COMPUTING FOR GRAVITATIONAL-WAVE RESEARCH TOWARDS O4

Stefano Bagnasco | INFN Torino

On behalf of the LVK Collaboration

Burst sources:

- CBC: Compact Binary Coalescence
  - Coalescing Compact Binary Systems (Neutron Star-NS, Black Hole-NS, BH-BH): Strong emitters, well modelled for much of the parameter space
- Burst: Unmodeled transient bursts
  - Asymmetric Core Collapse Supernovae: weak emitters, not well-modelled ("bursts"), transient
  - Cosmic strings, soft gamma repeaters, pulsar glitches,...
  - Who knows?

Continuous sources:

- CW: Continuous waves
  - Spinning neutron stars (known waveform, long/continuous duration)
  - All-sky and targeted searches
- SGWB: Continuous stochastic background
  - Cosmological stochastic background (residue of the Big Bang, cosmic GW background, long duration)
  - Astrophysical stochastic background

GW SIGNALS

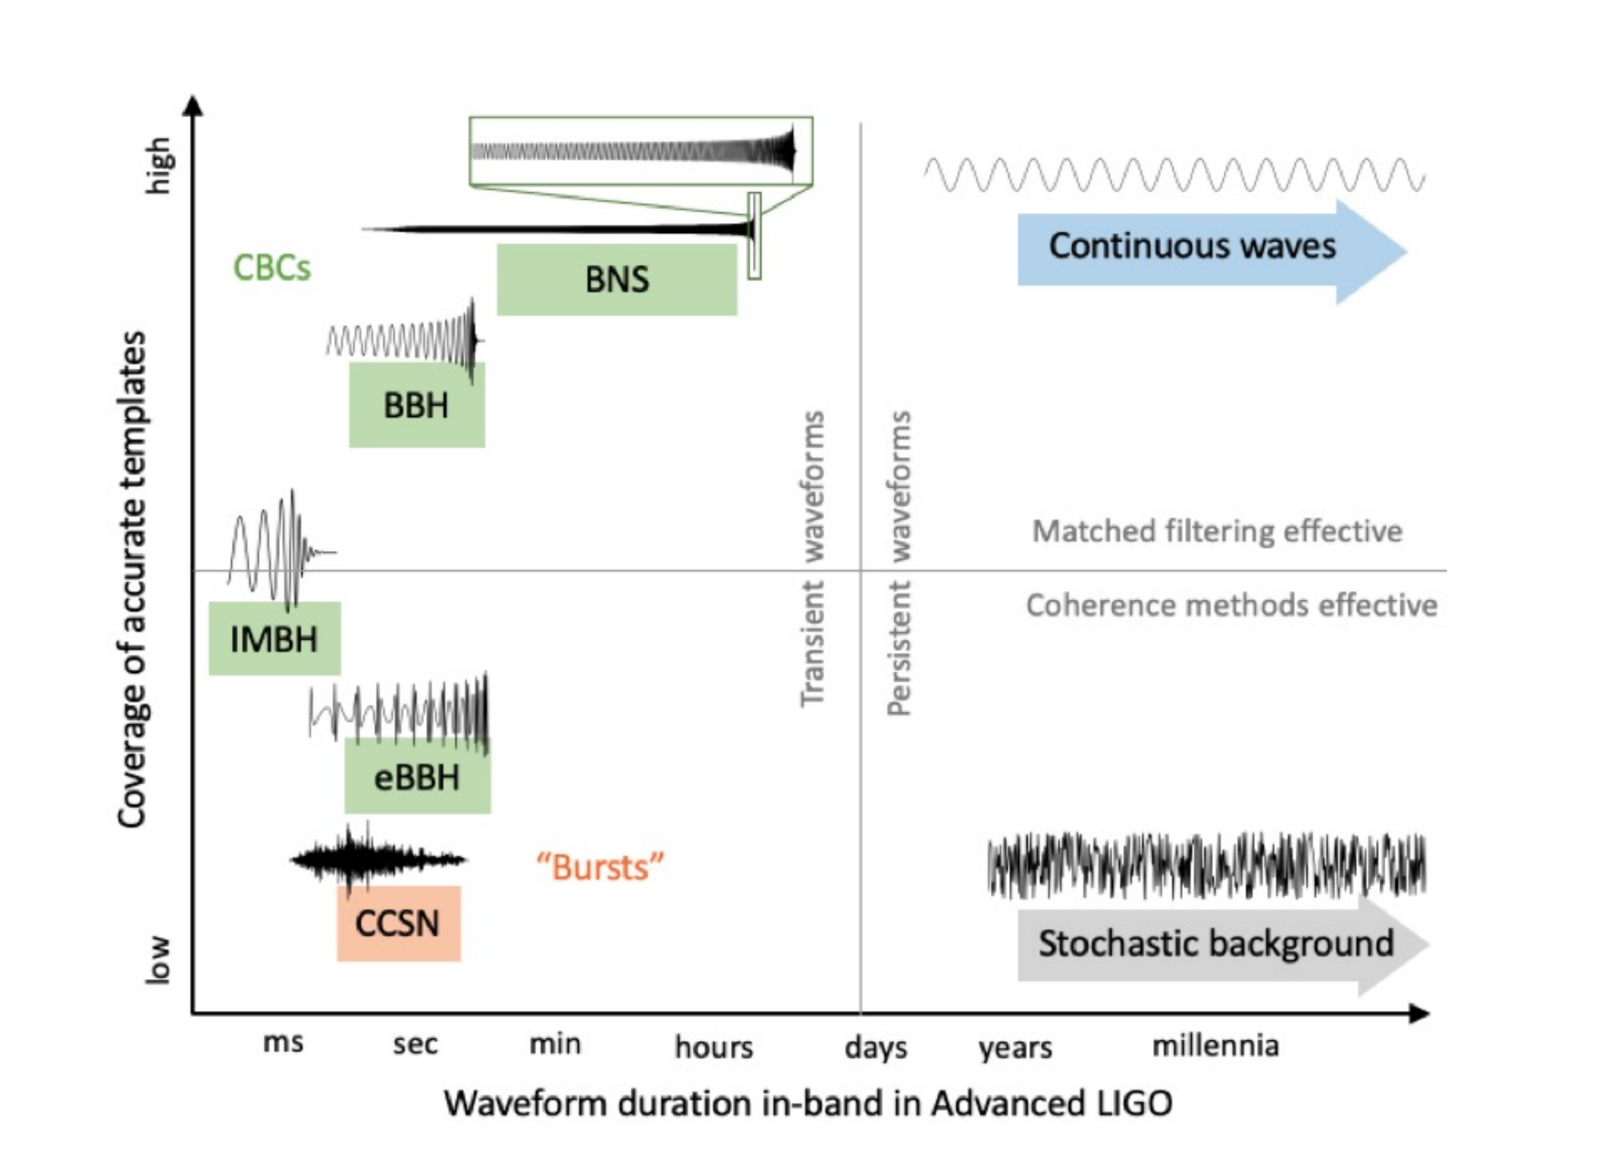

Jess McIver, D. H. Shoemaker, *Discovering Gravitational Waves with Advanced Ligo*, LIGO Document P200530-v1 (2000)

European Gravitational Observatory, Cascina (Italy)

