for GRB and GW superevent coincidences. This only runs when skymaps from both triggers are available to download.
- Add human vetting for initial GCN notices. For each new superevent that passes state vector checks, the `ADVREQ` label is applied. Rapid response team users should set their GraceDb notification preferences to alert them on `ADVREQ` labels. If a user sets the `ADVOK` label, then an initial notice is issued. If a user sets the `ADVNO` label, then a retraction notice is issued.
- Update the LVAlert host for gracedb-playground.ligo.org.
- Add experimental integration with Sentry for log aggregation and error reporting.
- Track API and LVAlert schema changes in ligo-gracedb 2.0.0.

9.16 0.0.31 (2018-09-04)

- Refactor external trigger handling to separate it from the orchestrator.
- Fixed a bug in the VOEvent broker to only issue "iamalive" messages after sending the first VOEvent.
- Pass group argument to set time windows appropriately when performing raven coincidence searches. Search in the [-600, 60]s range and [-5, 1]s range around external triggers for Burst events and CBC events respectively. Similarly, search in the [-60, 600]s and [-1, 5]s range around Burst and CBC events for external triggers.
- Compute and upload FAR for GRB external trigger/superevent coincidence upon receipt of the EM_COINC label application to a superevent.
- Add continuous integration testing for Python 3.7, and run test suite against all supported Python versions (3.6, 3.7).
- Update ligo.skymap to 0.0.15.

9.17 0.0.30 (2018-08-02)

- Manage superevents for production, test, and MDC events separately.

- Add some more validation of LIGO/Virgo VOEvents from GCN.
- Remove now-unused task `gwcelery.tasks.orchestartor.continue_if`.
- Add `check_vectors` run for external triggers.
- Change the preferred event selection criteria for burst events to be FAR instead of SNR.
- Add `gwcelery nagios` subcommand for Nagios monitoring.
- Incorporate Virgo DQ veto streams into `check_vectors`
- Update ligo-raven to 1.3 and ligo-followup-advocate to 0.0.11.

9.18 0.0.29 (2018-07-31)

- Add a workflow graph to superevents module documentation.
- Add `gwcelery condor resubmit` as a shortcut for `gwcelery condor rm; gwcelery condor submit`.
- Fix deprecation warning due to renaming of `ligo.gracedb.rest.Gracedb.createTag` to `ligo.gracedb.rest.Gracedb.addTag`.
- Update ligo-gracedb to 2.0.0.dev1.

9.19 0.0.28 (2018-07-25)

- Add injection checks to `check_vector`.
- Bitmasks are now defined symbolically in `detchar`.
- Refactor configuration so that it is possible to customize settings through an environment variable.

9.20 0.0.27 (2018-07-22)

- The preferred event for superevents is now decided based on higher SNR value instead of lower FAR in the case of a tie between groups.
- A check for the existence of the gstlal trigger database is performed so that compute_p_astro does not return None.

9.21 0.0.26 (2018-07-20)

- Fix spelling of the label that is applied to events after p_astro finishes, changed from `P_ASTRO_READY` to `PASTRO_READY`.
- Run p_astro calculation for mock events.
- Overhaul preliminary alert pipeline so that it is mostly feature complete for both CBC and Burst events, and uses a common code path for both types. Sky map annotations now occur for both CBC and Burst localizations.

- Switch to using the pre-registered port 8096 for receiving proprietary LIGO/Virgo alerts on emfollow.ligo.caltech.edu. This means that the capability to receive GCNs requires setting up a site configuration in advance with Scott Barthelmey.

  Once we switch to sending public alerts exclusively, then we can switch back to using port 8099 for anonymous access, requiring no prior site configuration.

9.22 0.0.25 (2018-07-19)

- Reintroduce pipeline-dependent pre/post peeks for `check_vector` after fixing issue where pipeline information was being looked for in the wrong dictionary.
- `check_vector` checks all detectors regardless of instruments used, but only appends labels based on active instruments.
- Fix a few issues in the GCN broker:
  - Decrease the frequency of keepalive ("iamalive" in VOEvent Transport Protocol parlance) packets from once a second to once a minute at the request of Scott Barthelmey.
  - Fix a possible race condition that might have caused queued VOEvents to be thrown away unsent shortly after a scheduled keepalive packet.
  - Consume and ignore all keepalive and ack packets from the client so that the receive buffer does not overrun.
- Add `p_astro` computation for `gstlal` pipeline. The copmutation is launched for all cbc_gstlal triggers.

