VIRGO & THE WLCG

European Gravitational Observatory, Cascina (Italy)

A WORLDWIDE NETWORK

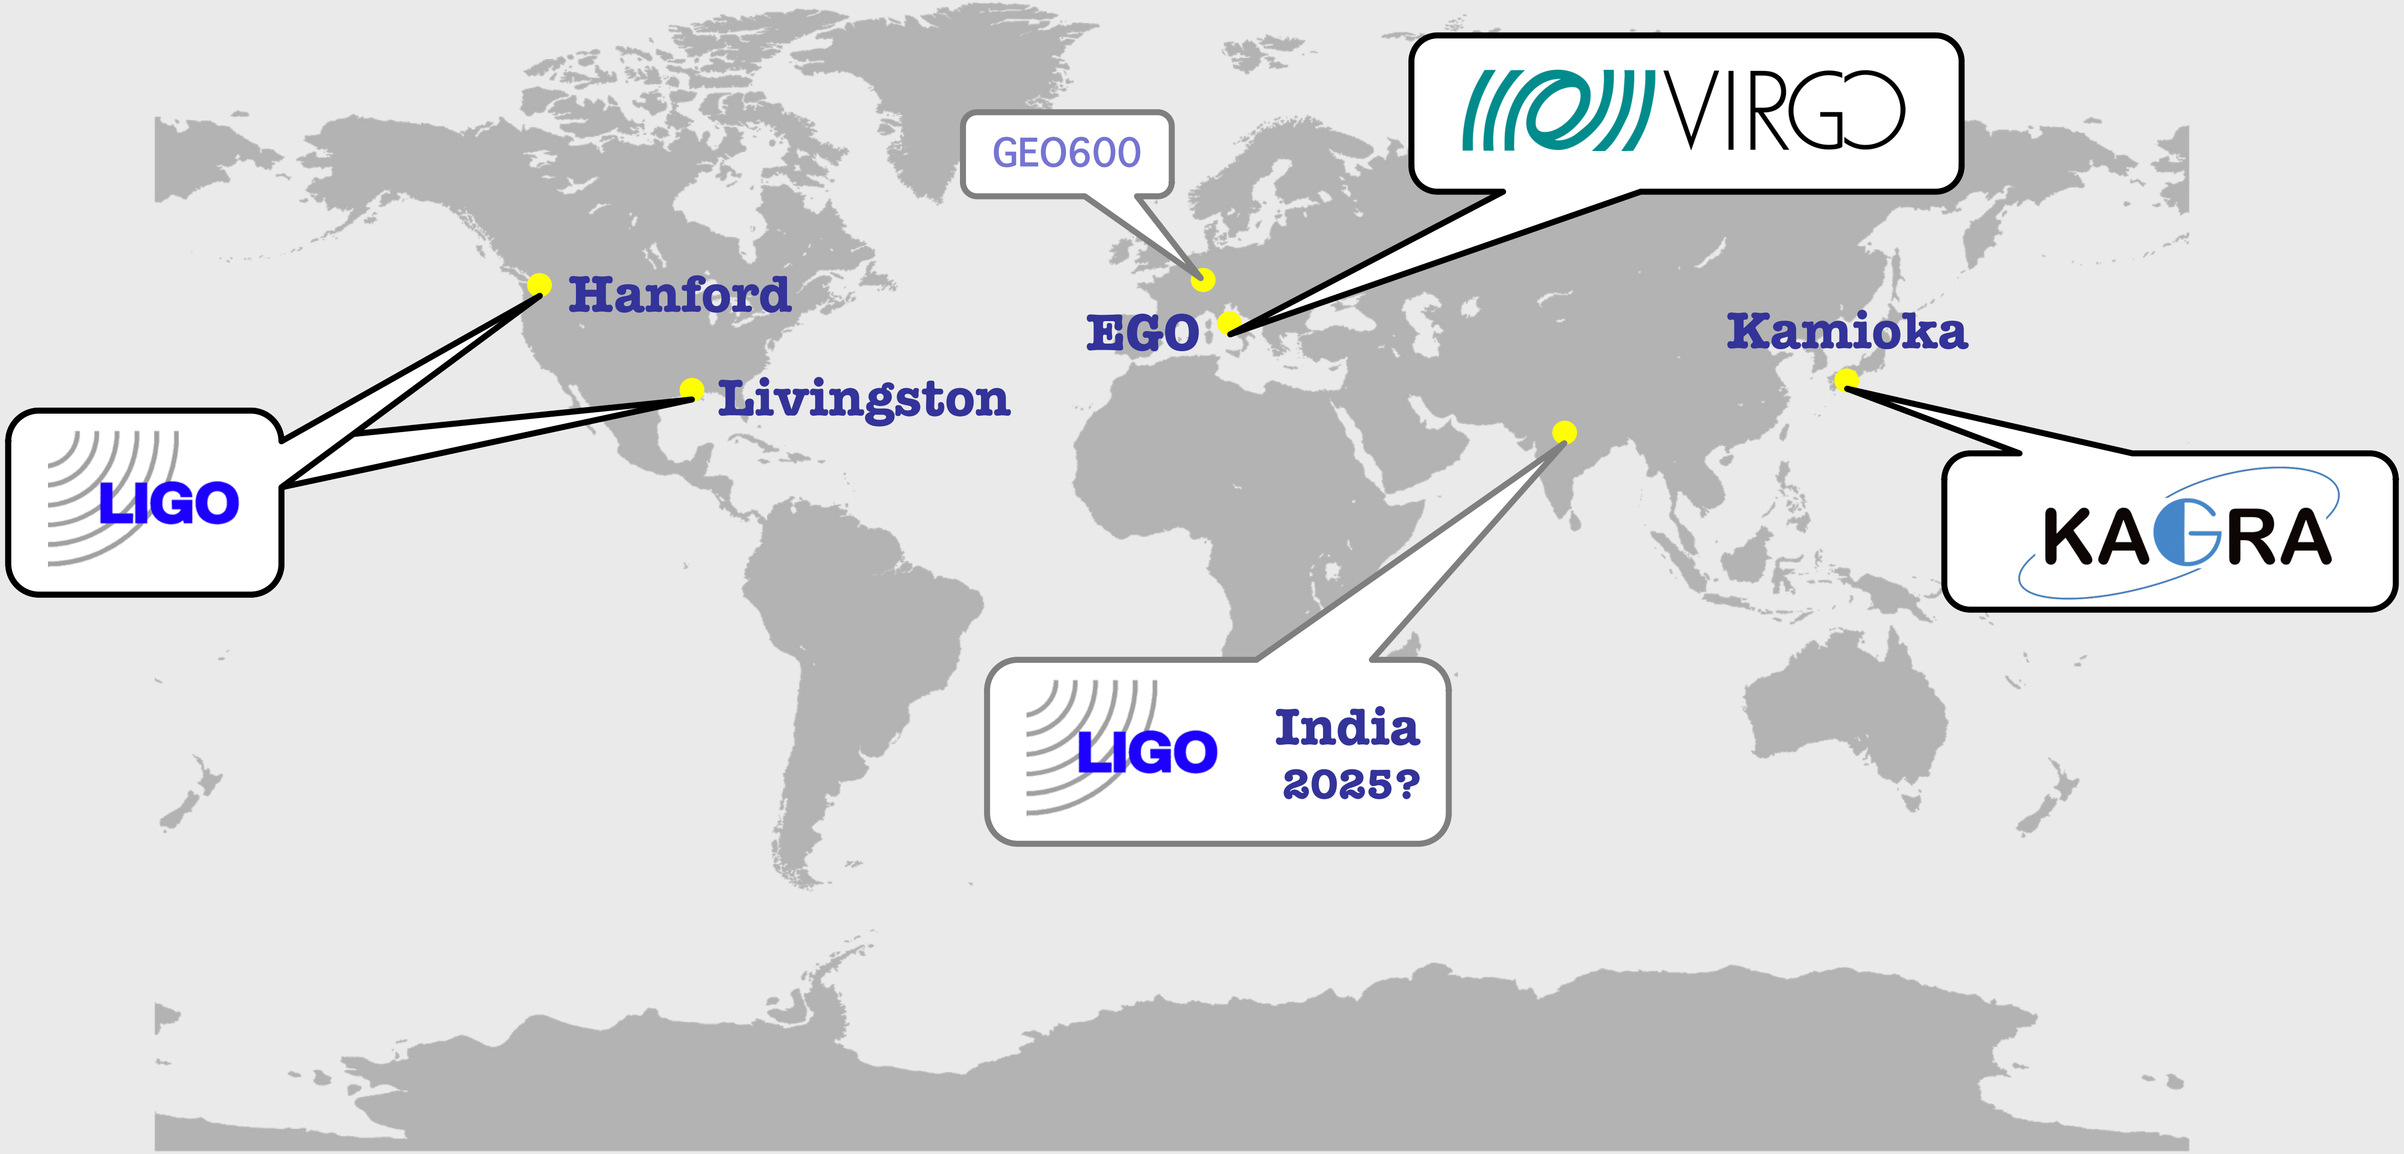

THREE COMPUTING DOMAINS

