IGWN: VIRGO computing towards O4

and Low Latency computing for GW science

CCR Workshop May 24-28 2021

VIRGO

S. Bagnasco, S. Vallero

A worldwide network

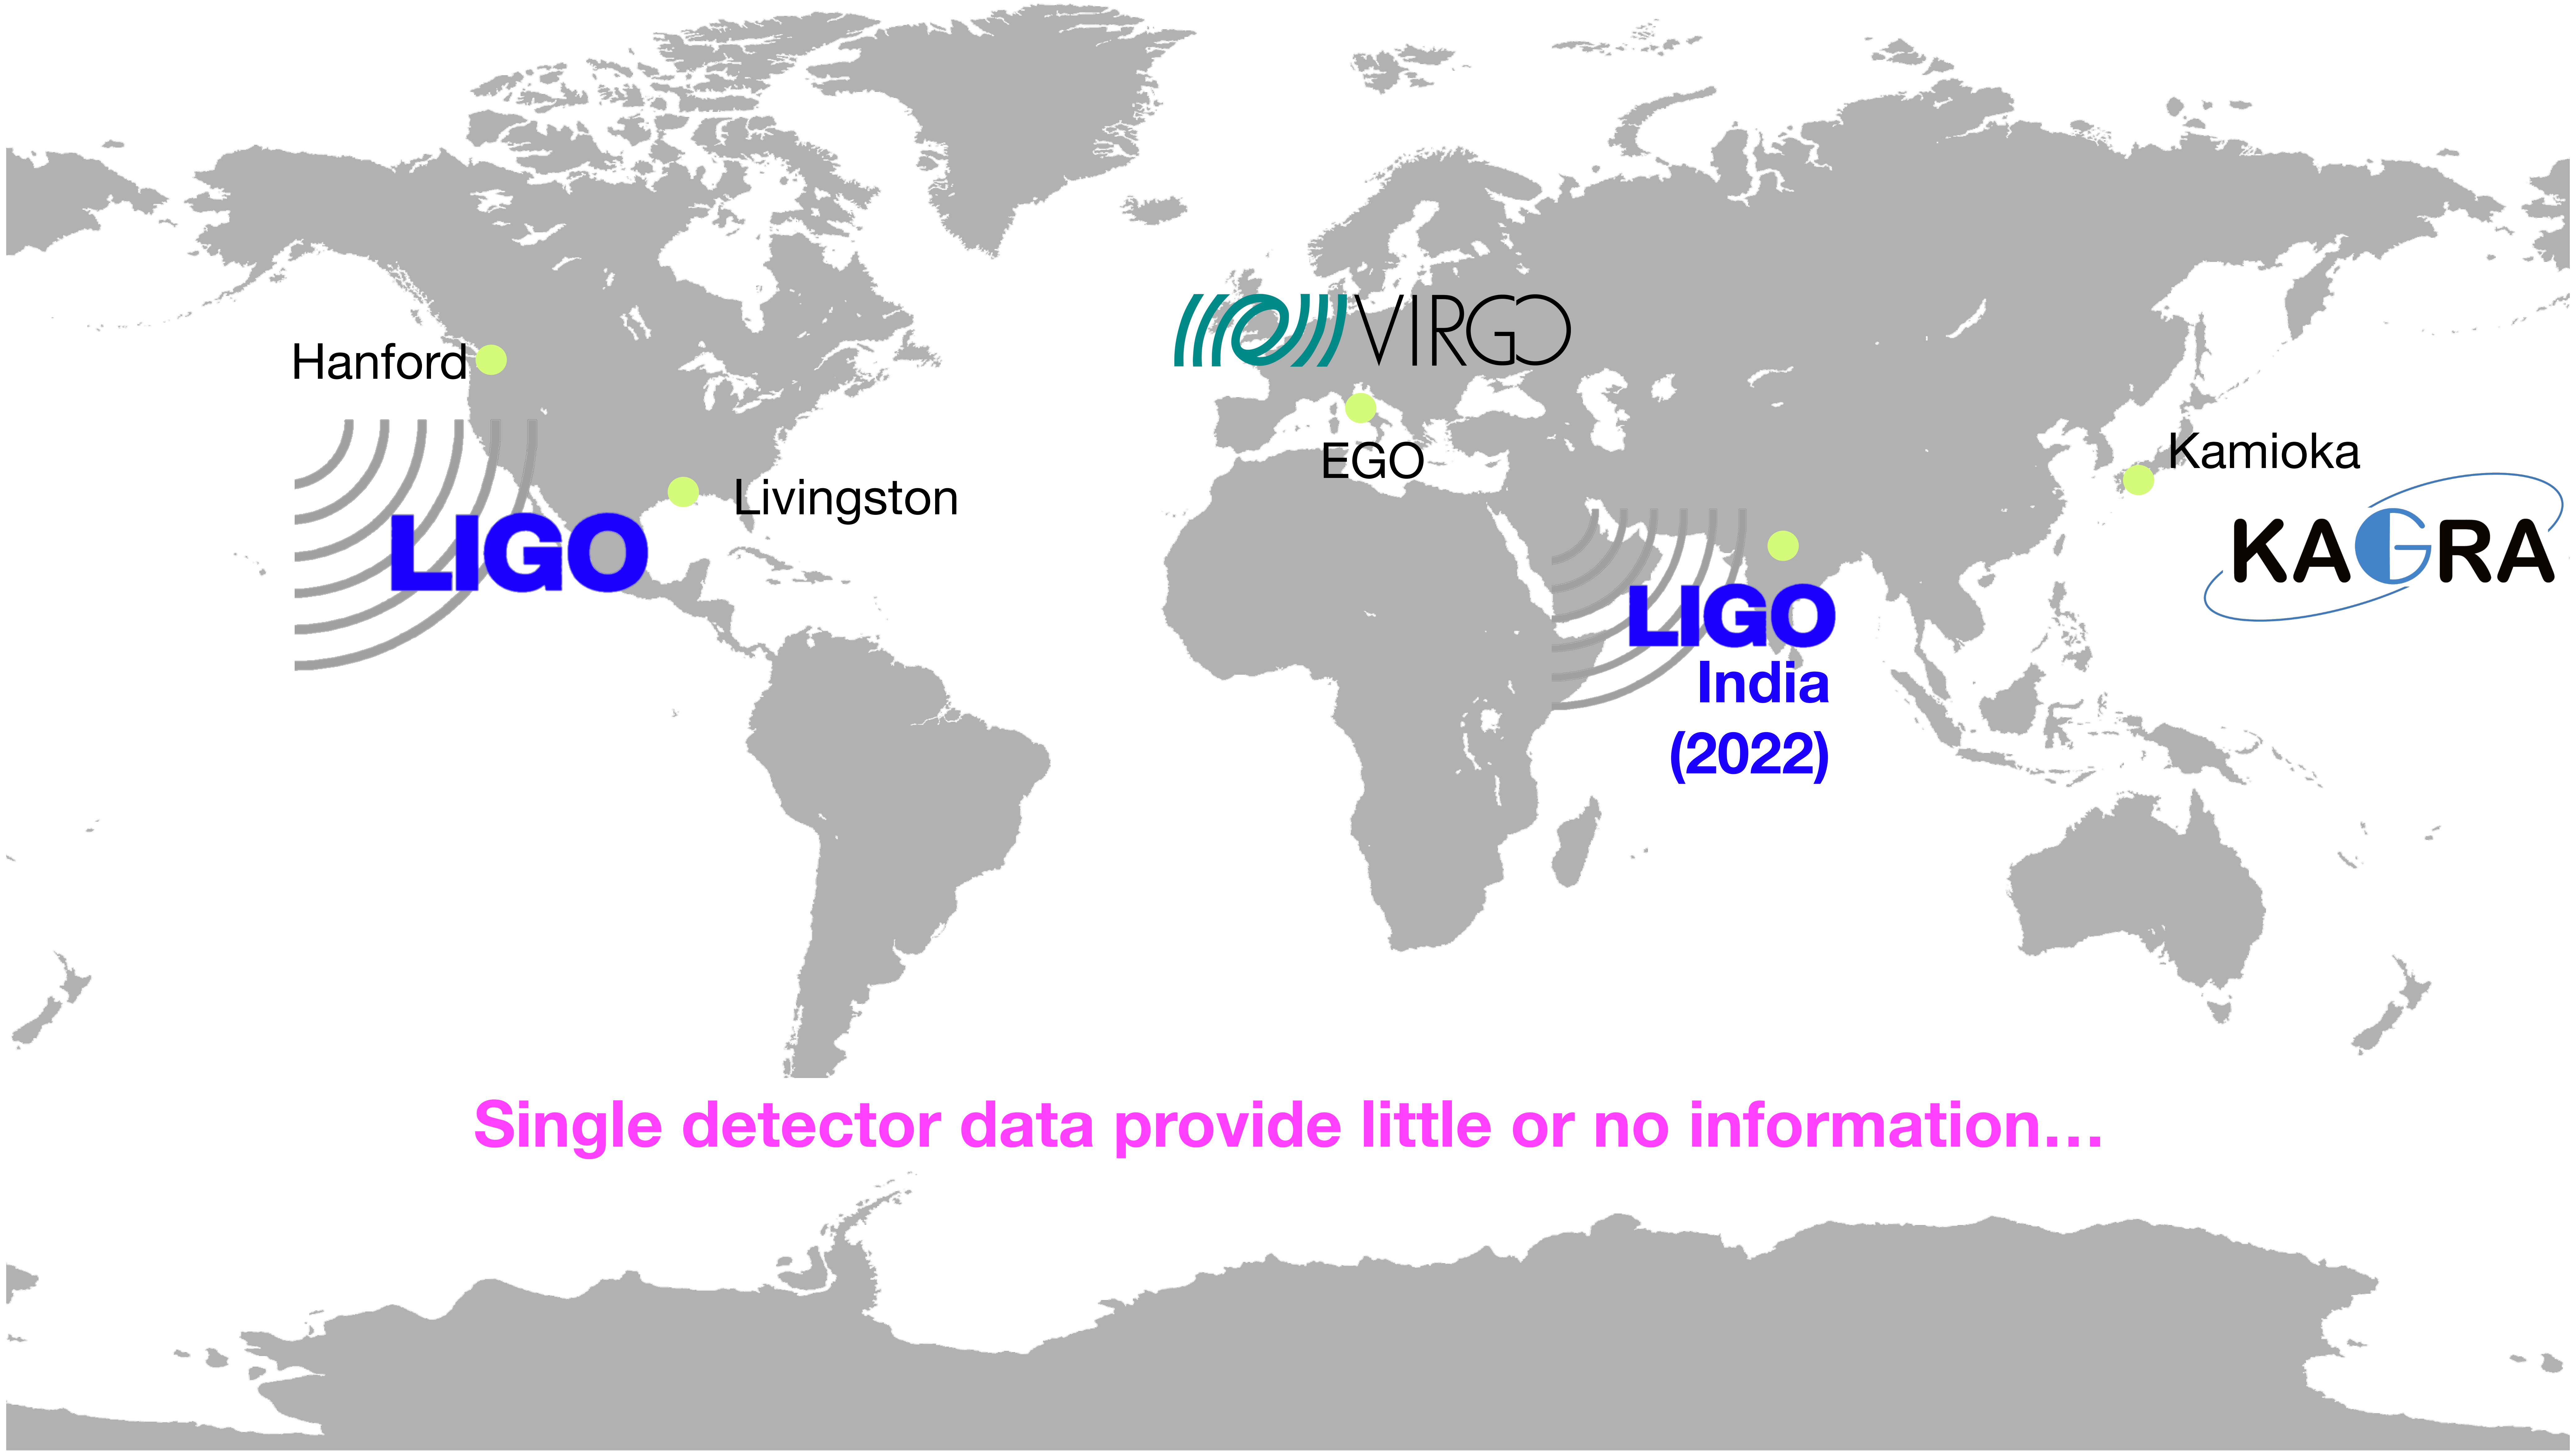

Single detector data provide little or no information…

Array of detectors to localize the GW source (triangulation) and rule out noise.

S. Bagnasco, S. Vallero

What are we looking for?

Burst sources