9.23 0.0.24 (2018-07-18)

- Revert pipeline-dependent pre/post peeks for `check_vector` because they introduced a regression: it caused the orchestrator failed without running any annotations.

9.24 0.0.23 (2018-07-18)

- Add timeout and keepalive messages to GCN broker.
- Update ligo-gracedb to 2.0.0.dev0 and ligo.skymap to 0.0.12.
- Add superevent duration for gstlal-spiir pipeline.
- Fix fallback for determining superevent duration for unknown pipelines.
- Make `check_vector` pre/post peeks pipeline dependent.

9.25 0.0.22 (2018-07-11)

- Process gstlal-spiir events.
- Create combined LVC-Fermi skymap in case of coincident triggers and upload to GraceDb superevent page. Also upload the original external trigger sky map to the external trigger GraceDb page.
- Generalize conditional processing of complex canvases by replacing the `continue_if_group_is()` task with a more general task that can be used like `continue_if(group='CBC')`.

- Add a `check_vector_prepost` configuration variable to control how much padding is added around an event for querying the state vector time series.

  This should have the beneficial side effect of fixing some crashes for burst events, for which the bare duration of the superevent segment was less than one sample.

9.26 0.0.21 (2018-07-10)

- MBTA events in GraceDb leave the `search` field blank. Work around this in `gwcelery.tasks.detchar.check_vectors` where we expected the field to be present.
- Track change in GraceDb JSON response for VOEvent creation.

9.27 0.0.20 (2018-07-09)

- After fixing some minor bugs in code that had not yet been tested live, sending VOEvents to GCN now works.

9.28 0.0.19 (2018-07-09)

- Rewrite the GCN broker so that it does not require a dedicated worker.
- Send VOEvents for preliminary alerts to GCN.
- Only perform state vector checks for detectors that were online, according to the preferred event.
- Exclude mock data challenge events from state vector checks.

9.29 0.0.18 (2018-07-06)

- Add detector state vector checks to the preliminary alert workflow.

9.30 0.0.17 (2018-07-05)

- Undo accidental configuration change in last version.

9.31 0.0.16 (2018-07-05)

- Stop listening for three unnecessary GCN notice types: `SWIFT_BAT_ALARM_LONG`, `SWIFT_BAT_ALARM_SHORT`, and `SWIFT_BAT_KNOWN_SRC`.
- Switch to SleekXMPP for the LVAlert client, instead of PyXMPP2. Because SleekXMPP has first-class support for publish-subscribe, the LVAlert listener can now automatically subscribe to all LVAlert nodes for which our code has handlers. Most of the client code now lives in a new external package, sleek-lvalert.

9.32 0.0.15 (2018-06-29)

- Change superevent threshold and mock event rate to once per hour.
- Add `gracedb.create_label` task.
- Always upload external triggers to the 'External' group.
- Add rudimentary burst event workflow to orchestrator: it just generates VOEvents and circulars.
- Create a label in GraceDb whenever `em_bright` or `bayestar` completes.

9.33 0.0.14 (2018-06-28)

- Fix typo that was causing a task to fail.
- Decrease orchestrator timeout to 15 seconds.

9.34 0.0.13 (2018-06-28)

- Change FAR threshold for creation of superevents to 1 per day.
- Update ligo-followup-advocate to >= 0.0.10. Re-enable automatic generation of GCN circulars.
- Add "EM bright" classification. This is rudimentary and based only on the point mass estimates from the search pipeline because some of the EM bright classifier's dependencies are not yet ready for Python 3.
- Added logic to select CBC events as preferred event over Burst. FAR acts as tie breaker when groups for preferred event and new event match.
- BAYESTAR now adds GraceDb URLs of events to FITS headers.

9.35 0.0.12 (2018-06-28)

- Prevent receiving duplicate copies of LVAlert messages by unregistering redundant LVAlert message types.
- Update to ligo-followup-advocate >= 0.0.9 to update GCN Circular text for superevents. Unfortunately, circulars are still disabled due to a regression in ligo-gracedb (see https://git.ligo.org/lscsoft/gracedb-client/issues/7).
- Upload BAYESTAR sky maps and annotations to superevents.
- Create (but do not send) preliminary VOEvents for all superevents. No vetting is performed yet.