Caltech MIT/Ligo Laboratory

A WORLDWIDE NETWORK

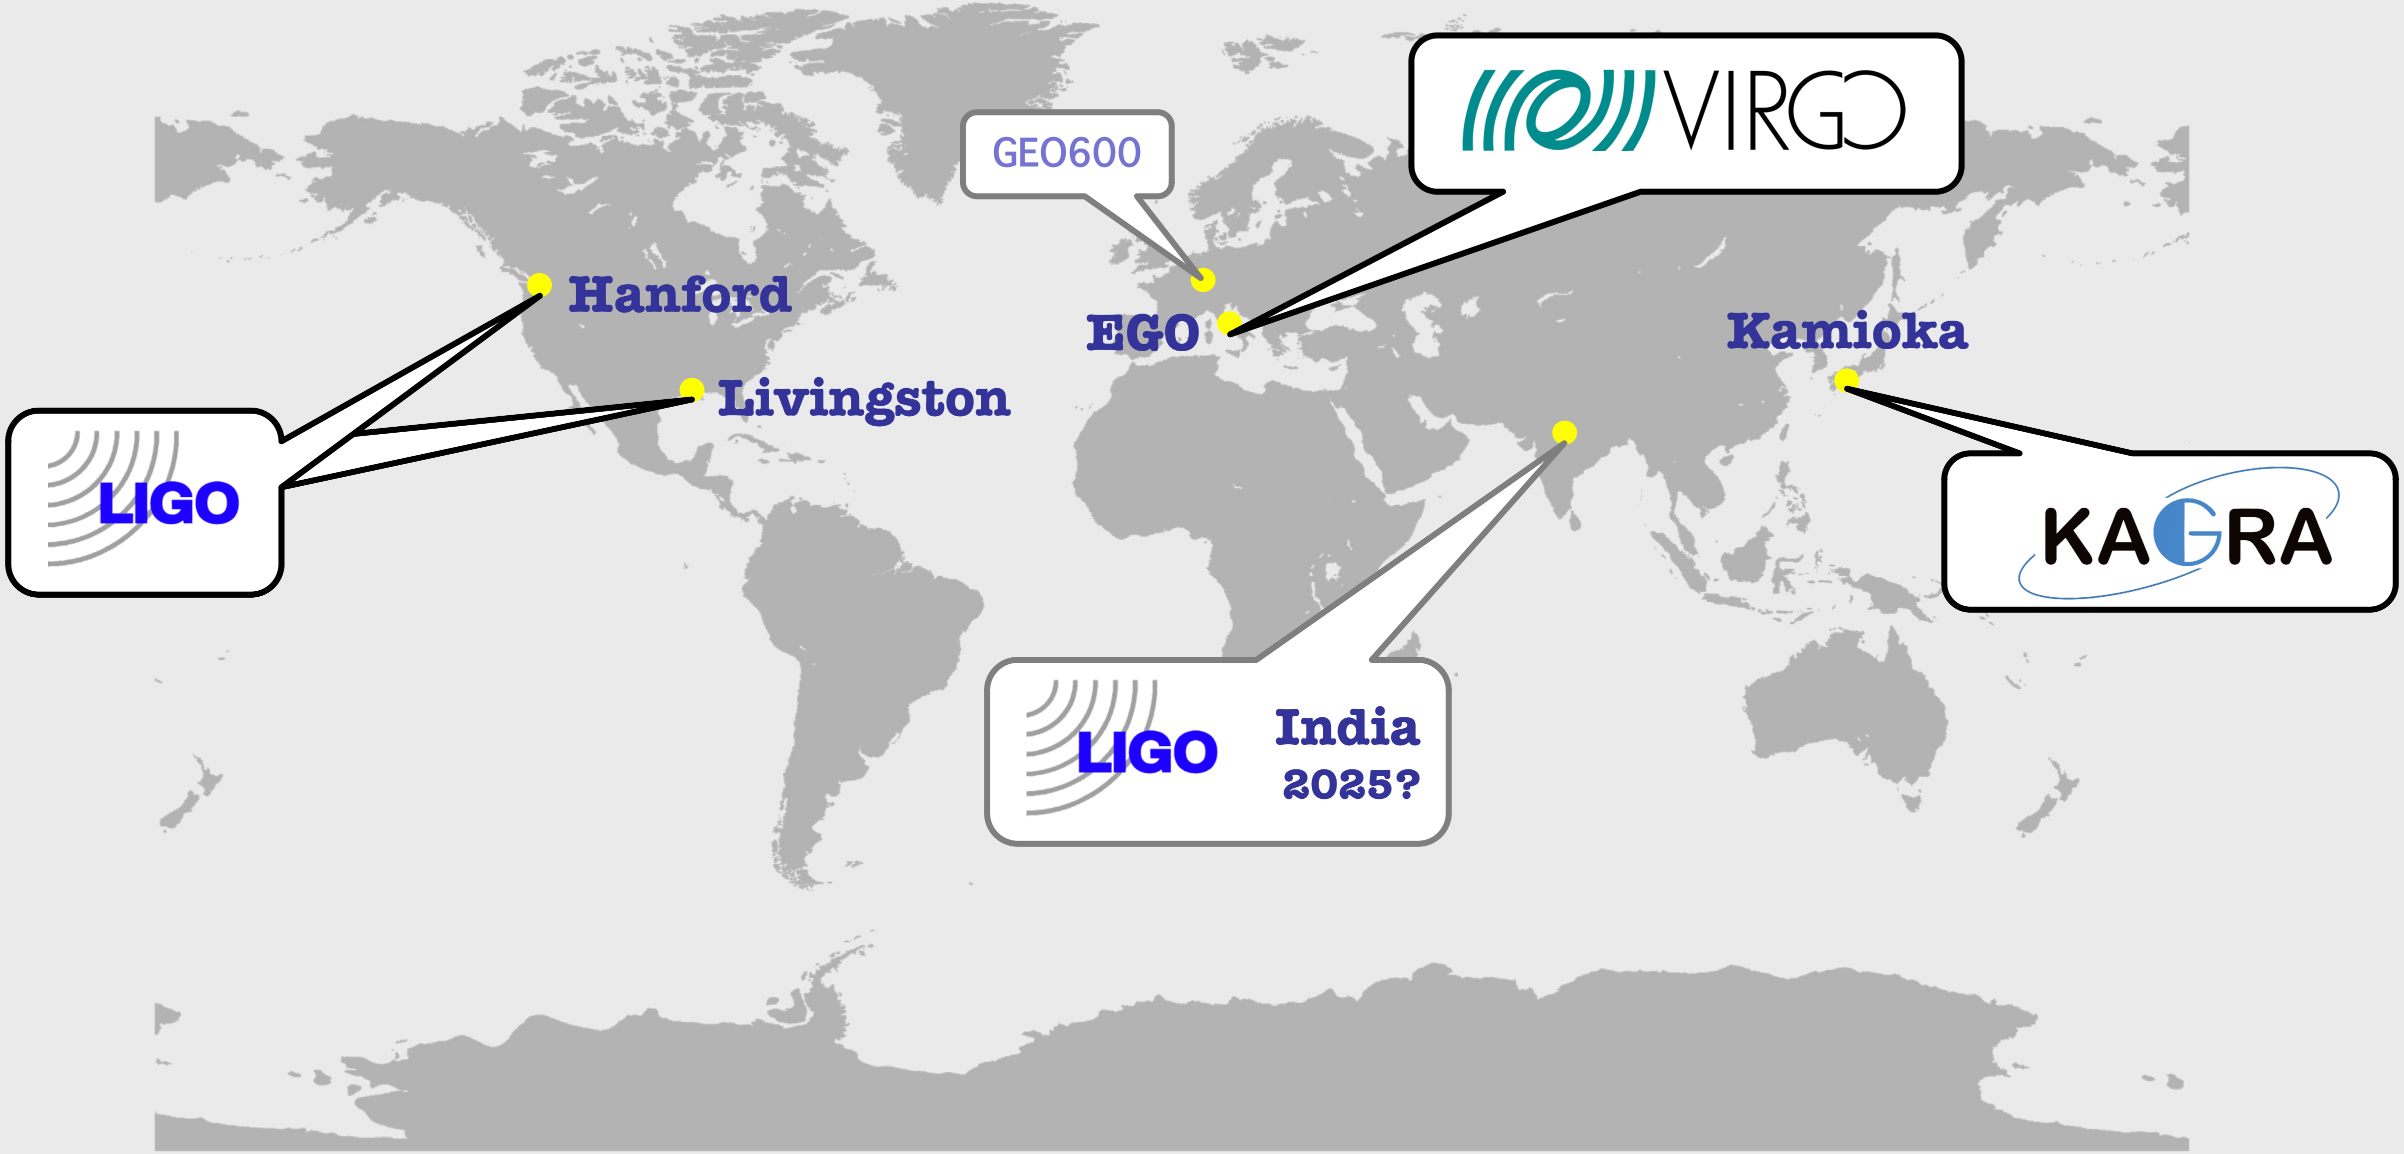

On-site infrastructure

Online
- Data acquisition and pre-processing
- Instrument control
- Environmental monitoring
- …

Plain old HTC (and some HPC)

Offline
- Deep searches
- Offline parameter estimation
- Detector Characterization (DetChar)
- (Template bank generation)

Here's the fun

Low-latency
- Candidate search
- Sky localization
- LL parameter estimation
- Alert generation and distribution