CBC: Compact Binary Coalescence

- Coalescing Compact Binary Systems (NS-NS, BH-NS, BH-BH): strong emitters, well modeled

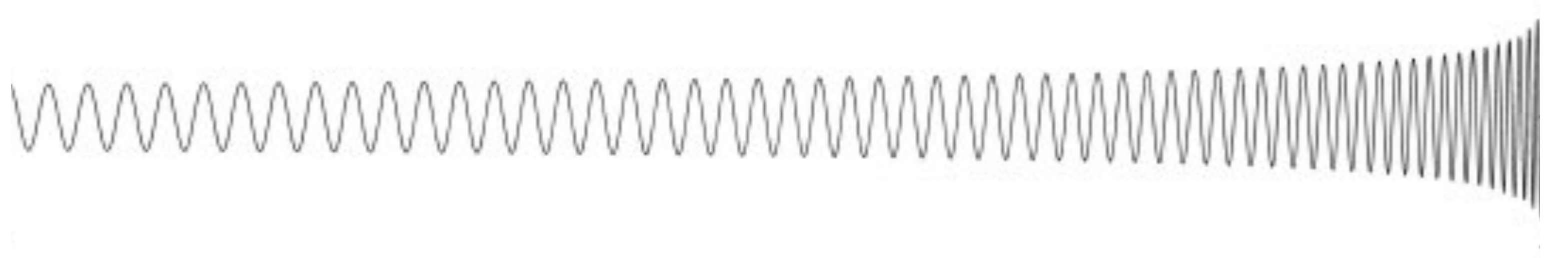

Low Latency analysis targets

Burst: Unmodeled transient bursts

- Asymmetric Core Collapse Supernovae: weak emitters, not well-modeled (‘bursts’), transient
- Cosmic strings, soft gamma repeaters, pulsar glitches,...
- Who knows?