9.36 0.0.11 (2018-06-27)

- Submit handler tasks to Celery as a single group.
- Retry GraceDb tasks that raise a `TimeoutError` exception.
- The superevent handler now skips LVAlert messages that do not affect the false alarm rate of an event (e.g. simple log messages).

(Note that the false alarm rate in GraceDb is set by the initial event upload and can be updated by replacing the event; however replacing the event does not produce an LVAlert message at all, so there is no way to intercept it.)

- Added a query kwarg to superevents method to reduce latency in fetching the superevents from gracedb.
- Refactored getting event information for update type events so that gracedb is polled only once to get the information needed for superevent manager.
- Renamed the `set_preferred_event` task in gracedb.py to `update_superevent` to be a full wrapper around the `updateSuperevent` client function. Now it can be used to set preferred event and also update superevent time windows.
- Many `cwb` (extra) attributes, which should be floating point numbers, are present in lvalert packet as strings. Casting them to avoid embarassing TypeErrors.
- Reverted back the typecasting of far, gpstime into float. This is fixed in https://git.ligo.org/lscsoft/gracedb/issues/10
- CBC `t_start` and `t_end` values are changed to 1 sec interval.
- Added ligo-raven to run on external trigger and superevent creation lvalerts to search for coincidences. In case of coincidence, EM_COINC label is applied to the superevent and external trigger page and the external trigger is added to the list of em_events in superevent object dictionary.
- `cwb` and `lib` nodes added to superevent handler.
- Events are treated as finite segment window, initial superevent creation with preferred event window. Addition of events to superevents may change the superevent window and also the preferred event.
- Change default GraceDb server to https://gracedb-playground.ligo.org/ for open public alert challenge.
- Update to ligo-gracedb >= 1.29dev1.
- Rename the `get_superevent` task to `get_superevents` and add a new `get_superevent` task that is a trivial wrapper around `ligo.gracedb.rest.GraceDb.superevent()`.

9.37 0.0.10 (2018-06-13)

- Model the time extent of events and superevents using the `glue.segments` module.
- Replace GraceDb.get with GraceDb.superevents from the recent dev release of gracedb-client.
- Fix possible false positive matches between GCNs for unrelated GRBs by matching on both TrigID (which is generally the mission elapsed time) and mission name.
- Add the configuration variable `superevent_far_threshold` to limit the maximum false alarm rate of events that are included in superevents.
- LVAlert handlers are now passed the actual alert data structure rather than the JSON text, so handlers are no longer responsible for calling `json.loads`. It is a little bit more convenient and possibly also faster for Celery to deserialize the alert messages.
- Introduce `Production`, `Development`, `Test`, and `Playground` application configuration objects in order to facilitate quickly switching between GraceDb servers.
- Pipeline specific start and end times for superevent segments. These values are controlled via configuration variables.

9.38 0.0.9 (2018-06-06)

- Add missing LVAlert message types to superevent handler.

9.39 0.0.8 (2018-06-06)

- Add some logging to the GCN and LVAlert dispatch code in order to diagnose missed messages.

9.40 0.0.7 (2018-05-31)

- Ingest Swift, Fermi, and SNEWS GCN notices and save them in GraceDb.
- Depend on the pre-release version of the GraceDb client, ligo-gracedb 1.29.dev0, because this is the only version that supports superevents at the moment.

9.41 0.0.6 (2018-05-26)

- Generate GCN Circular drafts using ligo-followup-advocate.
- In the continuous integration pipeline, validate PEP8 naming conventions using pep8-naming.
- Add instructions for measuring test coverage and running the linter locally to the contributing guide.
- Rename `gwcelery.tasks.voevent` to `gwcelery.tasks.gcn` to make it clear that this submodule contains functionality related to GCN notices, rather than VOEvents in general.
- Rename `gwcelery.tasks.dispatch` to `gwcelery.tasks.orchestrator` to make it clear that this module encapsulates the behavior associated with the "orchestrator" in the O3 low-latency design document.
- Mock up calls to BAYESTAR in test suite to speed it up.
- Unify dispatch of LVAlert and GCN messages using decorators. GCN notice handlers are declared like this:

```
import lxml.etree
from gwcelery.tasks import gcn

@gcn.handler(gcn.NoticeType.FERMI_GBM_GND_POS,
             gcn.NoticeType.FERMI_GBM_FIN_POS)
def handle_fermi(payload):
    root = lxml.etree.fromstring(payload)
    # do work here...
```

  LVAlert message handlers are declared like this:

```
import json
from gwcelery.tasks import lvalert

@lvalert.handler('cbc_gstlal',
                 'cbc_pycbc',
                 'cbc_mbta')
def handle_cbc(alert_content):
    alert = json.loads(alert_content)
    # do work here...
```

- Instead of carrying around the GraceDb service URL in tasks, store the GraceDb host name in the Celery application config.
- Create superevents by simple clustering in time. Currently this is only supported by the `gracedb-dev1` host.