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

- 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

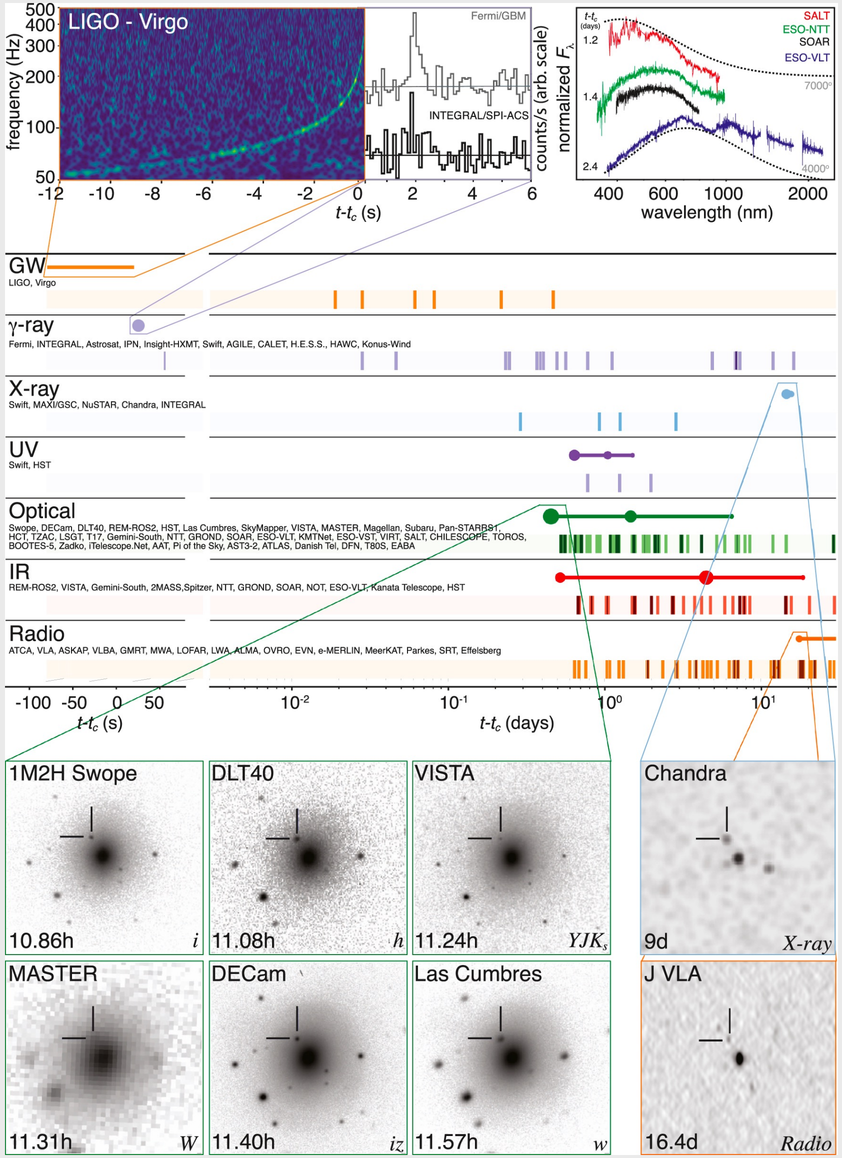

“Multi-messenger Observations of a Binary Neutron Star Merger”
B. P. Abbott *et al.* 2017 *ApJL* 848 L12
doi:10.3847/2041-8213/aa91c9

