

---

# **GWCelery Documentation**

***Release 0.9.0***

**Leo Singer**

**Nov 23, 2019**



---

## Contents:

---

<b>1</b>	<b>Quick start</b>	<b>3</b>
1.1	To install . . . . .	3
1.2	To test . . . . .	3
1.3	To start . . . . .	4
<b>2</b>	<b>Design and anatomy of GWCelery</b>	<b>7</b>
2.1	Conceptual overview . . . . .	7
2.2	Block diagram . . . . .	8
2.3	Processes . . . . .	10
2.4	Eternal tasks . . . . .	11
2.5	Handlers . . . . .	11
<b>3</b>	<b>Configuration</b>	<b>13</b>
3.1	Authentication . . . . .	13
<b>4</b>	<b>Running under HTCondor</b>	<b>15</b>
4.1	Shortcuts . . . . .	15
4.2	Managing multiple deployments . . . . .	16
4.3	Job accounting . . . . .	16
<b>5</b>	<b>Monitoring and Management</b>	<b>17</b>
5.1	Flower . . . . .	17
5.2	Sentry . . . . .	17
5.3	Flask . . . . .	18
5.4	Nagios . . . . .	18
5.5	Command-Line Tools . . . . .	19
<b>6</b>	<b>API Reference</b>	<b>21</b>
6.1	gwcclery.conf module . . . . .	21
6.2	gwcclery.email module . . . . .	25
6.3	gwcclery.lvalert module . . . . .	26
6.4	gwcclery.sentry module . . . . .	27
6.5	gwcclery.tasks module . . . . .	27
6.6	gwcclery.tools module . . . . .	57
6.7	gwcclery.util module . . . . .	59
6.8	gwcclery.voevent module . . . . .	60

<b>7</b>	<b>Contributing</b>	<b>63</b>
7.1	Development model . . . . .	63
7.2	Where new code should go . . . . .	63
7.3	Guidelines for tasks . . . . .	63
7.4	Unit tests . . . . .	64
7.5	Code style . . . . .	64
7.6	Documentation . . . . .	64
<b>8</b>	<b>Deployment</b>	<b>67</b>
8.1	Continuous deployment . . . . .	67
8.2	Making a new release . . . . .	67
<b>9</b>	<b>Changelog</b>	<b>71</b>
9.1	0.9.0 (2019-11-23) . . . . .	71
9.2	0.8.7 (2019-09-14) . . . . .	73
9.3	0.8.5.1 (2019-09-04) . . . . .	73
9.4	0.8.6 (2019-09-01) . . . . .	73
9.5	0.8.5 (2019-08-23) . . . . .	73
9.6	0.8.4 (2019-08-16) . . . . .	74
9.7	0.8.3 (2019-08-09) . . . . .	74
9.8	0.8.2 (2019-08-02) . . . . .	74
9.9	0.8.1 (2019-07-29) . . . . .	75
9.10	0.8.0 (2019-07-26) . . . . .	75
9.11	0.7.1 (2019-07-12) . . . . .	75
9.12	0.7.0 (2019-06-21) . . . . .	76
9.13	0.6.3 (2019-06-14) . . . . .	76
9.14	0.6.2 (2019-06-07) . . . . .	76
9.15	0.6.1 (2019-05-24) . . . . .	77
9.16	0.6.0 (2019-05-20) . . . . .	77
9.17	0.5.7 (2019-05-13) . . . . .	77
9.18	0.5.6 (2018-05-08) . . . . .	78
9.19	0.5.5 (2019-05-03) . . . . .	78
9.20	0.5.4 (2019-05-01) . . . . .	79
9.21	0.5.3 (2019-04-17) . . . . .	79
9.22	0.5.2 (2019-04-15) . . . . .	79
9.23	0.5.1 (2019-04-12) . . . . .	79
9.24	0.5.0 (2019-04-12) . . . . .	79
9.25	0.4.3 (2019-04-05) . . . . .	80
9.26	0.4.2 (2019-04-05) . . . . .	80
9.27	0.4.1 (2019-04-02) . . . . .	80
9.28	0.4.0 (2019-03-29) . . . . .	81
9.29	0.3.1 (2019-03-18) . . . . .	81
9.30	0.3.0 (2019-03-01) . . . . .	82
9.31	0.2.6 (2019-02-12) . . . . .	82
9.32	0.2.5 (2019-02-01) . . . . .	82
9.33	0.2.4 (2018-12-17) . . . . .	83
9.34	0.2.3 (2018-12-16) . . . . .	83
9.35	0.2.2 (2018-12-14) . . . . .	83
9.36	0.2.1 (2018-12-14) . . . . .	83
9.37	0.2.0 (2018-12-14) . . . . .	83
9.38	0.1.7 (2018-11-27) . . . . .	84
9.39	0.1.6 (2018-11-14) . . . . .	84
9.40	0.1.5 (2018-11-13) . . . . .	84
9.41	0.1.4 (2018-10-29) . . . . .	84

9.42	0.1.3 (2018-10-26)	85
9.43	0.1.2 (2018-10-11)	85
9.44	0.1.1 (2018-10-04)	85
9.45	0.1.0 (2018-09-26)	86
9.46	0.0.31 (2018-09-04)	86
9.47	0.0.30 (2018-08-02)	86
9.48	0.0.29 (2018-07-31)	87
9.49	0.0.28 (2018-07-25)	87
9.50	0.0.27 (2018-07-22)	87
9.51	0.0.26 (2018-07-20)	87
9.52	0.0.25 (2018-07-19)	88
9.53	0.0.24 (2018-07-18)	88
9.54	0.0.23 (2018-07-18)	88
9.55	0.0.22 (2018-07-11)	88
9.56	0.0.21 (2018-07-10)	89
9.57	0.0.20 (2018-07-09)	89
9.58	0.0.19 (2018-07-09)	89
9.59	0.0.18 (2018-07-06)	89
9.60	0.0.17 (2018-07-05)	89
9.61	0.0.16 (2018-07-05)	89
9.62	0.0.15 (2018-06-29)	90
9.63	0.0.14 (2018-06-28)	90
9.64	0.0.13 (2018-06-28)	90
9.65	0.0.12 (2018-06-28)	90
9.66	0.0.11 (2018-06-27)	90
9.67	0.0.10 (2018-06-13)	91
9.68	0.0.9 (2018-06-06)	92
9.69	0.0.8 (2018-06-06)	92
9.70	0.0.7 (2018-05-31)	92
9.71	0.0.6 (2018-05-26)	92
9.72	0.0.5 (2018-05-08)	93
9.73	0.0.4 (2018-04-28)	93
9.74	0.0.3 (2018-04-27)	93
9.75	0.0.2 (2018-04-27)	93
9.76	0.0.1 (2018-04-27)	93
<b>10 License</b>		<b>95</b>
<b>11 Indices and tables</b>		<b>103</b>
<b>Bibliography</b>		<b>105</b>
<b>Python Module Index</b>		<b>107</b>
<b>Index</b>		<b>109</b>



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  $\geq 3.7$ .

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 LVAAlert 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 LVAAlert, first complete the [LVAAlert Account Activation](#) form once for each LVAAlert 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 five [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 --lvalert
$ gwcelery worker -l info -n gwcelery-extttrig-worker -Q extttrig -c 1
$ 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-voevent-worker -Q voevent -P solo
$ 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](#).

---



---

## 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 LVAAlert, 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.

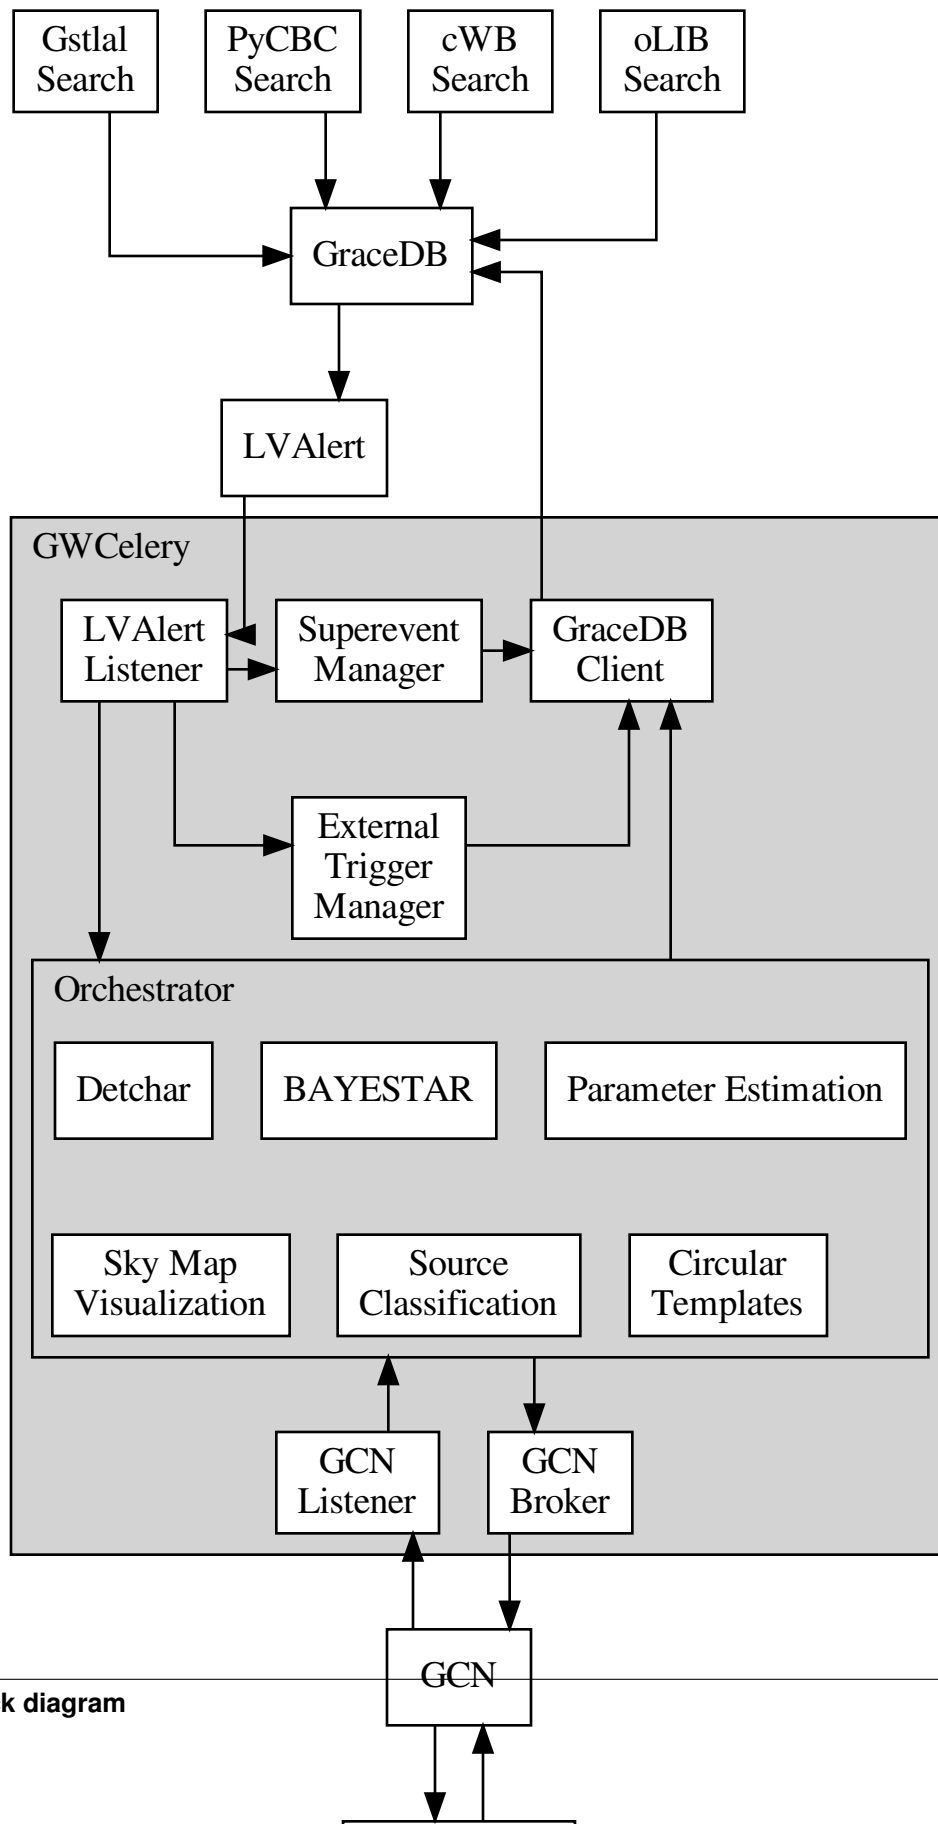
GWCelery interacts with GraceDB by listening for LVAAlert messages and making REST API requests through the GraceDB client. GWCelery interacts with GCN by listening for and sending GCN Notices using the Comet VOEvent broker.

The major subsystems of GWCelery are:

- the LVAAlert 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.



## 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. VOEvent Worker

A Celery worker that is dedicated to sending and receiving VOEvents. It runs an embedded instance of the [Comet](#) VOEvent broker, which is started and stopped using a set of custom [Celery bootsteps](#). Note that the VOEvent worker must be started with the `--pool=solo` option so that tasks are executed in the same Python process that is running the VOEvent broker.

### 8. General-Purpose Worker

A Celery worker that accepts all other tasks.

### 9. 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.lvalert.listen()`

These tasks 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_mbttaonline')
def handle_cbc(alert):
    # do work here...
```



---

## Configuration

---

Like any Celery application, GWCelery's configuration options are stored at run time in a global configuration object, `gwcelery.app.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 LVAAlert 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 LVAAlert.

#### 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`.

## LVAAlert

You must provide a valid username and password for LVAAlert. You can request an account using the [LVAAlert Account Activation](#) form. The LVAAlert username and password should be stored in your [netrc file](#).

---

## 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
```

## 4.3 Job accounting

When GWCelery is started using `gwcelery condor submit` or `gwcelery condor resubmit`, the HTCondor accounting group is set based on which GWCelery configuration you are using:

- `ligo.prod.o3.cbc.pe.bayestar` for production
- `ligo.dev.o3.cbc.pe.bayestar` for all other configurations, including playground

---

## Monitoring and Management

---

GW Celery 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
```

### 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 GW Celery 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:

```
$ gwcclery flask run
```

## 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](http://dashboard.ligo.org) and [monitor.ligo.org](http://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 `gwcclery 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
```

```

celery events 4.2.1 (windowlicker)

  UUID                                WORKER                                TASK                                TIME  STATE
  ----                                -
  34e0c7be-18b7-4378-98c4-6130d78a31d7  celery@gwcelery-worker                gwcelery.tasks.skymaps.fits_header  02:28:20 SUCCESS
  aa9960ea-edb4-4879-a9de-eaab64cfd15  celery@gwcelery-worker                gwcelery.tasks.gracedb.upload       02:28:20 STARTED
  5871398c-ec68-4b3d-b8fd-d3378851d5ae  celery@gwcelery-worker                gwcelery.tasks.skymaps.plot_allsky   02:28:20 STARTED
  f9f66149-d35a-4dc8-b09f-a2403196acaf  celery@gwcelery-worker                gwcelery.tasks.skymaps.annotate_fits_volume 02:28:20 SUCCESS
  b751ebc5-3758-4546-93e7-001a6c901ac3  celery@gwcelery-worker                gwcelery.tasks.orchestrator.handle_cbc_event 02:28:20 SUCCESS
  2382918c-321f-4398-a78d-3277ee10c5a  celery@gwcelery-worker                gwcelery.tasks.orchestrator.download       02:28:20 SUCCESS
  f55aa19d-3108-473e-9379-e481ab098660  celery@gwcelery-worker                gwcelery.tasks.gracedb.upload       02:28:20 STARTED
  e6f8c1b9-80c5-4ea0-9787-edd58305ff84  celery@gwcelery-superevent-worker      gwcelery.tasks.superevents.handle       02:28:20 SUCCESS
  d620c02c-cd51-4641-a77e-d4a52608a60d  celery@gwcelery-worker                gwcelery.tasks.orchestrator.handle_cbc_event 02:28:19 SUCCESS
  28e591de-7b21-479e-8b54-0063f40e75c3  celery@gwcelery-superevent-worker      gwcelery.tasks.superevents.handle       02:28:19 SUCCESS
  7be64365-fbb7-4234-8130-baf866528097  celery@gwcelery-worker                gwcelery.tasks.gracedb.create_label     02:28:18 SUCCESS

Selected: args=('SKYMAP_READY', 'M324973') kwargs={} retries=0 result=None runtime=0.76 root_id=d5343836-893f-4aa7-865f-0e31624e08e0 parent_id=b0f4c4e3-2ec9-4b13-ac5e-af5c4c14d9f0
Workers online: celery@gwcelery-xttrig-worker, celery@gwcelery-openmp-worker, celery@gwcelery-superevent-worker, celery@gwcelery-worker
Info: events: 2077 tasks: 629 workers: 4/4
Keys: j:down k:up i:info t:traceback r:result c:revoke ^c: quit

```

Celery application initialization.

```
gwcclery.app = <Celery gwcclery>
    Celery application object.
```

## 6.1 gwcclery.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:

- `gwcclery.conf.development`
- `gwcclery.conf.playground` (the default)
- `gwcclery.conf.production`
- `gwcclery.conf.test`

```
gwcclery.conf.condor_accounting_group = 'ligo.dev.o3.cbc.pe.bayestar'
    HTCondor accounting group for Celery workers launched with condor_submit.
```

```
gwcclery.conf.expose_to_public = False
    Set to True if events meeting the public alert threshold really should be exposed to the public.
```

```
gwcclery.conf.lvalert_host = 'lvalert-playground.cgca.uwm.edu'
    LVA alert host.
```

```
gwcclery.conf.gracedb_host = 'gracedb-playground.ligo.org'
    GraceDB host.
```

`gwcelery.conf.voevent_broadcaster_address = ':5342'`

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.voevent_broadcaster_whitelist = []`

List of hosts from which the broker will accept connections. If empty, then completely disable the broker's broadcast capability.