9.42 0.0.5 (2018-05-08)

- Disable socket access during most unit tests. This adds some extra assurance that we don't accidentally interact with production servers during the unit tests.
- Ignore BAYESTAR jobs that raise a `DetectorDisabled` error. These exceptions are used for control flow and do not constitute a real error. Ignoring these jobs avoids polluting logs and the Flower monitor.

9.43 0.0.4 (2018-04-28)

- FITS history and comment entries are now displayed in a monospaced font.
- Adjust error reporting for some tasks.
- Depend on newer version of `ligo.skymap`.
- Add unit tests for the `gwcelery condor submit` subcommand.

9.44 0.0.3 (2018-04-27)

- Fix some compatibility issues between the `gwcelery condor submit` subcommand and the format of `condor_q -totals -xml` with older versions of HTCondor.

9.45 0.0.2 (2018-04-27)

- Add `gwcelery condor submit` and related subcommands as shortcuts for managing GWCelery running under HTCondor.

9.46 0.0.1 (2018-04-27)

- This is the initial release. It provides rapid sky localization with BAYESTAR, sky map annotation, and sending mock alerts.
- By default, GWCelery is configured to listen to the test LVAlert server.
- Sending VOEvents to GCN/TAN is disabled for now.

CHAPTER 10

License

The GWCelery Logo is a composite of Celery2 by Tiia Monto and Lorentzian Wormhole by Kes47 from Wikimedia Commons (CC BY-SA 3.0).

```
                    GNU GENERAL PUBLIC LICENSE
                       Version 2, June 1991

 Copyright (C) 1989, 1991 Free Software Foundation, Inc.,
 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA
 Everyone is permitted to copy and distribute verbatim copies
 of this license document, but changing it is not allowed.

                            Preamble

  The licenses for most software are designed to take away your
freedom to share and change it.  By contrast, the GNU General Public
License is intended to guarantee your freedom to share and change free
software--to make sure the software is free for all its users.  This
General Public License applies to most of the Free Software
Foundation's software and to any other program whose authors commit to
using it.  (Some other Free Software Foundation software is covered by
the GNU Lesser General Public License instead.)  You can apply it to
your programs, too.

  When we speak of free software, we are referring to freedom, not
price.  Our General Public Licenses are designed to make sure that you
have the freedom to distribute copies of free software (and charge for
this service if you wish), that you receive source code or can get it
if you want it, that you can change the software or use pieces of it
in new free programs; and that you know you can do these things.

  To protect your rights, we need to make restrictions that forbid
anyone to deny you these rights or to ask you to surrender the rights.
These restrictions translate to certain responsibilities for you if you
distribute copies of the software, or if you modify it.
```

(continues on next page)

(continued from previous page)

```
  For example, if you distribute copies of such a program, whether
gratis or for a fee, you must give the recipients all the rights that
you have.  You must make sure that they, too, receive or can get the
source code.  And you must show them these terms so they know their
rights.

  We protect your rights with two steps: (1) copyright the software, and
(2) offer you this license which gives you legal permission to copy,
distribute and/or modify the software.

  Also, for each author's protection and ours, we want to make certain
that everyone understands that there is no warranty for this free
software.  If the software is modified by someone else and passed on, we
want its recipients to know that what they have is not the original, so
that any problems introduced by others will not reflect on the original
authors' reputations.

  Finally, any free program is threatened constantly by software
patents.  We wish to avoid the danger that redistributors of a free
program will individually obtain patent licenses, in effect making the
program proprietary.  To prevent this, we have made it clear that any
patent must be licensed for everyone's free use or not licensed at all.

  The precise terms and conditions for copying, distribution and
modification follow.

                    GNU GENERAL PUBLIC LICENSE
   TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

  0. This License applies to any program or other work which contains
a notice placed by the copyright holder saying it may be distributed
under the terms of this General Public License.  The "Program", below,
refers to any such program or work, and a "work based on the Program"
means either the Program or any derivative work under copyright law:
that is to say, a work containing the Program or a portion of it,
either verbatim or with modifications and/or translated into another
language.  (Hereinafter, translation is included without limitation in
the term "modification".)  Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not
covered by this License; they are outside its scope.  The act of
running the Program is not restricted, and the output from the Program
is covered only if its contents constitute a work based on the
Program (independent of having been made by running the Program).
Whether that is true depends on what the Program does.

  1. You may copy and distribute verbatim copies of the Program's
source code as you receive it, in any medium, provided that you
conspicuously and appropriately publish on each copy an appropriate
copyright notice and disclaimer of warranty; keep intact all the
notices that refer to this License and to the absence of any warranty;
and give any other recipients of the Program a copy of this License
along with the Program.

You may charge a fee for the physical act of transferring a copy, and
you may at your option offer warranty protection in exchange for a fee.
```

