
GWCelery Documentation

Release 0.1.1

Leo Singer

Nov 16, 2018

Contents:

1	Quick start	3
1.1	To install	3
1.2	To test	3
1.3	To start	3
2	Configuration	5
3	Monitoring and Management	7
3.1	Flower	7
3.2	Command-Line Tools	9
3.3	Nagios	9
4	Running under HTCondor	11
4.1	Shortcuts	11
5	Contributing	13
5.1	Development model	13
5.2	Where new code should go	13
5.3	Guidelines for tasks	13
5.4	Unit tests	14
5.5	Code style	14
5.6	Documentation	14
6	Design and anatomy of GWCelery	15
6.1	Processes	15
6.2	Eternal tasks	16
6.3	Handlers	16
7	API Reference	19
7.1	gwcclery.conf module	20
7.2	gwcclery.sentry module	20
7.3	gwcclery.tasks module	20
7.4	gwcclery.util module	27
8	Changelog	29
8.1	0.1.1 (2018-10-04)	29
8.2	0.1.0 (2018-09-26)	29

8.3	0.0.31 (2018-09-04)	30
8.4	0.0.30 (2018-08-02)	30
8.5	0.0.29 (2018-07-31)	30
8.6	0.0.28 (2018-07-25)	31
8.7	0.0.27 (2018-07-22)	31
8.8	0.0.26 (2018-07-20)	31
8.9	0.0.25 (2018-07-19)	31
8.10	0.0.24 (2018-07-18)	32
8.11	0.0.23 (2018-07-18)	32
8.12	0.0.22 (2018-07-11)	32
8.13	0.0.21 (2018-07-10)	32
8.14	0.0.20 (2018-07-09)	32
8.15	0.0.19 (2018-07-09)	33
8.16	0.0.18 (2018-07-06)	33
8.17	0.0.17 (2018-07-05)	33
8.18	0.0.16 (2018-07-05)	33
8.19	0.0.15 (2018-06-29)	33
8.20	0.0.14 (2018-06-28)	33
8.21	0.0.13 (2018-06-28)	34
8.22	0.0.12 (2018-06-28)	34
8.23	0.0.11 (2018-06-27)	34
8.24	0.0.10 (2018-06-13)	35
8.25	0.0.9 (2018-06-06)	35
8.26	0.0.8 (2018-06-06)	35
8.27	0.0.7 (2018-05-31)	35
8.28	0.0.6 (2018-05-26)	36
8.29	0.0.5 (2018-05-08)	36
8.30	0.0.4 (2018-04-28)	36
8.31	0.0.3 (2018-04-27)	37
8.32	0.0.2 (2018-04-27)	37
8.33	0.0.1 (2018-04-27)	37
9	License	39
9.1	GNU GENERAL PUBLIC LICENSE	39
9.2	Preamble	39
9.3	TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION	40
9.4	END OF TERMS AND CONDITIONS	43
9.5	How to Apply These Terms to Your New Programs	43
10	Indices and tables	45

GWCelery is a simple and reliable package for annotating and orchestrating LIGO/Virgo alerts, built from widely used open source components.

1.1 To install

GWCelery requires Python ≥ 3.6 .

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

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

- **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 (gwcelery/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

NOTE that GWCelery requires [redis](#). Your package manager (`apt`, `yum`, `macports`) should be able to install, configure, and automatically launch a suitable redis server, but otherwise you can use the [Redis Quick Start](#) instructions to build redis 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
```

GWCelery itself consists of four workers:

```
$ gwcclery worker -l info -n gwcclery-worker -Q celery -B
$ gwcclery worker -l info -n gwcclery-openmp-worker -Q openmp -c 1
$ gwcclery worker -l info -n gwcclery-superevent-worker -Q superevent -c 1
    $ gwcclery worker -l info -n gwcclery-extttrig-worker -Q extttrig -c 1
```


CHAPTER 2

Configuration

By default, GWCelery will talk to the playground GraceDb server, `gracedb-playground.ligo.org`. To switch to using the production GraceDb server, `gracedb.ligo.org`, set the following environment variable before starting GWCelery:

```
CELERY_CONFIG_MODULE=gwcelery.conf.production
```

For further customization, see the API documentation for the `gwcelery.conf` module.

Monitoring and Management

Like all Celery applications, GWCelery supports a rich selection of management and monitoring tools. Here is an introduction to a few of them.

3.1 Flower

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