Continuous sources

SGWB: Continuous stochastic background

- Cosmological stochastic background (residue of the Big Bang, cosmic GW background, long duration)
- Astrophysical stochastic background

CW: Continuous waves

- Spinning neutron stars (known waveform, long/continuous duration)
- All-sky and targeted searches

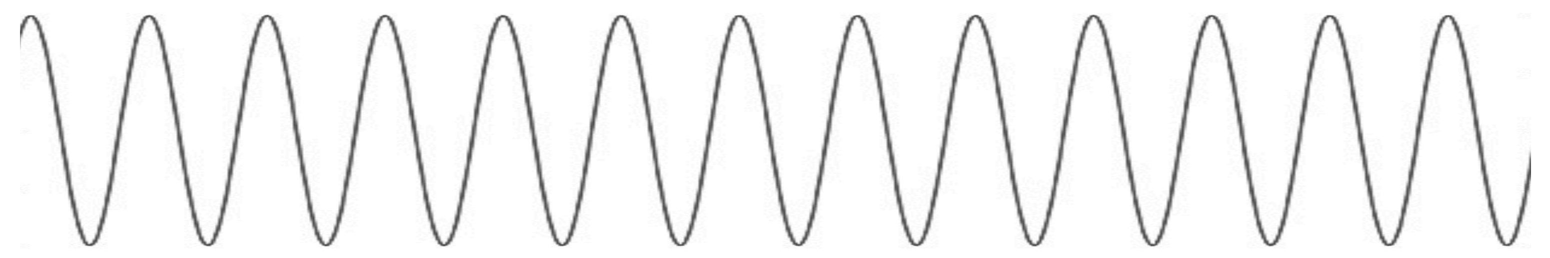

So many waveforms

BNS: Binary Neutron Stars

BBH: Binary Black Holes

IMBH: Intermediate Mass Black Holes