(continues on next page)

(continued from previous page)

```
  2. You may modify your copy or copies of the Program or any portion
of it, thus forming a work based on the Program, and copy and
distribute such modifications or work under the terms of Section 1
above, provided that you also meet all of these conditions:

    a) You must cause the modified files to carry prominent notices
    stating that you changed the files and the date of any change.

    b) You must cause any work that you distribute or publish, that in
    whole or in part contains or is derived from the Program or any
    part thereof, to be licensed as a whole at no charge to all third
    parties under the terms of this License.

    c) If the modified program normally reads commands interactively
    when run, you must cause it, when started running for such
    interactive use in the most ordinary way, to print or display an
    announcement including an appropriate copyright notice and a
    notice that there is no warranty (or else, saying that you provide
    a warranty) and that users may redistribute the program under
    these conditions, and telling the user how to view a copy of this
    License.  (Exception: if the Program itself is interactive but
    does not normally print such an announcement, your work based on
    the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole.  If
identifiable sections of that work are not derived from the Program,
and can be reasonably considered independent and separate works in
themselves, then this License, and its terms, do not apply to those
sections when you distribute them as separate works.  But when you
distribute the same sections as part of a whole which is a work based
on the Program, the distribution of the whole must be on the terms of
this License, whose permissions for other licensees extend to the
entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest
your rights to work written entirely by you; rather, the intent is to
exercise the right to control the distribution of derivative or
collective works based on the Program.

In addition, mere aggregation of another work not based on the Program
with the Program (or with a work based on the Program) on a volume of
a storage or distribution medium does not bring the other work under
the scope of this License.

  3. You may copy and distribute the Program (or a work based on it,
under Section 2) in object code or executable form under the terms of
Sections 1 and 2 above provided that you also do one of the following:

    a) Accompany it with the complete corresponding machine-readable
    source code, which must be distributed under the terms of Sections
    1 and 2 above on a medium customarily used for software interchange; or,

    b) Accompany it with a written offer, valid for at least three
    years, to give any third party, for a charge no more than your
    cost of physically performing source distribution, a complete
    machine-readable copy of the corresponding source code, to be
```

(continues on next page)

(continued from previous page)

```
    distributed under the terms of Sections 1 and 2 above on a medium
    customarily used for software interchange; or,

    c) Accompany it with the information you received as to the offer
    to distribute corresponding source code.  (This alternative is
    allowed only for noncommercial distribution and only if you
    received the program in object code or executable form with such
    an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for
making modifications to it.  For an executable work, complete source
code means all the source code for all modules it contains, plus any
associated interface definition files, plus the scripts used to
control compilation and installation of the executable.  However, as a
special exception, the source code distributed need not include
anything that is normally distributed (in either source or binary
form) with the major components (compiler, kernel, and so on) of the
operating system on which the executable runs, unless that component
itself accompanies the executable.

If distribution of executable or object code is made by offering
access to copy from a designated place, then offering equivalent
access to copy the source code from the same place counts as
distribution of the source code, even though third parties are not
compelled to copy the source along with the object code.

  4. You may not copy, modify, sublicense, or distribute the Program
except as expressly provided under this License.  Any attempt
otherwise to copy, modify, sublicense or distribute the Program is
void, and will automatically terminate your rights under this License.
However, parties who have received copies, or rights, from you under
this License will not have their licenses terminated so long as such
parties remain in full compliance.

  5. You are not required to accept this License, since you have not
signed it.  However, nothing else grants you permission to modify or
distribute the Program or its derivative works.  These actions are
prohibited by law if you do not accept this License.  Therefore, by
modifying or distributing the Program (or any work based on the
Program), you indicate your acceptance of this License to do so, and
all its terms and conditions for copying, distributing or modifying
the Program or works based on it.

  6. Each time you redistribute the Program (or any work based on the
Program), the recipient automatically receives a license from the
original licensor to copy, distribute or modify the Program subject to
these terms and conditions.  You may not impose any further
restrictions on the recipients' exercise of the rights granted herein.
You are not responsible for enforcing compliance by third parties to
this License.

  7. If, as a consequence of a court judgment or allegation of patent
infringement or for any other reason (not limited to patent issues),
conditions are imposed on you (whether by court order, agreement or
otherwise) that contradict the conditions of this License, they do not
excuse you from the conditions of this License.  If you cannot
distribute so as to satisfy simultaneously your obligations under this
```