MULTIMESSENGER ASTRONOMY

- Only GW170817 so far...
- Public Alerts are “Triggers” for ground- and space-based EM observatories
- In O3 average latency was o(1/2 hour)
- Target latency for O4 is o(minute)
- However, for some events “early warning” alerts with negative latency are possible

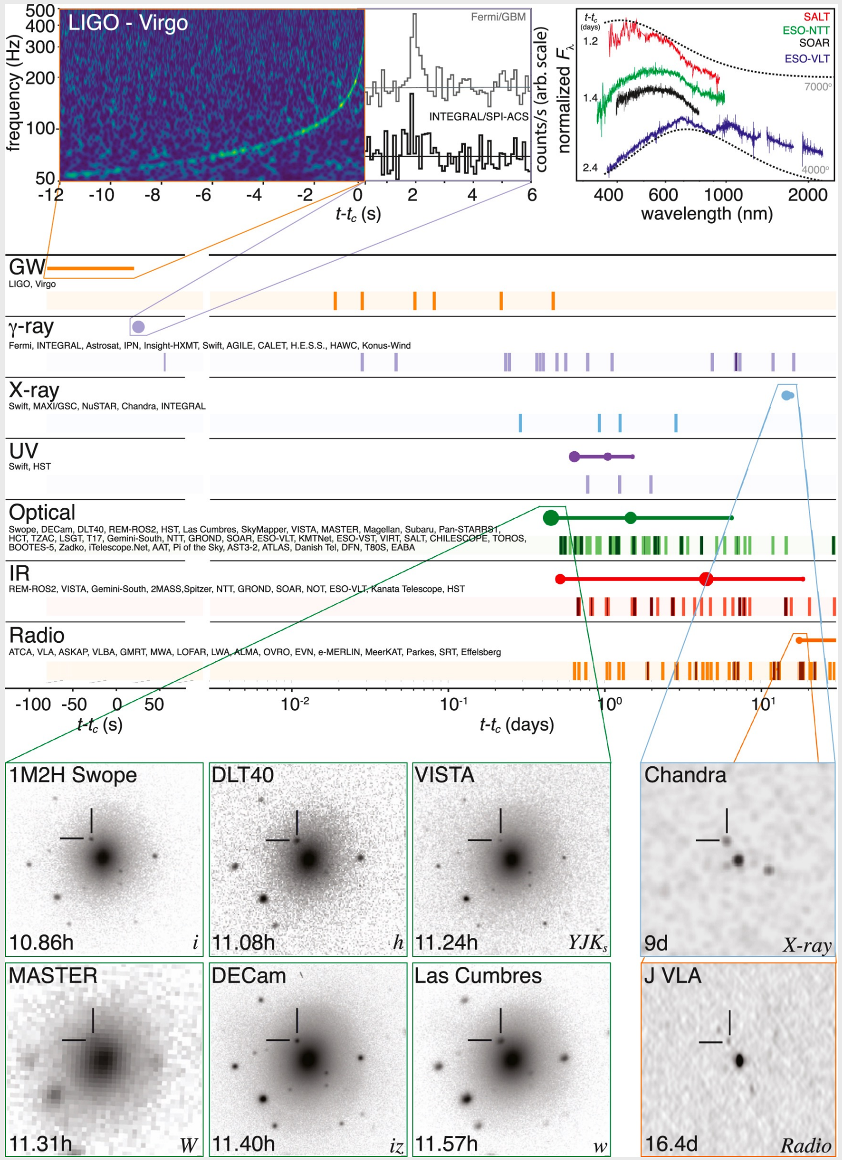

“Multi-messenger Observations of a Binary Neutron Star Merger”
B. P. Abbott *et al.* 2017 *ApJL* 848 L12
doi:10.3847/2041-8213/aa91c9