```
$ gwcelery flower
```


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

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

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

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

Some additional firewall configuration may be required.



Worker:

celery@gwcelery-worker

[Pool](#) [Broker](#) [Queues](#) **Tasks** [Limits](#) [Config](#) [System](#)

Refresh

Processed number of completed tasks

gwcelery.tasks.lvalert.lvalert_listen	1
gwcelery.tasks.skymaps.plot_allsky	2
gwcelery.tasks.skymaps.fits_header	2
gwcelery.tasks.gracedb.upload	8
gwcelery.tasks.gracedb.download	4
gwcelery.tasks.dispatch.dispatch	34
gwcelery.tasks.skymaps.is_3d_fits_file	2

Active currently executing tasks

Name	UUID	Ack	PID	args	kwargs
gwcelery.tasks.lvalert.lvalert_listen	b786eb15-0fc2-400c-8772-d6c79d566a8d	True	2725763	()	{}

Scheduled scheduled (eta/countdown/retry) tasks

Name	UUID	args	kwargs
------	------	------	--------

Reserved tasks that have been received, but are still waiting to be

3.2 Command-Line Tools

All Celery application provide [command-line monitoring and management utilities](#), including the following:

- `gwcclery 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:

```
$ gwcclery 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()
```

- `gwcclery 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:

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

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

```
$ gwcclery 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')
```

3.3 Nagios

This tool is specific to GWCelery.

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

GWCelery provides the command `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:

```
$ gwcclery 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](#)

CHAPTER 4

Running under HTCondor

The recommended way to start and stop GWCelery on the LIGO Data Grid cluster is using [HTCondor](#). See the example HTCondor submit file [gwcclery.sub](#). This submit file will start up Redis, the worker processes, 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 ~/gwcclery/var && cd ~/gwcclery/var
```

Then run the submit file as follows:

```
$ gwcclery.sub
Submitting job(s).....
6 job(s) submitted to cluster 293497.
```

To stop GWCelery, run the `condor_hold` command:

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

To restart GWCelery, run `condor_release`:

```
$ condor_release -constraint 'JobBatchName == "gwcclery"'
All jobs matching constraint (JobBatchName == "gwcclery") 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
```


Contributors may familiarize themselves with Celery itself by going through the [First Steps with Celery](#) tutorial.

5.1 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](#).

5.2 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.

5.3 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 `install_requires` section in `setup.cfg` so that they are installed automatically when GW Celery is installed with `pip`.

5.4 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`.

5.5 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 .
```

5.6 Documentation

Documentation strings should be written in the [Numpydoc style](#).

Design and anatomy of GWCelery

6.1 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:

- `gwcclery.tasks.bayestar.localize()`
- `gwcclery.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:

- `gwcclery.tasks.superevents.handle()`

6. External Trigger Worker

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

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

7. General-Purpose Worker

A Celery worker that accepts all other tasks.

6.2 Eternal tasks

GWCelery has a couple 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_etalernal.EternalTask` or `celery_etalernal.EternalProcessTask`.

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

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

6.3 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 LVAAlert message handlers.

6.3.1 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...
```

6.3.2 LVAAlert messages

LVAAlert 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...
```