(continues on next page)

(continued from previous page)

```
License and any other pertinent obligations, then as a consequence you
may not distribute the Program at all.  For example, if a patent
license would not permit royalty-free redistribution of the Program by
all those who receive copies directly or indirectly through you, then
the only way you could satisfy both it and this License would be to
refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under
any particular circumstance, the balance of the section is intended to
apply and the section as a whole is intended to apply in other
circumstances.

It is not the purpose of this section to induce you to infringe any
patents or other property right claims or to contest validity of any
such claims; this section has the sole purpose of protecting the
integrity of the free software distribution system, which is
implemented by public license practices.  Many people have made
generous contributions to the wide range of software distributed
through that system in reliance on consistent application of that
system; it is up to the author/donor to decide if he or she is willing
to distribute software through any other system and a licensee cannot
impose that choice.

This section is intended to make thoroughly clear what is believed to
be a consequence of the rest of this License.

  8. If the distribution and/or use of the Program is restricted in
certain countries either by patents or by copyrighted interfaces, the
original copyright holder who places the Program under this License
may add an explicit geographical distribution limitation excluding
those countries, so that distribution is permitted only in or among
countries not thus excluded.  In such case, this License incorporates
the limitation as if written in the body of this License.

  9. The Free Software Foundation may publish revised and/or new versions
of the General Public License from time to time.  Such new versions will
be similar in spirit to the present version, but may differ in detail to
address new problems or concerns.

Each version is given a distinguishing version number.  If the Program
specifies a version number of this License which applies to it and "any
later version", you have the option of following the terms and conditions
either of that version or of any later version published by the Free
Software Foundation.  If the Program does not specify a version number of
this License, you may choose any version ever published by the Free Software
Foundation.

  10. If you wish to incorporate parts of the Program into other free
programs whose distribution conditions are different, write to the author
to ask for permission.  For software which is copyrighted by the Free
Software Foundation, write to the Free Software Foundation; we sometimes
make exceptions for this.  Our decision will be guided by the two goals
of preserving the free status of all derivatives of our free software and
of promoting the sharing and reuse of software generally.

                            NO WARRANTY
```

(continues on next page)

(continued from previous page)

```
  11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY
FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW.  EXCEPT WHEN
OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES
PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED
OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF
MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.  THE ENTIRE RISK AS
TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU.  SHOULD THE
PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING,
REPAIR OR CORRECTION.

  12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING
WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR
REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES,
INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING
OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED
TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY
YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER
PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE
POSSIBILITY OF SUCH DAMAGES.

                     END OF TERMS AND CONDITIONS

            How to Apply These Terms to Your New Programs

  If you develop a new program, and you want it to be of the greatest
possible use to the public, the best way to achieve this is to make it
free software which everyone can redistribute and change under these terms.

  To do so, attach the following notices to the program.  It is safest
to attach them to the start of each source file to most effectively
convey the exclusion of warranty; and each file should have at least
the "copyright" line and a pointer to where the full notice is found.

    <one line to give the program's name and a brief idea of what it does.>
    Copyright (C) <year>  <name of author>

    This program is free software; you can redistribute it and/or modify
    it under the terms of the GNU General Public License as published by
    the Free Software Foundation; either version 2 of the License, or
    (at your option) any later version.

    This program is distributed in the hope that it will be useful,
    but WITHOUT ANY WARRANTY; without even the implied warranty of
    MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See the
    GNU General Public License for more details.

    You should have received a copy of the GNU General Public License along
    with this program; if not, write to the Free Software Foundation, Inc.,
    51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA.

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this
when it starts in an interactive mode:

    Gnomovision version 69, Copyright (C) year name of author
    Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type `show w'.
```