`gwcelery.conf.voevent_receiver_address = '68.169.57.253:8099'`

The VOEvent listener will connect to this address to receive GCNs. For options, see [GCN's list of available VOEvent servers](#). If this is an empty string, then completely disable the GCN listener.

`gwcelery.conf.email_host = 'imap.gmail.com'`

IMAP hostname to receive the GCN e-mail notice formats.

`gwcelery.conf.superevent_d_t_start = {'gstlal': 1.0, 'mbtaonline': 1.0, 'pycbc': 1.0, 'sp': 1.0}`

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, 'sp': 1.0}`

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 = 0.00027777777777777778`

Maximum false alarm rate to consider events superevents.

`gwcelery.conf.preliminary_alert_timeout = 30.0`

Wait this many seconds for the preferred event to stabilize before issuing a preliminary alert.

`gwcelery.conf.preliminary_alert_far_threshold = {'burst': 3.1709791983764586e-08, 'cbc': 1.0}`

Group specific maximum false alarm rate to consider sending preliminary alerts.

`gwcelery.conf.preliminary_alert_trials_factor = {'burst': 3.0, 'cbc': 4.0}`

Trials factor corresponding to trigger categories. For CBC and Burst, trials factor is the number of pipelines. CBC pipelines are gstlal, pycbc, mbtaonline and spiir. Burst searches are cwb.allsky, cwb.bbh and cwb.imbh.

`gwcelery.conf.snews_gw_far_threshold = 1.1574074074074073e-05`

Maximum false alarm rate for a superevent to send out a coincidence alert between an external SNEWS alert and the superevent.

`gwcelery.conf.superevent_clean_up_timeout = 270.0`

The orchestrator will wait this many seconds from the time of the application of the GCN\_PRELIM\_SENT label to revise the preferred event out of the accumulated events.

`gwcelery.conf.subthreshold_annotation_timeout = 300.0`

The orchestrator will wait this many seconds from the time of the creation of a new superevent to the time that subthreshold superevents are annotated. It is expected that the timeout is long enough such that there are no more G events being added to the superevent.

`gwcelery.conf.pe_timeout = 345.0`

The orchestrator will wait this many seconds from the time of the creation of a new superevent to the time that parameter estimation begins, in case the preferred event is updated with high latency.

`gwcelery.conf.check_vector_prepost` = {'CWB': [1.5, 1.5], 'Fermi': [2, 2], 'HardwareInjection': False}  
 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/kafka/{detector}\_O2/\*.gwf'  
 File glob for playground low-latency h(t) frames. Currently points to O2 replay data.

`gwcelery.conf.llhoft_channels` = {'H1:DMT-DQ\_VECTOR': 'dmt\_dq\_vector\_bits', 'H1:GDS-CALIB\_STATE\_VECTOR': 'gds\_calib\_state\_vector\_bits'}  
 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\_OVL\_32\_2048', 'L1:IDQ-PGLITCH\_OVL\_32\_2048']  
 Low-latency iDQ p(glitch) channel names from O2 replay.

`gwcelery.conf.idq_pglitch_thresh` = 0.95  
 If P(Glitch) is above this threshold, and `idq_veto` for the pipeline is true, DQV will be labeled for the event.

`gwcelery.conf.idq_veto` = {'CWB': False, 'Fermi': False, 'HardwareInjection': False, 'LIBRA': False}  
 If true for a pipeline, iDQ values above the threshold defined in `thresh` will cause DQV to be labeled. Currently all False, pending iDQ review (should be done before O3).

`gwcelery.conf.low_latency_frame_types` = {'H1': 'H1\_O2\_llhoft', 'L1': 'L1\_O2\_llhoft', 'V1': 'V1\_O2\_llhoft'}  
 Types of low latency frames used in Parameter Estimation (see `gwcelery.tasks.inference`) and in cache creation for detchar checks (see `gwcelery.tasks.detchar`).

`gwcelery.conf.high_latency_frame_types` = {'H1': None, 'L1': None, 'V1': None}  
 Types of high latency frames used in Parameter Estimation and in cache creation for detchar checks. They do not exist for O2Replay data. (see `gwcelery.tasks.inference` and `gwcelery.tasks.detchar`)

`gwcelery.conf.strain_channel_names` = {'H1': 'H1:GDS-CALIB\_STRAIN\_O2Replay', 'L1': 'L1:GDS-CALIB\_STRAIN\_O2Replay'}  
 Names of h(t) channels used in Parameter Estimation (see `gwcelery.tasks.inference`) and in detchar omegascan creation (see `gwcelery.tasks.detchar`).

`gwcelery.conf.state_vector_channel_names` = {'H1': 'H1:GDS-CALIB\_STATE\_VECTOR', 'L1': 'L1:GDS-CALIB\_STATE\_VECTOR'}  
 Names of state vector channels used in Parameter Estimation (see `gwcelery.tasks.inference`)

`gwcelery.conf.detchar_bit_definitions` = {'dmt\_dq\_vector\_bits': {'bits': {1: 'NO\_OMC\_DCPI'}}}  
 Bit definitions for detchar checks

`gwcelery.conf.omegascan_durations` = [0.5, 2.0, 10.0]  
 Durations for omegascans, symmetric about t0

`gwcelery.conf.pe_threshold` = 4.133597883597884e-07  
 FAR threshold in Hz for Parameter Estimation. PE group now applies 1/(4 weeks) as a threshold. 86400 seconds = 1 day and 28 days = 4 weeks.

`gwcelery.conf.pe_results_path` = '/home/docs/public\_html/online\_pe'  
 Path to the results of Parameter Estimation (see `gwcelery.tasks.inference`)

`gwcelery.conf.pe_results_url` = 'https://ldas-jobs.ligo.caltech.edu/~docs/online\_pe/'  
 URL of page where all the results of Parameter Estimation are outputted (see `gwcelery.tasks.inference`)

`gwcelery.conf.raven_coincidence_windows` = {'GRB\_Burst': [-600, 60], 'GRB\_CBC': [-5, 1], 'GRB\_OTHER': [-5, 1]}  
 Time coincidence windows passed to ligo-raven. External events and superevents of the appropriate type are considered to be coincident if within time window of each other.

`gwcclery.conf.mock_events_simulate_multiple_uploads = False`

If True, then upload each mock event several times in rapid succession with random jitter in order to simulate multiple pipeline uploads.

### 6.1.1 gwcclery.conf.development module

Application configuration for `gracedb-dev1.ligo.org`.

Inherits all settings from `gwcclery.conf.test`, with the exceptions below.

`gwcclery.conf.development.gracedb_host = 'gracedb-dev1.ligo.org'`

GraceDB host.

`gwcclery.conf.development.sentry_environment = 'development'`

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

### 6.1.2 gwcclery.conf.playground module

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

`gwcclery.conf.playground.sentry_environment = 'playground'`

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

`gwcclery.conf.playground.mock_events_simulate_multiple_uploads = True`

If True, then upload each mock event several times in rapid succession with random jitter in order to simulate multiple pipeline uploads.

### 6.1.3 gwcclery.conf.production module

Application configuration for `gracedb.ligo.org`.

Inherits all settings from `gwcclery.conf.playground`, with the exceptions below.

`gwcclery.conf.production.condor_accounting_group = 'ligo.prod.o3.cbc.pe.bayestar'`

HTCondor accounting group for Celery workers launched with `condor_submit`.

`gwcclery.conf.production.expose_to_public = True`

Set to True if events meeting the public alert threshold really should be exposed to the public.

`gwcclery.conf.production.lvalert_host = 'lvalert.cgca.uwm.edu'`

LVAalert host.

`gwcclery.conf.production.gracedb_host = 'gracedb.ligo.org'`

GraceDB host.

`gwcclery.conf.production.voevent_broadcaster_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.

`gwcclery.conf.production.voevent_broadcaster_whitelist = ['capella2.gsfc.nasa.gov']`

List of hosts from which the broker will accept connections. If empty, then completely disable the broker's broadcast capability.

`gwcclery.conf.production.llhoft_glob = '/dev/shm/kafka/{detector}/*.gwf'`

File glob for low-latency h(t) frames.

```

gwcclery.conf.production.low_latency_frame_types = {'H1': 'H1_llhoft', 'L1': 'L1_llhoft'}
    Types of frames used in Parameter Estimation (see gwcclery.tasks.inference) and in cache creation
    for detchar checks (see gwcclery.tasks.detchar).

gwcclery.conf.production.high_latency_frame_types = {'H1': 'H1_HOFT_C00', 'L1': 'L1_HOFT_C00'}
    Types of high latency frames used in Parameter Estimation (see gwcclery.tasks.inference) and in
    cache creation for detchar checks (see gwcclery.tasks.detchar).

gwcclery.conf.production.idq_channels = ['H1:IDQ-PGLITCH_OVL_16_4096', 'L1:IDQ-PGLITCH_OVL_16_4096']
    Low-latency iDQ p(glitch) channel names from live O3 frames

gwcclery.conf.production.strain_channel_names = {'H1': 'H1:GDS-CALIB_STRAIN_CLEAN', 'L1': 'L1:GDS-CALIB_STRAIN_CLEAN'}
    Names of h(t) channels used in Parameter Estimation (see gwcclery.tasks.inference)

gwcclery.conf.production.sentry_environment = 'production'
    Record this environment tag <https://docs.sentry.io/enriching-error-data/environments/> in Sentry log mes-
    sages.

```

### 6.1.4 gwcclery.conf.test module

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

Inherits all settings from *gwcclery.conf.playground*, with the exceptions below.

```

gwcclery.conf.test.lvalert_host = 'lvalert-test.cgca.uwm.edu'
    LVAlerter host.

gwcclery.conf.test.gracedb_host = 'gracedb-test.ligo.org'
    GraceDB host.

gwcclery.conf.test.sentry_environment = 'test'
    Record this environment tag <https://docs.sentry.io/enriching-error-data/environments/> in Sentry log mes-
    sages.

```

## 6.2 gwcclery.email module

Embed an IMAP email client into a Celery worker by [extending Celery with bootsteps](#).

```

gwcclery.email.add_worker_arguments(parser)

gwcclery.email.install(app)
    Register the email client subsystem in the application boot steps.

```

### 6.2.1 gwcclery.email.bootsteps module

```

class gwcclery.email.bootsteps.Receiver(consumer, email=False, **kwargs)
    Bases: gwcclery.email.bootsteps.EmailBootStep

    Run the global email receiver in background thread.

    name = 'email client'

    create(consumer)
        Create the step.

    start(consumer)

    stop(consumer)

```

## 6.2.2 gwcelery.email.signals module

Definitions of custom Celery signals related to emails.

These signals allow us to keep the VOEvent validation code decoupled from the email client itself.

`gwcelery.email.signals.email_received = <Signal: email_received providing_args={'rfc822'}>`

Fired whenever an email message is received.

**Parameters** `rfc822` (*bytes*) – The **RFC 822** contents of the message.

### Examples

Register an email listener like this:

```
import email
import email.policy

@email_received.connect
def on_email_received(rfc822, **kwargs):
    # Parse the RFC822 email.
    message = email.message_from_bytes(rfc822, policy=email.policy.default)
    # Print some of the message headers.
    print('Subject:', message['Subject'])
    print('From:', message['From'])
    # Print the plain-text message body.
    body = message.get_body(['plain']).get_content()
    print(body)
```

## 6.3 gwcelery.lvalert module

Embed a Comet LVAAlert listener into a Celery worker by extending Celery with bootsteps.

`gwcelery.lvalert.add_worker_arguments(parser)`

`gwcelery.lvalert.install(app)`

Register the LVAAlert subsystem in the application boot steps.

### 6.3.1 gwcelery.lvalert.bootsteps module

**class** `gwcelery.lvalert.bootsteps.Receiver` (*consumer, lvalert=False, \*\*kwargs*)

Bases: `gwcelery.lvalert.bootsteps.LVAAlertBootStep`

Run the global LVAAlert receiver in background thread.

**name** = 'LVAAlert client'

**create** (*consumer*)

Create the step.

**start** (*consumer*)

**stop** (*consumer*)

**info** (*consumer*)



### 6.3.2 gwcelery.lvalert.signals module

Definitions of custom Celery signals related to VOEvents.

These signals allow us to keep the VOEvent broker code decoupled from any GCN-specific logic. Notably, it allows us to keep all of the details of the GCN-specific “Notice Type” concept out of `gwcelery.voevent`.

`gwcelery.lvalert.signals.lvalert_received = <Signal: lvalert_received providing_args={'notice_type'}`  
 Fired whenever a VOEvent is received.

**Parameters** `xml_document` (`comet.utility.xml.xml_document`) – The XML document that was received. The raw file contents are available as `xml_document.raw_bytes`. The lxml.etree representation of the document is available as `xml_document.element`.

## 6.4 gwcelery.sentry module

Error telemetry for Sentry.

`gwcelery.sentry.DSN = 'https://sentry.io/1425216'`  
 Sentry data source name (DSN).

`gwcelery.sentry.configure()`  
 Configure Sentry logging integration for Celery.  
 See the [official instructions for Celery integration](#).

### Notes

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

## 6.5 gwcelery.tasks module

All Celery tasks are declared in submodules of this module.

### 6.5.1 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.

## 6.5.2 gwcelery.tasks.circulars module

Generate and upload automated circulars.

**(task)** `gwcelery.tasks.circulars.create_initial_circular(graceid)`

Create and return circular txt.

**(task)** `gwcelery.tasks.circulars.create_emcoinc_circular(graceid)`

Create and return the em\_coinc circular txt.

**(task)** `gwcelery.tasks.circulars.create_update_circular(graceid, update_types=['sky_localization', 'em_bright', 'p_astro'])`

Create and return update circular txt.

**(task)** `gwcelery.tasks.circulars.create_retraction_circular(graceid)`

Create and return retraction circular txt.

## 6.5.3 gwcelery.tasks.condor module

Submit and monitor HTCCondor jobs<sup>1</sup>.

### Notes

Internally, we use the XML condor log format<sup>2</sup> for easier parsing.

### References

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

Bases: `Exception`

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

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

Bases: `Exception`

Raised if an HTCCondor job is still running.

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

Bases: `subprocess.CalledProcessError`

Raised if an HTCCondor job fails.

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

Submit a job using HTCCondor.

---

<sup>1</sup> [http://research.cs.wisc.edu/htcondor/manual/latest/condor\\_submit.html](http://research.cs.wisc.edu/htcondor/manual/latest/condor_submit.html)

<sup>2</sup> <http://research.cs.wisc.edu/htcondor/classad/refman/node3.html>

#### 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)** `gwcclery.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')
```

## 6.5.4 gwcclery.tasks.detchar module

### Flow Chart

The flow chart below shows the decision process for the application of DQOK and DQV labels.



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.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>,
 ...