eBBH: Eccentric Binary Black Holes

CCSN: Core Collapse Supernovae

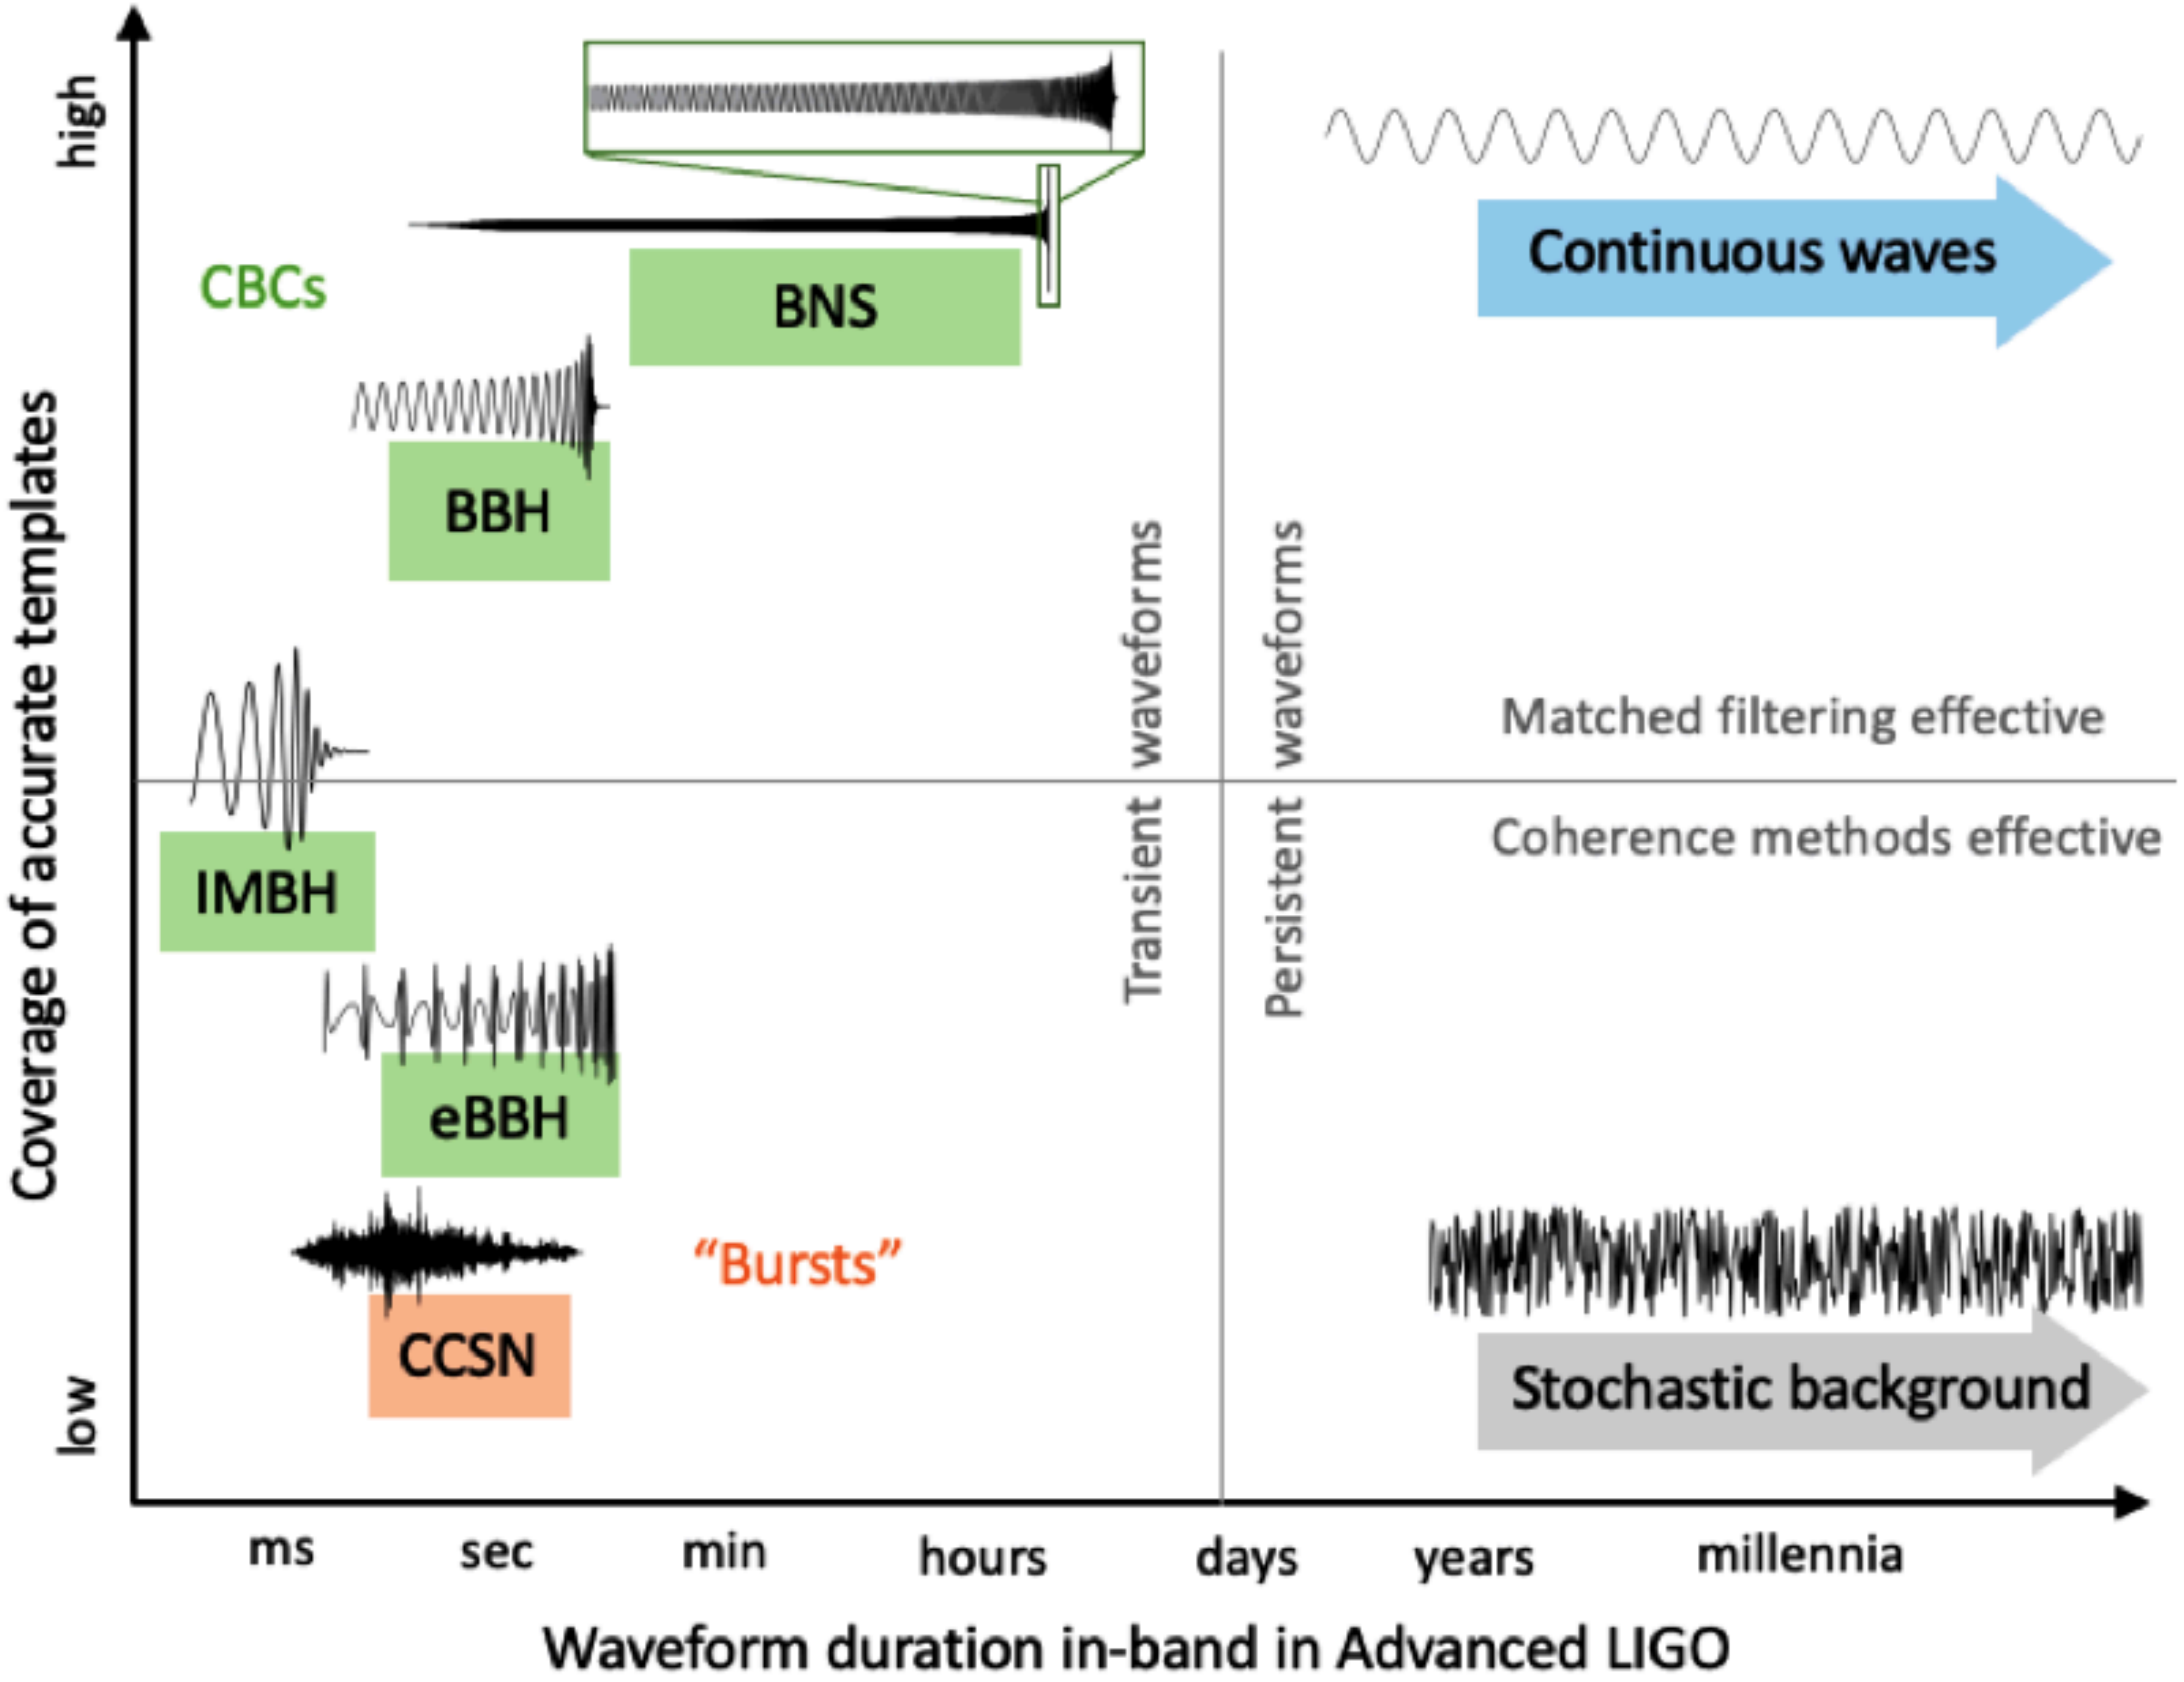

Matched filtering:

Compare detector data to predicted GW signals from general relativity (templates). Cross correlation evaluated as signal-to-noise ratio (SNR).

Coherence methods:

Correlate data from different detectors.

Jess McIver, D. H. Shoemaker, *Discovering Gravitational Waves with Advanced Ligo,* LIGO Document P2000530-v1 (2000)