CHAPTER 7

API Reference

7.1 gwcelery.conf module

7.1.1 gwcelery.conf.development module

7.1.2 gwcelery.conf.playground module

7.1.3 gwcelery.conf.production module

7.1.4 gwcelery.conf.test module

7.2 gwcelery.sentry module

7.3 gwcelery.tasks module

7.3.1 gwcelery.tasks.bayestar module

7.3.2 gwcelery.tasks.circulars module

7.3.3 gwcelery.tasks.condor module

7.3.4 gwcelery.tasks.detchar module

7.3.5 gwcelery.tasks.em_bright module

7.3.6 gwcelery.tasks.external_triggers module

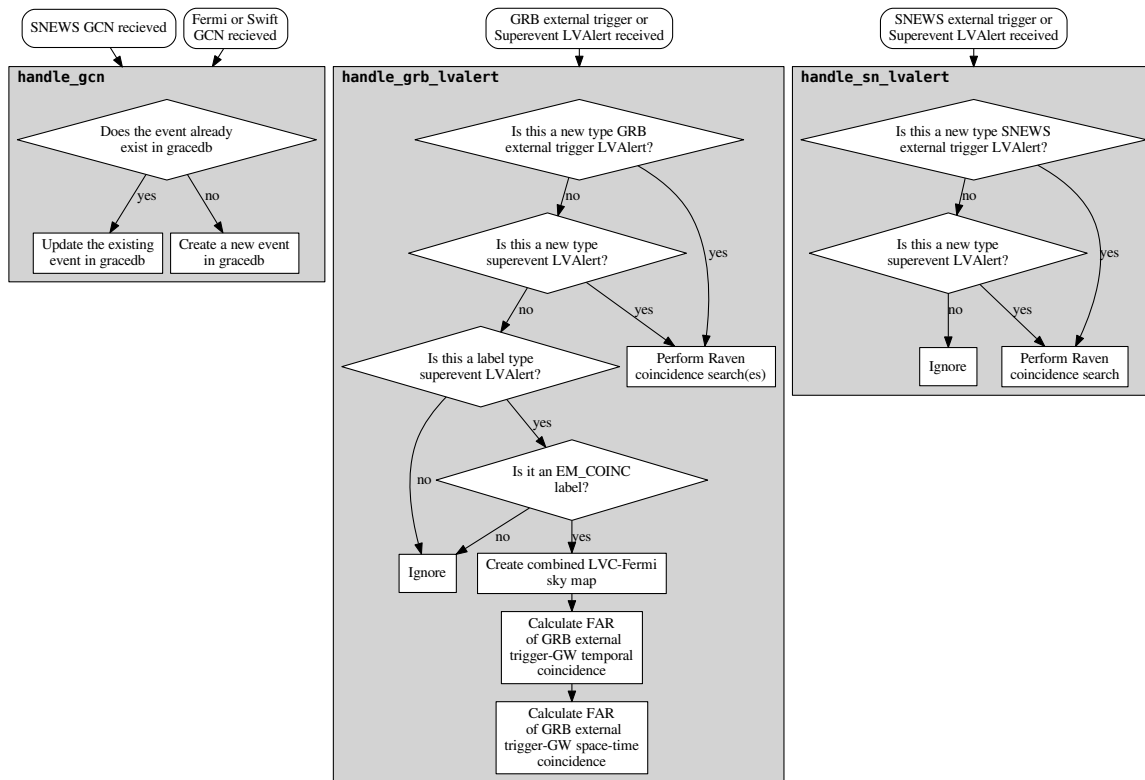
This module listens to the *GCNs* from SNEWS and the Fermi and Swift missions. It is also responsible for carrying out tasks related to external trigger-gravitational wave coincidences, including looking for temporal coincidences, creating

combined GRB-GW sky localization probability maps, and computing their joint temporal and spatio-temporal false alarm rates.

