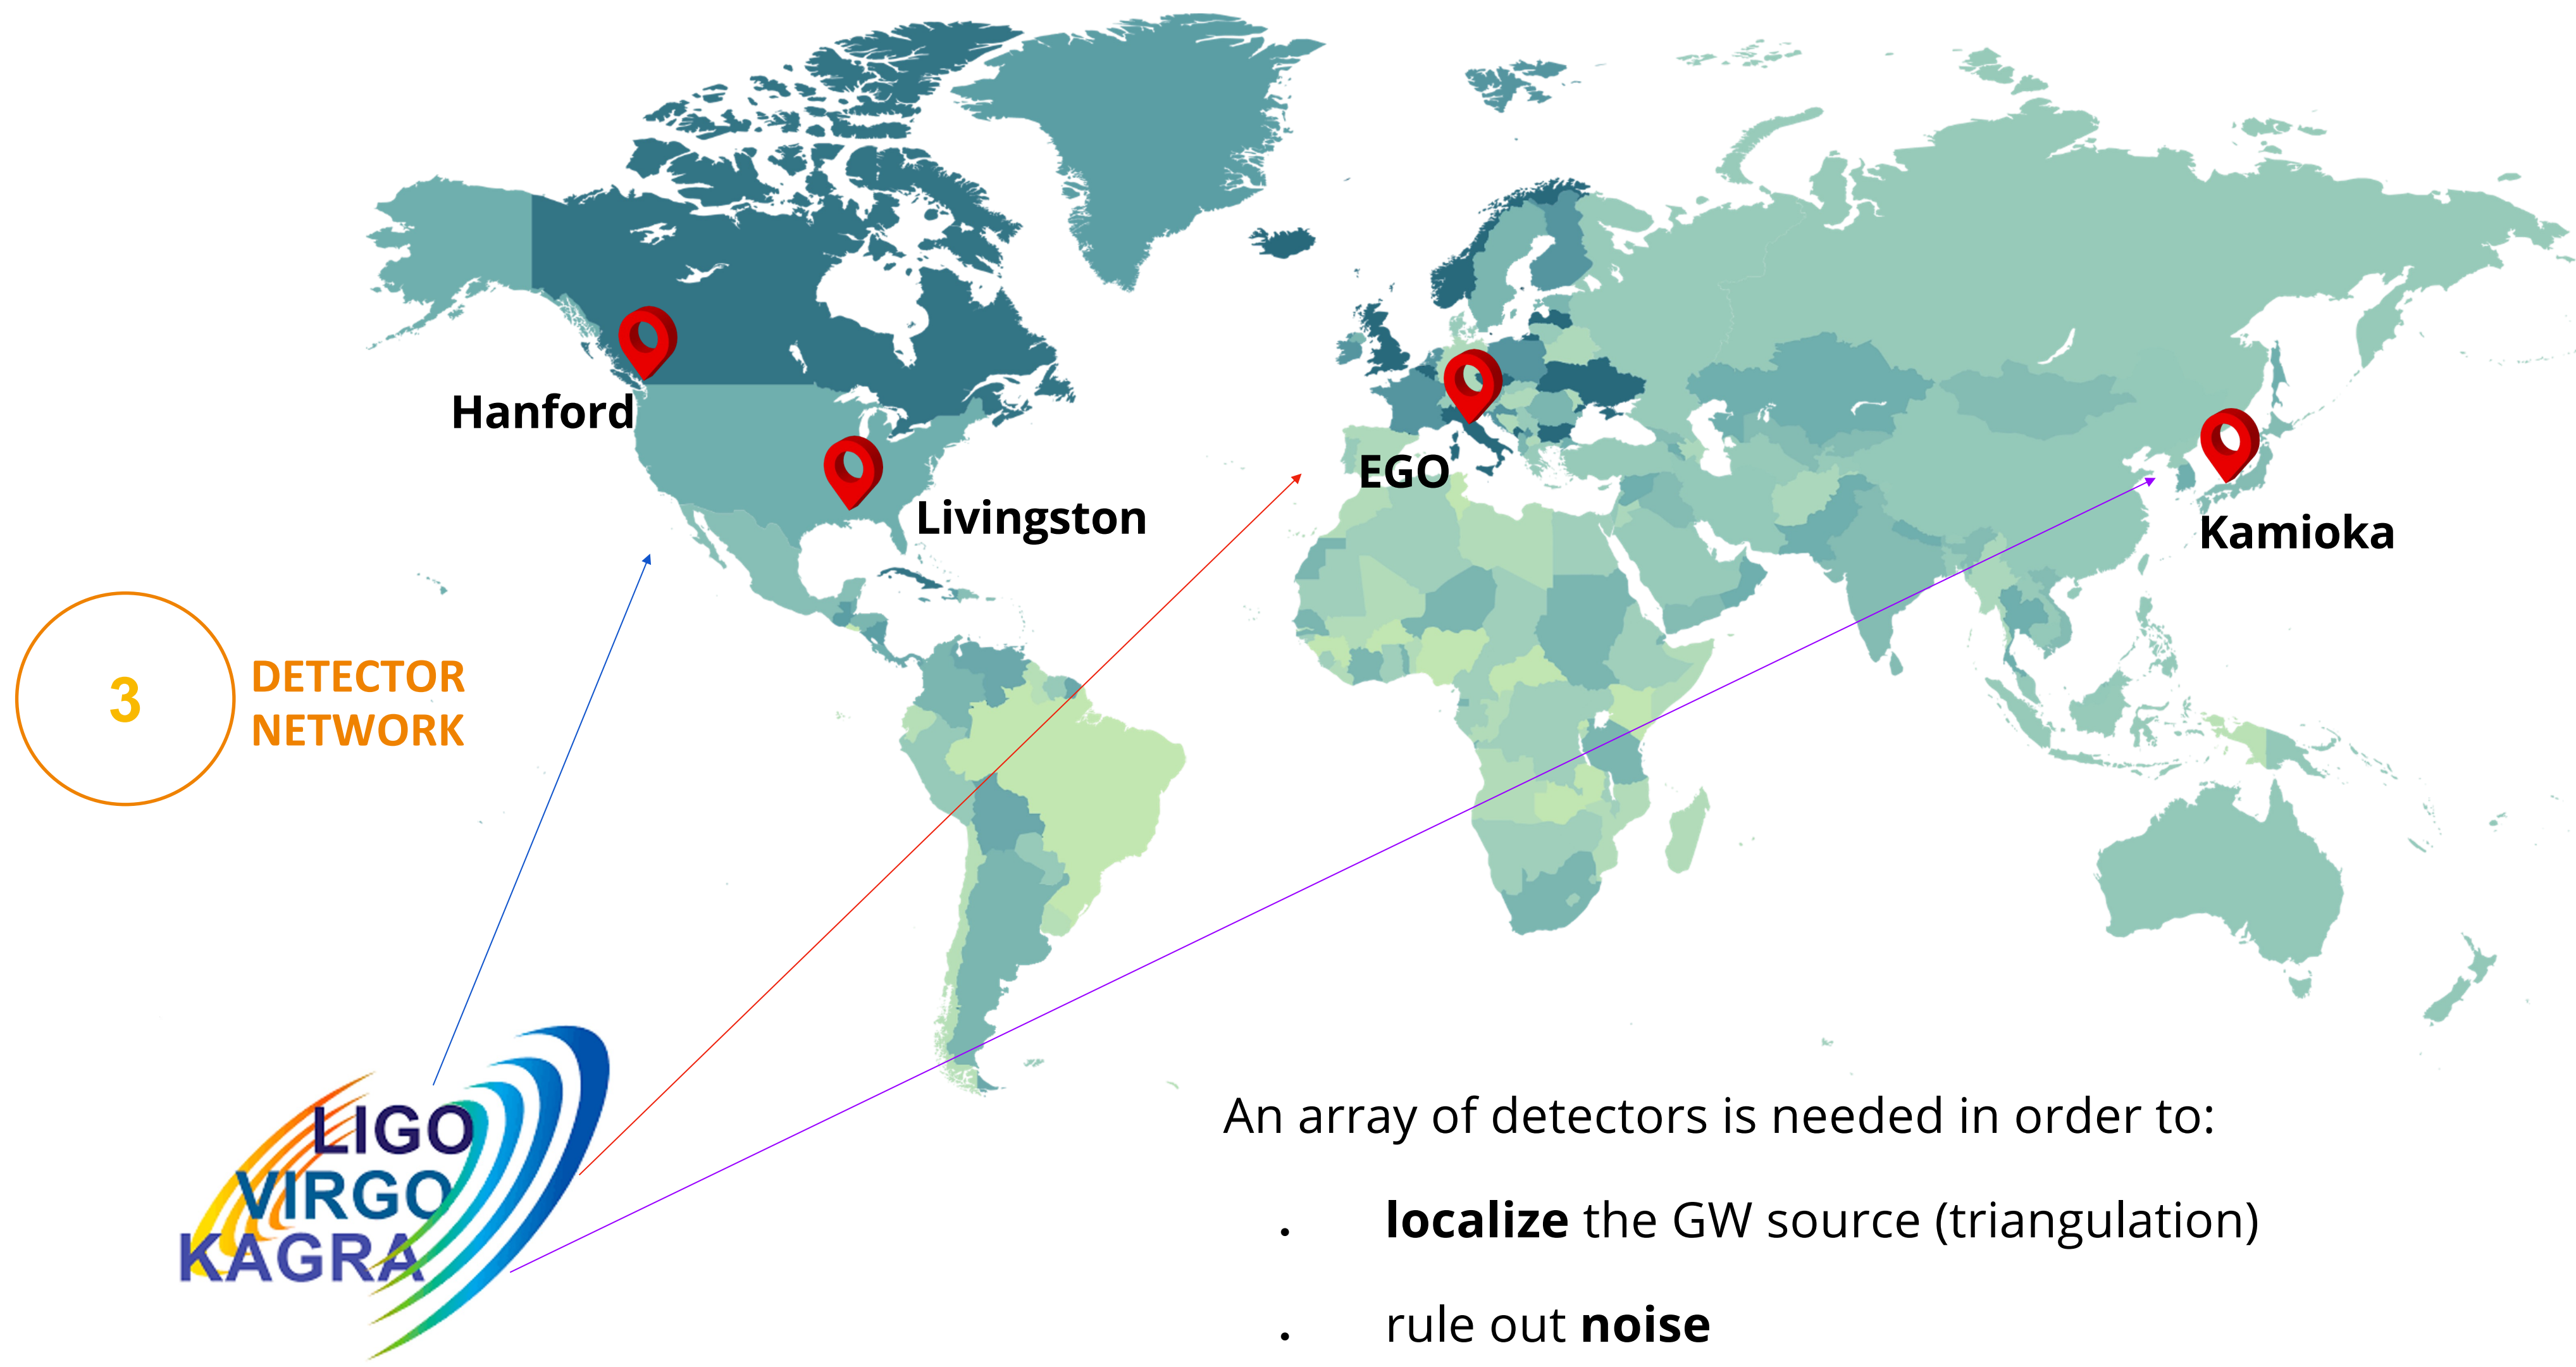


The Ligo/Virgo/Kagra Low Latency Alert Generation Infrastructure

Einstein Telescope E-Infrastructure Board Workshop
Aachen, March 9-10 2023




S. Vallero - INFN Torino



An array of detectors is needed in order to:

- **localize** the GW source (triangulation)
- rule out **noise**

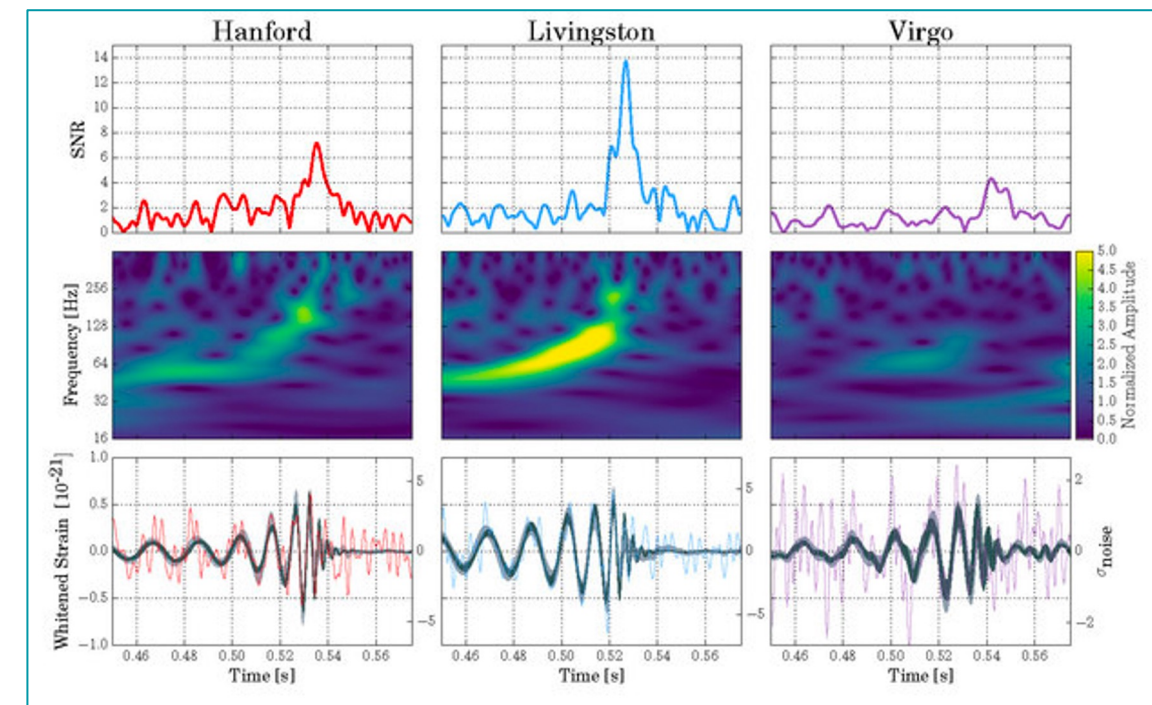




LOW LATENCY SEARCH PIPELINES

4 LOW LATENCY SEARCHES

GW CANDIDATE

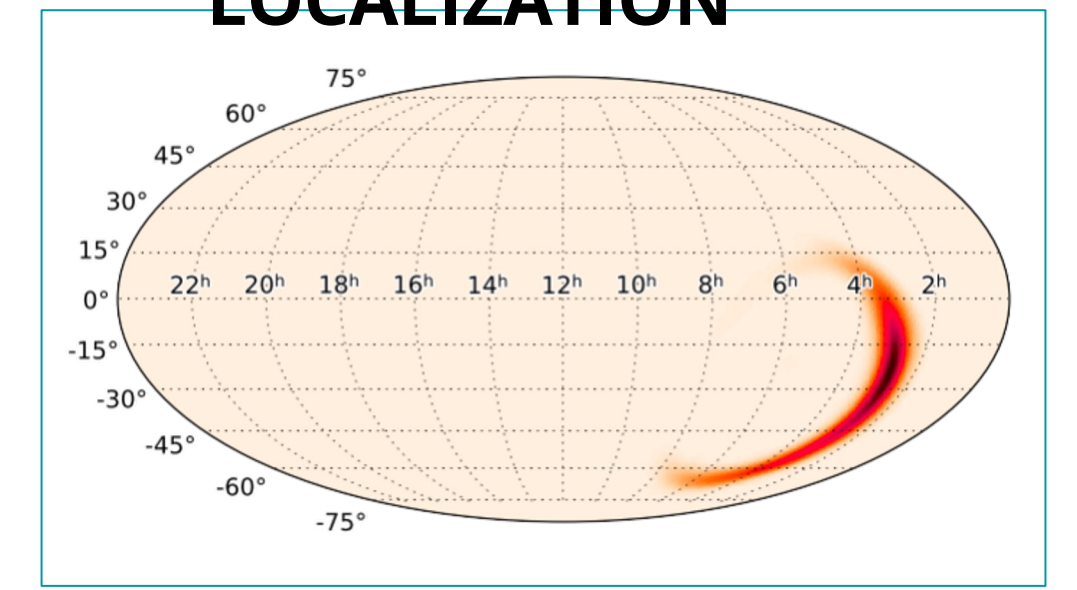


~ 1 minute



- Detector sanity
- Data quality
- Localization
- ...

SKY LOCALIZATION



→ Validation



PUBLIC ALERT

Send the first alert notice within **10 minutes** for EM follow-up (telescopes):

- improved source localization
- multi-messenger astronomy

```


TITLE: GCN CIRCULAR
NUMBER: 21505
SUBJECT: Fermi GBM trigger 524666471/170817529: LIGO/Virgo Identification of a possible
gravitational-wave counterpart
DATE: 17/08/17 13:21:42 GMT
FROM: Reed Clasey Essick at MIT <ressick@mit.edu>

The LIGO Scientific Collaboration and the Virgo Collaboration report:

The online CBC pipeline (gstlal) has made a preliminary
identification of a GW candidate associated with the time of Fermi GBM trigger
524666471/170817529 at gps time 1187008884.47 (Thu Aug 17 12:41:06 GMT 2017) with
RA=186.62deg Dec=-48.84deg and an error radius of 17.45deg.

The candidate is consistent with a neutron star binary coalescence with
False Alarm Rate of ~1/10,000 years.

An offline analysis is ongoing. Any significant updates will be provided
by a new Circular.
    
```