COMPLEX OVERALL DATA FLOWS

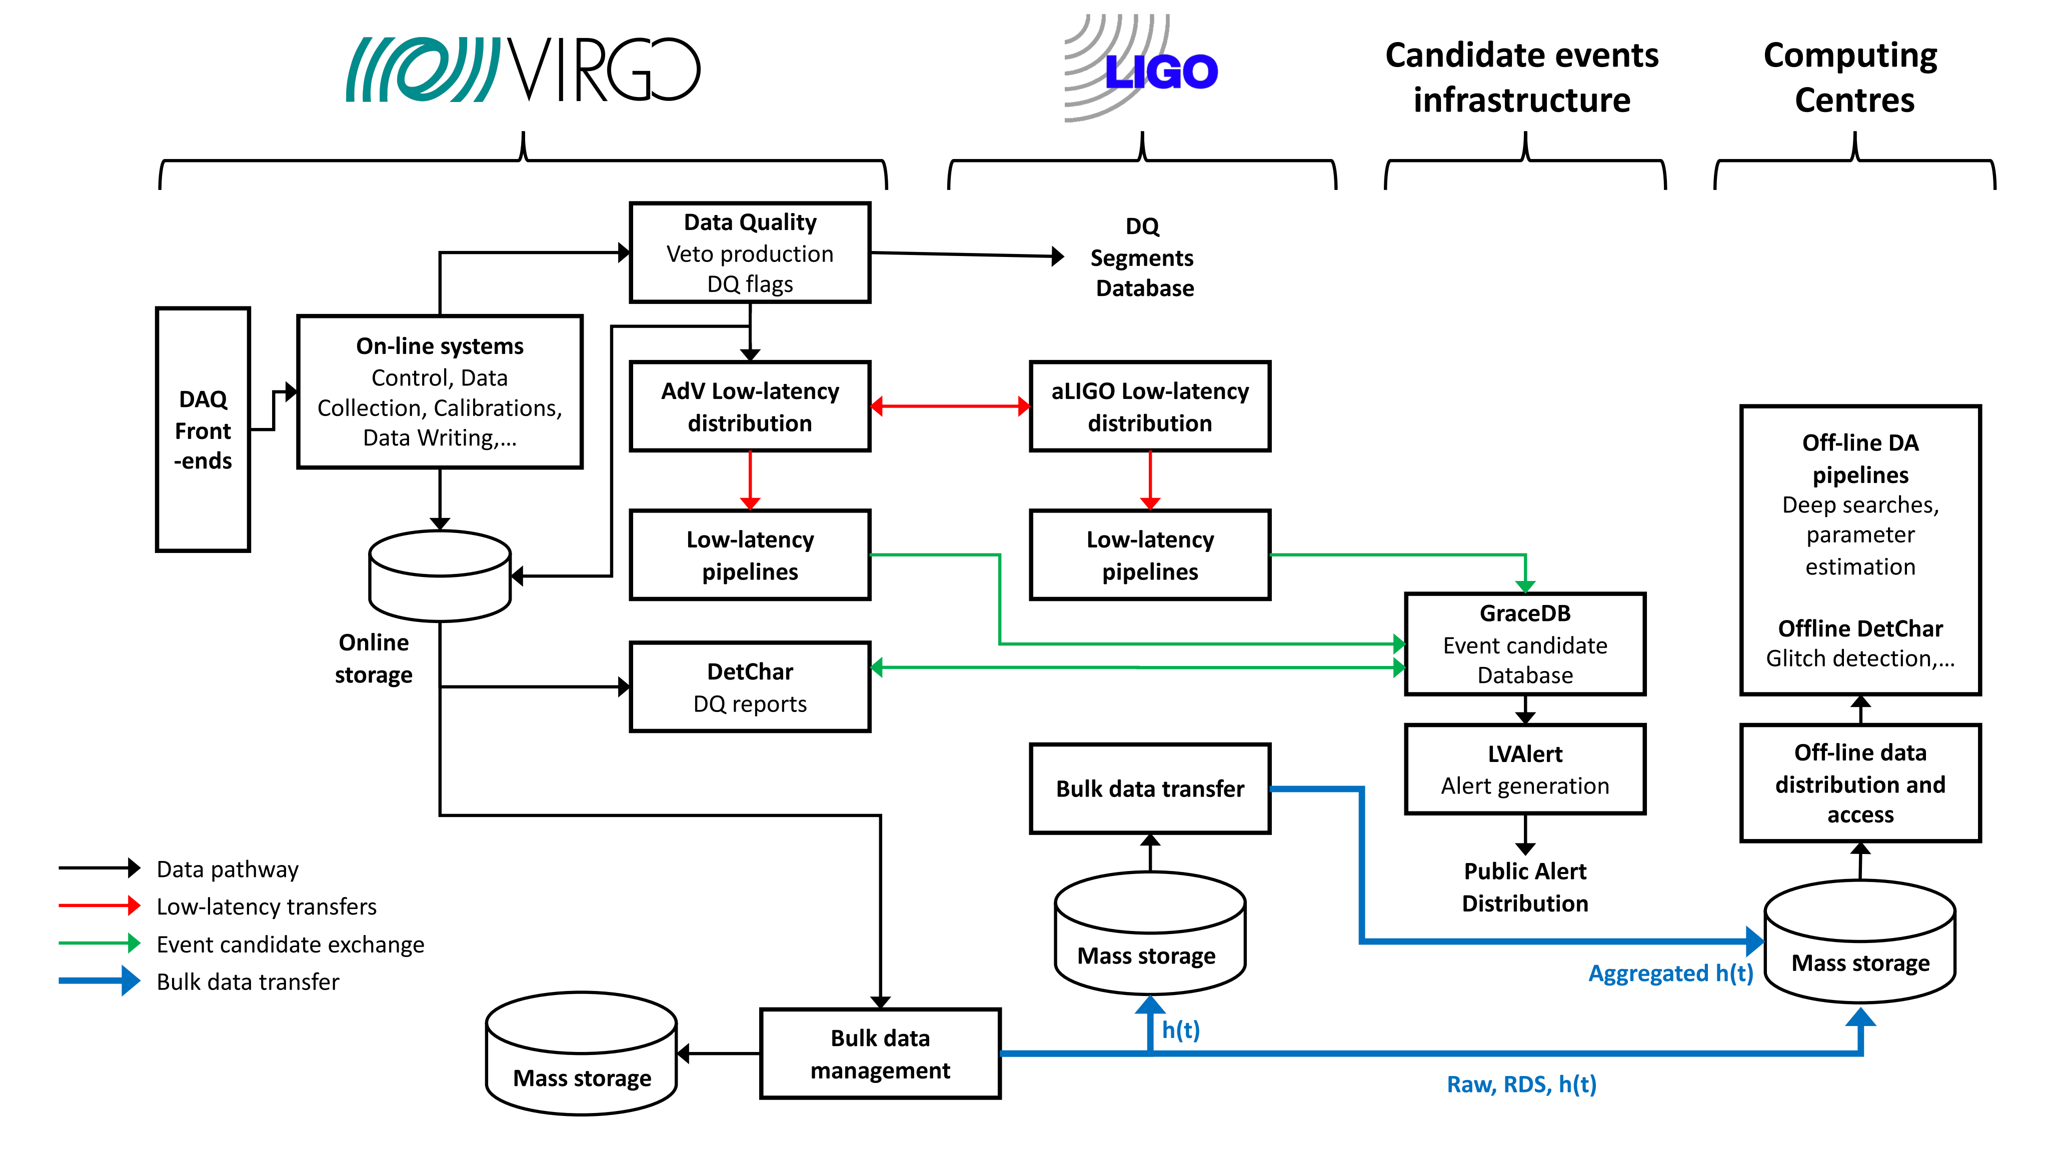

GRACEDB

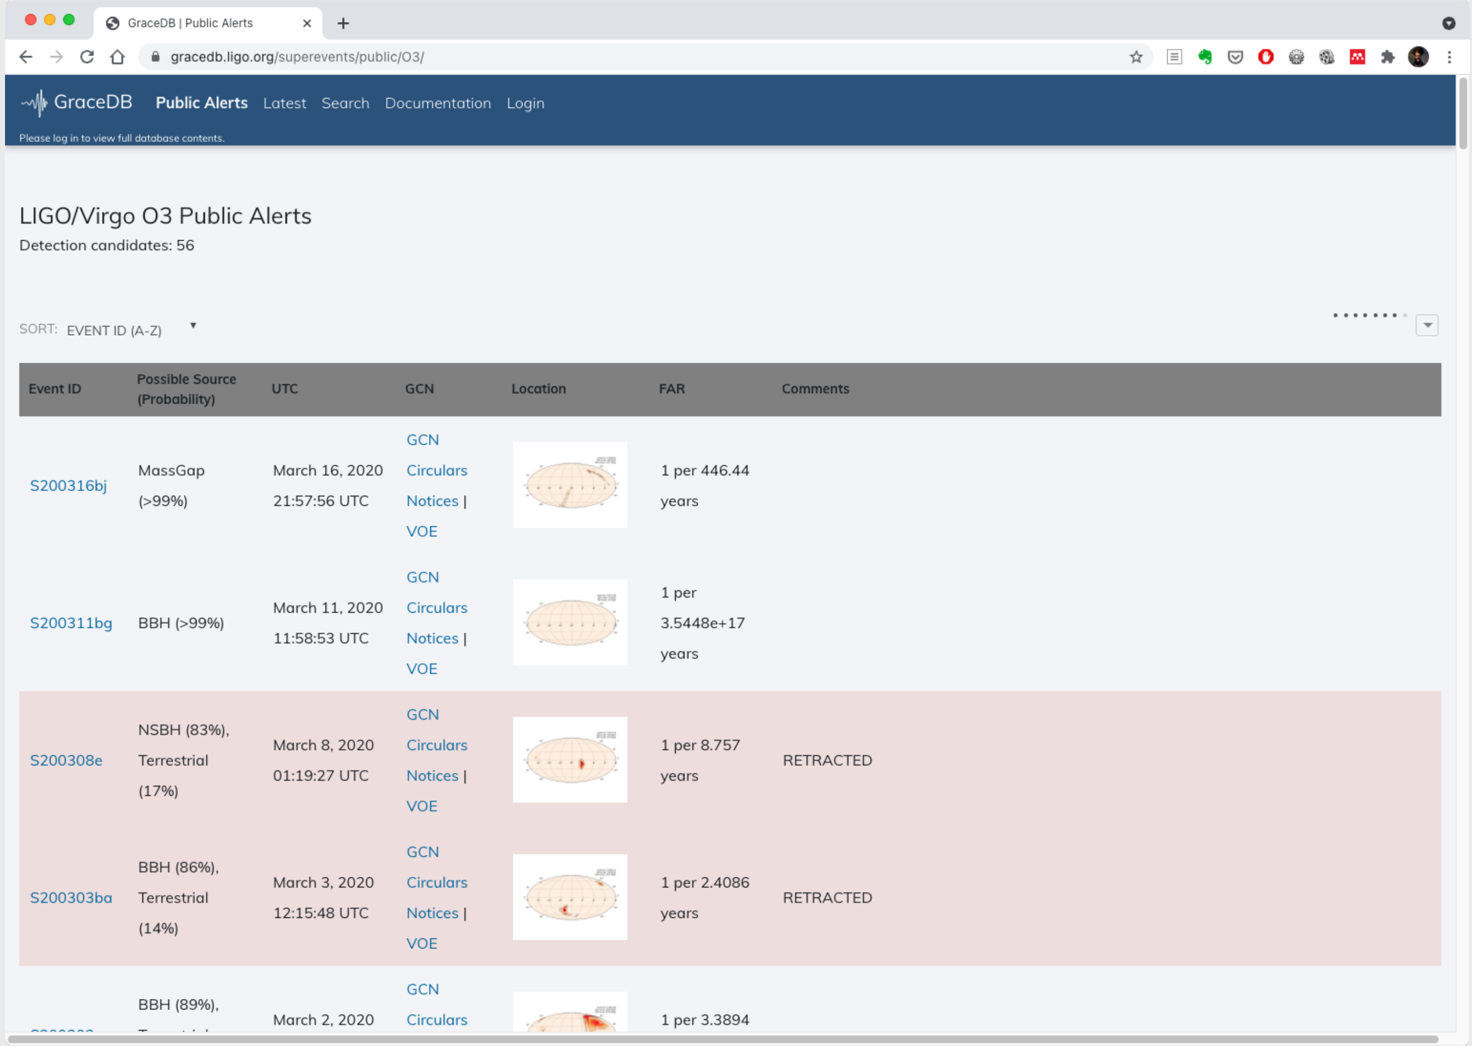

- Centralized aggregation point for information about candidate gravitational-wave events
- Provides web interface and programmatic API
- Works together with a notification service (XMPP-based LVAlert, being replaced by HOPSKOTCH/Kafka-based igwn-alert for O4)
- GWCelery: a package for annotating and orchestrating LIGO/Virgo alerts, built on the Celery distributed task queue.