```

(continues on next page)

(continued from previous page)

```
<glue.lal.CacheEntry at 0x7fbae6b15080>,
<glue.lal.CacheEntry at 0x7fbae6b15828>]
```

**(task)** `gwcclery.tasks.detchar.make_omegascan(ifo, t0, durs)`

Helper function to create a single omegascan image, with multiple durations.

**Parameters**

- **ifo** (*str*) – ‘H1’, ‘L1’, or ‘V1’
- **t0** (*int or float*) – Central time of the omegascan.
- **durs** (*list of floats/ints*) – List of three durations which will be scanned symmetrically about t0. Example: [0.5, 2, 10]

**Returns** bytes of png of the omegascan, or None if no omegascan created.

**Return type** `bytes` or `None`

**(task)** `gwcclery.tasks.detchar.omegascan(t0, graceid)`

Create omegascan for a certain event.

**Parameters**

- **t0** (*float*) – Central event time.
- **graceid** (*str*) – GraceDB ID to which to upload the omegascan.

`gwcclery.tasks.detchar.generate_table(title, high_bit_list, low_bit_list, unknown_bit_list)`

Make a nice table which shows the status of the bits checked.

**Parameters**

- **title** (*str*) – Title of the table.
- **high\_bit\_list** (*list*) – List of bit names which are high.
- **low\_bit\_list** (*list*) – List of bit names which are low.
- **unknown\_bit\_list** (*list*) – List of bit names which are unknown.

**Returns** HTML string of the table.

**Return type** `str`

`gwcclery.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`

`gwcclery.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, ligo_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.

**Returns** **event** – Details of the event, reflecting any labels that were added.

**Return type** *dict*

### 6.5.5 gwcelery.tasks.em\_bright module

This module computes the probabilities that there is a neutron star in the binary, and that the coalescence event resulted in creation of tidally disrupted matter.

The result is returned in the form of a JSON file:

```
{'HasNS': 1.0, 'HasRemnant': 1.0}
```

- **HasNS: The probability that at least one of the component masses** in the binary is a neutron star. The definition of a neutron star in this context simply means an object with mass less than 3.0 solar mass.
- **HasRemnant: The probability that the binary system can produce** tidally disrupted matter during coalescence. This is computed using the fitting formula in [arXiv:1807.00011](https://arxiv.org/abs/1807.00011) We are currently using an extremely stiff equation of state (2H) to compute the compactness of the neutron star. This results in a higher chance of labelling a systems with non-zero HasRemnant value.

Qualitative source classification for CBC events.

**(task)** `gwcelery.tasks.em_bright.em_bright_posterior_samples` (*posterior\_file\_content*)

Returns the probability of having a NS component and remnant using LALInference posterior samples.

**Parameters** **posterior\_file\_content** (*hdf5 posterior file content*) –

**Returns** JSON formatted string storing HasNS and HasRemnant probabilities

**Return type** *str*

#### Example

```
>>> em_bright_posterior_samples(GraceDb().files('S190930s',
... 'LALInference.posterior_samples.hdf5').read())
{'HasNS': 0.014904901243599122, 'HasRemnant': 0.0}
```

**(task)** `gwcelery.tasks.em_bright.classifier_other` (*args, graceid*)

Returns the boolean probability of having a NS component and the probability of having non-zero disk mass. This method is used for pipelines that do not provide the data products necessary for computation of the source properties probabilities.

#### Parameters

- **args** (*tuple*) – Tuple containing (m1, m2, spin1z, spin2z, snr)
- **graceid** (*str*) – The graceid of the event

**Returns** JSON formatted string storing HasNS and HasRemnant probabilities

**Return type** `str`

### Example

```
>>> em_bright.classifier_other((2.0, 1.0, 0.0, 0.0, 10.), 'S123456')
'{"HasNS": 1.0, "HasRemnant": 1.0}'
```

**(task)** `gwcelery.tasks.em_bright.classifier_gstlal` (*args*, *graceid*)

Returns the probability of having a NS component and the probability of having non-zero disk mass in the detected event. This method will be using the data products obtained from the weekly supervised learning runs for injections campaigns. The data products are in pickle formatted RandomForestClassifier objects. The method `predict_proba` of these objects provides us the probabilities of the coalescence being EM-Bright and existence of neutron star in the binary.

#### Parameters

- **args** (*tuple*) – Tuple containing (m1, m2, spin1z, spin2z, snr)
- **graceid** (*str*) – The graceid of the event

**Returns** JSON formatted string storing HasNS and HasRemnant probabilities

**Return type** `str`

### Notes

This task would only work from within the CIT cluster.

## 6.5.6 gwcelery.tasks.external\_triggers module

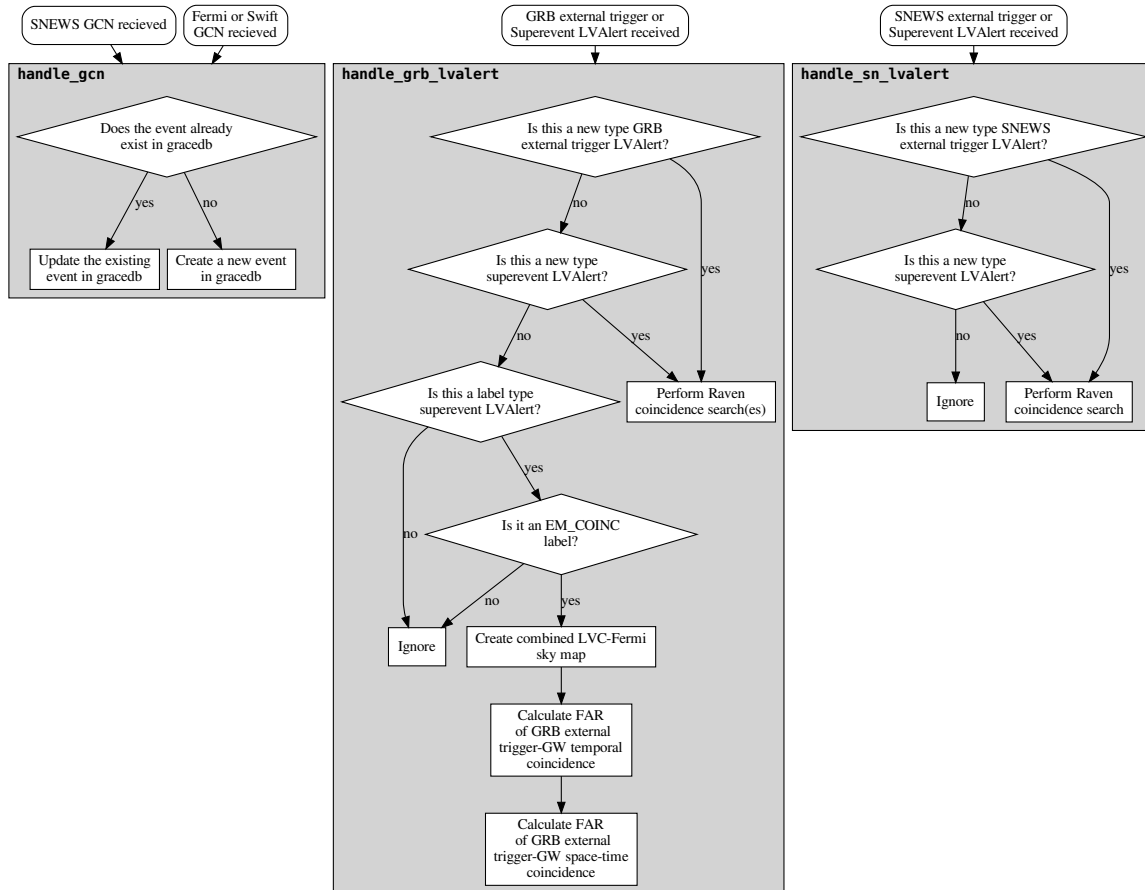
This module listens to the *GCNs* from SNEWS and the Fermi and Swift missions. It is also responsible for carrying out tasks related to external trigger-gravitational wave coincidences, including looking for temporal coincidences, creating combined GRB-GW sky localization probability maps, and computing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

There are two GCN and two LVAAlert 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 LVAAlert.
- `handle_grb_lvalert()` is called for each Fermi and Swift external trigger and superevent LVAAlert.



## Flow Chart



## Tasks

**(task)** `gwcelery.tasks.external_triggers.handle_snews_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_snews_lvalert(alert)`  
 Parse an LVAAlert message related to superevents/SN external triggers and dispatch it to other tasks.

## Notes

This LVAAlert 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()`.

## 6.5.7 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._vet_event(superevents)`

**(task)** `gwcelery.tasks.first2years._upload_psd(graceid)`

**(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.

## 6.5.8 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.get_external_skymap_filename(graceid)`  
 Get the external 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.

**(task)** `gwcelery.tasks.ligo_fermi_skymaps.get_upload_external_skymap(graceid)`  
If a Fermi sky map is not uploaded yet, tries to download one and upload to external event. If sky map is not available, passes so that this can be re-run the next time an update GCN notice is received.

## 6.5.9 gwcelery.tasks.gcn module

Tasks to send and receive Gamma-ray Coordinates Network [GCN] notices.

### References

`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...
```

**exception** `gwcelery.tasks.gcn.SendingError`

Bases: `RuntimeError`

A generic error associated with sending VOEvents.

**(task)** `gwcelery.tasks.gcn.send(message)`  
Send a VOEvent to GCN.

This task will be retried several times if the VOEvent cannot be sent. See the Raises section below for circumstances that cause a retry.

**Parameters** `message` (*bytes*) – The raw VOEvent file contents.

**Raises** `SendingError` – If the VOEvent could not be sent because there were no network peers connected to the VOEvent broadcaster.

**(task)** `gwcelery.tasks.gcn.validate(payload)`  
Validate LIGO/Virgo GCN notices.

Check that the contents of a public LIGO/Virgo GCN matches the original VOEvent in GraceDB.

## Notes

If the VOEvent broadcaster is disabled by setting `voevent_broadcaster_whitelist` to an empty list, then this task becomes a no-op.

### 6.5.10 gwcelery.tasks.gracedb module

Communication with GraceDB.

**class** gwcelery.tasks.gracedb.**RetryableHTTPError** (\*args, \*\*kwargs)  
 Bases: sphinx.ext.autodoc.importer.\_MockObject

Exception class for server-side HTTP errors that we should retry.

gwcelery.tasks.gracedb.**catch\_retryable\_http\_errors** (f)  
 Decorator to capture server-side errors that we should retry.

We retry HTTP status 502 (Bad Gateway), 503 (Service Unavailable), and 504 (Gateway Timeout).

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

**(task)** gwcelery.tasks.gracedb.**create\_event** (\*args, \*\*kwargs)  
 Create an event in GraceDB.

**(task)** gwcelery.tasks.gracedb.**create\_label** (\*args, \*\*kwargs)  
 Create a label in GraceDB.

**(task)** gwcelery.tasks.gracedb.**remove\_label** (\*args, \*\*kwargs)  
 Create a label in GraceDB.

**(task)** gwcelery.tasks.gracedb.**create\_signoff** (\*args, \*\*kwargs)  
 Create a label in GraceDB.

**(task)** gwcelery.tasks.gracedb.**create\_tag** (\*args, \*\*kwargs)  
 Create a tag in GraceDB.

**(task)** gwcelery.tasks.gracedb.**create\_voevent** (\*args, \*\*kwargs)  
 Create a VOEvent.

**Returns** The filename of the new VOEvent.

**Return type** str

**(task)** gwcelery.tasks.gracedb.**download** (\*args, \*\*kwargs)  
 Download a file from GraceDB.

**(task)** gwcelery.tasks.gracedb.**expose** (\*args, \*\*kwargs)  
 Expose an event to the public.

## Notes

If `expose_to_public` is False, then this because a no-op.

**(task)** gwcelery.tasks.gracedb.**get\_events** (\*args, \*\*kwargs)  
 Get events from GraceDB.

**(task)** gwcelery.tasks.gracedb.**get\_event** (\*args, \*\*kwargs)  
 Retrieve an event from GraceDB.

**(task)** gwcelery.tasks.gracedb.**get\_search** (\*args, \*\*kwargs)  
 Retrieve the search field of an event from GraceDB.

**(task)** gwcelery.tasks.gracedb.get\_labels(\*args, \*\*kwargs)  
Get all labels for an event in GraceDB.

**(task)** gwcelery.tasks.gracedb.get\_log(\*args, \*\*kwargs)  
Get all log messages for an event in GraceDB.

**(task)** gwcelery.tasks.gracedb.get\_superevent(\*args, \*\*kwargs)  
Retrieve a superevent from GraceDB.

**(task)** gwcelery.tasks.gracedb.replace\_event(\*args, \*\*kwargs)  
Get an event from GraceDB.

**(task)** gwcelery.tasks.gracedb.upload(\*args, \*\*kwargs)  
Upload a file to GraceDB.

**(task)** gwcelery.tasks.gracedb.get\_superevents(\*args, \*\*kwargs)  
List matching superevents in gracedb.

#### Parameters

- **\*args** – arguments passed to GraceDb.superevents()
- **\*\*kwargs** – keyword arguments passed to GraceDb.superevents()

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

**Return type** list

**(task)** gwcelery.tasks.gracedb.update\_superevent(\*args, \*\*kwargs)  
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(\*args, \*\*kwargs)  
Create new superevent in GraceDB with graceid

#### Parameters

- **graceid**(str) – graceid with which superevent is created.
- **t0**(float) – t\_0 parameter of superevent
- **t\_start**(float) – t\_start parameter of superevent
- **t\_end**(float) – t\_end parameter of superevent
- **category**(str) – superevent category

**(task)** gwcelery.tasks.gracedb.add\_event\_to\_superevent(\*args, \*\*kwargs)  
Add an event to a superevent in GraceDB.

## 6.5.11 gwcelery.tasks.inference module

Source Parameter Estimation with LALInference and Bilby.

**exception** `gwcclery.tasks.inference.NotEnoughData`

Bases: `Exception`

Raised if found data is not enough due to the latency of data transfer

**(task)** `gwcclery.tasks.inference.query_data` (*trigtime*)

Continues to query data until it is found with `gwdatafind` and return frametypes for the data. If data is not found in 86400 seconds = 1 day, raise `NotEnoughData`.

**(task)** `gwcclery.tasks.inference.upload_no_frame_files` (*request, exc, traceback, superevent\_id*)

Upload notification when no frame files are found.

#### Parameters

- **request** (*Context (placeholder)*) – Task request variables
- **exc** (*Exception*) – Exception raised by `condor.submit`
- **traceback** (*str (placeholder)*) – Traceback message from a task
- **superevent\_id** (*str*) – The GraceDB ID of a target superevent

**(task)** `gwcclery.tasks.inference.prepare_ini` (*frametype\_dict, event, superevent\_id=None*)

Determine an appropriate PE settings for the target event and return ini file content for LALInference pipeline

`gwcclery.tasks.inference.pre_pe_tasks` (*event, superevent\_id*)

Return canvas of tasks executed before parameter estimation starts

**(task)** `gwcclery.tasks.inference._setup_dag_for_lalinference` (*coinc\_psd, ini\_contents, rundir, superevent\_id*)

Create DAG for a lalinference run and return the path to DAG.

#### Parameters

- **coinc\_psd** (*tuple of byte contents*) – Tuple of the byte contents of `coinc.xml` and `psd.xml.gz`
- **ini\_contents** (*str*) – The content of `online_lalinference_pe.ini`
- **rundir** (*str*) – The path to a run directory where the DAG file exits
- **superevent\_id** (*str*) – The GraceDB ID of a target superevent

**Returns** `path_to_dag` – The path to the .dag file

**Return type** `str`

**(task)** `gwcclery.tasks.inference._setup_dag_for_bilby` (*event, rundir, preferred\_event\_id, superevent\_id*)

Create DAG for a bilby run and return the path to DAG.

#### Parameters

- **event** (*json contents*) – The json contents retrieved from `gracedb.get_event()`
- **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** `path_to_dag` – The path to the .dag file

**Return type** `str`

**(task)** `gwcelery.tasks.inference._condor_no_submit` (*path\_to\_dag*)

Run 'condor\_submit\_dag -no\_submit' and return the path to .sub file.

**(task)** `gwcelery.tasks.inference.dag_prepare_task` (*rundir*, *superevent\_id*, *preferred\_event\_id*, *pe\_pipeline*, *ini\_contents=None*)

Return a canvas of tasks to prepare DAG.

#### Parameters

- **rundir** (*str*) – The path to a run directory where the DAG file exits
- **superevent\_id** (*str*) – The GraceDB ID of a target superevent
- **preferred\_event\_id** (*str*) – The GraceDB ID of a target preferred event
- **pe\_pipeline** (*str*) – The parameter estimation pipeline used Either 'lalinference' OR 'bilby'
- **ini\_contents** (*str*) – The content of online\_lalinference\_pe.ini Required if `pe_pipeline == 'lalinference'`

**Returns** `canvas` – The canvas of tasks to prepare DAG

**Return type** `canvas` of tasks

**(task)** `gwcelery.tasks.inference.job_error_notification` (*request*, *exc*, *traceback*, *superevent\_id*, *rundir*, *pe\_pipeline*)

Upload notification when condor.submit terminates unexpectedly.

#### Parameters

- **request** (*Context (placeholder)*) – Task request variables
- **exc** (*Exception*) – Exception raised by condor.submit
- **traceback** (*str (placeholder)*) – Traceback message from a task
- **superevent\_id** (*str*) – The GraceDB ID of a target superevent
- **rundir** (*str*) – The run directory for PE
- **pe\_pipeline** (*str*) – The parameter estimation pipeline used Either lalinference OR bilby

**(task)** `gwcelery.tasks.inference._upload_url` (*pe\_results\_path*, *graceid*, *pe\_pipeline*)

Upload url of a page containing all of the plots.

**(task)** `gwcelery.tasks.inference._get_result_contents` (*pe\_results\_path*, *filename*)

Return the contents of a PE results file by reading it from the local filesystem.

**(task)** `gwcelery.tasks.inference.clean_up` (*rundir*)

Clean up a run directory.

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

**(task)** `gwcelery.tasks.inference.dag_finished` (*rundir*, *preferred\_event\_id*, *superevent\_id*, *pe\_pipeline*)

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
- **pe\_pipeline** (*str*) – The parameter estimation pipeline used Either lalinference OR bilby

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

**Return type** `canvas`

**(task)** `gwcelery.tasks.inference._download_psd(gid)`

Download `psd.xml.gz` and return its content. If that file does not exist, return `None`.

**(task)** `gwcelery.tasks.inference.start_pe(ini_contents, preferred_event_id, superevent_id, pe_pipeline)`

Run Parameter Estimation on a given event.

**Parameters**

- **ini\_contents** (*str*) – The content of `online_lalinference_pe.ini`
- **preferred\_event\_id** (*str*) – The GraceDB ID of a target preferred event
- **superevent\_id** (*str*) – The GraceDB ID of a target superevent
- **pe\_pipeline** (*str*) – The parameter estimation pipeline used lalinference OR bilby

## 6.5.12 gwcelery.tasks.lvalert module

LVAAlert client.

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

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

**Parameters**

- **\*keys** – List of LVAAlert 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_mbttaonline')
def handle_cbc(alert_content):
    # do work here...
```