Low Latency searches today

CCR Workshop May 2021
S. Bagnasco, S. Vallero

GW CANDIDATE

SKY LOCALIZATION

PUBLIC ALERT

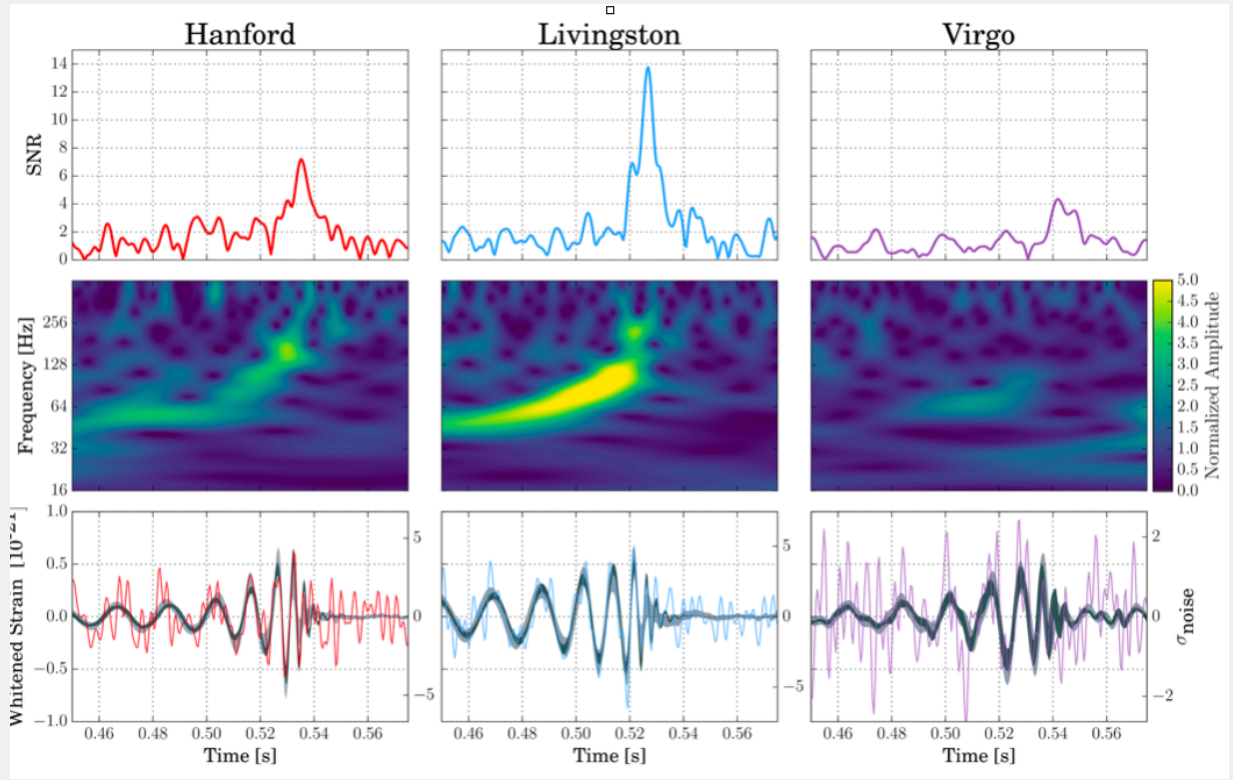

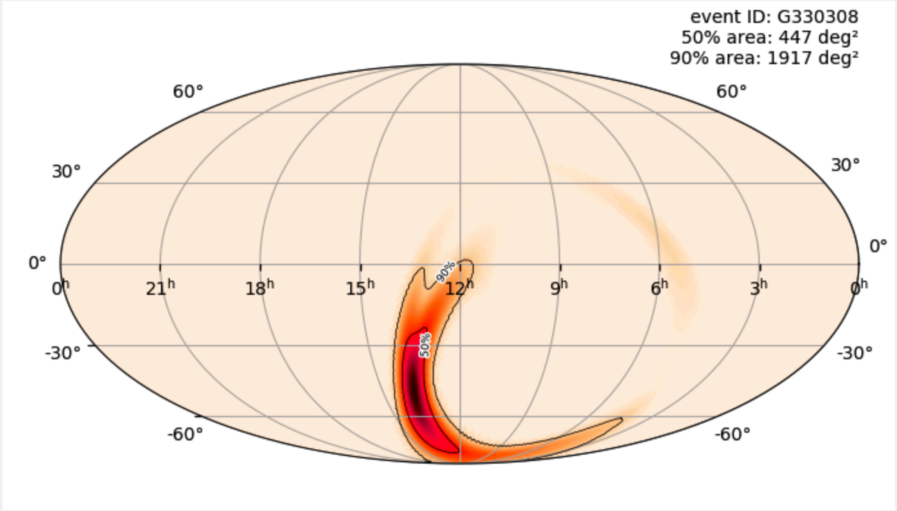

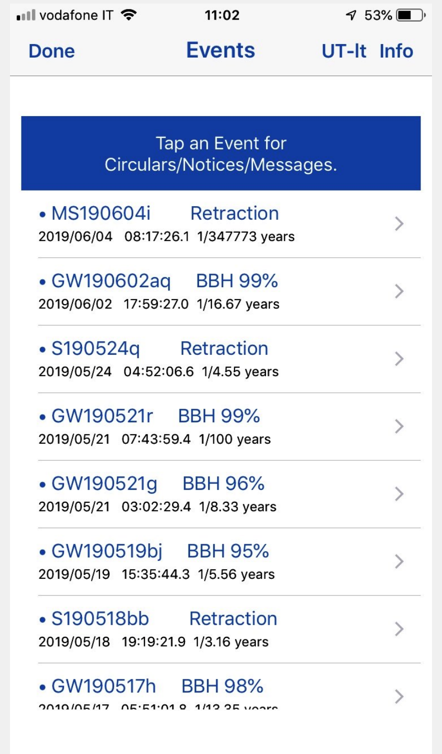

Low Latency searches

- Detector sanity
- Data quality
- Localization
- …

Candidate event validation

On Site

A few minutes → 1/2 hour

Off Site

- Parameter estimation
- GW Candidate update

Hours, days

Multi-messenger astronomy

MMA is the study of astronomical sources using different types of *messenger* particles:

- photons
- neutrinos (2013)
- cosmic rays
- gravitational waves (2015)

It's a difficult task… no other event with EM counterpart observed after the one in 2017

→ provided **useful information on heavy elements formation**.