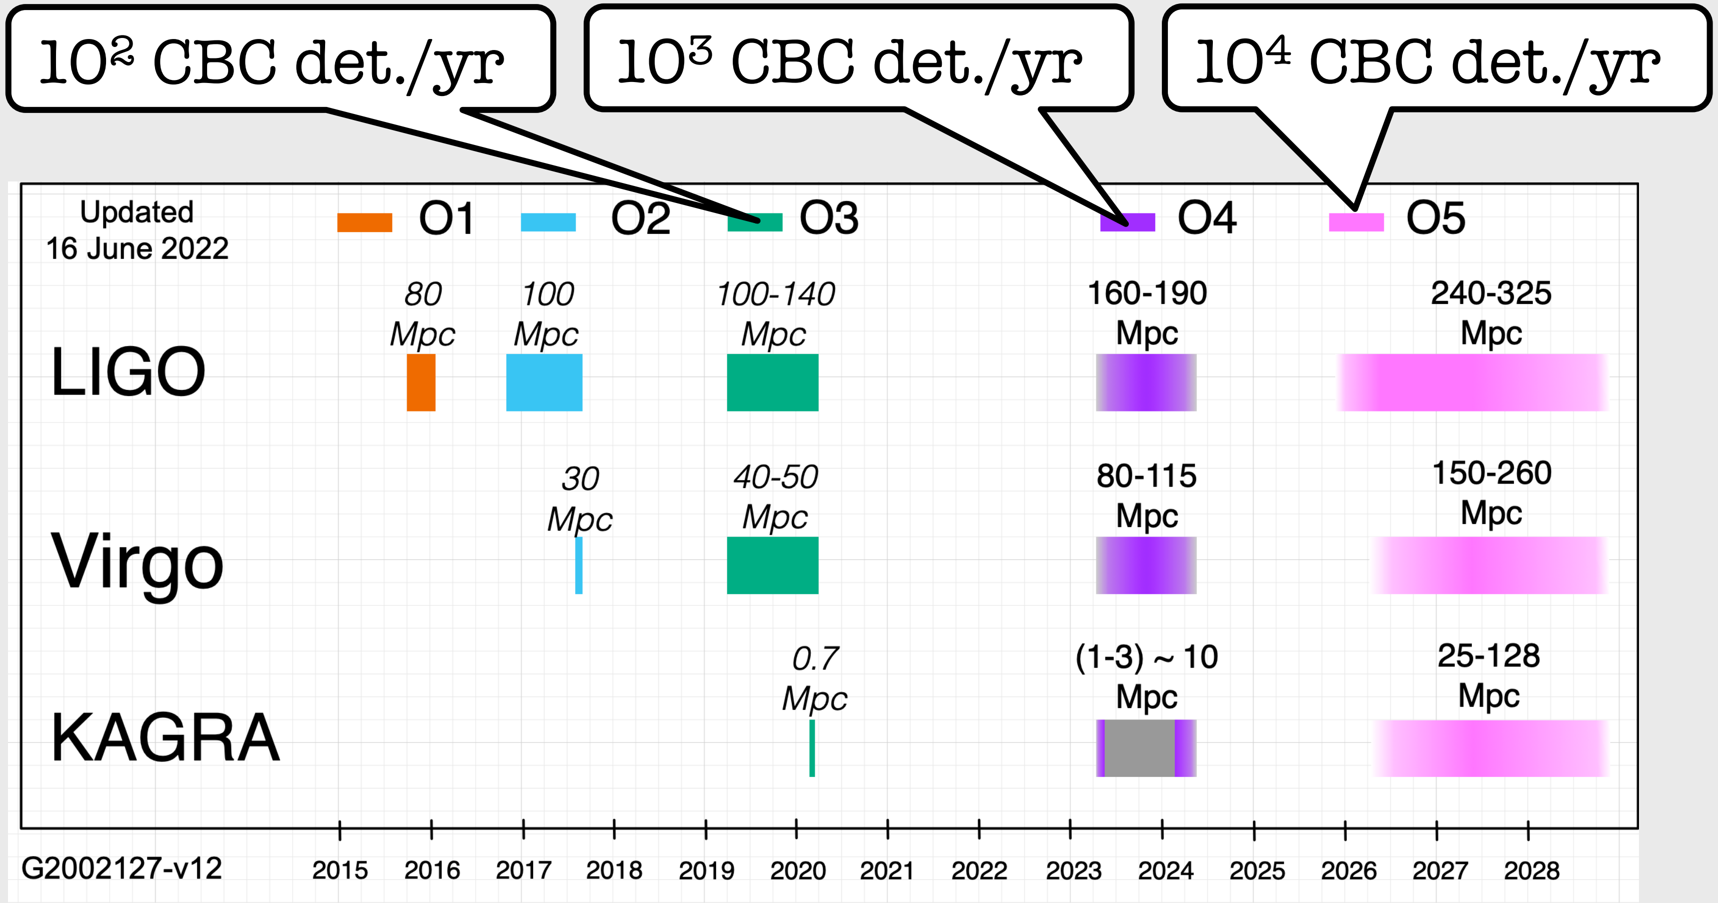

Then:

- "Post-O5" plans being prepared
- Projects for Third-generation interferometers being proposed (Einstein Telescope, Cosmic Explorer)
- Growth in "computing size" (relatively) gradual

OVERALL COMPUTING RESOURCES USAGE

- Including offline, low-latency and DetChar
- Overall CPU for O3 was ≈ 7000 MHS06 Hours
- Large peaks after end of observation periods (frequency-domain analyses)
- As expected, interest for older data wanes as new data become available
- Required computing power for O4 about 1.5÷2 × O3
- Overall: about 10% of an LHC experiment