## 6.5.13 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 LVAAlert 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.




## 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.

The orchestrator is 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 performs data quality checks with `gwcelery.tasks.detach.check_vectors()` and calls `preliminary_alert()` to send a preliminary GCN notice.

**(task)** `gwcelery.tasks.orchestrator.handle_cbc_event(alert)`  
Perform annotations for CBC events that depend on pipeline-specific matched-filter parameter estimates.

## Notes

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

File	Task
<code>bayestar.multiorder.fits</code>	<code>gwcelery.tasks.bayestar.localize()</code>
<code>em_bright.json</code>	<code>gwcelery.tasks.em_bright.classifier()</code>
<code>p_astro.json.json</code>	<code>gwcelery.tasks.p_astro.compute_p_astro()</code>

**(task)** `gwcelery.tasks.orchestrator.handle_posterior_samples(alert)`  
Generate multi-resolution and flat-resolution fits files and skymaps from an uploaded HDF5 file containing posterior samples.

**(task)** `gwcelery.tasks.orchestrator._update_if_dqok(superevent_id, event_id)`  
Update `preferred_event` of `superevent_id` to `event_id` if `DQOK` label has been applied.

**(task)** `gwcelery.tasks.orchestrator._get_preferred_event(superevent_id)`  
Determine preferred event for a superevent by querying GraceDB.

This works just like `gwcelery.tasks.gracedb.get_superevent()`, except that it returns only the preferred event, and not the entire GraceDB JSON response.

**(task)** `gwcelery.tasks.orchestrator._create_voevent(classification, *args, **kwargs)`  
Create a VOEvent record from an EM bright JSON file.

### Parameters

- **classification** (*tuple*, *None*) – A collection of JSON strings, generated by `gwcelery.tasks.em_bright.classifier()` and `gwcelery.tasks.p_astro.compute_p_astro()` or content of `p_astro.json` uploaded by `gstlal` respectively; or *None*
- **\*args** – Additional positional arguments passed to `gwcelery.tasks.gracedb.create_voevent()`.
- **\*\*kwargs** – Additional keyword arguments passed to `gwcelery.tasks.gracedb.create_voevent()`.

**Returns** The filename of the newly created VOEvent.

**Return type** `str`

**(task)** `gwcelery.tasks.orchestrator._create_label_and_return_filename(filename, label, graceid)`

**(task)** `gwcelery.tasks.orchestrator._leave_log_message_and_return_event_dict` (*event*, *superevent\_id*, *message*, *initiate\_voevent=True*, *\*\*kwargs*)

Wrapper around `gracedb.update_superevent()` that returns the event dictionary.

**(task)** `gwcelery.tasks.orchestrator._update_superevent_and_return_event_dict` (*event*, *superevent\_id*)

Wrapper around `gracedb.update_superevent()` that returns the event dictionary.

**(task)** `gwcelery.tasks.orchestrator._proceed_if_no_advocate_action` (*filenames*, *superevent\_id*)

Return filenames in case the superevent does not have labels indicating advocate action.

**(task)** `gwcelery.tasks.orchestrator.preliminary_alert` (*event*, *superevent\_id*, *annotation\_prefix=""*, *initiate\_voevent=True*)

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._get_lowest_far` (*superevent\_id*)

Obtain the lowest FAR of the events in the target superevent.

**(task)** `gwcelery.tasks.orchestrator.parameter_estimation` (*far\_event*, *superevent\_id*)

Tasks for Parameter Estimation Followup with LALInference or Bilby

For LALInference, this consists of the following steps:

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

For Bilby, this consists of the following steps:

1. Start Parameter Estimation if FAR is smaller than the PE threshold.
2. Upload of ini file during Parameter Estimation

**(task)** `gwcelery.tasks.orchestrator.preliminary_initial_update_alert` (*filenames*, *superevent\_id*, *alert\_type*, *labels=[]*)

Create and send a preliminary, initial, or update GCN notice.

#### Parameters

- **filenames** (*tuple*) – A list of the sky map, em\_bright, and p\_astro filenames.
- **superevent\_id** (*str*) – The superevent ID.
- **alert\_type** ({'preliminary', 'initial', 'update'}) – The alert type.
- **labels** (*list*) – A list of labels applied to superevent.

### Notes

This function is decorated with `gwcclery.tasks.gracedb.task` rather than `gwcclery.app.task` so that a synchronous call to `gwcclery.tasks.gracedb.get_log()` is retried in the event of GraceDB API failures. If `EM_COINC` in labels will create a RAVEN circular.

**(task)** `gwcclery.tasks.orchestrator.initial_alert` (*filenames*, *superevent\_id*, *labels=[]*)  
Produce an initial alert.

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

#### Parameters

- **filenames** (*tuple*) – A list of the sky map, em\_bright, and p\_astro filenames.
- **superevent\_id** (*str*) – The superevent ID.
- **labels** (*list*) – A list of labels applied to superevent.

### Notes

This function is decorated with `gwcclery.tasks.gracedb.task` rather than `gwcclery.app.task` so that a synchronous call to `gwcclery.tasks.gracedb.get_log()` is retried in the event of GraceDB API failures.

**(task)** `gwcclery.tasks.orchestrator.update_alert` (*filenames*, *superevent\_id*)  
Produce an update alert.

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

#### Parameters

- **filenames** (*tuple*) – A list of the sky map, em\_bright, and p\_astro filenames.
- **superevent\_id** (*str*) – The superevent ID.

### Notes

This function is decorated with `gwcclery.tasks.gracedb.task` rather than `gwcclery.app.task` so that a synchronous call to `gwcclery.tasks.gracedb.get_log()` is retried in the event of GraceDB API failures.

**(task)** `gwcclery.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.

## 6.5.14 gwcelery.tasks.p\_astro module

Computation of `p_astro` by source category and utilities related to `p_astro.json` source classification files. See Kapadia et al (2019), arXiv:1903.06881, for details.

**(task)** `gwcelery.tasks.p_astro.compute_p_astro` (*snr, far, mass1, mass2, pipeline, instruments*)

Task to compute `p_astro` by source category.

### Parameters

- **snr** (*float*) – event’s SNR
- **far** (*float*) – event’s cfar
- **mass1** (*float*) – event’s mass1
- **mass2** (*float*) – event’s mass2
- **instruments** (*set*) – set of instruments that detected the event

**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, 'Terrestrial': 0.001}
```

**(task)** `gwcelery.tasks.p_astro.plot` (*contents*)

Make a visualization of the source classification.

**Parameters** `contents` (*str, bytes*) – The contents of the `p_astro.json` file.

**Returns** `png` – The contents of a PNG file.

**Return type** `bytes`

### Notes


The unusually small size of the plot (2.5 x 2 inches) is optimized for viewing in GraceDB’s image display widget.

### Examples

```
>>> from gwcelery.tasks import p_astro
>>> contents = '''
... {"Terrestrial": 0.001, "BNS": 0.65, "NSBH": 0.20,
...  "MassGap": 0.10, "BBH": 0.059}
... '''
>>> p_astro.plot(contents)
```

**(task)** `gwcelery.tasks.p_astro.handle` (*alert*)

LVAAlert handler to plot and upload a visualization of every `p_astro.json` that is added to a superevent.



### 6.5.15 gwcelery.tasks.raven module

Search for GRB-GW coincidences with ligo-raven.

**(task)** `gwcelery.tasks.raven.calculate_coincidence_far` (*superevent\_id*, *exttrig\_id*, *preferred\_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` (*search*, *se\_id*, *exttrig\_id*, *tl*, *th*, *incl\_sky=False*, *se\_fitsfile=None*)

Calculate FAR of GRB exttrig-GW coincidence

**(task)** `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*) – `Ivalert['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*) – `Ivalert['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.raven_pipeline` (*raven\_search\_results*, *gracedb\_id*, *alert\_object*, *gw\_group*)

Executes much of the full raven pipeline, including adding the external trigger to the superevent, calculating the coincidence false alarm rate, and applying ‘EM\_COINC’ to the appropriate events. Also a preliminary alert will be triggered if the coincidence passes threshold.

### Parameters

- **raven\_search\_results** (*list*) – list of dictionaries of each related gracedb trigger
- **gracedb\_id** (*str*) – ID of either a superevent or external trigger
- **alert\_object** (*dict*) – lvalert[‘object’], either a superevent or an external event
- **gw\_group** (*str*) – Burst or CBC

**(task)** `gwcelery.tasks.raven.preferred_superevent` (*raven\_search\_results*)

Chooses the superevent with the lowest far for an external event to be added to. This is to prevent errors from trying to add one external event to multiple superevents.

**Parameters** **raven\_search\_results** (*list*) – list of dictionaries of each related gracedb trigger

**(task)** `gwcelery.tasks.raven._get_coinc_far_try_raven_alert` (*superevent*, *ext\_event*, *gracedb\_id*, *superevent\_id*, *exttrig\_id*, *gw\_group*)

**(task)** `gwcelery.tasks.raven.trigger_raven_alert` (*coinc\_far\_json*, *superevent*, *gracedb\_id*, *ext\_event*, *gw\_group*)

Determine whether an event should be published as a public alert. If yes, then launches an alert by applying *EM\_COINC* to the preferred event.

All of the following conditions must be true for a public alert:

- The external event must be a threshold GRB or SNEWS event.
- If triggered on a SNEW event, the GW false alarm rate must pass *snews\_gw\_far\_threshold*.
- The event’s RAVEN coincidence false alarm rate, weighted by the group-specific trials factor as specified by the *preliminary\_alert\_trials\_factor* configuration setting, is less than or equal to *preliminary\_alert\_far\_threshold*.

### Parameters

- **coinc\_far\_json** (*dict*) – Dictionary containing coincidence false alarm rate results from RAVEN
- **superevent** (*dict*) – superevent dictionary
- **gracedb\_id** (*str*) – ID of the trigger that launched RAVEN
- **ext\_event** (*dict*) – external event dictionary
- **gw\_group** (*str*) – Burst or CBC

### 6.5.16 gwcelery.tasks.skymaps module

Annotations for sky maps.

**(task) gwcelery.tasks.skymaps.annotate\_fits** (*filecontents*, *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](#).

**(task) gwcelery.tasks.skymaps.skymap\_from\_samples** (*samplefilecontents*)

Generate multi-resolution fits file from samples.

### 6.5.17 gwcelery.tasks.superevents module

Superevents are an abstraction to unify gravitational-wave candidates from multiple search pipelines. 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 time. At any given time, one event belonging to the superevent is identified as the *preferred event*.

This module provides the Superevent Manager, an LVAAlert handler that creates and updates superevents whenever new events are uploaded to GraceDB. It also checks whether the superevent qualifies to be sent out as a LIGO-Virgo public alert.

Event candidates are only considered for membership in a superevent if their false alarm rate is less than or equal to the value of the [superevent\\_far\\_threshold](#) configuration setting.

Each superevent has a time window described by a central time  $t_0$ , a start time  $t_{\text{start}}$ , and a end time  $t_{\text{end}}$ . The central time  $t_0$  is just the time of the preferred event. The start and end time are extended to encompass all of the events that belong to the superevent (see [get\\_ts\(\)](#)).

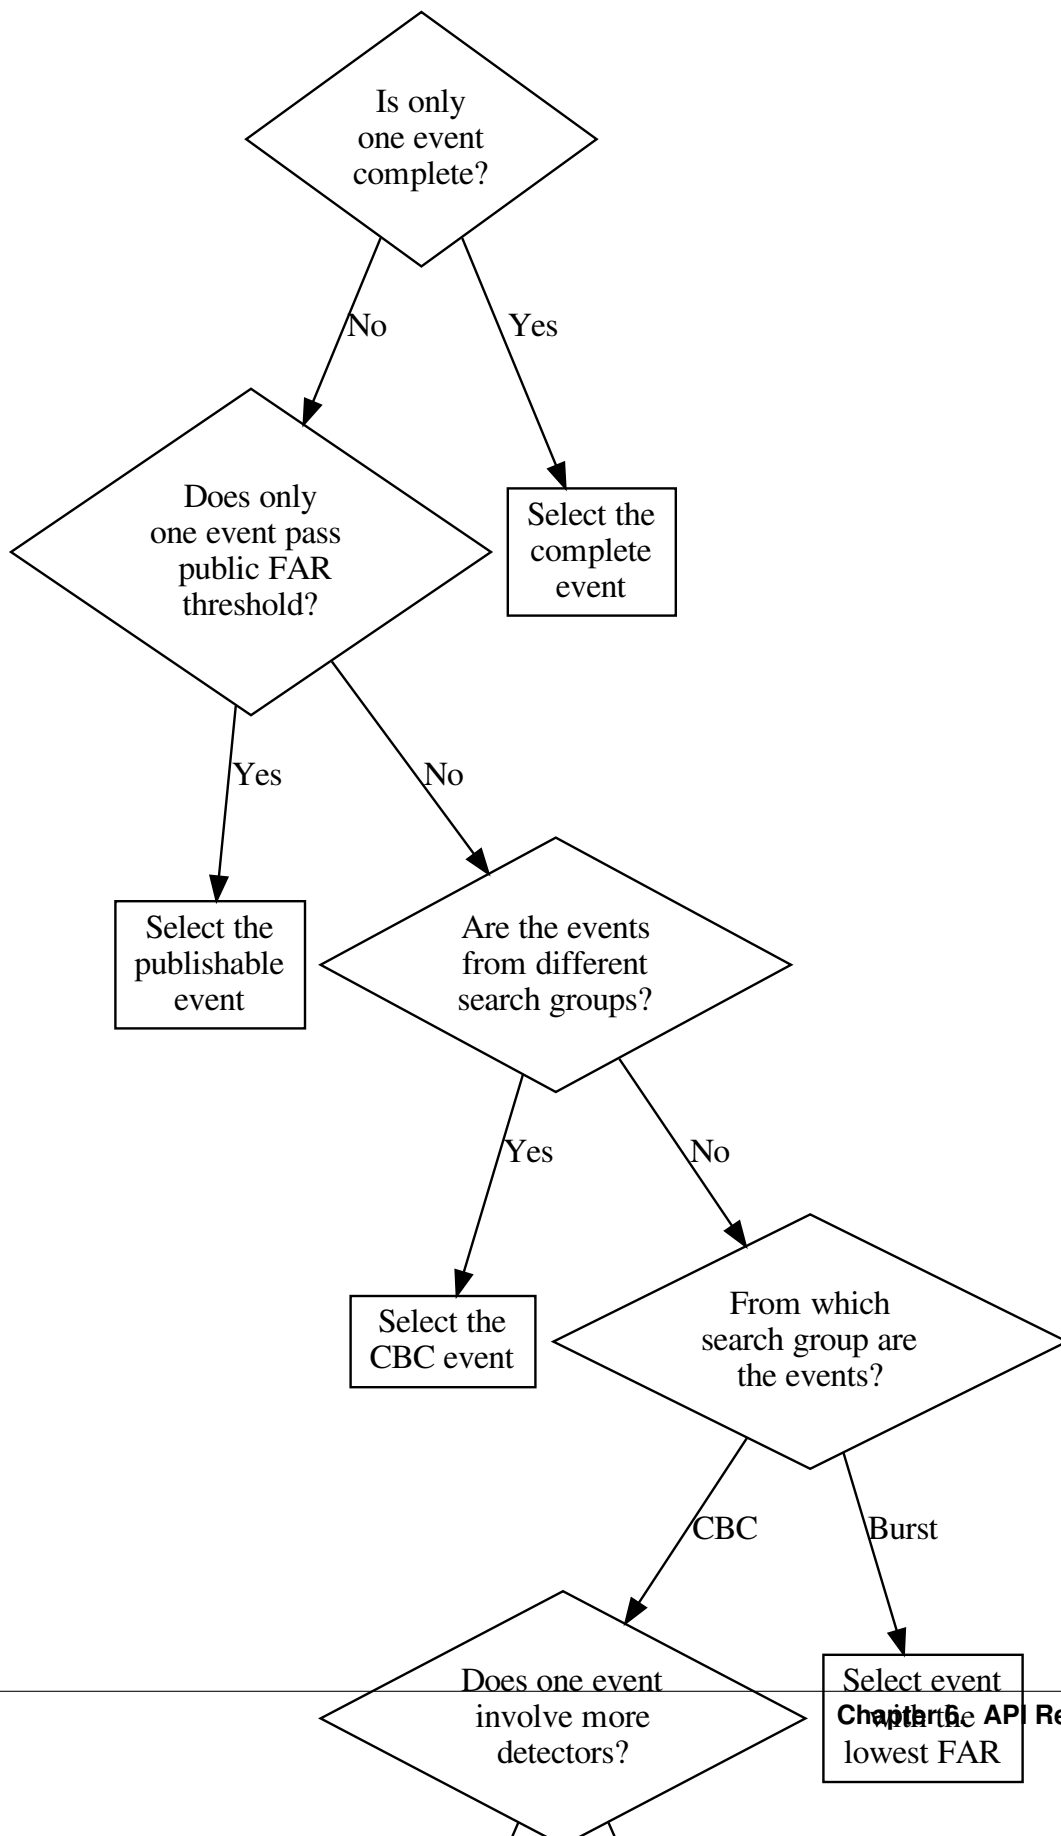
The first candidate reported from a search pipeline creates a superevent, with it being the preferred event. Subsequent candidate additions to the superevent may result in a change of the superevent time window. The preferred event may also be updated as more significant candidates are added. However, this will stop once a candidate passing the public false alarm rate threshold (mentioned in [preliminary\\_alert\\_far\\_threshold](#)) is added to the superevent. At this point, the preferred event is frozen and an automatically generated preliminary notice is sent with the data products of the preferred event. Triggers could however still be added to the superevent as the preliminary alert and the ensuing annotations are being processed. Once the preliminary alert is dispatched to the GCN broker, the preferred event would be revised after a wait time of [superevent\\_clean\\_up\\_timeout](#), following which a second automatic preliminary alert would be issued.

## Selection of the preferred event

When a new event is added to a superevent, it may or may not become the new preferred event. The preferred event is selected by considering the following factors in order to resolve any ties:

1. **Completeness:** Would the event be complete, as determined by the function `is_complete()`, for sending an automated public alert?
2. **Public FAR threshold:** Does the false alarm rate pass the false alarm rate threshold mentioned in `preliminary_alert_far_threshold`?
3. **Search group:** Is it a CBC event or a burst event? CBC events takes precedence.
4. **Number of detectors:** How many detectors contributed data to the event? For CBC events, events with triggers from more detectors take precedence.
5. **Significance:** For CBC events, which has the highest SNR? For burst events, which has the lowest FAR?

The selection of the preferred event from a pair of events is illustrated by the decision tree below.



**Note:** When a preferred event is assigned to a superevent, it may not be complete i.e., its data products may not have been computed yet. Once all the data products of the preferred event is ready, the `EM_READY` label is applied to the superevent.

The preferred event is frozen once an event candidate passing the public false alarm rate threshold is added to the superevent. This is denoted by the application of the `EM_Selected` label on the superevent.

When the preliminary alert has been dispatched to the GCN broker, the `GCN_PRELIM_SENT` label is applied to the superevent which is used to revise the preferred event and launch a second preliminary alert.

The second preliminary is sent even if the preferred event stays unchanged after the revision. In this case, it contains the same content as the first preliminary alert.

The application of `ADVNO` before the launching of the second preliminary alert stops the process. A retraction notice is sent instead.

## 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 `process()` function.