COMPLEX OVERALL DATA FLOWS

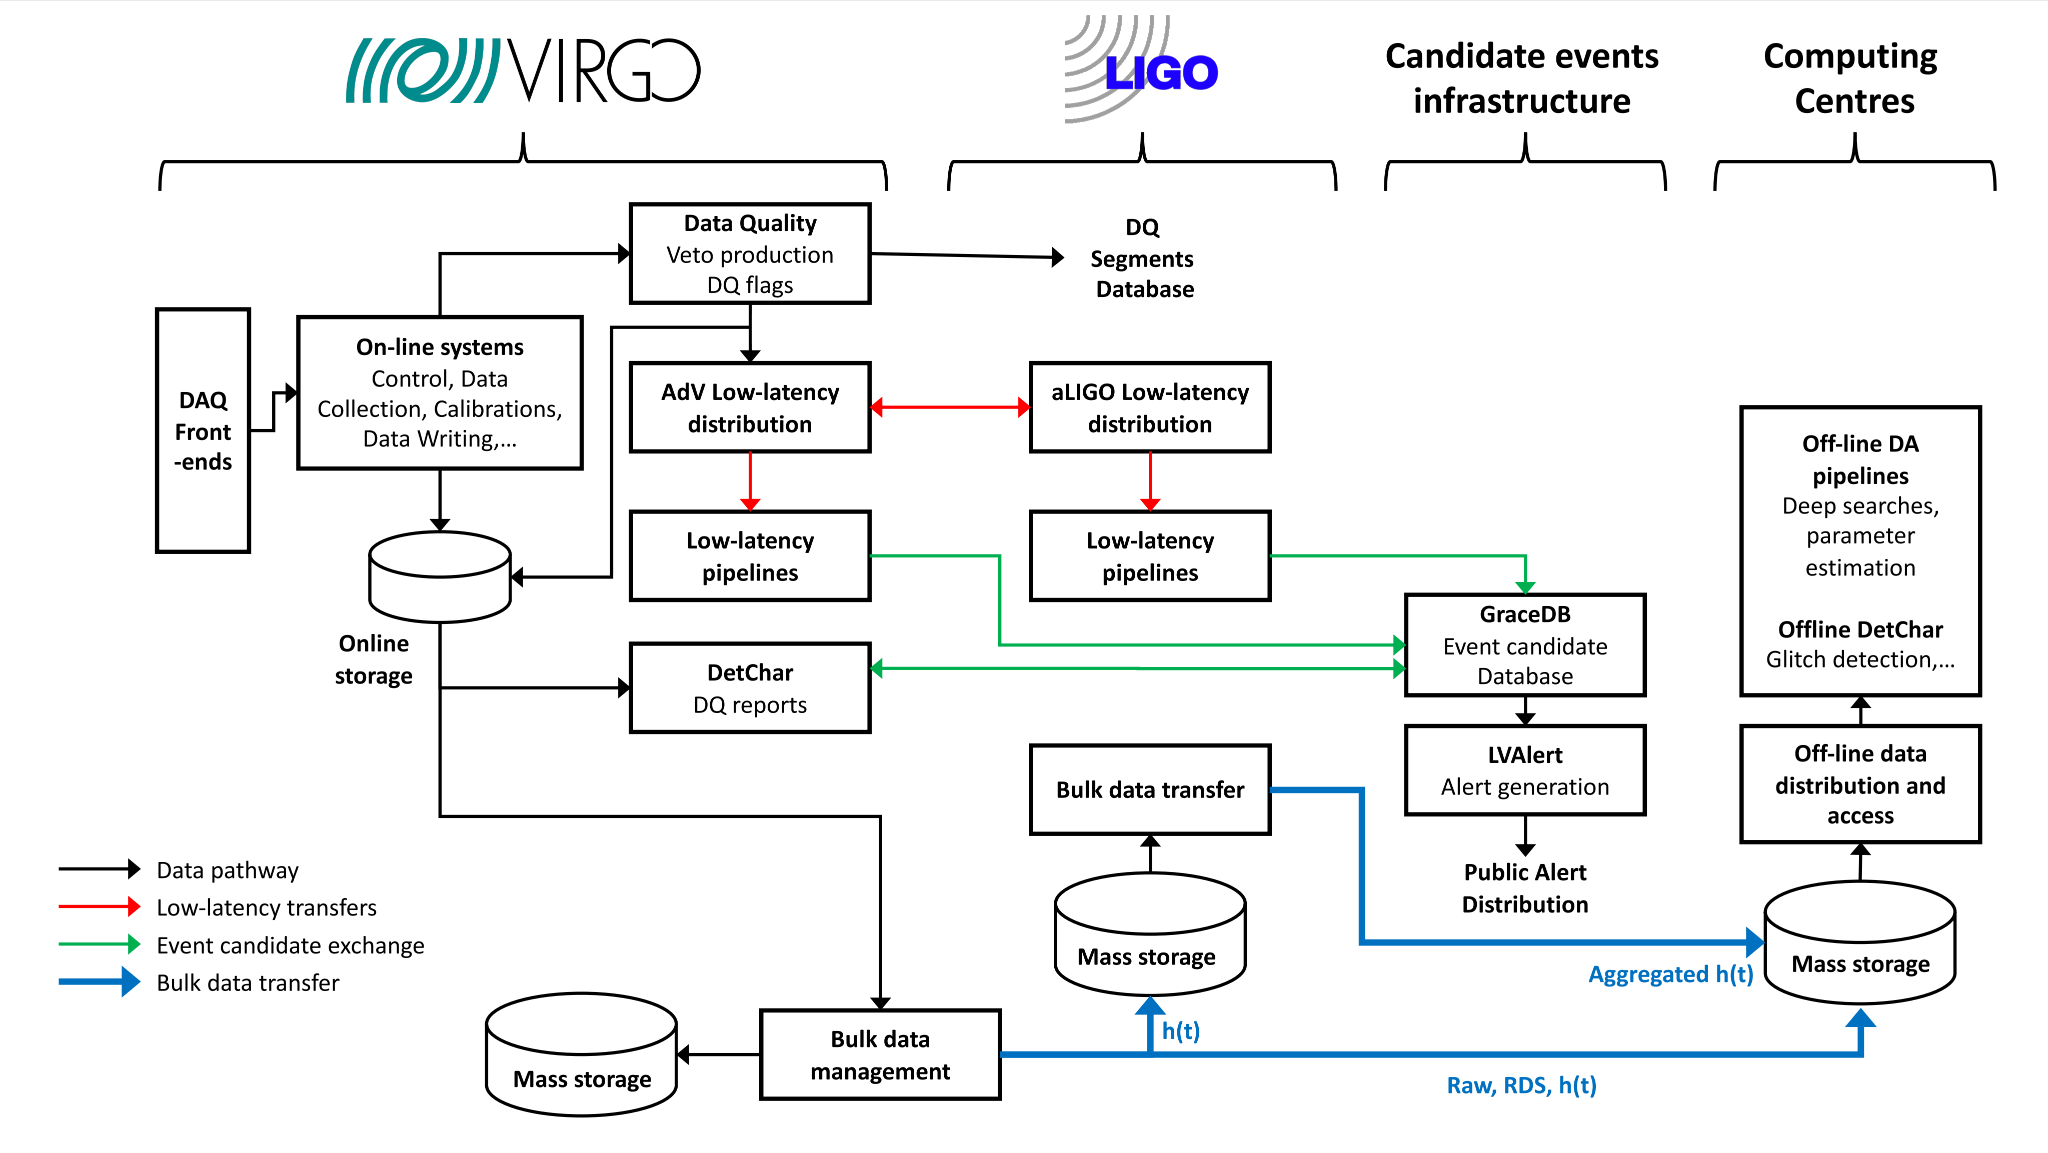

THE EU SIDE OF THE IGWN

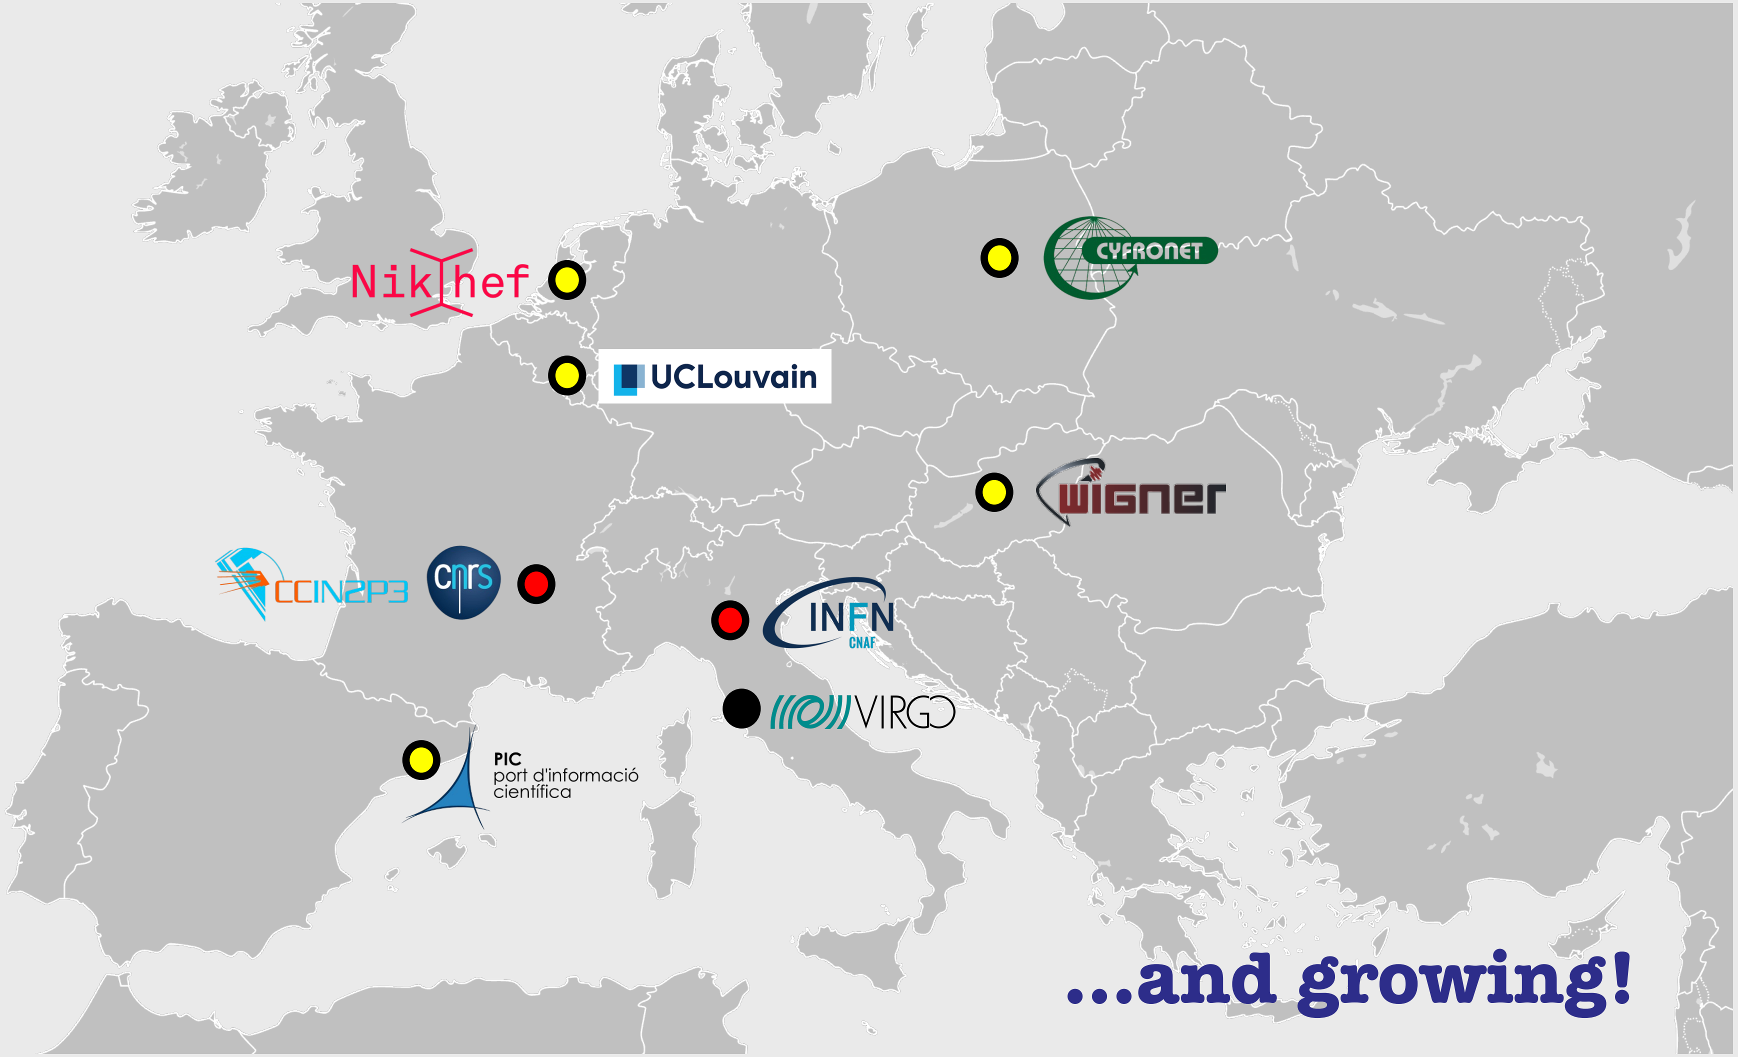