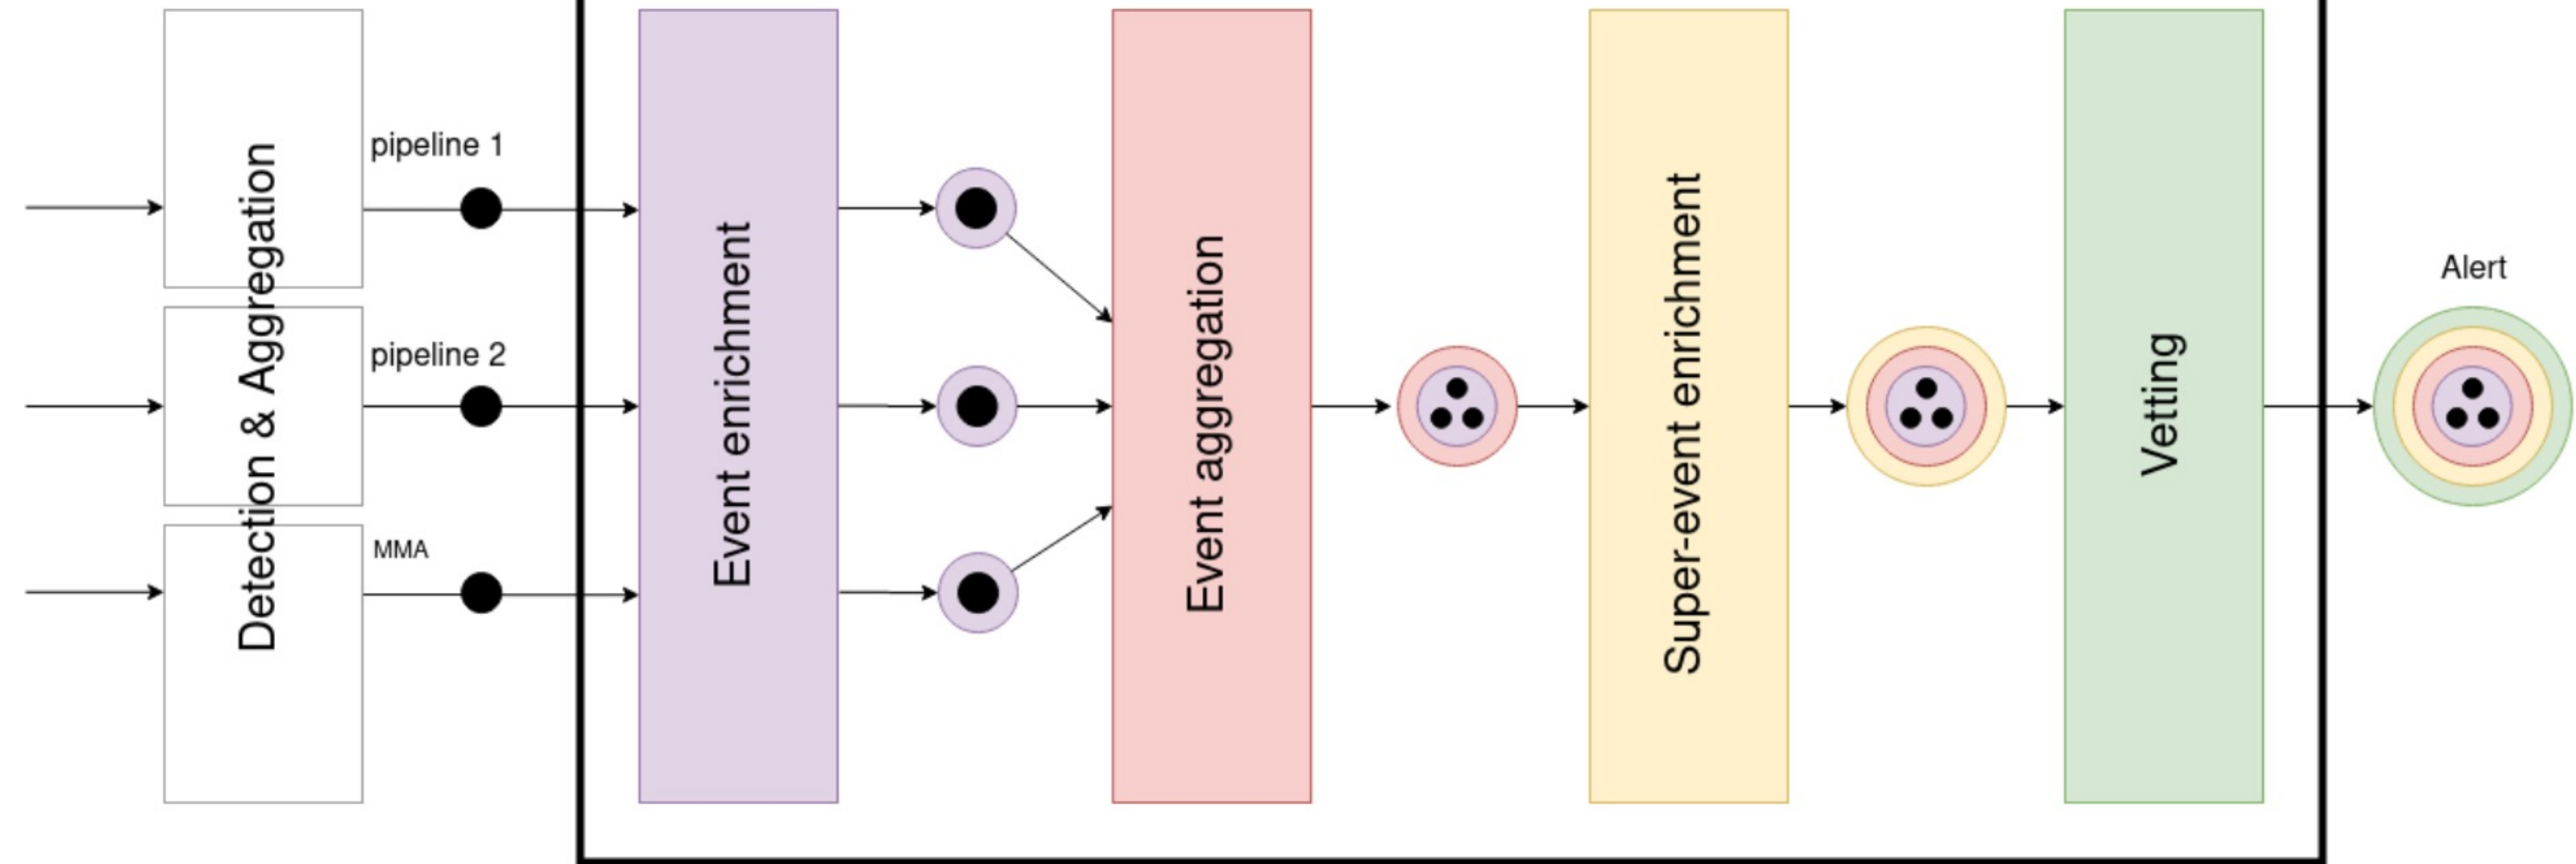
Ripples in spacetime: how we detect them and why

The purpose

1. Disseminate public alerts of transient GW (and MMA involving GW) detections.
2. Enable the discovery of EM and neutrino counterparts to GWs (and vice versa) and assist the (common) source characterisation.
3. Provide feedback to the instrument teams by facilitating the diagnosing of detector problems via real-time analyses.

Workflow

- GW data
- EM data
- Neutrino data



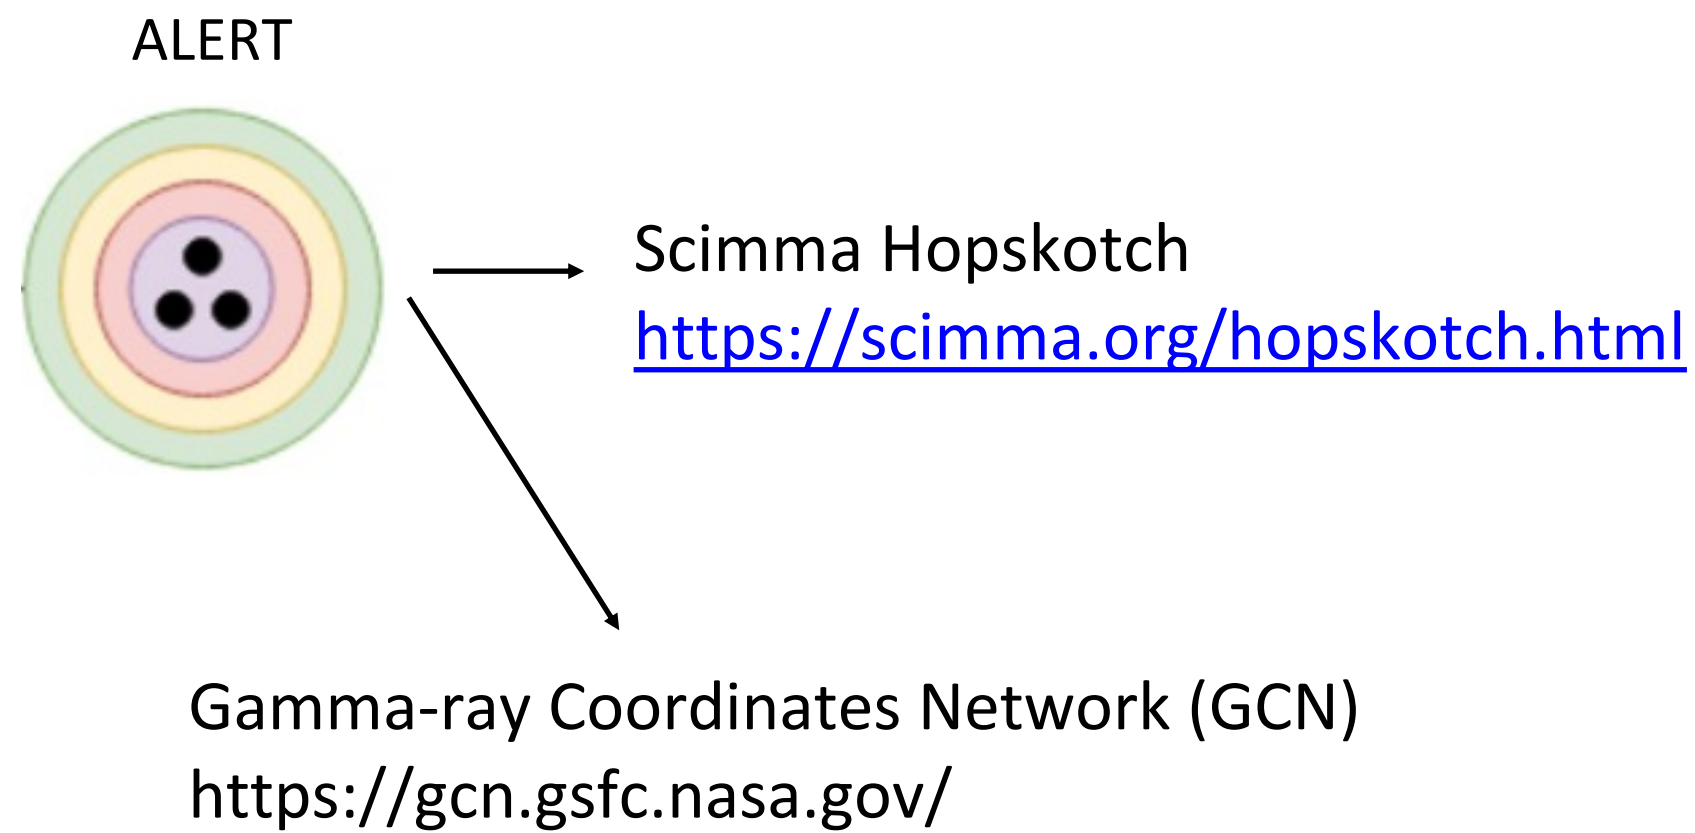
- All-sky searches:
 - no assumption on sky location or time of a transient
 - Compact Binary Coalescence: modelled, matched-filtering
 - Bursts: no assumption on signal morphology, time-frequency analysis
 - no information from non GW sources
 - also low significance alerts for early warning alerts (pre-merger)
- Multi-messenger searches:
 - integrate information from EM or neutrino external triggers
 - can be targeted to a region of the sky or a time identified by the external trigger (medium-latency)