```
gwcelery.tasks.superevents.REQUIRED_LABELS_BY_GROUP = {'burst': {'SKYMAP_READY'}, 'cbc':
```

These labels should be present on an event to consider it to be complete.

```
gwcelery.tasks.superevents.FROZEN_LABEL = 'EM_Selected'
```

This label indicates that the superevent manager should make no further changes to the preferred event.

```
gwcelery.tasks.superevents.READY_LABEL = 'EM_READY'
```

This label indicates that a preferred event has been assigned and it has all data products required to make it ready for annotations.

**(task)** `gwcelery.tasks.superevents.handle(payload)`  
Respond to lvalert nodes from low-latency search pipelines and delegate to `process()` for superevent management.

**(task)** `gwcelery.tasks.superevents.process(*args, **kwargs)`  
Respond to `payload` and serially processes them to create new superevents, add events to existing ones

**Parameters** `payload(dict)` – LVAAlert payload

```
gwcelery.tasks.superevents.get_category(event)
```

Get the superevent category for an event.

**Parameters** `event(dict)` – Event dictionary (e.g., the return value from `gwcelery.tasks.gracedb.get_event()`).

**Returns**

**Return type** {'mdc', 'test', 'production'}

```
gwcelery.tasks.superevents.get_ts(event)
```

Get time extent of an event, depending on pipeline-specific parameters.

- For CWB, use the event's `duration` field.
- For oLIB, use the ratio of the event's `quality_mean` and `frequency_mean` fields.

- For all other pipelines, use the `superevent_d_t_start` and `superevent_d_t_end` configuration values.

**Parameters** `event` (*dict*) – Event dictionary (e.g., the return value from `gwcelery.tasks.gracedb.get_event()`).

**Returns**

- `t_0` (*float*) – Segment center time in GPS seconds.
- `t_start` (*float*) – Segment start time in GPS seconds.
- `t_end` (*float*) – Segment end time in GPS seconds.

`gwcelery.tasks.superevents.get_snr(event)`

Get the SNR from the LVA alert packet.

Different groups and pipelines store the SNR in different fields.

**Parameters** `event` (*dict*) – Event dictionary (e.g., the return value from `gwcelery.tasks.gracedb.get_event()`).

**Returns** `snr` – The SNR.

**Return type** `float`

`gwcelery.tasks.superevents.get_instruments(event)`

Get the instruments that contributed data to an event.

**Parameters** `event` (*dict*) – Event dictionary (e.g., the return value from `gwcelery.tasks.gracedb.get_event()`).

**Returns** The set of instruments that contributed to the event.

**Return type** `set`

`gwcelery.tasks.superevents.get_instruments_in_ranking_statistic(event)`

Get the instruments that contribute to the false alarm rate.

**Parameters** `event` (*dict*) – Event dictionary (e.g., the return value from `gwcelery.tasks.gracedb.get_event()`).

**Returns** The set of instruments that contributed to the ranking statistic for the event.

**Return type** `set`

## Notes

The number of instruments that contributed *data* to an event is given by the `instruments` key of the GraceDB event JSON structure. However, some pipelines (e.g. `gstlal`) have a distinction between which instruments contributed *data* and which were considered in the *ranking* of the candidate. For such pipelines, we infer which pipelines contributed to the ranking by counting only the `SingleInspiral` records for which the `chi squared` field is non-empty.

For PyCBC Live in the O3 configuration, an empty `chi^2` field does not mean that the detector did not contribute to the ranking; in fact, *all* detectors listed in the `SingleInspiral` table contribute to the significance even if the `chi^2` is not computed for some of them. Hence PyCBC Live is handled as a special case.

**(task)** `gwcelery.tasks.superevents.select_preferred_event(events)`

Select the preferred event out of a list of events, typically contents of a `superevent`, based on `keyfunc()`.

**Parameters** `events` (*list*) – list of event dictionaries

`gwcelery.tasks.superevents.is_complete(event)`

Determine if a G event is complete in the sense of the event has its data products complete i.e. has PAS-TRO\_READY, SKYMAP\_READY, EMBRIGHT\_READY for CBC events and the SKYMAP\_READY label for the Burst events. Test events are not processed by low-latency infrastructure and are always labeled complete.

**Parameters** `event` (*dict*) – Event dictionary (e.g., the return value from `gwcelery.tasks.gracedb.get_event()`).

`gwcelery.tasks.superevents.should_publish(event)`

Determine whether an event should be published as a public alert.

All of the following conditions must be true for a public alert:

- The event's offline flag is not set.
- The event should be complete based on `is_complete()`.
- The event's false alarm rate, weighted by the group-specific trials factor as specified by the `preliminary_alert_trials_factor` configuration setting, is less than or equal to `preliminary_alert_far_threshold`.

**Parameters** `event` (*dict*) – Event dictionary (e.g., the return value from `gwcelery.tasks.gracedb.get_event()`).

**Returns** `should_publish` – `True` if the event meets the criteria for a public alert or `False` if it does not.

**Return type** `bool`

`gwcelery.tasks.superevents.keyfunc(event)`

Key function for selection of the preferred event.

Return a value suitable for identifying the preferred event. Given events `a` and `b`, `a` is preferred over `b` if `keyfunc(a) > keyfunc(b)`, else `b` is preferred.

**Parameters** `event` (*dict*) – Event dictionary (e.g., the return value from `gwcelery.tasks.gracedb.get_event()`).

**Returns** `key` – The comparison key.

**Return type** `tuple`

## Notes

Tuples are compared lexicographically in Python: they are compared element-wise until an unequal pair of elements is found.

## 6.6 gwcelery.tools module

Custom Celery subcommands.

The subcommand that is implemented in `gwcelery.tools.submodule` can be invoked from the command line as `gwcelery submodule`.

### 6.6.1 gwcelery.tools.condor module

Shortcuts for HTCondor commands to manage deployment of GWCelery on LIGO Data Grid clusters.

These commands apply to the GWCelery instance that is running in the current working directory.

`gwcelery.tools.condor.get_constraints()`

`gwcelery.tools.condor.run_exec(*args)`

`gwcelery.tools.condor.running()`

Determine if GWCelery is already running under HTCondor.

`gwcelery.tools.condor.submit(app)`

Submit all GWCelery jobs to HTCondor (if not already running).

`gwcelery.tools.condor.resubmit(app)`

Remove any running GWCelery jobs and resubmit to HTCondor.

`gwcelery.tools.condor.rm(app)`

Remove all GWCelery jobs.

`gwcelery.tools.condor.hold(app)`

Put all GWCelery jobs on hold.

`gwcelery.tools.condor.release(app)`

Release all GWCelery jobs from hold status.

`gwcelery.tools.condor.q(app)`

Show status of all GWCelery jobs.

**class** `gwcelery.tools.condor.CondorCommand`(*app=None, get\_app=None, no\_color=False, stdout=None, stderr=None, quiet=False, on\_error=None, on\_usage\_error=None*)

Bases: `celery.bin.base.Command`

Shortcuts for HTCondor commands to manage deployment of GWCelery on LIGO Data Grid clusters.

These commands apply to the GWCelery instance that is running in the current working directory.

**add\_arguments** (*parser*)

**run** (*func=None, \*\*kwargs*)

### 6.6.2 gwcelery.tools.flask module

Flask web application for manually triggering certain tasks.

`gwcelery.tools.flask.main()`

Flask web application for manually triggering certain tasks.

**class** `gwcelery.tools.flask.FlaskCommand`(*app=None, get\_app=None, no\_color=False, stdout=None, stderr=None, quiet=False, on\_error=None, on\_usage\_error=None*)

Bases: `celery.bin.base.Command`

Flask web application for manually triggering certain tasks.

**add\_arguments** (*parser*)

**run** (*\*args, flask\_args=(), \*\*kwargs*)



### 6.6.3 gwcelery.tools.nagios module

A Nagios plugin for monitoring GWCelery.

**class** gwcelery.tools.nagios.NagiosPluginStatus

Bases: `enum.IntEnum`

Nagios plugin status codes.

**OK** = 0

**WARNING** = 1

**CRITICAL** = 2

**UNKNOWN** = 3

**exception** gwcelery.tools.nagios.NagiosCriticalError

Bases: `Exception`

An exception that maps to a Nagios status of *CRITICAL*.

gwcelery.tools.nagios.get\_active\_queues(*inspector*)

gwcelery.tools.nagios.get\_active\_lvalert\_nodes(*inspector*)

gwcelery.tools.nagios.get\_expected\_queues(*app*)

gwcelery.tools.nagios.get\_expected\_lvalert\_nodes(*app*)

gwcelery.tools.nagios.get\_active\_voevent\_peers(*inspector*)

gwcelery.tools.nagios.check\_status(*app*)

**class** gwcelery.tools.nagios.NagiosCommand(*app=None, get\_app=None, no\_color=False, stdout=None, stderr=None, quiet=False, on\_error=None, on\_usage\_error=None*)

Bases: `celery.bin.base.Command`

A Nagios plugin for monitoring GWCelery.

**run** (*\*\*kwargs*)

## 6.7 gwcelery.util module

Miscellaneous utilities that are useful inside many different tasks.

gwcelery.util.handling\_system\_exit()

Catch any `SystemExit` and re-raise it as `RuntimeError`.

Some Celery tasks in this package call main functions of command-line tools from other packages. Those main functions may try to exit the Python interpreter (if, for example, the command-line arguments are not understood).

Catch any `SystemExit` exception. If the exit code is zero (signifying a normal exit status), then ignore the exception. If the exit code is nonzero (signifying an error exit status), then re-raise it as a `RuntimeError` so that the error is reported but the Celery worker is not killed.

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

Bases: `object`

`gwcclery.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()`.

## 6.8 gwcclery.voevent module

Embed a [Comet](#) VOEvent broker and subscriber into a Celery worker by [extending Celery with bootsteps](#).

`gwcclery.voevent.install (app)`

Register the VOEvent subsystem in the application boot steps.

### 6.8.1 gwcclery.voevent.bootsteps module

**class** `gwcclery.voevent.bootsteps.Reactor (consumer, **kwargs)`

Bases: `gwcclery.voevent.bootsteps.VOEventBootStep`

Run the global Twisted reactor in background thread.

The Twisted reactor is a global run loop that drives all Twisted services and operations. This boot step starts the Twisted reactor in a background thread when the Celery consumer starts, and stops the thread when the Consumer terminates.

**name** = 'Twisted reactor'

**create** (*consumer*)

Create the step.

**start** (*consumer*)

**stop** (*consumer*)

**class** `gwcclery.voevent.bootsteps.Broadcaster (consumer, **kwargs)`

Bases: `gwcclery.voevent.bootsteps.TwistedService`

Comet-based VOEvent broadcaster.

Run a Comet-based VOEvent broadcaster (`comet.protocol.broadcaster.VOEventBroadcasterFactory`). Starts after the [Reactor](#) bootstep.

A few [configuration options](#) are available:

- **voevent\_broadcaster\_address**: The address to bind to, in `host:port` format.
- **voevent\_broadcaster\_whitelist**: A list of hostnames, IP addresses, or CIDR address ranges from which to accept connections.

The list of active connections is made available [inspection](#) with the `gwcclery inspect stats` command under the `voevent-broker-peers` key.

**name** = 'VOEvent broadcaster'

**create\_service** (*consumer*)

**info** (*consumer*)

```
class gwcelery.voevent.bootsteps.Receiver (consumer, **kwargs)
```

Bases: gwcelery.voevent.bootsteps.TwistedService

VOEvent receiver.

Run a Comet-based VOEvent receiver (comet.protocol.subscriber.VOEventSubscriberFactory). Starts after the *Reactor* bootstrap.

A few *configuration options* are available:

- voevent\_receiver\_address: The address to connect to, in *host:port* format.

The list of active connections is made available *inspection* with the gwcelery inspect stats command under the voevent-receiver-peers key.