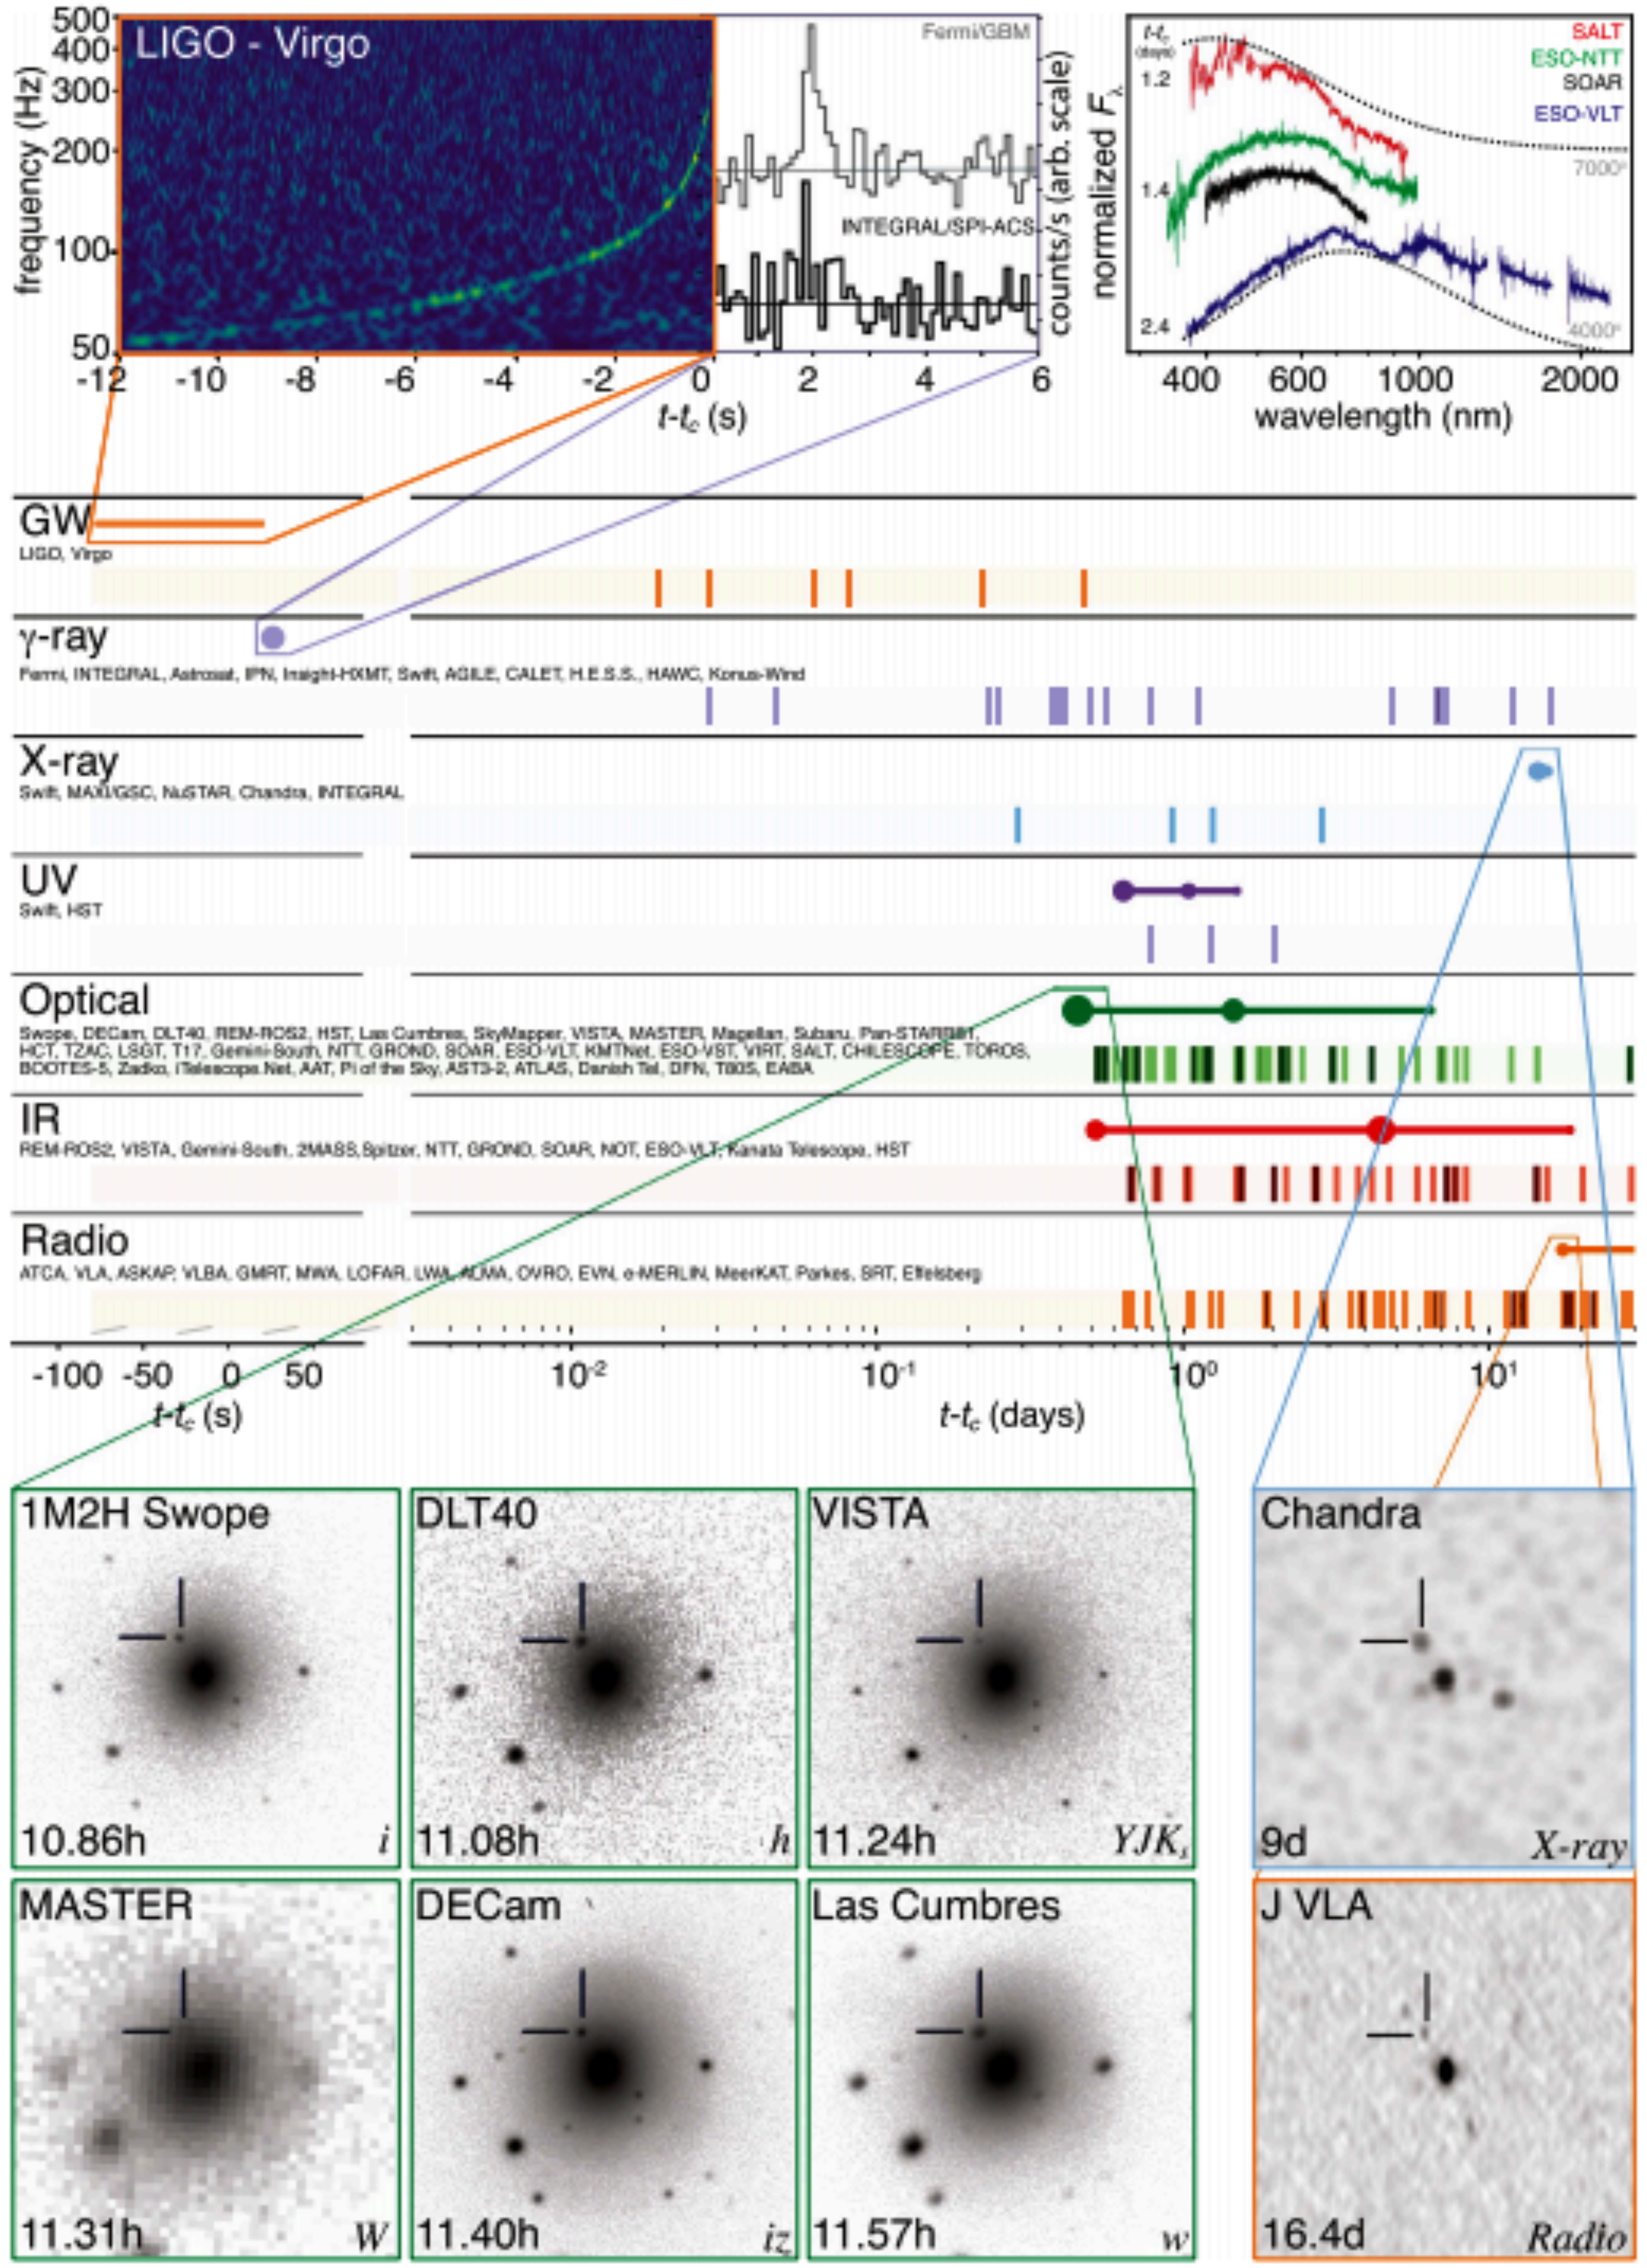

Study a broad energy window and provide complementary physics insight:

- **GW and neutrinos** carry information from the inner regions of the astrophysical sources
- **Photons** can give a precise localization, gaining sensitivity in spin and source mass measurements.
- The presence or absence of **EM emission** can be a tool to constrain the NS equation of state
- …

M. Branchesi, *Multi-messenger astronomy: gravitational waves, neutrinos, photons, and cosmic rays*
2016 J. Phys.: Conf. Ser. **718** 022004