- Event aggregation:
 - grouping of events related to the same astrophysical (or not) cause into a *superevent*
 - combining triggers across multiple pipelines but also from a given pipeline

- Superevent enrichment:
 - source classification
 - sky localisation
 - data quality
 - ...

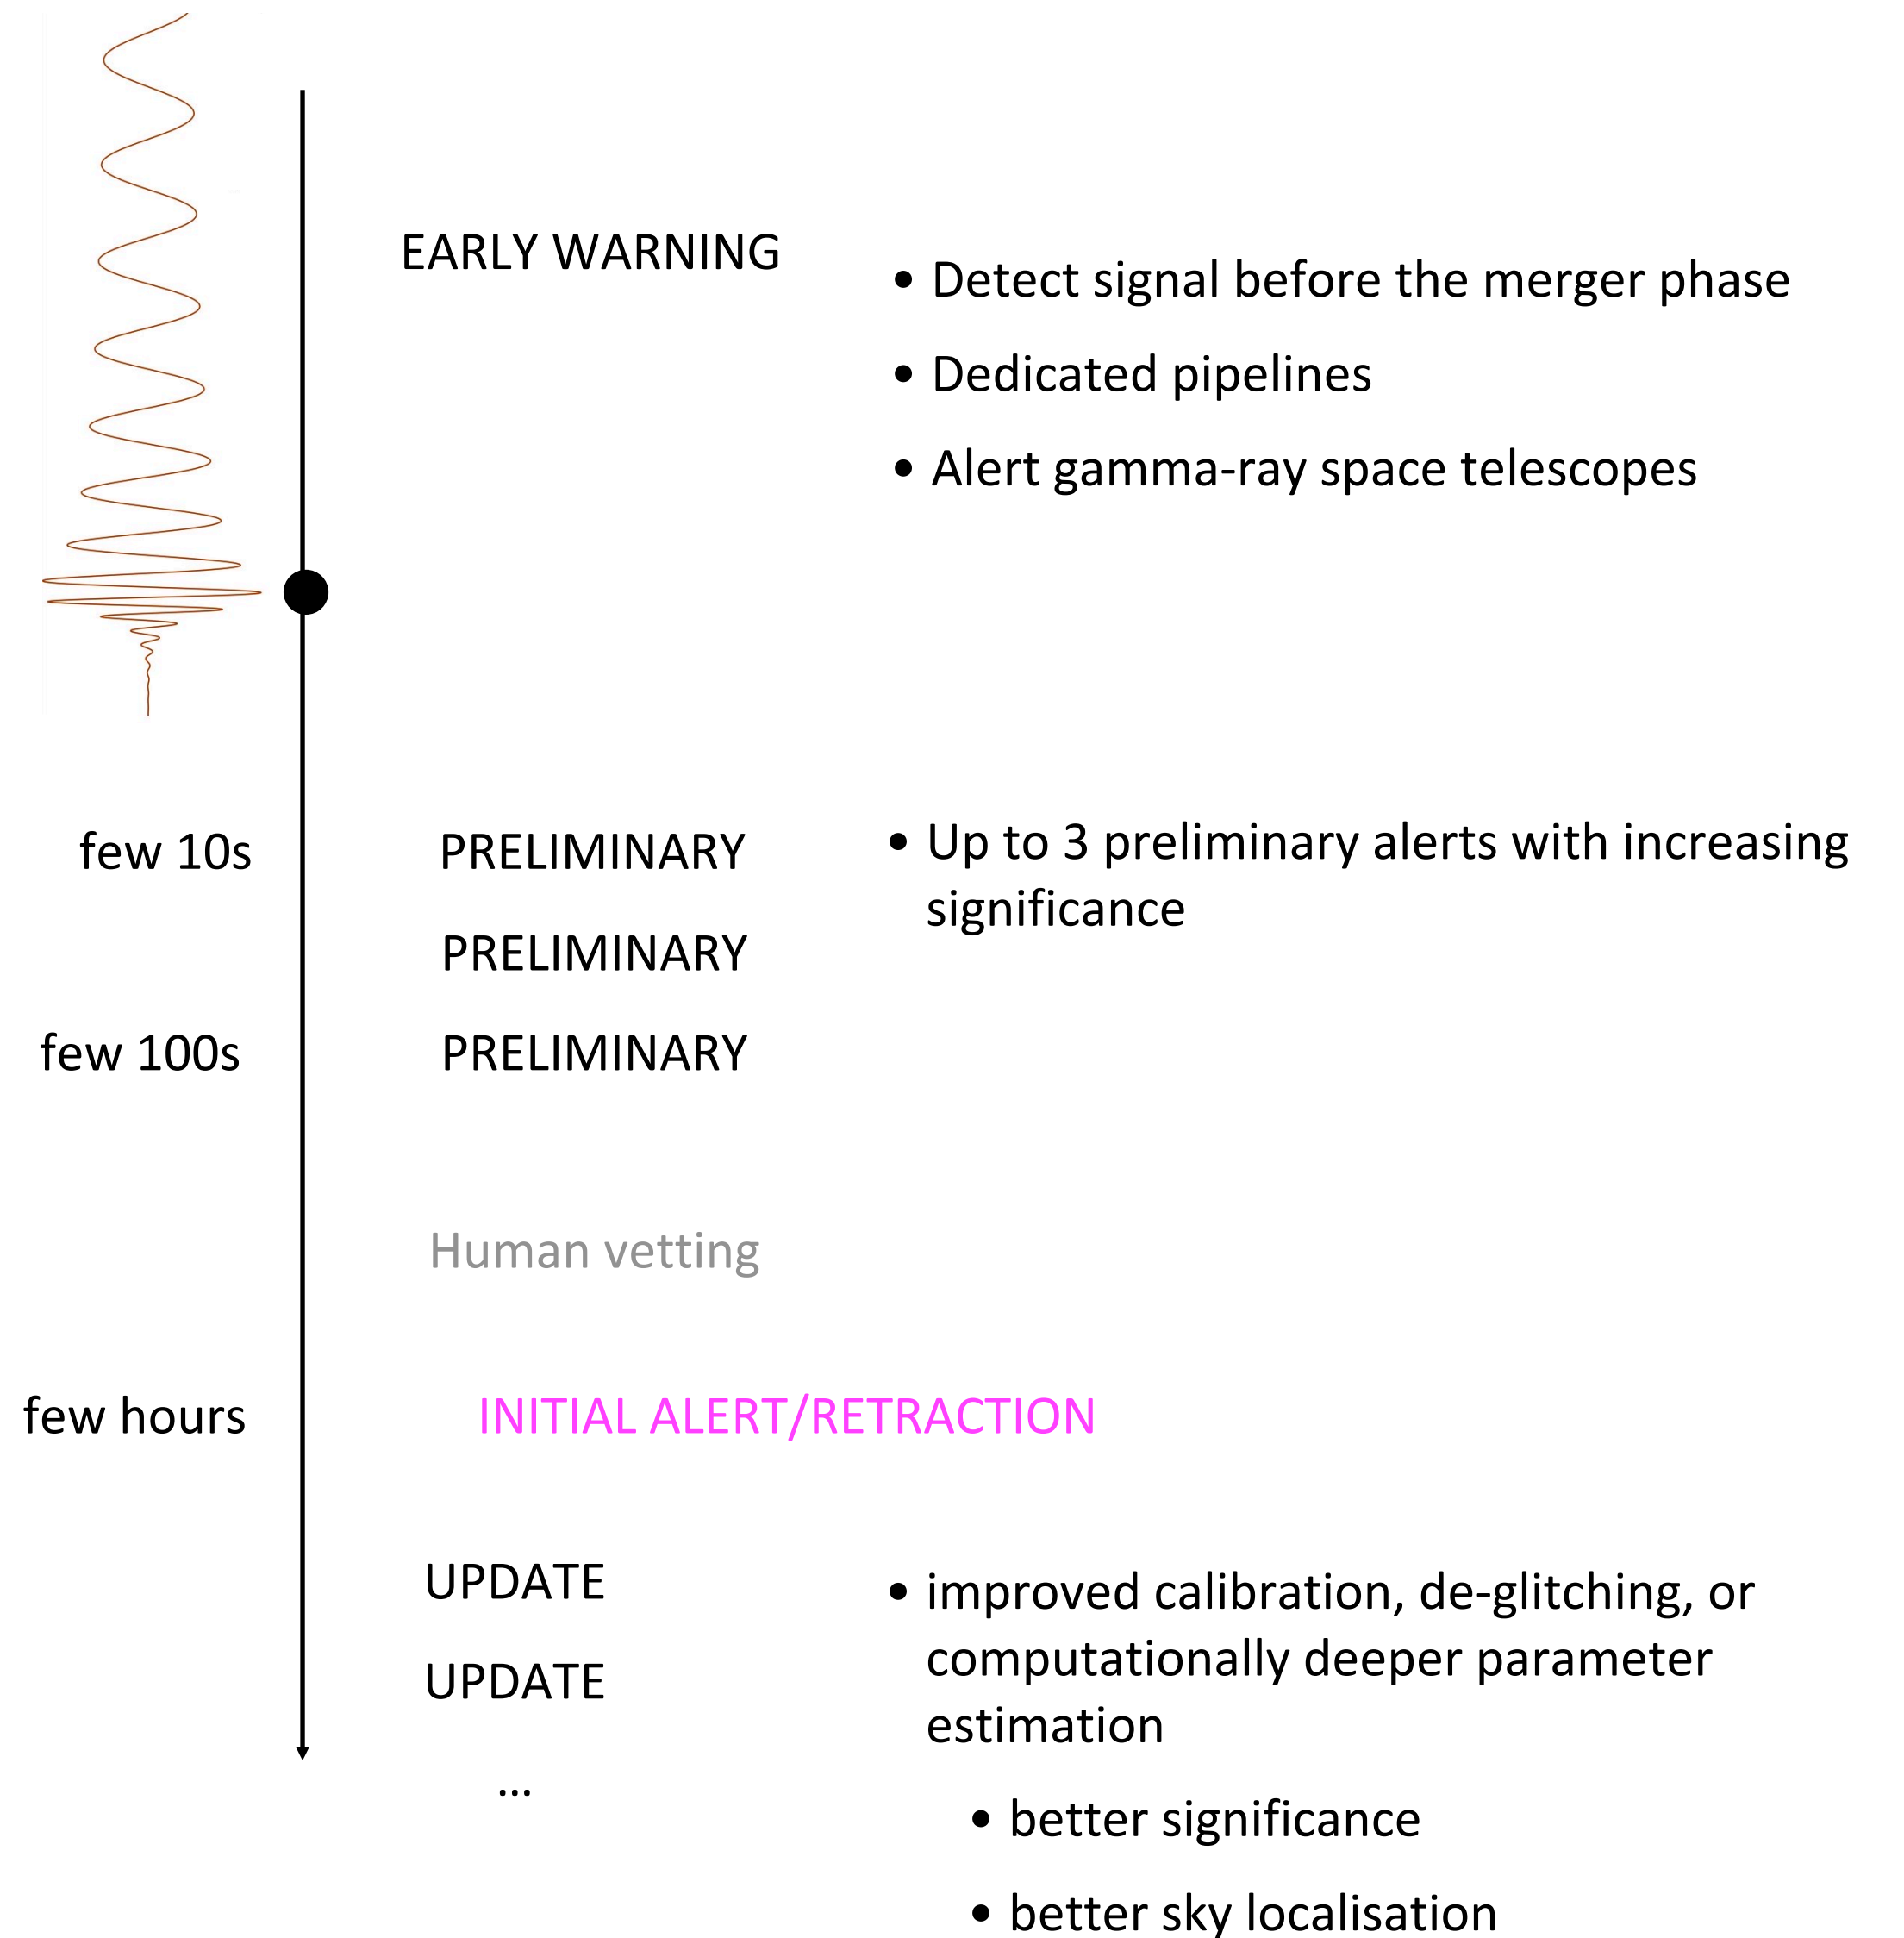
- Vetting:
 - human or automatic decision concerning the publication or retraction of the GW detection
 - based on the enriched information available for the *superevent*