```
name = 'VOEvent receiver'
```

```
requires = ('celery.worker.consumer:Connection', step:Twisted reactor{('celery.worker.
```

```
create_service (consumer)
```

```
info (consumer)
```

## 6.8.2 gwcelery.voevent.logging module

Integration between the Celery, Twisted, and Comet logging systems.

```
gwcelery.voevent.logging.after_setup_logger (logger, loglevel, **kwargs)
```

Celery *signal handler* to set up capturing of all log messages from Comet and Twisted.

- Celery uses the Python standard library's *logging* module. Twisted has its own separate logging facility. Use Twisted's *PythonLoggingObserver* to forward all Twisted log messages to the Python *logging* module.
- Comet uses the Twisted logging facility, but has its own separate management of log severity level (e.g., *info*, *debug*). Set Comet's log level to match Celery's.

## 6.8.3 gwcelery.voevent.util module

VOEvent-related utilities.

```
gwcelery.voevent.util.get_host_port (address)
```

Split a network address of the form *host:port*.

**Parameters** *network* (*str*) – The network address.

**Returns**

- **host** (*str*) – The hostname, or an empty string if missing.
- **port** (*int*, *None*) – The port number, or None if missing.

```
gwcelery.voevent.util.get_local_ivo (app)
```

Create an IVOID to identify this application in VOEvent Transport Protocol packets.

**Returns** A local IVOID composed of the machine's fully qualified domain name and the Celery application name (for example, *ivo://emfollow.ligo.caltech.edu/gwcelery*).

**Return type** *str*

```
gwcelery.voevent.util.get_network (address)
```

Find the IP network prefix for a hostname or CIDR notation.

**Parameters** `address` (*str*) – A hostname, such as `ligo.org`, or an IP address prefix in CIDR notation, such as `127.0.0.0/8`.

**Returns** An object representing the IP address prefix.

**Return type** `ipaddress.IPv4Network`

## 6.8.4 gwcelery.voevent.signals module

Definitions of custom Celery signals related to VOEvents.

These signals allow us to keep the VOEvent broker code decoupled from any GCN-specific logic. Notably, it allows us to keep all of the details of the GCN-specific “Notice Type” concept out of `gwcelery.voevent`.

`gwcelery.voevent.signals.voevent_received` = <Signal: `voevent_received` providing\_args={'xml\_document'}>  
 Fired whenever a VOEvent is received.

**Parameters** `xml_document` (`comet.utility.xml.xml_document`) – The XML document that was received. The raw file contents are available as `xml_document.raw_bytes`. The `lxml.etree` representation of the document is available as `xml_document.element`.

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 `gwc celery.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 GW Celery 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.

To build the documentation, run the following command in the top of your source directory:

```
$ python setup.py build_sphinx
```

Then to view the documentation, open the file `build/sphinx/html/index.html` in your favorite web browser.





## 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`.

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. **Complete the acceptance tests.** Our acceptance tests consist of a manual checklist for verifying that the pipeline satisfies certain requirements on the playground environment. The checklist is maintained as a [GitLab issue template](#) and is under version control in the special directory `.gitlab/issue_templates`.

Create a [new issue](#) in GitLab. Set the title to Release version `MAJOR.MINOR.PATCH`. In the Choose a template dropdown menu, select Create a Release. The description field will be automatically populated with the checklist. Submit the issue.

Complete the items in the checklist and check them off one by one on the release issue before proceeding to the next step.



4. **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; closes #N`, where `N` is the release issue's number.

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

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

5. **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`.

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

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

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

8. If desired, navigate to the GitLab project's [Environments](#) page and trigger a deployment to production.

### 9.1 0.9.0 (2019-11-23)

This is the initial release of GWCelery for O3b.

- Changes related to configuration settings
  - Use the Redis server that is provided by the operating system (e.g. as a systemd unit) rather than starting our own Redis server. This prevents a race condition between the shutdown of Redis and the shutdown of the workers that caused the workers to hang on shutdown.
  - Update HTCondor accounting group from O2 to O3.
  - Increase throughput for sky localization tasks by offloading processing of the `openmp` Celery queue to 40 workers that are launched via HTCondor on specially configured cluster nodes.
  - Use `mpich` as the MPI runtime for parameter estimation.
  - Use different HTCondor accounting groups for Celery workers depending on whether GWCelery is running in the playground environment (`ligo.dev.o3.cbc.pe.bayestar`) or the production environment (`ligo.prod.o3.cbc.pe.bayestar`).
  - Drop support for Python 3.6 so that we can use the `check_output` keyword argument that was added to `subprocess.run()` in Python 3.7.
  - Pin `gwp` to `<= 0.15.0` since the updated `gwp` fails to read Virgo's state vector.
  - Update `ligo-followup-advocate` version to 1.1.3.
- Changes related to superevent/orchestrator design
  - Add event completeness to publishability criterion. All three of `PASTRO_READY`, `SKYMAP_READY`, and `EMBRIGHT_READY` will be used to evaluate event completeness for CBC events. Only the `SKYMAP_READY` label will be used to evaluate completeness for burst events.
  - Use `EM_Selected` to freeze the preferred event of a superevent and launch a preliminary alert.
  - Make sub-threshold annotations independent of annotations for superevents which pass public alert threshold.

- Prevent second preliminary to be sent in the event of any advocate action. Previously, this was only being prevented for ADVNO.
  - Make skymaps from parameter estimation public automatically.
- Changes related to online parameter estimation
  - Move a comment attached to posterior samples to the corresponding skymap.
  - Add bilby online parameter inference workflow.
  - Fix approximant name used for automatic parameter estimation.
  - Start parameter estimation on mock events.
  - Add acceptance tests of parameter estimation.
  - Use nodes dedicated to online PE also for playground events so that the test runs do not get stuck due to the lack of resources.
  - Add spins in online PE on playground events so that embright probabilities are calculated based on the posterior samples without errors.
  - Remove skymap generation from PE DAG so that it will not be generated twice.
  - Notify which pe pipeline failed for the failure of pe condor jobs.
- Changes related to external coincidences
  - Create RAVEN circular if EM\_COINC label is applied to superevent.
  - Make coincidence FAR synchronous within RAVEN pipeline to fix race condition.
  - Remove redundant SNEWS handler key.
  - Remove generation of em\_coinc circular when EM\_COINC label is applied.
  - Apply EM\_COINC to preferred event when coincidence passes RAVEN publishing conditions.
  - Attempt fetching and uploading Fermi skymap upon receiving GCN notice.
- Changes related to skymap generation
  - Revert back to running BAYESTAR for all G events.
  - Pass the `-j` flag to `ligo-skymap-from-samples` to speed up skymap generation.
- Changes related to automated data quality checks
  - Create omegascans for all detectors upon creation of new superevent.
  - Run `check_vectors` upon the creation of a superevent. This will allow subthreshold superevents to be annotated with DQOK or DQV label.
- Changes to the Flask dashboard
  - Teach preliminary alert form in Flask dashboard to present a dropdown of events sorted by the preferred event criterion.
  - Incorporate update circular into flask app.
- Other changes
  - Add a task to `em_bright.py` to compute and upload source properties upon the upload of `LALInference.posterior_samples.hdf5`.

## 9.2 0.8.7 (2019-09-14)

- Update ligo-raven version to 1.17.

## 9.3 0.8.5.1 (2019-09-04)

This is a non-sequential bugfix release based on version 0.8.5.1 to fix the following issue:

- Fix a lethal bug in `em_bright.py` introduced in version 0.8.5. The bug would incorrectly use the `snr` as the maximum mass of the NS and therefore the source property estimation for pipelines apart from `gstlal` would be grossly incorrect.

## 9.4 0.8.6 (2019-09-01)

- Update ligo-raven version to 1.16.
- Fix a bug that prevented retrying of failed GraceDB API calls in the superevent manager.
- Add a retry for one more potential GraceDB API failure in the initial and update alert workflows.
- In the playground environment only, upload each mock event several times in rapid succession with random jitter in order to simulate multiple pipeline uploads.
- Expose events to the public prior to sending any kind of alert: preliminary, initial, update, or retraction. Previously this behavior only occurred for preliminary alerts, which created the unusual and undesirable possibility of a public GCN for an event that is not public. As before, events are only exposed to the public in the production environment, and not in the playground environment.
- Propagate sky map file extensions (as in `bayestar.fits.gz`, 1) to the URLs that are presented in GCN notices.
- Generate flattened FITS files and sky map visualizations for all superevents, even those that do not rise to the public alert threshold. Note that as a side effect all superevents will have the `EM_Selected` label applied, since it is used as a semaphore to trigger the annotations. The `ADVREQ` label used to serve double duty as the semaphore and also as the wake-up call for follow-up advocates, but now it only serves the latter purpose.

The feature of generating flattened FITS files and sky map plots for all superevents comes as a request from the Fermi and Swift sub-threshold searches.

- Delay running BAYESTAR until the superevent's preferred event has stabilized. BAYESTAR is the most computationally intensive postprocessing task and running it for all events belonging to a superevent was a bottleneck.
- For the playground environment only, decrease the timeout for stabilization of the preferred event from 5 minutes to 2 minutes, which is comparable to how long it has taken recent events to settle. This does not affect the configuration of the production environment.
- Changed `handle_cbc_event` handler to call `gstlal` trained ML based inference for source property computation for `gstlal` triggers.
- Apply `EM_COINC` to superevent and external event in parallel.

## 9.5 0.8.5 (2019-08-23)

- Made `raven.py` tests more robust and have increased coverage.

- Removed the feature of `p-astro` and `em_bright` reading mean counts, livetimes or ML classifiers from `emfollow/data`; moved them to `lscsoft/p_astro` as package data. Added back the `test_tasks_p_astro.py` that was accidentally taken out in `v0.8.0`. Pinned `p_astro == 0.8.0`.

## 9.6 0.8.4 (2019-08-16)

- Fixed a bug in `gracedb.create_tag` to handle the scenario when multiple log messages exist for the same filename. The tag is applied to the most recent log message.
- Retry GraceDB API calls that fail due to receiving incomplete or malformed HTTP responses, as indicated by `http.client.HTTPException` exceptions. This should work around the increased incidence of `RemoteDisconnected` exceptions that we have seen recently, and that caused a delay in sending out the preliminary alert for S190814bv.

## 9.7 0.8.3 (2019-08-09)

- Enable Redis integration for Sentry error reporting.
- Downgrade `lalsuite` to 6.54 since `lal inference_pipe` in 6.59 has a minor bug, which breaks automatic parameter estimation.
- Include the number of participating detectors in the preferred event selection criterion for compact binaries: 3-detector events should be preferred over 2-detector events, and 2-detector events over 1-detector events, on the basis of more accurate localization. Ties are still broken by SNR.
- Catch `SystemExit` exceptions from Python command line tools called in Celery tasks and re-raise them as `RuntimeError` exceptions so that they do not cause the workers to exit.

## 9.8 0.8.2 (2019-08-02)

- Apply the `public` tag to data products before sending out an update GCN notice. This will prevent human errors related to not exposing LALInference files before sending a GCN notice.
- Don't read the entire HTTP response from GraceDB POST requests. We only need the HTTP status code. This change might speed up GraceDB API calls a little bit.
- Increase preliminary alert timeout back to 5 minutes.
- Make `gracedb.create_superevent`, `gracedb.update_superevent` and `gracedb.add_event_to_superevent` idempotent by catching the `HTTPError` from GraceDB that occurs if the superevent has already been created.
- Fix bug where neither the space-time nor temporal coincidence far is calculated if external sky map is unavailable.
- Update `ligo.skymap` to 0.1.9. This version changes the data type of the multi-resolution HEALPix format's `UNIQ` column from an unsigned integer to a signed integer.  
Starting with this version, the Linux builds of `ligo.skymap` are compiled and optimized using the Intel C Compiler.
- Change the trials factor for CBC searches to 4, since SPIIR is performing a single search, and that for burst to 3, since oLIB is not currently in operation.



## 9.9 0.8.1 (2019-07-29)

- Downgrade lalsuite to 6.59.
- Revert change that tried to fix incorrect key for querying external events. The keys were correct before.

## 9.10 0.8.0 (2019-07-26)

- Assign `gwcelery.tasks.skymaps.plot_volume` tasks a reduced Celery priority as compared to `gwcelery.tasks.bayestar.localize` so that the latter are given preference. This ought to speed up the preparation of preliminary GCN notices because only the latter are required for GCNs but both kinds of tasks compete for slots in the resource-intensive OpenMP queue.
- Reduce priority for CBC annotation tasks for events that do not pass the public alert threshold.
- Update lalsuite to 6.60.
- Ensure `gracedb` calls to create and update superevents are retried in the event of transient GraceDB API errors.
- Update `ligo-raven` version to 1.15. Apply `EM_COINC` label in `raven.py` to give more control and prevent race conditions.
- Use the space-time coincidence FAR as the default for RAVEN, use the temporal coincidence FAR when sky maps are not available.
- Check if GRB is sub-threshold, set search to be 'SubGRB'. Pass search through external triggers pipeline and RAVEN.
- Tune Celery's `result_expires` setting from its default value of one day to five minutes. Since we pass large byte strings as task arguments and return values, one day is too long to keep task tombstones in the database. This adjustment should reduce the memory footprint of the Redis server during periods with very high rates of GraceDB uploads.

The downside is that task details will remain browsable in Flower for a much shorter period.

- Remove `p_astro_gstlal.py` module, corresponding test modules, and documentation; `p_astro` will be reported as a pipeline product from `gstlal`. The computation of `p_astro` for all other pipelines is unaffected.
- Fix `EM_COINC` bug where it is being over-applied to superevents.
- Fix bug where wrong key was called for querying external events.

## 9.11 0.7.1 (2019-07-12)

- The initial alert workflow will now consider only `*.fits.gz` sky maps and not `*.fits` sky maps for GCN Notices. It was an oversight that we did not exclude `*.fits` files from the list of extensions to consider when we updated the handling of multiresolution sky maps.
- Catch and retry HTTP 429 ("Too Many Requests") errors from GraceDB.
- Enable Sentry integration for Tornado in order to capture errors from the Flower console.
- Fix file extensions for LALInference sky map PNG files: they should be named `LALInference.png`, not `LALInference.multiorder.png`.
- Increase the Redis server's log verbosity in order to help diagnose Redis client connection dropouts.

- Run sky map plotting and annotation tasks asynchronously so that they do not block sending preliminary alerts. Their outputs are only for human consumption; they are not needed in order to prepare GCN Notices.

## 9.12 0.7.0 (2019-06-21)

- Trigger a preliminary alert for a superevent upon the first time that the preferred event is set to an event that meets the public alert criterion.

This fixes a longstanding issue that has prevented automated preliminary alerts from being sent so far. The preferred event *at the instant that the timeout ended* did not meet the public alert criterion, but a preferred event that was selected some tens of seconds later did.

- Decrease preliminary alert timeout to one minute.
- The combined effect of these changes should be to decrease the latency for producing preliminary alerts from 7 minutes to 2 minutes.

## 9.13 0.6.3 (2019-06-14)

- Work around a Celery canvas bug that prevented LALInference postprocessing from completing.
- Fix a copy-paste error that caused DQV and INJ labels to be ignored when determining whether to send a preliminary alert.
- Move RAVEN time coincidence windows to the application configuration.
- Document the acceptance tests checklist in the instructions for preparing a release.
- Update ligo-raven to version 1.14.

## 9.14 0.6.2 (2019-06-07)

- Add a dependency on dnspython to silence the following warning message from SleekXMPP:

```
DNS: dnspython not found. Can not use SRV lookup.
```

- Pin some recently updated dependencies of Celery that caused unit test failures: amqp <= 2.4.2, kombu <= 4.5.0, vine <= 1.3.0.
- Prevent subthreshold GRBs with low reliability from being processed as external events.
- Add a task in orchestrator.py to generate FITS files and sky map images automatically whenever an HDF5 posterior samples file is uploaded.
- Remove special-case handling of single-instrument events. Now, the eligibility of an event for a public alert is determined only on the basis of its false alarm rate.
- Run parameter estimation on nodes dedicated to online-PE.
- Emcoinc circular is triggered when RAVEN uploads a coincident FAR.
- Pin scipy since scipy>=1.3.0 removes an interpolation function which lalinference postprocessing requires.

## 9.15 0.6.1 (2019-05-24)

- Work around a bug in the Sentry Python SDK that caused excessive reporting of certain GraceDB exceptions that are listed in tasks' `autoretry_for` settings. See [getsentry/sentry-python#370](#).
- Change the name of BAYESTAR localization files to `bayestar.multiorder.fits` to distinguish them from flat-resolution HEALPix files, which are still named `bayestar.fits.gz`.
- Reimplement LVAalert listener as a Celery bootstrap to avoid needing to track a singleton task using a Redis lock, because Redis locks do not play nicely with Redis persistence. The `--lvalert` command line option must now be passed in order to enable the LVAalert listener.
- Turn on Redis database persistence so that Celery task state is preserved across restarts.
- Add `expose_to_public` setting to disable exposing GraceDB events to the public in all environments except for production.
- Update to the latest version of GWPy and un-pin Matplotlib because GWPy now supports Matplotlib 3.1.
- Pin LALSuite to version 6.54 because LALInference in LALSuite 6.55 is not compatible with Python 3.

## 9.16 0.6.0 (2019-05-20)

- Work around a bug in complex Celery canvases (see [celery/celery#5512](#)) that prevented initial GCN notices from being sent. As a side effect of this workaround, the initial, update, and retraction canvases will not automatically expose events to the public.

The preliminary alert canvas still *does* expose events to the public, so under normal circumstances, the follow-up advocate should not have to manually do that. However, if the event has not been exposed to the public for whatever reason, then the follow-up advocate should expose it to the public manually before applying the ADVOK label. See [emfollow/followup-advocate-guide!2](#).

- Reduce the false alarm rate threshold for parameter estimation to decrease cluster load.
- Remove redundant LVAalert subscription in `handle_lvalert_grb` to prevent double calls to RAVEN.
- Read template weights for `P_astro` from hdf5 file using `h5py` for speedup.
- Require `matplotlib < 3.1` because `matplotlib 3.1` breaks importing `gwpy`:

```
/usr/local/lib/python3.7/site-packages/gwpy/plot/rc.py:79: in <module>
    rcParams.get('text.latex.preamble', []) + tex.MACROS),
E   TypeError: can only concatenate str (not "list") to str
```

- Make `gwcelery.tasks.gracedb.get_superevents` and `gwcelery.tasks.gracedb.get_events` take any number of keyword arguments to be passed to corresponding client methods.
- Update the superevent `t_0` field whenever the preferred event changes.

## 9.17 0.5.7 (2019-05-13)

- If the VOEvent broker is disabled by setting `voevent_broker_whitelist` to an empty list, then suppress the normal error message that would occur when attempting to send a VOEvent when there are no broker connections.
- Rearrange preliminary alert workflow so that sky map plots are generated for the newly added FITS file rather than an older FITS file that coincidentally has the same name.

- Have `gwcelery.detchar.check_vectors` task apply all GraceDB log messages in order to increase robustness to recoverable GraceDB API errors.
- Port over majority of `P_astro` code from `gwcelery` to the `p-astro` package.
- Use cleaned data for parameter estimation.
- The `DQOK` and `DQV` labels should be mutually exclusive. When `gwcelery.tasks.detchar.check_vectors` adds one of the `DQOK` or `DQV` labels, it will now first remove the other label.
- Change exception in `VOEvent` parsing of Fermi subthreshold alerts to match real incoming alerts.
- Update Celery to 4.3.0.
- Automatically select the most up-to-date calibration uncertainties for parameter estimation.

## 9.18 0.5.6 (2018-05-08)

- Extend the `orchestrator_timeout` to 300s and the `pe_timeout` to 345s. The previous timeout was not sufficient for the online pipelines to upload all of their possible candidates, hence the extension.

## 9.19 0.5.5 (2019-05-03)

- Cycle through `llhoft`, high latency frames, and low latency frames in `detchar`'s cache creation.
- Add explanations on options in `online_pe.jinja2` for those who start parameter estimation based on the ini files uploaded to GraceDB.
- Calculate horizon distance with `psd.xml.gz` to determine the upper limit of distance prior for parameter estimation.
- Start parameter estimation when the lowest FAR of the events in a superevent is lower than the threshold.
- Update the calibration uncertainties used for parameter estimation.
- Handle an exception in `VOEvent` parsing of Fermi subthreshold alerts due to different param names.
- Stop uploading corner plots of intrinsic parameters.
- Connect to different GCN servers to receive alerts in the production and playground environments, because GCN does not support multiple receiver connections from the same client IP address to the same server.
- Change the preferred event assignment logic to not let accidental candidates like G330298 which have low FAR but high SNR values to become the preferred event. From now on, `superevents.should_publish` takes maximum precedence for selecting the preferred event. The same is also used by `orchestrator` to expose events.
- Update RAVEN coinc FAR task call which uses string params versus un-pickleable class object params.
- Make sure to consume the entire response from every GraceDB API request. This will ensure that GraceDB API call has completed before the pipeline continues, and will fix errors like we encountered with S190426c where the pipeline would march along before uploads had finished.
- Apply ADVREQ label earlier in the preliminary alert workflow.
- Update LALSuite to version 6.54. We are now using a stable version again instead of a nightly build.
- Add Nagios checks for GCN connectivity.
- Improve uploaded comments so that it is easily understood which event has triggered parameter estimation.

## 9.20 0.5.4 (2019-05-01)

- Provide a value for terrestrial count for  $P_{\text{astro}}$  for non-gstlal pipelines that is consistent with the FAR threshold used.

## 9.21 0.5.3 (2019-04-17)

- Update ligo-followup-advocate to 0.0.28.
- Stop using unreviewed cleaned data for parameter estimation.
- Update detchar check to analyze full template duration for CBC events.

## 9.22 0.5.2 (2019-04-15)

- Fix typo in `gracedb.get_instruments`: there was the attribute lookup `single.ifo`, which should have been the dictionary lookup `single[ifo]`.
- Fix `gwcelery.tasks.p_astro_other.choose_snr` for gstlal. This method did not previously expect to be called for gstlal, since it is typically only invoked for other pipelines. However, there is one case when `choose_snr` is invoked for gstlal, which is when the `ranking_data` file from gstlal is corrupted with NaNs, causing  $P_{\text{astro}}$  for gstlal to fail. Thus, `choose_snr` has now been fixed to also handle gstlal as a pipeline.

## 9.23 0.5.1 (2019-04-12)

- Changed default for `em-bright` from 2.83 to 3.0  $M_{\text{sun}}$  to be consistent with notices.

## 9.24 0.5.0 (2019-04-12)

- Give permissions to read the files under parameter estimation run directories to non-owner people so that rota people can check their progresses. The naming convention of the run directories changed.
- EM-Bright ML classification requires review. Until then, give answer based on low-latency estimates.
- Compute  $P_{\text{astro}}$  with mass-based template weighting. Template weights are now keyed on template parameters, rather than bin numbers. This should make  $P_{\text{astro}}$  immune to binning conventions.
- Add form to manually send a preliminary GCN Notice.
- Fix a typo in `gwcelery.sub` that caused the Flower dashboard to fail to start.
- Round `iDQ p(glitch)` to 3 decimal places in GraceDB log message.
- Switch log telemetry from the on-premise instance of Sentry at Caltech to a cloud-hosted subscription to `sentry.io`.
- In the playground configuration, the `gwcelery.tasks.gcn.validate` task was producing false alarms because the GCN receiver was receiving `VOEvents` from the production instance, which would certainly differ in content from `VOEvents` in the playground instance. Fix this by having `gwcelery.tasks.gcn.validate` discard all `VOEvents` if the `VOEvent` broadcaster is disabled.
- Update ligo-followup-advocate to 0.0.27.

- Wait for 1 minute before parameter estimation in case the preferred event is updated with high latency.
- Ensure that `P_astro` accounts for very loud MBTA and PyCBC events, whose FAR saturate at certain low values depending on instrument combination, but whose SNRs can increase indefinitely.
- When a user triggers a Preliminary or Update alert through the Flask interface, create a GraceDB log message to record the username.
- The Flask interface will now show a confirmation dialog before sending any alerts.
- Add a terrifying warning to the Flask interface to make it clear that the interface is live.

## 9.25 0.4.3 (2019-04-05)

- Now that LIGO/Virgo alerts are public, switch the GCN listener that we use to confirm receipt of our own GCN Notices from a managed, private connection to an anonymous, public connection.
- Migrate the Flask and Flower dashboards from `ldas-jobs.ligo.caltech.edu` to `emfollow.ligo.caltech.edu`. The new URLs are:
  - <https://emfollow.ligo.caltech.edu/gwcelery>
  - <https://emfollow.ligo.caltech.edu/flower>
  - <https://emfollow.ligo.caltech.edu/playground/gwcelery>
  - <https://emfollow.ligo.caltech.edu/playground/flower>

Remove the `htaccess` file from our `public_html` directory, since the reverse proxy configuration is now the responsibility of system administrators.

- Display the GWCelery version number in the Flask application.
- Add visualizations for `p_astro.json` source classification files.

## 9.26 0.4.2 (2019-04-05)

- Calculation of number of instruments is now unified across superevent manager and orchestrator using `gracedb` method `get_number_of_instruments`.
- Enable automated preliminary alerts for all pipelines because disabling them in the orchestrator introduced some issues due to the criteria for releasing a public alert drifting away from the definition of a the preferred event of a superevent. We will instead trust pipelines that are still under review will upload events to the playground rather than the production environment.

## 9.27 0.4.1 (2019-04-02)

- Fixed normalization issues with `p_astro_gstlal.py`; normalization was being applied in the wrong places during Bayes factor computation.
- Require `celery < 4.3.0` because that version breaks the nagios unit tests.
- Update false alarm rate trials factors for preliminary alerts.
- Enable sending GCN notices for fully automated preliminary alerts.
- Add `threshold_snr` option in `online_pe.jinja2`, which is used to determine the upper limit of distance prior.

- Use the same criteria to decide whether to expose an event publicly in GraceDB as we use to decide whether to issue a public alert.
- Do not issue public alerts for single-instrument GW events.
- Disable automated preliminary alerts for all pipelines but gstlal and cWB due to outstanding review items for the other pipelines.

## 9.28 0.4.0 (2019-03-29)

- This is the penultimate release before LIGO/Virgo observing run 3 (O3).
- Make detchar results easier to read by formatting as HTML table.
- Allow iDQ to label DQV onto events based on p(glitch). Adjustable by pipeline.
- Move functions in tasks/lalinference.py to lalinference\_pipe.py in lalsuite.
- Take into account calibration errors in automatic Parameter Estimation.
- Do not use margphi option for automatic Parameter Estimation with ROQ waveform since that option is not compatible with ROQ likelihood.
- Adjust WSGI middleware configuration to adapt to a change in Werkzeug 0.15.0 that broke redirects on form submission in the Flask app. See <https://github.com/pallets/werkzeug/pull/1303>.
- Use the new `ligo.lw` module for reading gstlal's `ranking_data.psd.xml.gz` files, because these files are now written using the new LIGO-LW format that uses integer row IDs.
- Use clean data for parameter estimation.
- Use production accounting group for PE runs on gracedb events.
- Change threshold from log-likelihood equals 6 to a dynamic threshold that ensures that all gstlal events uploaded to gracedb get assigned a `P_astro` value.

## 9.29 0.3.1 (2019-03-18)

- Fix a bug in translating keys from `source_classification.json` to keyword arguments for `GraceDB.createVOEvent` that caused `VOEvents` to be missing the `HasNS` and `HasRemnant` fields.
- FAR threshold for sending preliminary notices for CBC is changed to 1 per 2 months.
- Upload log files when LALInference parameter estimation jobs fail or are aborted.
- Changed the filename `source_classification.json` to `em_bright.json`.
- Change condor log directory from `/var/tmp` to `~/cache/condor` since gwcelery workers have separate `/var/tmp` when they are running as condor jobs and that causes problems when gwcelery tries to read log files.
- Limit the maximum version of gwpy to 0.14.0 in order to work around a unit test failure that started with gwpy 0.14.1. See <https://git.ligo.org/emfollow/gwcelery/issues/95>.
- Upload a diff whenever a LIGO/Virgo `VOEvent` that we receive from GCN does not match the original that we sent.
- Wait for low-latency or high-latency frame files being transferred to the cluster before parameter estimation starts.

## 9.30 0.3.0 (2019-03-01)

- Fixed exponent in the expression of foreground count in `p_astro_other` task.
- Run the sky map postprocessing and add the `PE_READY` tag when LALInference finishes.
- Include `EM_COINC` triggered circulars to upload to the superevent page.
- `p-astro` reads mean values from a file on CIT, new mass-gap category added. Removed redundant functions from `p_astro_gstlal` module.
- Continuous deployment on the Caltech cluster now uses a robot keytab and `gsissh` instead of SSH keys and vanilla `ssh` because the new `my.ligo.org` SSH key management does not support scripted access.
- Improve the isolation between the production and playground instances of GWCelery by deploying them under two separate user accounts on the Caltech cluster.
- Add functionality for `em_bright` task to query `emfollow/data` for trained machine learning classifier and report probabilities based on it.

## 9.31 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 `gwcclery 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.

## 9.32 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 LVAAlert 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.



### 9.33 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.

### 9.34 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`.

### 9.35 0.2.2 (2018-12-14)

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

### 9.36 0.2.1 (2018-12-14)

- Update detach module configuration for ER13.

### 9.37 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 instead 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.

### 9.38 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](#).

### 9.39 0.1.6 (2018-11-14)

- Update URL for static example event.

### 9.40 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.

### 9.41 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.

## 9.42 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.

## 9.43 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.

## 9.44 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  $\geq 2.0.1$  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.

## 9.45 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 LVAAlert host for `gracedb-playground.ligo.org`.
- Add experimental integration with [Sentry](#) for log aggregation and error reporting.
- Track API and LVAAlert schema changes in ligo-gracedb 2.0.0.

## 9.46 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.

## 9.47 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 `gwcclery.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 `gwcclery 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.

## 9.48 0.0.29 (2018-07-31)

- Add a workflow graph to superevents module documentation.
- Add `gwcclery condor resubmit` as a shortcut for `gwcclery condor rm; gwcclery 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.

## 9.49 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.

## 9.50 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`.

## 9.51 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.

## 9.52 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 computation is launched for all `cbc_gstlal` triggers.

## 9.53 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.

## 9.54 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.

## 9.55 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.

## 9.56 0.0.21 (2018-07-10)

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

## 9.57 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.

## 9.58 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.

## 9.59 0.0.18 (2018-07-06)

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

## 9.60 0.0.17 (2018-07-05)

- Undo accidental configuration change in last version.

## 9.61 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 LAlert client, instead of [PyXMPP2](#). Because SleekXMPP has first-class support for publish-subscribe, the LAlert listener can now automatically subscribe to all LAlert nodes for which our code has handlers. Most of the client code now lives in a new external package, [sleek-lvalert](#).

## 9.62 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.

## 9.63 0.0.14 (2018-06-28)

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

## 9.64 0.0.13 (2018-06-28)

- Change FAR threshold for creation of superevents to 1 per day.
- Update ligo-followup-advocate to  $\geq 0.0.10$ . 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.

## 9.65 0.0.12 (2018-06-28)

- Prevent receiving duplicate copies of LVAAlert messages by unregistering redundant LVAAlert message types.
- Update to ligo-followup-advocate  $\geq 0.0.9$  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.

## 9.66 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 LVAAlert 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 LAlert 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 embarrassing `TypeError`s.
- 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()`.

## 9.67 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.
- LAlert 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.

## 9.68 0.0.9 (2018-06-06)

- Add missing LVAalert message types to superevent handler.

## 9.69 0.0.8 (2018-06-06)

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

## 9.70 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.

## 9.71 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 LVAalert 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...
```