(continues on next page)

(continued from previous page)

```
    This is free software, and you are welcome to redistribute it
    under certain conditions; type `show c' for details.

The hypothetical commands `show w' and `show c' should show the appropriate
parts of the General Public License.  Of course, the commands you use may
be called something other than `show w' and `show c'; they could even be
mouse-clicks or menu items--whatever suits your program.

You should also get your employer (if you work as a programmer) or your
school, if any, to sign a "copyright disclaimer" for the program, if
necessary.  Here is a sample; alter the names:

  Yoyodyne, Inc., hereby disclaims all copyright interest in the program
  `Gnomovision' (which makes passes at compilers) written by James Hacker.

  <signature of Ty Coon>, 1 April 1989
  Ty Coon, President of Vice

This General Public License does not permit incorporating your program into
proprietary programs.  If your program is a subroutine library, you may
consider it more useful to permit linking proprietary applications with the
library.  If this is what you want to do, use the GNU Lesser General
Public License instead of this License.
```

CHAPTER 11

Indices and tables

- genindex
- modindex
- search

GWCelery is open source and is licensed under the *GNU General Public License v2 or later*.

Bibliography

[LIGO] https://wiki.ligo.org/Calibration/TDCalibReview

[Virgo] https://dcc.ligo.org/G1801125/

[DMT] https://wiki.ligo.org/DetChar/DmtDqVector

[GCN] https://gcn.gsfc.nasa.gov

Python Module Index

g

gwcelery, 23
gwcelery.conf, 23
gwcelery.conf.development, 25
gwcelery.conf.playground, 25
gwcelery.conf.production, 25
gwcelery.conf.test, 26
gwcelery.sentry, 26
gwcelery.tasks, 26
gwcelery.tasks.bayestar, 27
gwcelery.tasks.circulars, 27
gwcelery.tasks.condor, 27
gwcelery.tasks.detchar, 29
gwcelery.tasks.em_bright, 31
gwcelery.tasks.external_triggers, 32
gwcelery.tasks.first2years, 33
gwcelery.tasks.gcn, 33
gwcelery.tasks.gcn.validate, 34
gwcelery.tasks.gracedb, 34
gwcelery.tasks.lalinference, 36
gwcelery.tasks.ligo_fermi_skymaps, 33
gwcelery.tasks.lvalert, 37
gwcelery.tasks.orchestrator, 40
gwcelery.tasks.p_astro_gstlal, 42
gwcelery.tasks.raven, 43
gwcelery.tasks.skymaps, 44
gwcelery.tasks.superevents, 47
gwcelery.util, 47

Index

A

annotate_fits() (*in module gwcelery.tasks.skymaps*), 44
app (*in module gwcelery*), 23

C

calculate_coincidence_far() (*in module gwcelery.tasks.raven*), 43
calculate_spacetime_coincidence_far() (*in module gwcelery.tasks.raven*), 43
check_idq() (*in module gwcelery.tasks.detchar*), 30
check_vector() (*in module gwcelery.tasks.detchar*), 30
check_vector_prepost (*in module gwcelery.conf*), 24
coincidence_search() (*in module gwcelery.tasks.raven*), 43
configure() (*in module gwcelery.sentry*), 26
create_cache() (*in module gwcelery.tasks.detchar*), 29
create_combined_skymap() (*in module gwcelery.tasks.ligo_fermi_skymaps*), 33

D

dag_finished() (*in module gwcelery.tasks.lalinference*), 36
dmt_dq_vector_bits (*in module gwcelery.tasks.detchar*), 29
dqr_json() (*in module gwcelery.tasks.detchar*), 29
DSN (*in module gwcelery.sentry*), 26