Public alerts

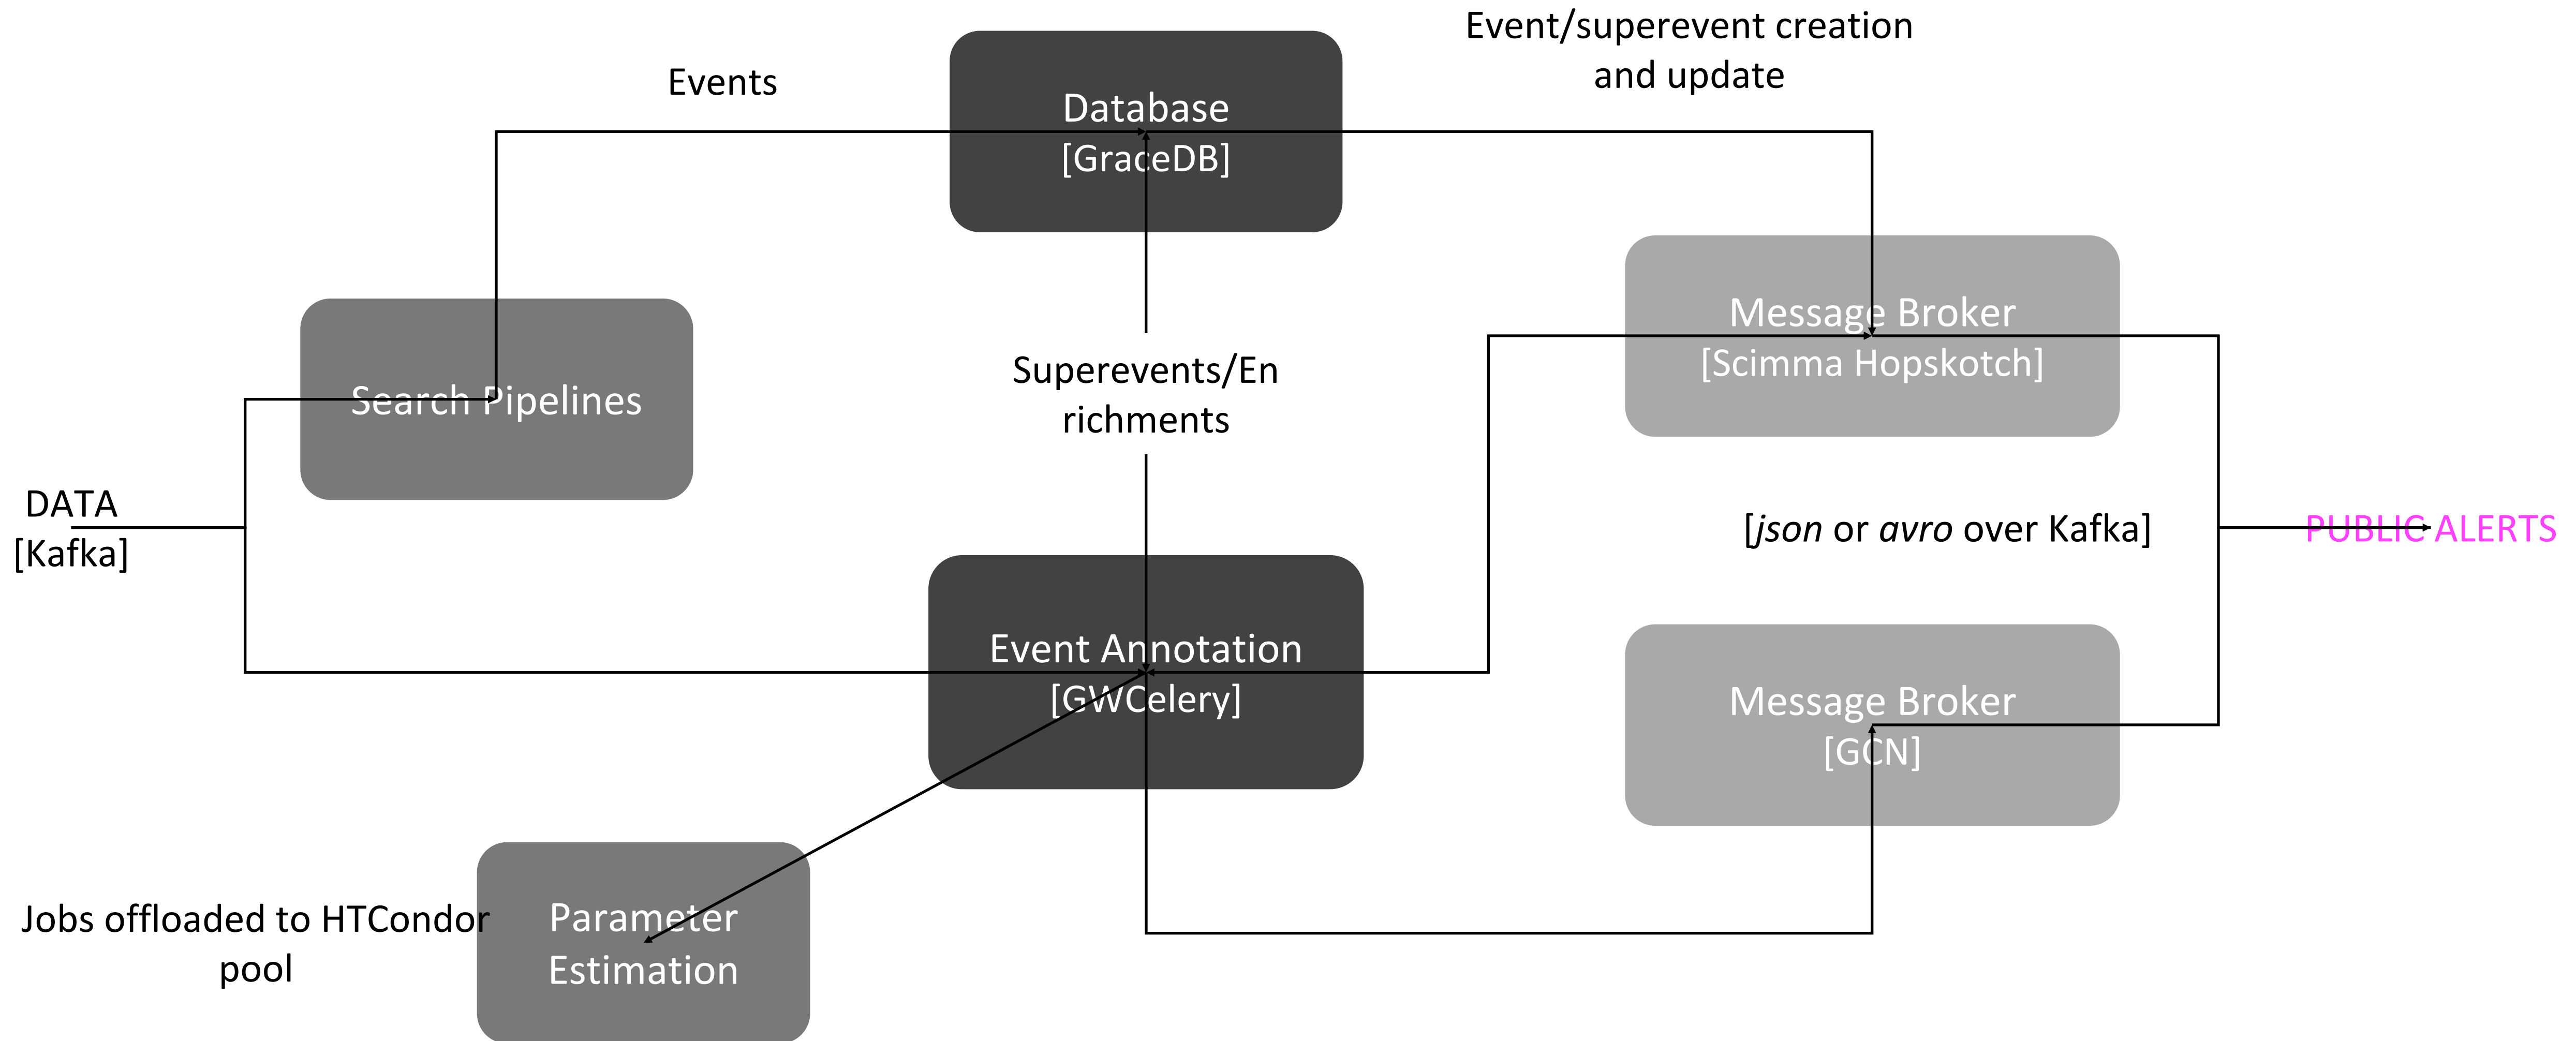


```
{
  "alert_type": "PRELIMINARY",
  "time_created": "2018-11-01T22:34:49Z",
  "superevent_id": "MS181101ab",
  "urls": { "gracedb": "https://example.org/superevents/MS181101ab/view/" },
  "event": {
    "time": "2018-11-01T22:22:46.654Z",
    "far": 9.11069936486e-14,
    "significance": "low significance" # 2/day > FAR > (1/month CBC and 1/year BURST)
    "significant" # (1/month CBC and 1/year BURST) > FAR
    "instruments": [ "H1", "L1", "V1" ],
    "group": "CBC",
    "pipeline": "gstlal",
    "search": "MDC",
    "classification": { "BNS": 0.95, "NSBH": 0.01, "BBH": 0.03, "Terrestrial": 0.01 },
    "properties": { "HasNS": 0.95, "HasRemnant": 0.91, "HasMassGap": 0.01 },
    "skymap": "U01NUEXFICA9ICAgICAgICAgICAgICAgICBUIC8gY29uZm..."
  },
  "external_coinc": null }

```



High-level architecture



GraceDB: the database

- Different deployments (tiers):
 - Production
 - Playground: designed for users to develop and test their own applications. It mimics the production instance, but events and associated data are not preserved indefinitely. Used in Mock Data Challenges
 - Test: designed for Quality Assurance (QA) testing and validation for GraceDB and electromagnetic follow-up developers. Software should meet QA milestones on the test instance before being moved to Playground or Production.
- Web application:
 - provides long-term storage for its inputs, outputs and by-products
 - accessible via GUI or API
 - relational database and local storage for persistence
 - deployed in replica 3 (each tier)

GraceDB (TEST) Public Alerts Latest Search Alerts Pipelines Documentation Admin Docs Logout							
Authenticated as: Sara Vallero							
MS230309I	EM_COINC EM_Selected EM_READY PASTRO_READY EMBRIGHT_READY SKYMAP_READY GCN_PRELIM_SENT DQR_REQUEST ADVOK	9.110699364861297e-14	M30761	M30767 M30766 M30765 M30764 M30763 M30762 M30761 M30760 M30759 M30758 M30757 M30756 M30755 M30754 M30753 M30752 M30751	1362397199.376	2023-03-09 11:44:49 UTC	Sara Vallero
MS230309k	EM_Selected EM_READY EMBRIGHT_READY PASTRO_READY SKYMAP_READY GCN_PRELIM_SENT DQR_REQUEST ADVNO	5.176954877338635e-16	M30736	M30750 M30749 M30748 M30747 M30746 M30745 M30744 M30743 M30742 M30741 M30740 M30739 M30738 M30737 M30736 M30735	1362393223.659	2023-03-09 10:44:48 UTC	Sara Vallero

GraceDB [Public Alerts](#) [Latest](#) [Search](#) [Alerts](#) [Pipelines](#) [Documentation](#) [Logout](#)

Authenticated as: Sara Vallero

S191112q

- [Log Messages](#)
- [Data Quality](#)
- [External Coincidence](#)
- [Full Event Log](#)

Data Quality

Log Image

Omegascans of L1:GDS-CALIB_STRAIN_CLEAN at 1257596016.362291

L1 omegascan
— Submitted by LIGO/Virgo EM Follow-Up on March 8, 2022 19:09:01 UTC