- Tier-1
- Tier-2

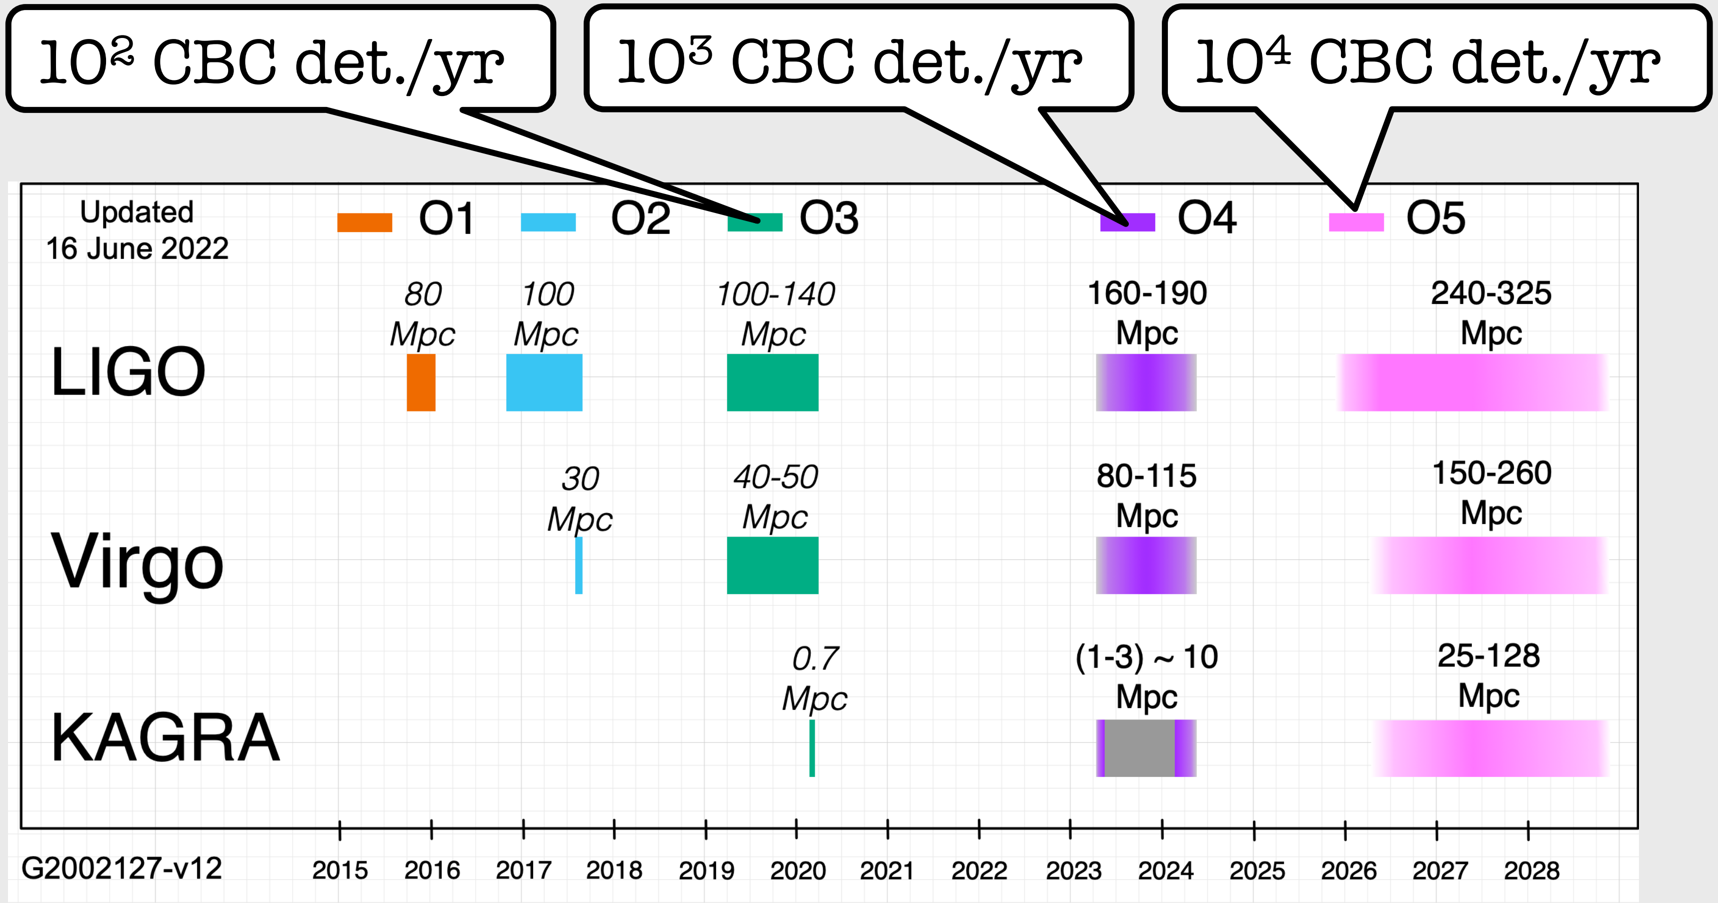

Then:

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

DATA AND MORE DATA

- 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

OVERALL COMPUTING RESOURCES USAGE

- Including offline and low-latency
- Overall CPU for O3 was ≈ 7000 MHS06 Hours
- Large peaks after end of observation period (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