There are two GCN and two 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

7.3.7 gwcelery.tasks.ligo_fermi_skymaps module

7.3.8 gwcelery.tasks.gcn module

7.3.9 gwcelery.tasks.gcn.validate module

7.3.10 gwcelery.tasks.gracedb module

7.3.11 gwcelery.tasks.lvalert module

7.3.12 gwcelery.tasks.orchestrator module

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

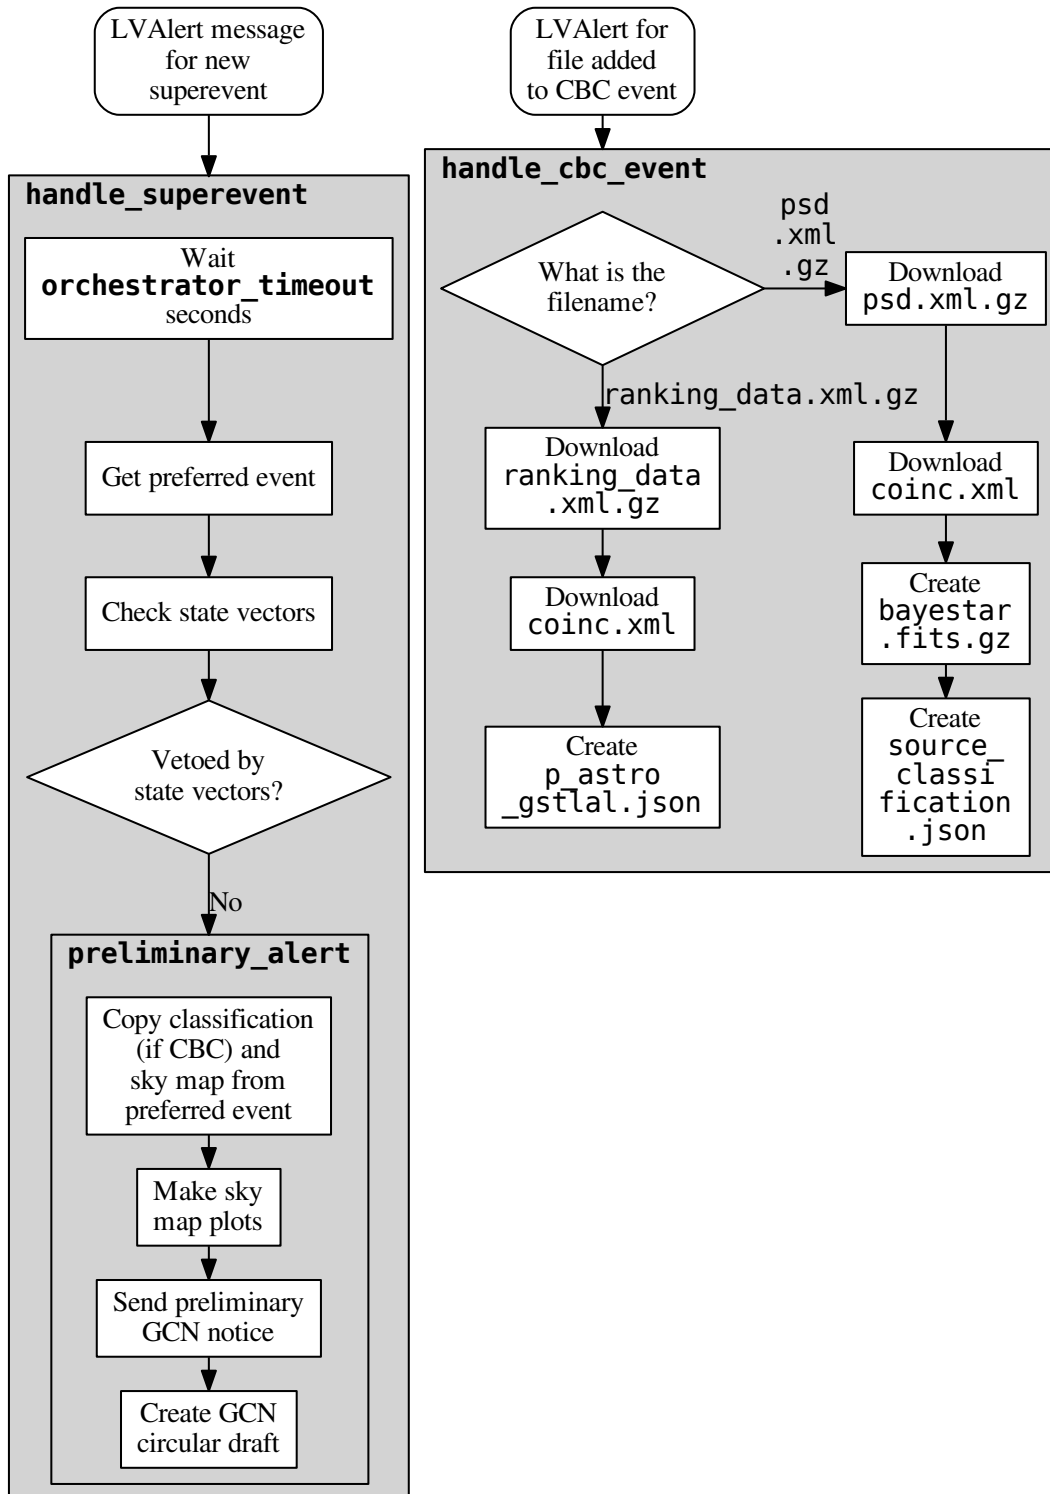
The orchestrator consists of two LVAlert message handlers:

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

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


Preliminary Alerts

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




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

7.3.13 `gwcelery.tasks.p_astro_gstlal` module

7.3.14 `gwcelery.tasks.raven` module


7.3.15 `gwcelery.tasks.skymaps` module

7.3.16 `gwcelery.tasks.superevents` module

Superevents are a new abstraction of gravitational-wave candidates introduced in the third LIGO/Virgo observing (O3). Each superevent is intended to represent a single astrophysical event. A superevent consists of one or more event candidates, possibly from different pipelines, that are neighbors in `gpstime`. One event belonging to the superevent is identified as the preferred event.

Flow Chart

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



Tasks

7.4 gwcelery.util module

8.1 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.

8.2 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`.

8.3 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.

8.4 0.0.30 (2018-08-02)

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

8.5 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.

8.6 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.

8.7 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`.

8.8 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.

8.9 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.

8.10 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.

8.11 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.

8.12 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.

8.13 0.0.21 (2018-07-10)

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

8.14 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.

8.15 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.

8.16 0.0.18 (2018-07-06)

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

8.17 0.0.17 (2018-07-05)

- Undo accidental configuration change in last version.

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

8.19 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.

8.20 0.0.14 (2018-06-28)

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

8.21 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.

8.22 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.

8.23 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 LVAAlert 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` \geq 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()`.

8.24 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.
- LVAAlert 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.

8.25 0.0.9 (2018-06-06)

- Add missing LVAAlert message types to superevent handler.

8.26 0.0.8 (2018-06-06)

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

8.27 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.

8.28 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.

8.29 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.

8.30 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 `gwc celery condor submit` subcommand.

8.31 0.0.3 (2018-04-27)

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

8.32 0.0.2 (2018-04-27)

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

8.33 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.

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

9.1 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.

9.2 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.

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.

9.3 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.

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 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 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

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.

9.4 END OF TERMS AND CONDITIONS

9.5 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 an idea of what it does.
Copyright (C) yyyy  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'.  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 10

Indices and tables

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

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