G

gcn_broker_accept_addresses (*in module gwcelery.conf*), 23
gcn_broker_accept_addresses (*in module gwcelery.conf.production*), 25
gcn_broker_address (*in module gwcelery.conf*), 23
gcn_client_address (*in module gwcelery.conf*), 23
gcn_client_address (*in module gwcelery.conf.production*), 25
gracedb_host (*in module gwcelery.conf*), 23
gracedb_host (*in module gwcelery.conf.development*), 25
gracedb_host (*in module gwcelery.conf.production*), 25
gracedb_host (*in module gwcelery.conf.test*), 26
gwcelery (*module*), 23
gwcelery.conf (*module*), 23
gwcelery.conf.development (*module*), 25
gwcelery.conf.playground (*module*), 25
gwcelery.conf.production (*module*), 25
gwcelery.conf.test (*module*), 26
gwcelery.sentry (*module*), 26
gwcelery.tasks (*module*), 26
gwcelery.tasks.bayestar (*module*), 27
gwcelery.tasks.circulars (*module*), 27
gwcelery.tasks.condor (*module*), 27
gwcelery.tasks.detchar (*module*), 29
gwcelery.tasks.em_bright (*module*), 31
gwcelery.tasks.external_triggers (*module*), 32
gwcelery.tasks.first2years (*module*), 33
gwcelery.tasks.gcn (*module*), 33
gwcelery.tasks.gcn.validate (*module*), 34
gwcelery.tasks.gracedb (*module*), 34
gwcelery.tasks.lalinference (*module*), 36
gwcelery.tasks.ligo_fermi_skymaps (*module*), 33
gwcelery.tasks.lvalert (*module*), 37
gwcelery.tasks.orchestrator (*module*), 40
gwcelery.tasks.p_astro_gstlal (*module*), 42
gwcelery.tasks.raven (*module*), 43
gwcelery.tasks.skymaps (*module*), 44
gwcelery.tasks.superevents (*module*), 47
gwcelery.util (*module*), 47

H

handler (*in module gwcelery.tasks.gcn*), 34

handler (*in module gwcelery.tasks.lvalert*), 37
high_latency_frame_types (*in module gwcelery.conf*), 25
high_latency_frame_types (*in module gwcelery.conf.production*), 26

I

idq_channels (*in module gwcelery.conf*), 24
idq_pglitch_thresh (*in module gwcelery.conf*), 24
is_3d_fits_file() (*in module gwcelery.tasks.skymaps*), 44

J

JobAborted, 27
JobFailed, 28
JobRunning, 27

L

llhoft_channels (*in module gwcelery.conf*), 24
llhoft_glob (*in module gwcelery.conf*), 24
low_latency_frame_types (*in module gwcelery.conf*), 24
low_latency_frame_types (*in module gwcelery.conf.production*), 26
lvalert_host (*in module gwcelery.conf*), 23
lvalert_host (*in module gwcelery.conf.production*), 25
lvalert_host (*in module gwcelery.conf.test*), 26

N

NamedTemporaryFile() (*in module gwcelery.util*), 47
no_dq_veto_pycbc_bits (*in module gwcelery.tasks.detchar*), 29

O

orchestrator_timeout (*in module gwcelery.conf*), 24

P

p_astro_gstlal_ln_likelihood_threshold (*in module gwcelery.conf*), 24
p_astro_update() (*in module gwcelery.tasks.p_astro_gstlal*), 42
pe_threshold (*in module gwcelery.conf*), 25
preliminary_alert_far_threshold (*in module gwcelery.conf*), 24
preliminary_alert_trials_factor (*in module gwcelery.conf*), 24
PromiseProxy (*class in gwcelery.util*), 47
Python Enhancement Proposals
  PEP 8, 50

S

sentry_environment (*in module gwcelery.conf.development*), 25
sentry_environment (*in module gwcelery.conf.playground*), 25
sentry_environment (*in module gwcelery.conf.production*), 26
sentry_environment (*in module gwcelery.conf.test*), 26
source_classification() (*in module gwcelery.tasks.em_bright*), 31
state_vector_bits (*in module gwcelery.tasks.detchar*), 29
state_vector_channel_names (*in module gwcelery.conf*), 25
strain_channel_names (*in module gwcelery.conf*), 25
strain_channel_names (*in module gwcelery.conf.production*), 26
superevent_d_t_end (*in module gwcelery.conf*), 24
superevent_d_t_start (*in module gwcelery.conf*), 24
superevent_default_d_t_end (*in module gwcelery.conf*), 24
superevent_default_d_t_start (*in module gwcelery.conf*), 24
superevent_far_threshold (*in module gwcelery.conf*), 24
superevent_query_d_t_end (*in module gwcelery.conf*), 24
superevent_query_d_t_start (*in module gwcelery.conf*), 24

T

task() (*in module gwcelery.tasks.gracedb*), 34

U

uses_gatedhoft (*in module gwcelery.conf*), 24