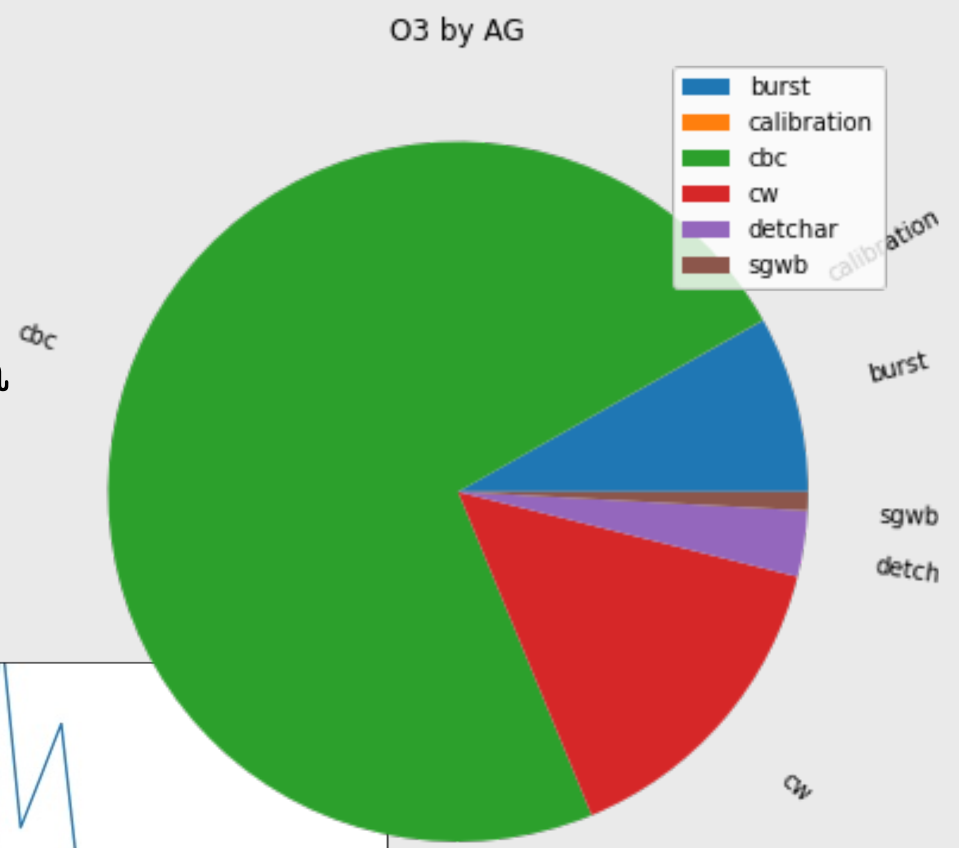

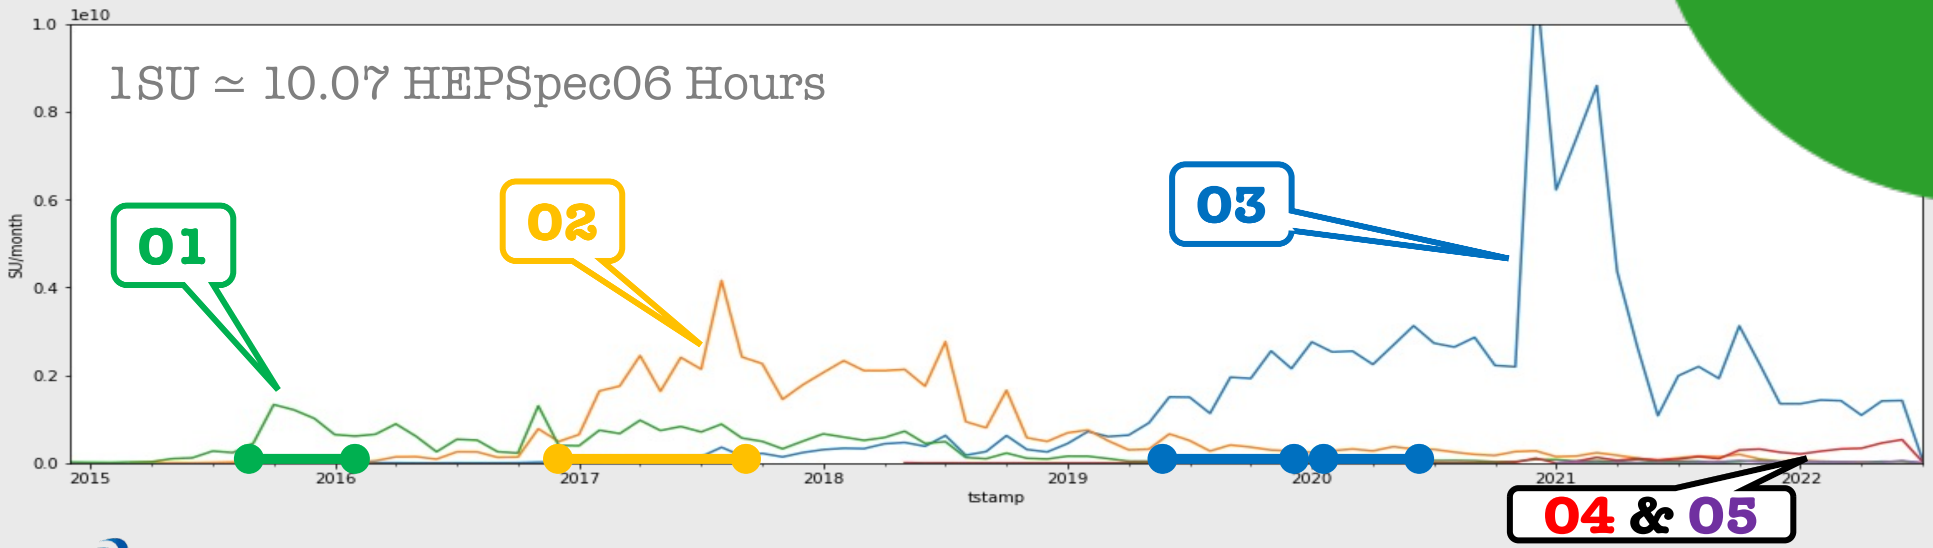

- Raw interferometer data don't grow much with increasing instrument sensitivity
  - In O3 we were writing o(1PB) per year of raw data per detector
  - h(t) (or "strain", the physics channel) + o(10^5) control channels
  - Pre-processed data for final user analysis is more than 1 order of magnitude smaller
  - 1.5 × expected in O4 (more control channels!)
- What grows is the amount of useful scientific information embedded in the data
  - And the computing power needed to wring it out

Need to define and deploy a common and sustainable GW computing environment

- Provide a uniform runtime environment for offline pipelines
- Full interoperability between Virgo, LIGO (and KAGRA)
- Provide scalability and the opportunity to exploit heterogeneous resources
- Adopt mainstream, widely used tools, leveraging upon HEP experience

Enter **IGWN** – the International Gravitational Waves observatories Network

- A coordination effort aimed at jointly discussing the computing policy, management, and architecture issues of LIGO, Virgo, and KAGRA.

So, for example…

- Use Apache Kafka for low-latency data transfers
- Move towards a fully distributed offline computing model
  - Based on HTCondor for workload management and CVMFS+StashCache for data distribution
  - Strong suport from OSG and HTCondor community
- Gradually adopt Rucio for data management tasks
  - E.g. transfer of h(t) data to StashCache Origin servers for distribution
- Streamline code management
  - GitLab, CI, Conda Environments, Containers,…
- …