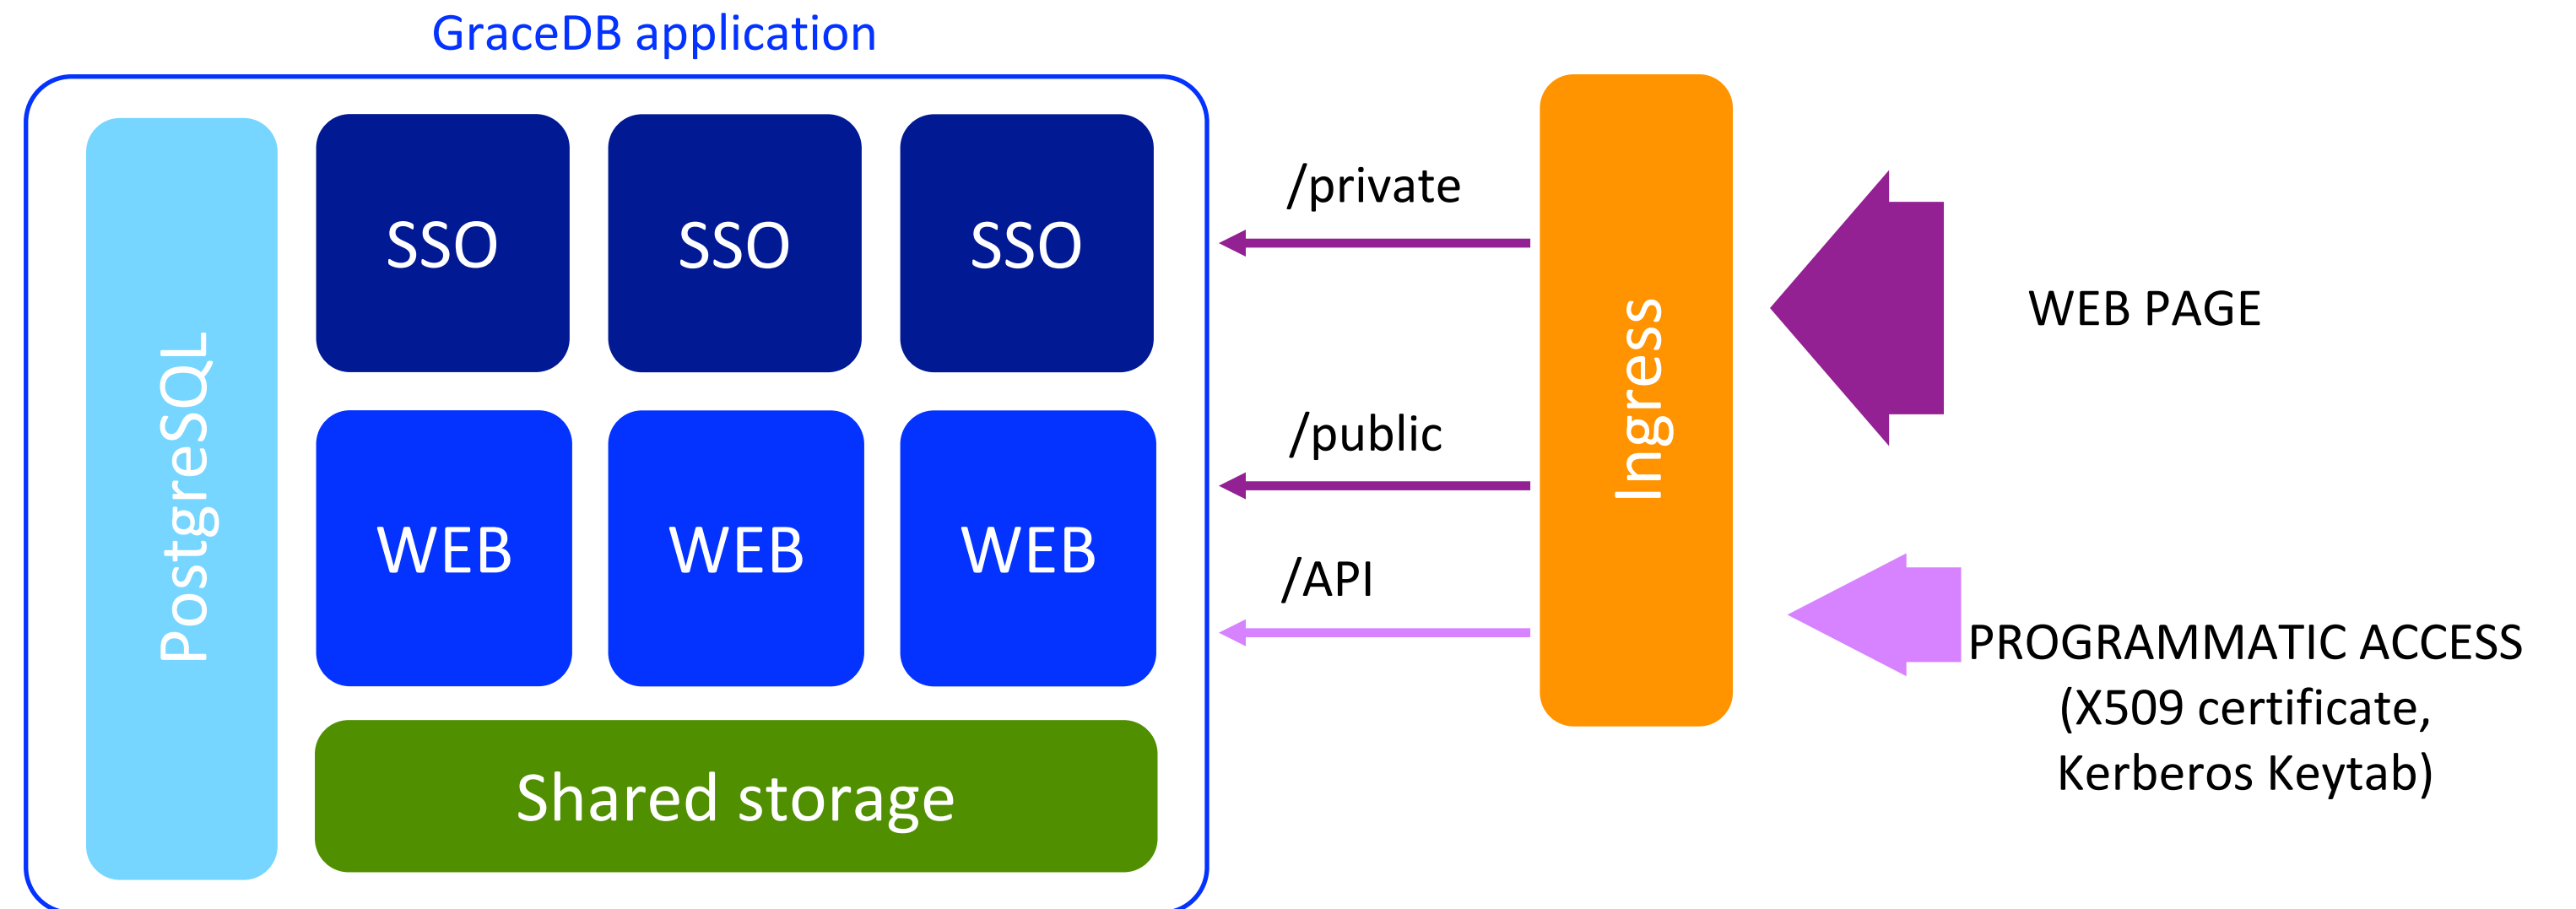
Log Image

Omegascans of H1:GDS-CALIB_STRAIN_CLEAN at 1257596016.362291

GraceDB: the database

Primary Authentication (Shibboleth):

- managing federated identities and providing a single sign-on (SSO) portal
- uses a metadata provider to collect user attributes from an attribute authority and put them into the user's session



Web application:

- Django application: GraceDB is written in Python and is constructed around the Django web framework.
- Backend Webserver (Gunicorn): Gunicorn is a lightweight Python webserver which interfaces directly with the Django service via the WSGI protocol.
- Frontend Webserver (Apache): used in concert with Gunicorn as an interface with Shibboleth. It is configured as a reverse proxy which gets authentication information from Shibboleth, sets that information in the headers, and then passes it on to Gunicorn.
- Igwn-alert Overseer: registers new events in Scimma

Shared storage:

- save event enrichment files (i.e. skymaps)
- relies on Amazon Elastic Block Store

Ingress:

- Traefik
- Redirects to different application components (Linux containers) according to url path

Deployed with
Docker Swarm on
AWS.

GWCelery: the event annotation service

<https://rtd.igwn.org/projects/gwcelery/en/latest/index.html>

- package for annotating and orchestrating LIGO/Virgo alerts
- built on the Celery distributed task queue and widely used opensource components
- different deployments (tiers), connected to corresponding GraceDB instances
- uses Redis to route and distribute Celery task messages and to store task results for later retrieval
- runs on dedicated VMs that also host Redis and offloads computing intensive tasks to a HTCondor cluster

- GWCelery's responsibilities include:
 - merging related candidates from multiple online LIGO/Virgo transient searches into *superevents*
 - correlating LIGO/Virgo events with gamma-ray bursts, neutrinos, and supernovae
 - launching automated follow-up analyses including data quality checks, rapid sky localisation, automated parameter estimation, and source classification
 - generating and sending preliminary and updated public Notices (machine readable)
 - automatically composing public Circulars (human readable)