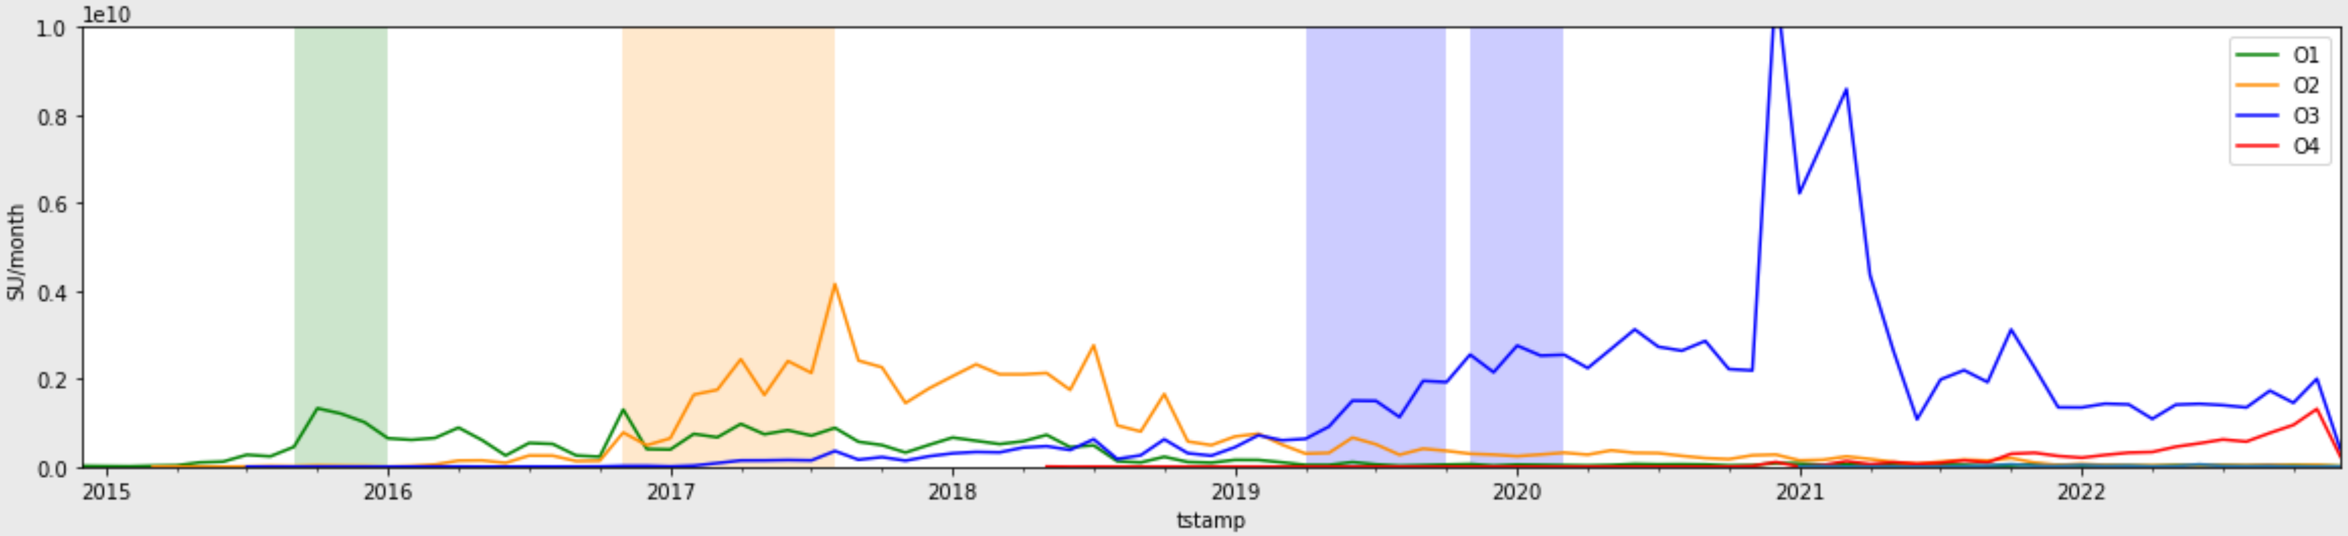

Need to define and deploy a common and sustainable GW computing environment

- Provide a uniform runtime environment for offline pipelines
- Full interoperability between Virgo, LIGO (and KAGRA)
- Provide scalability and move to distributed 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…

- Move towards a fully distributed offline computing model
  - Based on HTCondor for workload management and CVMFS+OSDF for data distribution
  - Strong suport by OSG and HTCondor community, hope for support also by WLCG ☺
- Gradually adopt Rucio for data management tasks
  - E.g. transfer of h(t) data to OSDF Origin server at UCLouvain for distribution
- Accept computing cycles from centres as in-kind contributions
  - Dicussion and update of MoUs underway

CONCLUSIONS

- GW computing started small but is coming of age
- IGWN is pushing the community towards “standard” (= HEP!) tools and architectures, most notably distributed computing
- We keep expanding our computing infratsructure to include more and more WLCG centres
- Einstein Telscope will be next player in the same timescale of HL-LHC
- Will need coordination with many WLCG sites, more practical to coordinate with WLCG itself!