LVAalert message handlers are declared like this:

```
import json
from gwcelery.tasks import lvalert

@lvalert.handler('cbc_gstlal',
                 'cbc_pycbc',
                 'cbc_mbt')
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.

## 9.72 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.

## 9.73 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.

## 9.74 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.

## 9.75 0.0.2 (2018-04-27)

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

## 9.76 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 LVAalert 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`

GWCCelery 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>





### C

`gwcelery.conf`, 21  
`gwcelery.conf.development`, 24  
`gwcelery.conf.playground`, 24  
`gwcelery.conf.production`, 24  
`gwcelery.conf.test`, 25

### e

`gwcelery.email`, 25  
`gwcelery.email.bootsteps`, 25  
`gwcelery.email.signals`, 26

### g

`gwcelery`, 21

### l

`gwcelery.lvalert`, 26  
`gwcelery.lvalert.bootsteps`, 26  
`gwcelery.lvalert.signals`, 27

### s

`gwcelery.sentry`, 27

### t

`gwcelery.tasks`, 27  
`gwcelery.tasks.bayestar`, 27  
`gwcelery.tasks.circulars`, 28  
`gwcelery.tasks.condor`, 28  
`gwcelery.tasks.detchar`, 30  
`gwcelery.tasks.em_bright`, 33  
`gwcelery.tasks.external_triggers`, 35  
`gwcelery.tasks.first2years`, 36  
`gwcelery.tasks.gcn`, 37  
`gwcelery.tasks.gracedb`, 38  
`gwcelery.tasks.inference`, 39  
`gwcelery.tasks.ligo_fermi_skymaps`, 36  
`gwcelery.tasks.lvalert`, 42  
`gwcelery.tasks.orchestrator`, 45  
`gwcelery.tasks.p_astro`, 49

`gwcelery.tasks.raven`, 50  
`gwcelery.tasks.skymaps`, 52  
`gwcelery.tasks.superevents`, 55  
`gwcelery.tools`, 57  
`gwcelery.tools.condor`, 58  
`gwcelery.tools.flask`, 58  
`gwcelery.tools.nagios`, 59

### u

`gwcelery.util`, 59

### v

`gwcelery.voevent`, 60  
`gwcelery.voevent.bootsteps`, 60  
`gwcelery.voevent.logging`, 61  
`gwcelery.voevent.signals`, 62  
`gwcelery.voevent.util`, 61



## A

`add_arguments()` (gwcelery.tools.condor.CondorCommand method), 58  
`add_arguments()` (gwcelery.tools.flask.FlaskCommand method), 58  
`add_worker_arguments()` (in module gwcelery.email), 25  
`add_worker_arguments()` (in module gwcelery.lvalert), 26  
`after_setup_logger()` (in module gwcelery.voevent.logging), 61  
`app` (in module gwcelery), 21

## B

`Broadcaster` (class in gwcelery.voevent.bootsteps), 60

## C

`catch_retryable_http_errors()` (in module gwcelery.tasks.gracedb), 38  
`check_idq()` (in module gwcelery.tasks.detchar), 31  
`check_status()` (in module gwcelery.tools.nagios), 59  
`check_vector()` (in module gwcelery.tasks.detchar), 32  
`check_vector_prepost` (in module gwcelery.conf), 22  
`condor_accounting_group` (in module gwcelery.conf), 21  
`condor_accounting_group` (in module gwcelery.conf.production), 24  
`CondorCommand` (class in gwcelery.tools.condor), 58  
`configure()` (in module gwcelery.sentry), 27  
`create()` (gwcelery.email.bootsteps.Receiver method), 25  
`create()` (gwcelery.lvalert.bootsteps.Receiver method), 26

`create()` (gwcelery.voevent.bootsteps.Reactor method), 60  
`create_cache()` (in module gwcelery.tasks.detchar), 30  
`create_combined_skymap()` (in module gwcelery.tasks.ligo\_fermi\_skymaps), 36  
`create_service()` (gwcelery.voevent.bootsteps.Broadcaster method), 60  
`create_service()` (gwcelery.voevent.bootsteps.Receiver method), 61  
`CRITICAL` (gwcelery.tools.nagios.NagiosPluginStatus attribute), 59

## D

`detchar_bit_definitions` (in module gwcelery.conf), 23  
`dqr_json()` (in module gwcelery.tasks.detchar), 31  
`DSN` (in module gwcelery.sentry), 27

## E

`email_host` (in module gwcelery.conf), 22  
`email_received` (in module gwcelery.email.signals), 26  
`expose_to_public` (in module gwcelery.conf), 21  
`expose_to_public` (in module gwcelery.conf.production), 24

## F

`FlaskCommand` (class in gwcelery.tools.flask), 58  
`FROZEN_LABEL` (in module gwcelery.tasks.superevents), 55

## G

`generate_table()` (in module gwcelery.tasks.detchar), 31  
`get_active_lvalert_nodes()` (in module gwcelery.tools.nagios), 59

- `get_active_queues()` (in module `gwcelery.tools.nagios`), 59
  - `get_active_voevent_peers()` (in module `gwcelery.tools.nagios`), 59
  - `get_category()` (in module `gwcelery.tasks.superevents`), 55
  - `get_constraints()` (in module `gwcelery.tools.condor`), 58
  - `get_expected_lvalert_nodes()` (in module `gwcelery.tools.nagios`), 59
  - `get_expected_queues()` (in module `gwcelery.tools.nagios`), 59
  - `get_host_port()` (in module `gwcelery.voevent.util`), 61
  - `get_instruments()` (in module `gwcelery.tasks.superevents`), 56
  - `get_instruments_in_ranking_statistic()` (in module `gwcelery.tasks.superevents`), 56
  - `get_local_ivo()` (in module `gwcelery.voevent.util`), 61
  - `get_network()` (in module `gwcelery.voevent.util`), 61
  - `get_snr()` (in module `gwcelery.tasks.superevents`), 56
  - `get_ts()` (in module `gwcelery.tasks.superevents`), 55
  - `gracedb_host` (in module `gwcelery.conf`), 21
  - `gracedb_host` (in module `gwcelery.conf.development`), 24
  - `gracedb_host` (in module `gwcelery.conf.production`), 24
  - `gracedb_host` (in module `gwcelery.conf.test`), 25
  - `gwcelery` (module), 21
  - `gwcelery.conf` (module), 21
  - `gwcelery.conf.development` (module), 24
  - `gwcelery.conf.playground` (module), 24
  - `gwcelery.conf.production` (module), 24
  - `gwcelery.conf.test` (module), 25
  - `gwcelery.email` (module), 25
  - `gwcelery.email.bootsteps` (module), 25
  - `gwcelery.email.signals` (module), 26
  - `gwcelery.lvalert` (module), 26
  - `gwcelery.lvalert.bootsteps` (module), 26
  - `gwcelery.lvalert.signals` (module), 27
  - `gwcelery.sentry` (module), 27
  - `gwcelery.tasks` (module), 27
  - `gwcelery.tasks.bayestar` (module), 27
  - `gwcelery.tasks.circulars` (module), 28
  - `gwcelery.tasks.condor` (module), 28
  - `gwcelery.tasks.detchar` (module), 30
  - `gwcelery.tasks.em_bright` (module), 33
  - `gwcelery.tasks.external_triggers` (module), 35
  - `gwcelery.tasks.first2years` (module), 36
  - `gwcelery.tasks.gcn` (module), 37
  - `gwcelery.tasks.gracedb` (module), 38
  - `gwcelery.tasks.inference` (module), 39
  - `gwcelery.tasks.ligo_fermi_skymaps` (module), 36
  - `gwcelery.tasks.lvalert` (module), 42
  - `gwcelery.tasks.orchestrator` (module), 45
  - `gwcelery.tasks.p_astro` (module), 49
  - `gwcelery.tasks.raven` (module), 50
  - `gwcelery.tasks.skymaps` (module), 52
  - `gwcelery.tasks.superevents` (module), 55
  - `gwcelery.tools` (module), 57
  - `gwcelery.tools.condor` (module), 58
  - `gwcelery.tools.flask` (module), 58
  - `gwcelery.tools.nagios` (module), 59
  - `gwcelery.util` (module), 59
  - `gwcelery.voevent` (module), 60
  - `gwcelery.voevent.bootsteps` (module), 60
  - `gwcelery.voevent.logging` (module), 61
  - `gwcelery.voevent.signals` (module), 62
  - `gwcelery.voevent.util` (module), 61
- ## H
- `handler` (in module `gwcelery.tasks.gcn`), 37
  - `handler` (in module `gwcelery.tasks.lvalert`), 42
  - `handling_system_exit()` (in module `gwcelery.util`), 59
  - `high_latency_frame_types` (in module `gwcelery.conf`), 23
  - `high_latency_frame_types` (in module `gwcelery.conf.production`), 25
  - `hold()` (in module `gwcelery.tools.condor`), 58
- ## I
- `idq_channels` (in module `gwcelery.conf`), 23
  - `idq_channels` (in module `gwcelery.conf.production`), 25
  - `idq_pglitch_thresh` (in module `gwcelery.conf`), 23
  - `idq_veto` (in module `gwcelery.conf`), 23
  - `info()` (`gwcelery.lvalert.bootsteps.Receiver` method), 26
  - `info()` (`gwcelery.voevent.bootsteps.Broadcaster` method), 60
  - `info()` (`gwcelery.voevent.bootsteps.Receiver` method), 61
  - `install()` (in module `gwcelery.email`), 25
  - `install()` (in module `gwcelery.lvalert`), 26
  - `install()` (in module `gwcelery.voevent`), 60
  - `is_3d_fits_file()` (in module `gwcelery.tasks.skymaps`), 52
  - `is_complete()` (in module `gwcelery.tasks.superevents`), 56
- ## J
- `JobAborted`, 28
  - `JobFailed`, 28
  - `JobRunning`, 28

## K

`keyfunc()` (in module `gwcclery.tasks.superevents`), 57

## L

`llhoft_channels` (in module `gwcclery.conf`), 23

`llhoft_glob` (in module `gwcclery.conf`), 23

`llhoft_glob` (in module `gwcclery.conf.production`), 24

`low_latency_frame_types` (in module `gwcclery.conf`), 23

`low_latency_frame_types` (in module `gwcclery.conf.production`), 24

`lvalert_host` (in module `gwcclery.conf`), 21

`lvalert_host` (in module `gwcclery.conf.production`), 24

`lvalert_host` (in module `gwcclery.conf.test`), 25

`lvalert_received` (in module `gwcclery.lvalert.signals`), 27

## M

`main()` (in module `gwcclery.tools.flask`), 58

`mock_events_simulate_multiple_uploads` (in module `gwcclery.conf`), 23

`mock_events_simulate_multiple_uploads` (in module `gwcclery.conf.playground`), 24

## N

`NagiosCommand` (class in `gwcclery.tools.nagios`), 59

`NagiosCriticalError`, 59

`NagiosPluginStatus` (class in `gwcclery.tools.nagios`), 59

`name` (`gwcclery.email.bootsteps.Receiver` attribute), 25

`name` (`gwcclery.lvalert.bootsteps.Receiver` attribute), 26

`name` (`gwcclery.voevent.bootsteps.Broadcaster` attribute), 60

`name` (`gwcclery.voevent.bootsteps.Reactor` attribute), 60

`name` (`gwcclery.voevent.bootsteps.Receiver` attribute), 61

`NamedTemporaryFile()` (in module `gwcclery.util`), 59

`NotEnoughData`, 39

## O

`OK` (`gwcclery.tools.nagios.NagiosPluginStatus` attribute), 59

`omegascan_durations` (in module `gwcclery.conf`), 23

## P

`pe_results_path` (in module `gwcclery.conf`), 23

`pe_results_url` (in module `gwcclery.conf`), 23

`pe_threshold` (in module `gwcclery.conf`), 23

`pe_timeout` (in module `gwcclery.conf`), 22

`pre_pe_tasks()` (in module `gwcclery.tasks.inference`), 40

`preliminary_alert_far_threshold` (in module `gwcclery.conf`), 22

`preliminary_alert_timeout` (in module `gwcclery.conf`), 22

`preliminary_alert_trials_factor` (in module `gwcclery.conf`), 22

`PromiseProxy` (class in `gwcclery.util`), 59

Python Enhancement Proposals  
PEP 8, 64

## Q

`q()` (in module `gwcclery.tools.condor`), 58

## R

`raven_coincidence_windows` (in module `gwcclery.conf`), 23

`Reactor` (class in `gwcclery.voevent.bootsteps`), 60

`READY_LABEL` (in module `gwcclery.tasks.superevents`), 55

`Receiver` (class in `gwcclery.email.bootsteps`), 25

`Receiver` (class in `gwcclery.lvalert.bootsteps`), 26

`Receiver` (class in `gwcclery.voevent.bootsteps`), 60

`release()` (in module `gwcclery.tools.condor`), 58

`REQUIRED_LABELS_BY_GROUP` (in module `gwcclery.tasks.superevents`), 55

`requires` (`gwcclery.voevent.bootsteps.Receiver` attribute), 61

`resubmit()` (in module `gwcclery.tools.condor`), 58

`RetryableHTTPError` (class in `gwcclery.tasks.gracedb`), 38

RFC

RFC 822, 26

`rm()` (in module `gwcclery.tools.condor`), 58

`run()` (`gwcclery.tools.condor.CondorCommand` method), 58

`run()` (`gwcclery.tools.flask.FlaskCommand` method), 58

`run()` (`gwcclery.tools.nagios.NagiosCommand` method), 59

`run_exec()` (in module `gwcclery.tools.condor`), 58

`running()` (in module `gwcclery.tools.condor`), 58

## S

`SendingError`, 37

`sentry_environment` (in module `gwcclery.conf.development`), 24

`sentry_environment` (in module `gwcclery.conf.playground`), 24

`sentry_environment` (in module `gwcclery.conf.production`), 25

`sentry_environment` (in module `gwcclery.conf.test`), 25

[should\\_publish\(\)](#) (in module `gwcelery.tasks.superevents`), 57  
[snews\\_gw\\_far\\_threshold](#) (in module `gwcelery.conf`), 22  
[start\(\)](#) (`gwcelery.email.bootsteps.Receiver` method), 25  
[start\(\)](#) (`gwcelery.lvalert.bootsteps.Receiver` method), 26  
[start\(\)](#) (`gwcelery.voevent.bootsteps.Reactor` method), 60  
[state\\_vector\\_channel\\_names](#) (in module `gwcelery.conf`), 23  
[stop\(\)](#) (`gwcelery.email.bootsteps.Receiver` method), 25  
[stop\(\)](#) (`gwcelery.lvalert.bootsteps.Receiver` method), 26  
[stop\(\)](#) (`gwcelery.voevent.bootsteps.Reactor` method), 60  
[strain\\_channel\\_names](#) (in module `gwcelery.conf`), 23  
[strain\\_channel\\_names](#) (in module `gwcelery.conf.production`), 25  
[submit\(\)](#) (in module `gwcelery.tools.condor`), 58  
[subthreshold\\_annotation\\_timeout](#) (in module `gwcelery.conf`), 22  
[superevent\\_clean\\_up\\_timeout](#) (in module `gwcelery.conf`), 22  
[superevent\\_d\\_t\\_end](#) (in module `gwcelery.conf`), 22  
[superevent\\_d\\_t\\_start](#) (in module `gwcelery.conf`), 22  
[superevent\\_default\\_d\\_t\\_end](#) (in module `gwcelery.conf`), 22  
[superevent\\_default\\_d\\_t\\_start](#) (in module `gwcelery.conf`), 22  
[superevent\\_far\\_threshold](#) (in module `gwcelery.conf`), 22  
[superevent\\_query\\_d\\_t\\_end](#) (in module `gwcelery.conf`), 22  
[superevent\\_query\\_d\\_t\\_start](#) (in module `gwcelery.conf`), 22

## T

[task\(\)](#) (in module `gwcelery.tasks.gracedb`), 38

## U

[UNKNOWN](#) (`gwcelery.tools.nagios.NagiosPluginStatus` attribute), 59  
[uses\\_gatedhoft](#) (in module `gwcelery.conf`), 23

## V

[voevent\\_broadcaster\\_address](#) (in module `gwcelery.conf`), 21  
[voevent\\_broadcaster\\_address](#) (in module `gwcelery.conf.production`), 24

[voevent\\_broadcaster\\_whitelist](#) (in module `gwcelery.conf`), 22  
[voevent\\_broadcaster\\_whitelist](#) (in module `gwcelery.conf.production`), 24  
[voevent\\_received](#) (in module `gwcelery.voevent.signals`), 62  
[voevent\\_receiver\\_address](#) (in module `gwcelery.conf`), 22

## W

[WARNING](#) (`gwcelery.tools.nagios.NagiosPluginStatus` attribute), 59