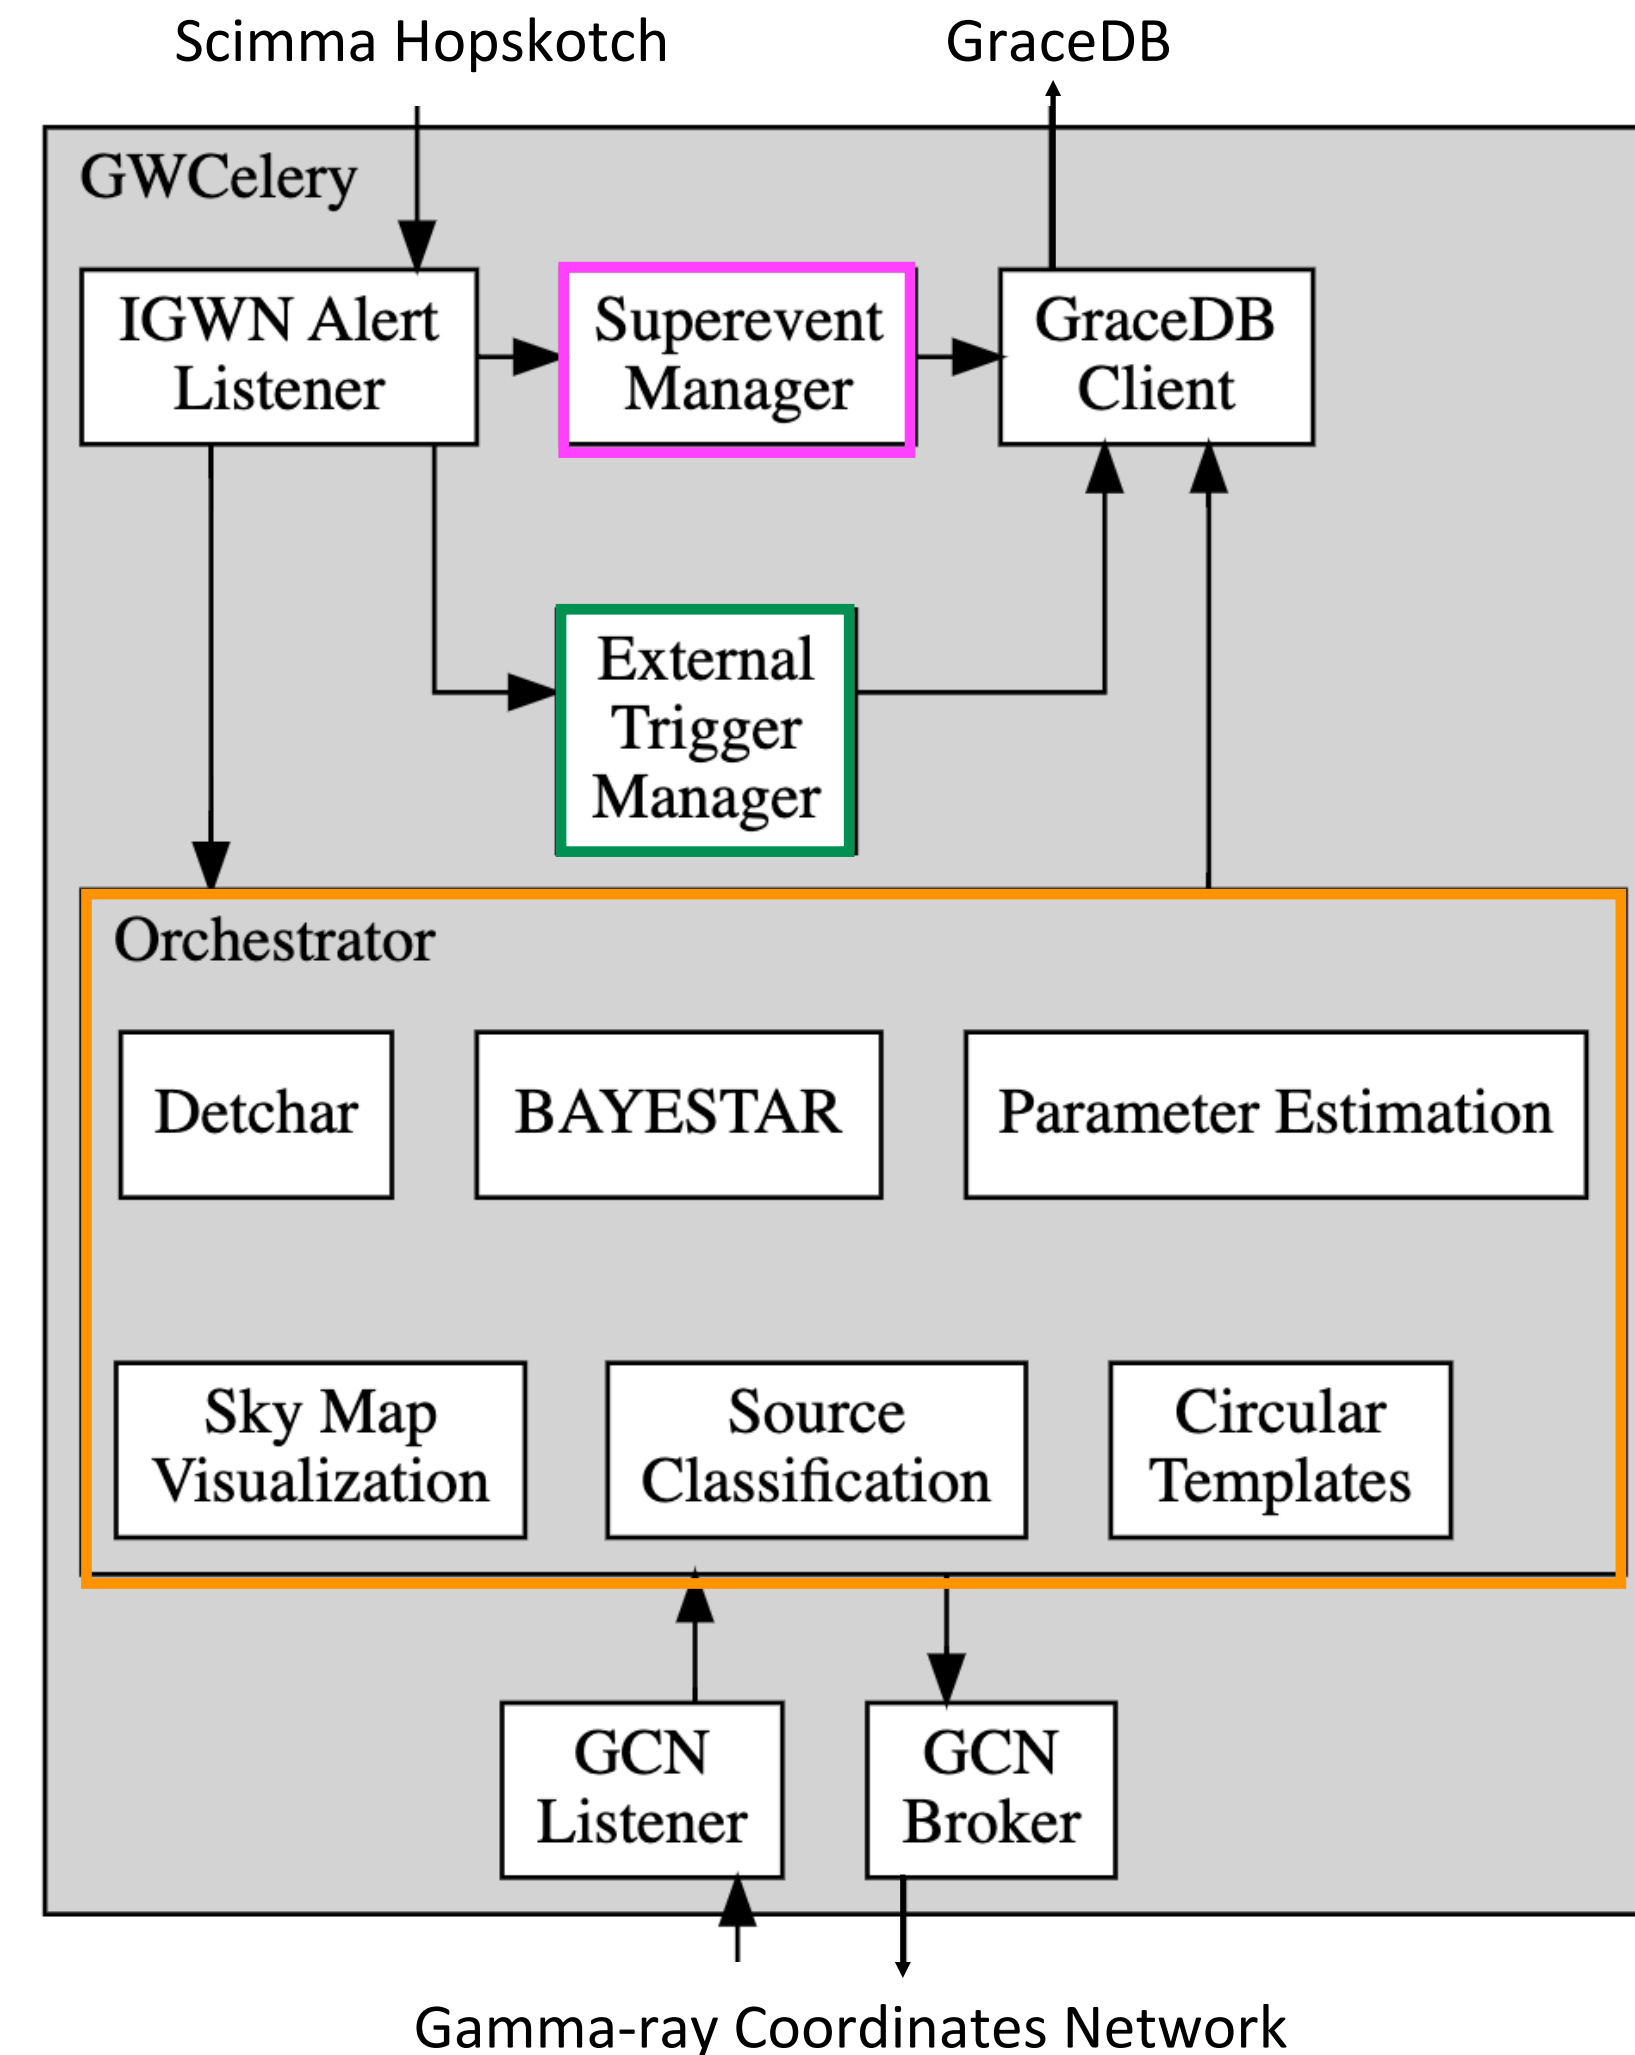
GW Celery: the event annotation service

Superevent Manager: clusters and merges related candidates into *superevents*

External Trigger Manager: correlates gravitational-wave events with GRB, neutrino, and supernova events

Orchestrator: executes the per-event annotation workflow

- these functional components are roughly mapped to Celery workers
- 1 Celery worker has been configured to accept only computationally intensive tasks that use OpenMP parallelism (HTCondor cluster)



Some considerations

- Elements in the critical path for latency:
 - Database
 - Annotation service
 - Data transfers (i.e. need data from other interferometers)
 - Search pipelines
- Latency is never “low enough”: consider early warning (negative latency) alerts
- Uniform running environment across experiments and within online/offline is desirable
- Asynchronous and distributed architecture of services is good