THE EU SIDE OF THE IGWN

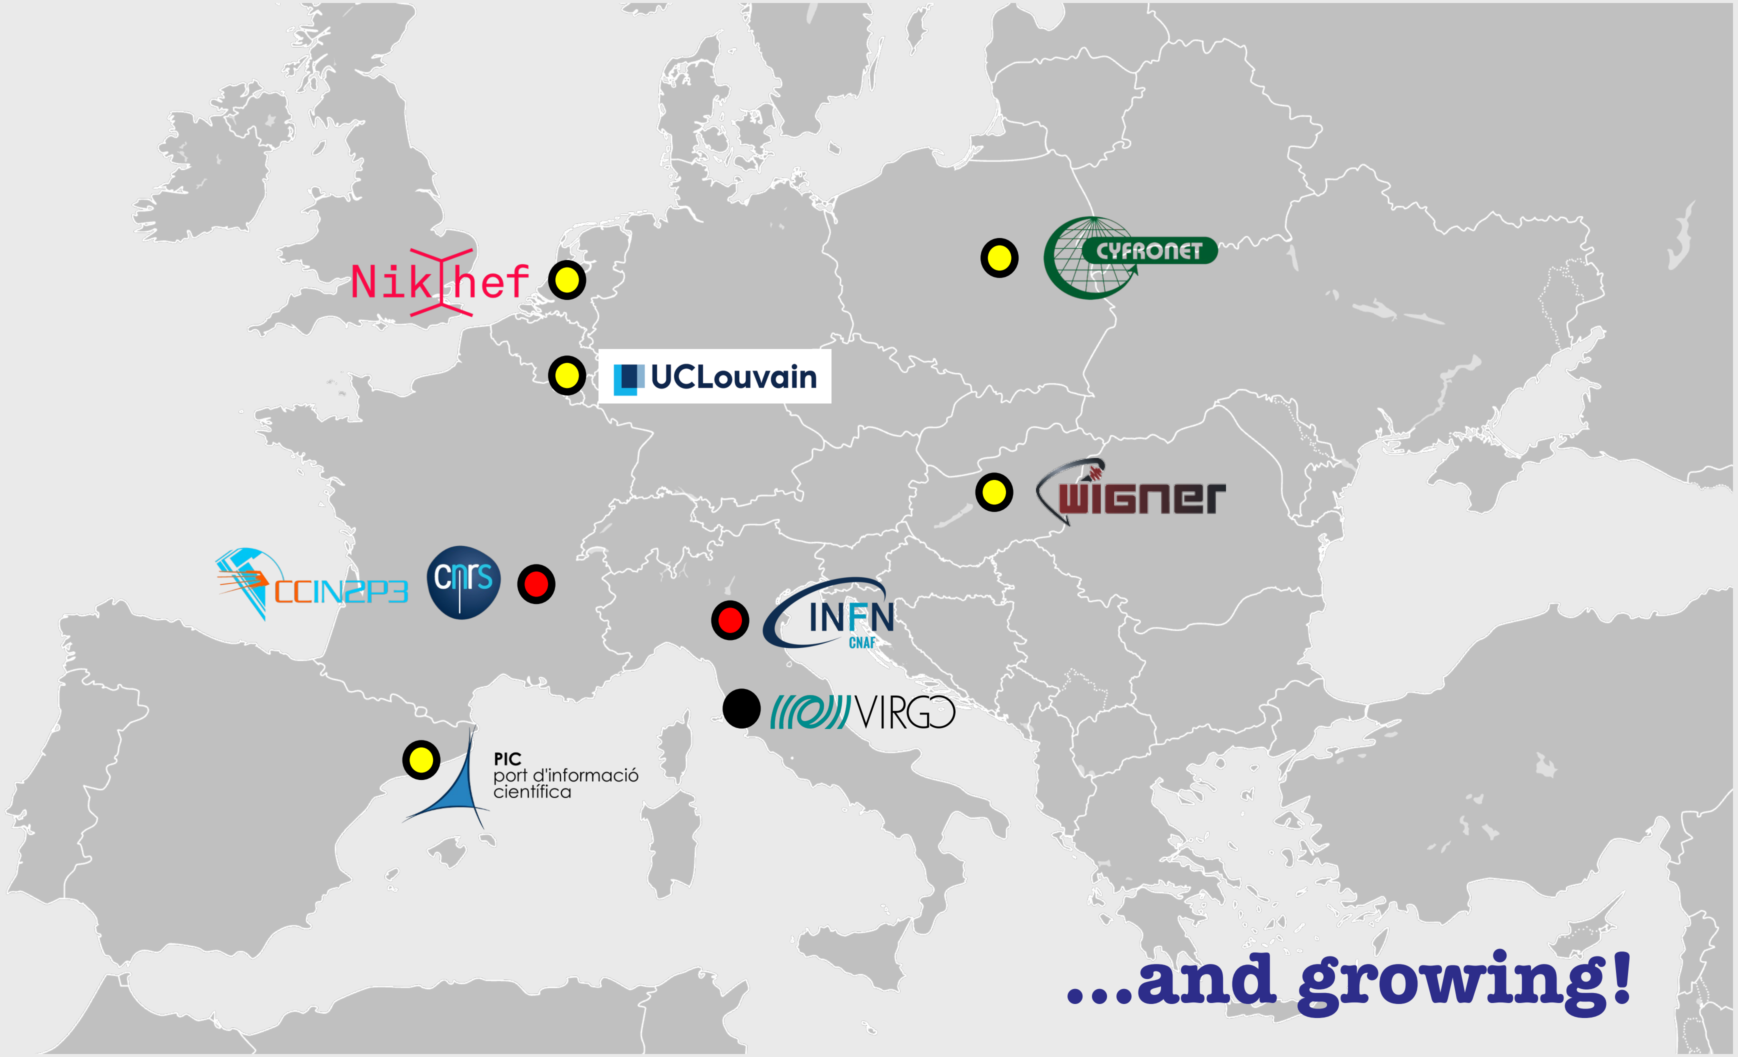

- Tier-1
- Tier-2

...and growing!

MODERN LOW-LATENCY SERVICES DEPLOYMENT

- Upgrade of main services (GWCelery and GraceDB)
  - And exploring a complete redesign
- Deployment on Kubernetes in INFN-Cloud at CNAF, INFN's Tier-1 centre in Bologna
  - With the idea of also having on-demand CPU resources for running jobs (not yet there)
  - Production and Development instances running at CNAF
  - Sandboxed configuration option for self-contained installation on Minikube
  - GitLab CI for deployment and configuration
  - Monitoring via Prometheus/Grafana
- Get ready for exploitation of upcoming "Scientific Cloud" resources

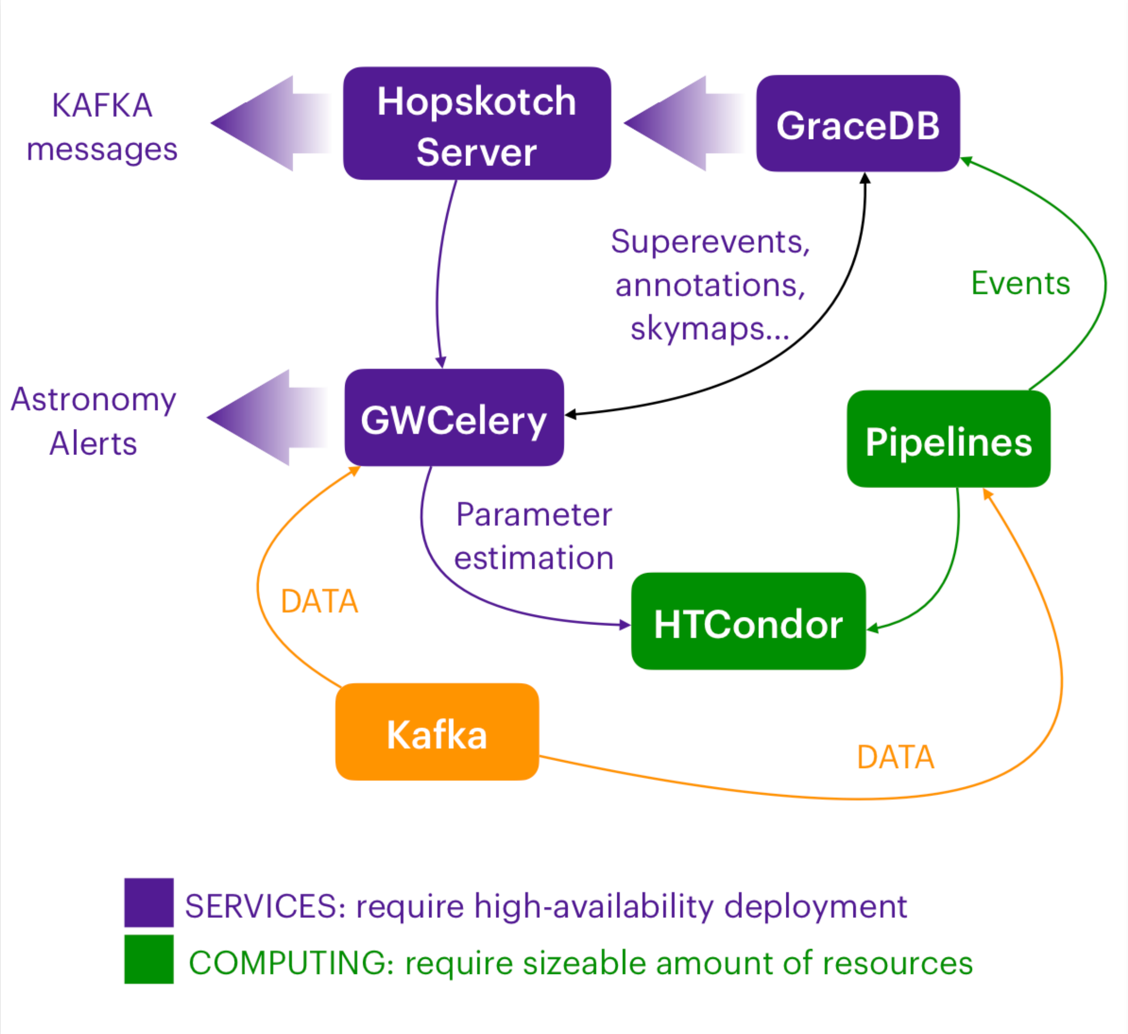

Sara.Vallero@to.infn.it

HOWEVER...

- Still a very large number of computing tools and services to maintain
  - Even adopting mainstream technologies, some high-level services are unavoidably proprietary
- Pipelines are very heterogeneous and difficult to port to distributed environment
  - Python, C++, Matlab,...
  - Essentially no common framework
  - Many have been developed to run in specific environments (data access technology, runtime environment,...)
- Difficult to recruit (and keep) the right computing skills
  - ...which translates into a large personpower gap for computing

CONCLUSIONS

- GW computing started small but is coming of age
- Increased interferometer sensitivity does not imply proportional growth in data size…
- …but computing power requirements do grow.
- No quantum leaps O4 → O5 → “post-O5” → 3G
- A large effort is being made to adopt modern computing technologies and strategies, and to integrate into the wider physics computing community
- A large personpower gap is slowing down the transition, but most important targets for O4 are within reach