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

We need to be faster

CCR Workshop May 2021
S. Bagnasco, S. Vallero

If we catch the signal in this phase…

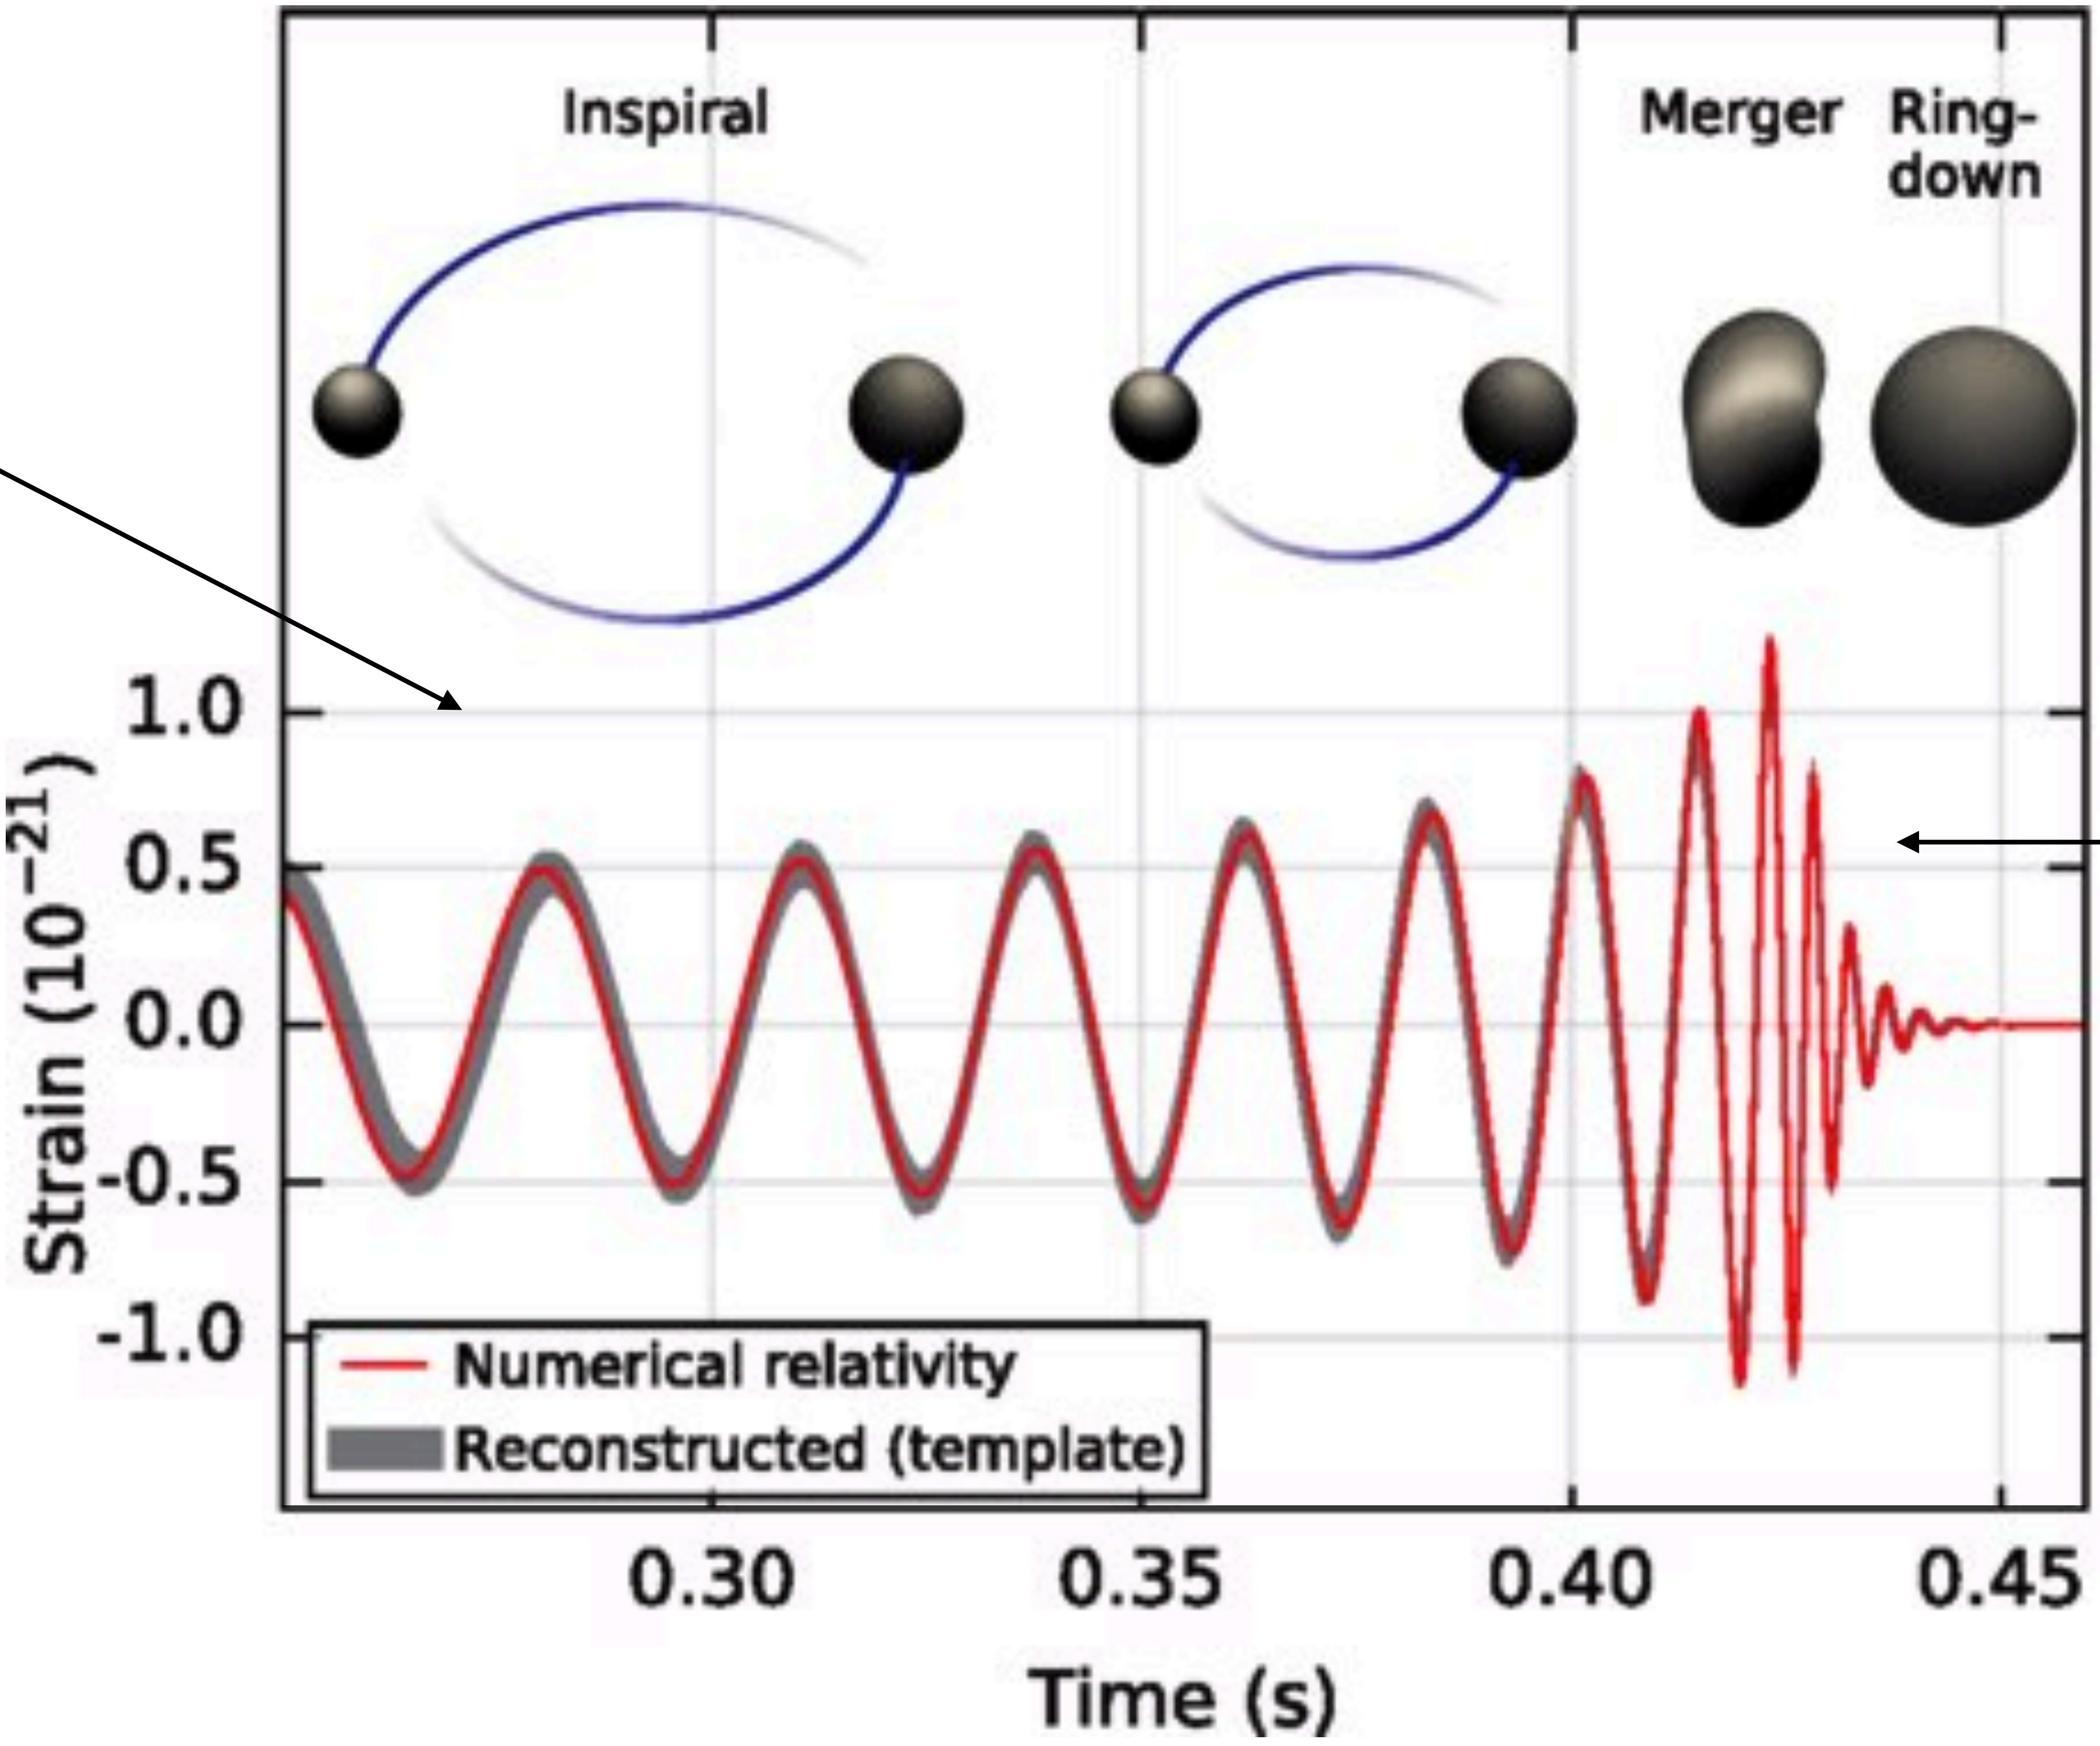

… we can provide a negative latency trigger for the merger phase.

B.P. Abbott et al. (LIGO Scientific Collaboration and Virgo Collaboration), *Observation of Gravitational Waves from a Binary Black Hole Merger,* Phys. Rev. Lett. **116**, 061102

Observing runs

AdV: o(10^2)/yr

AdV+ Phase I: o(10^3)/yr
(Quantum and technical noise reduction)

AdV+ Phase II: o(10^4)/yr
(Thermal noise reduction, mirrors replacement)

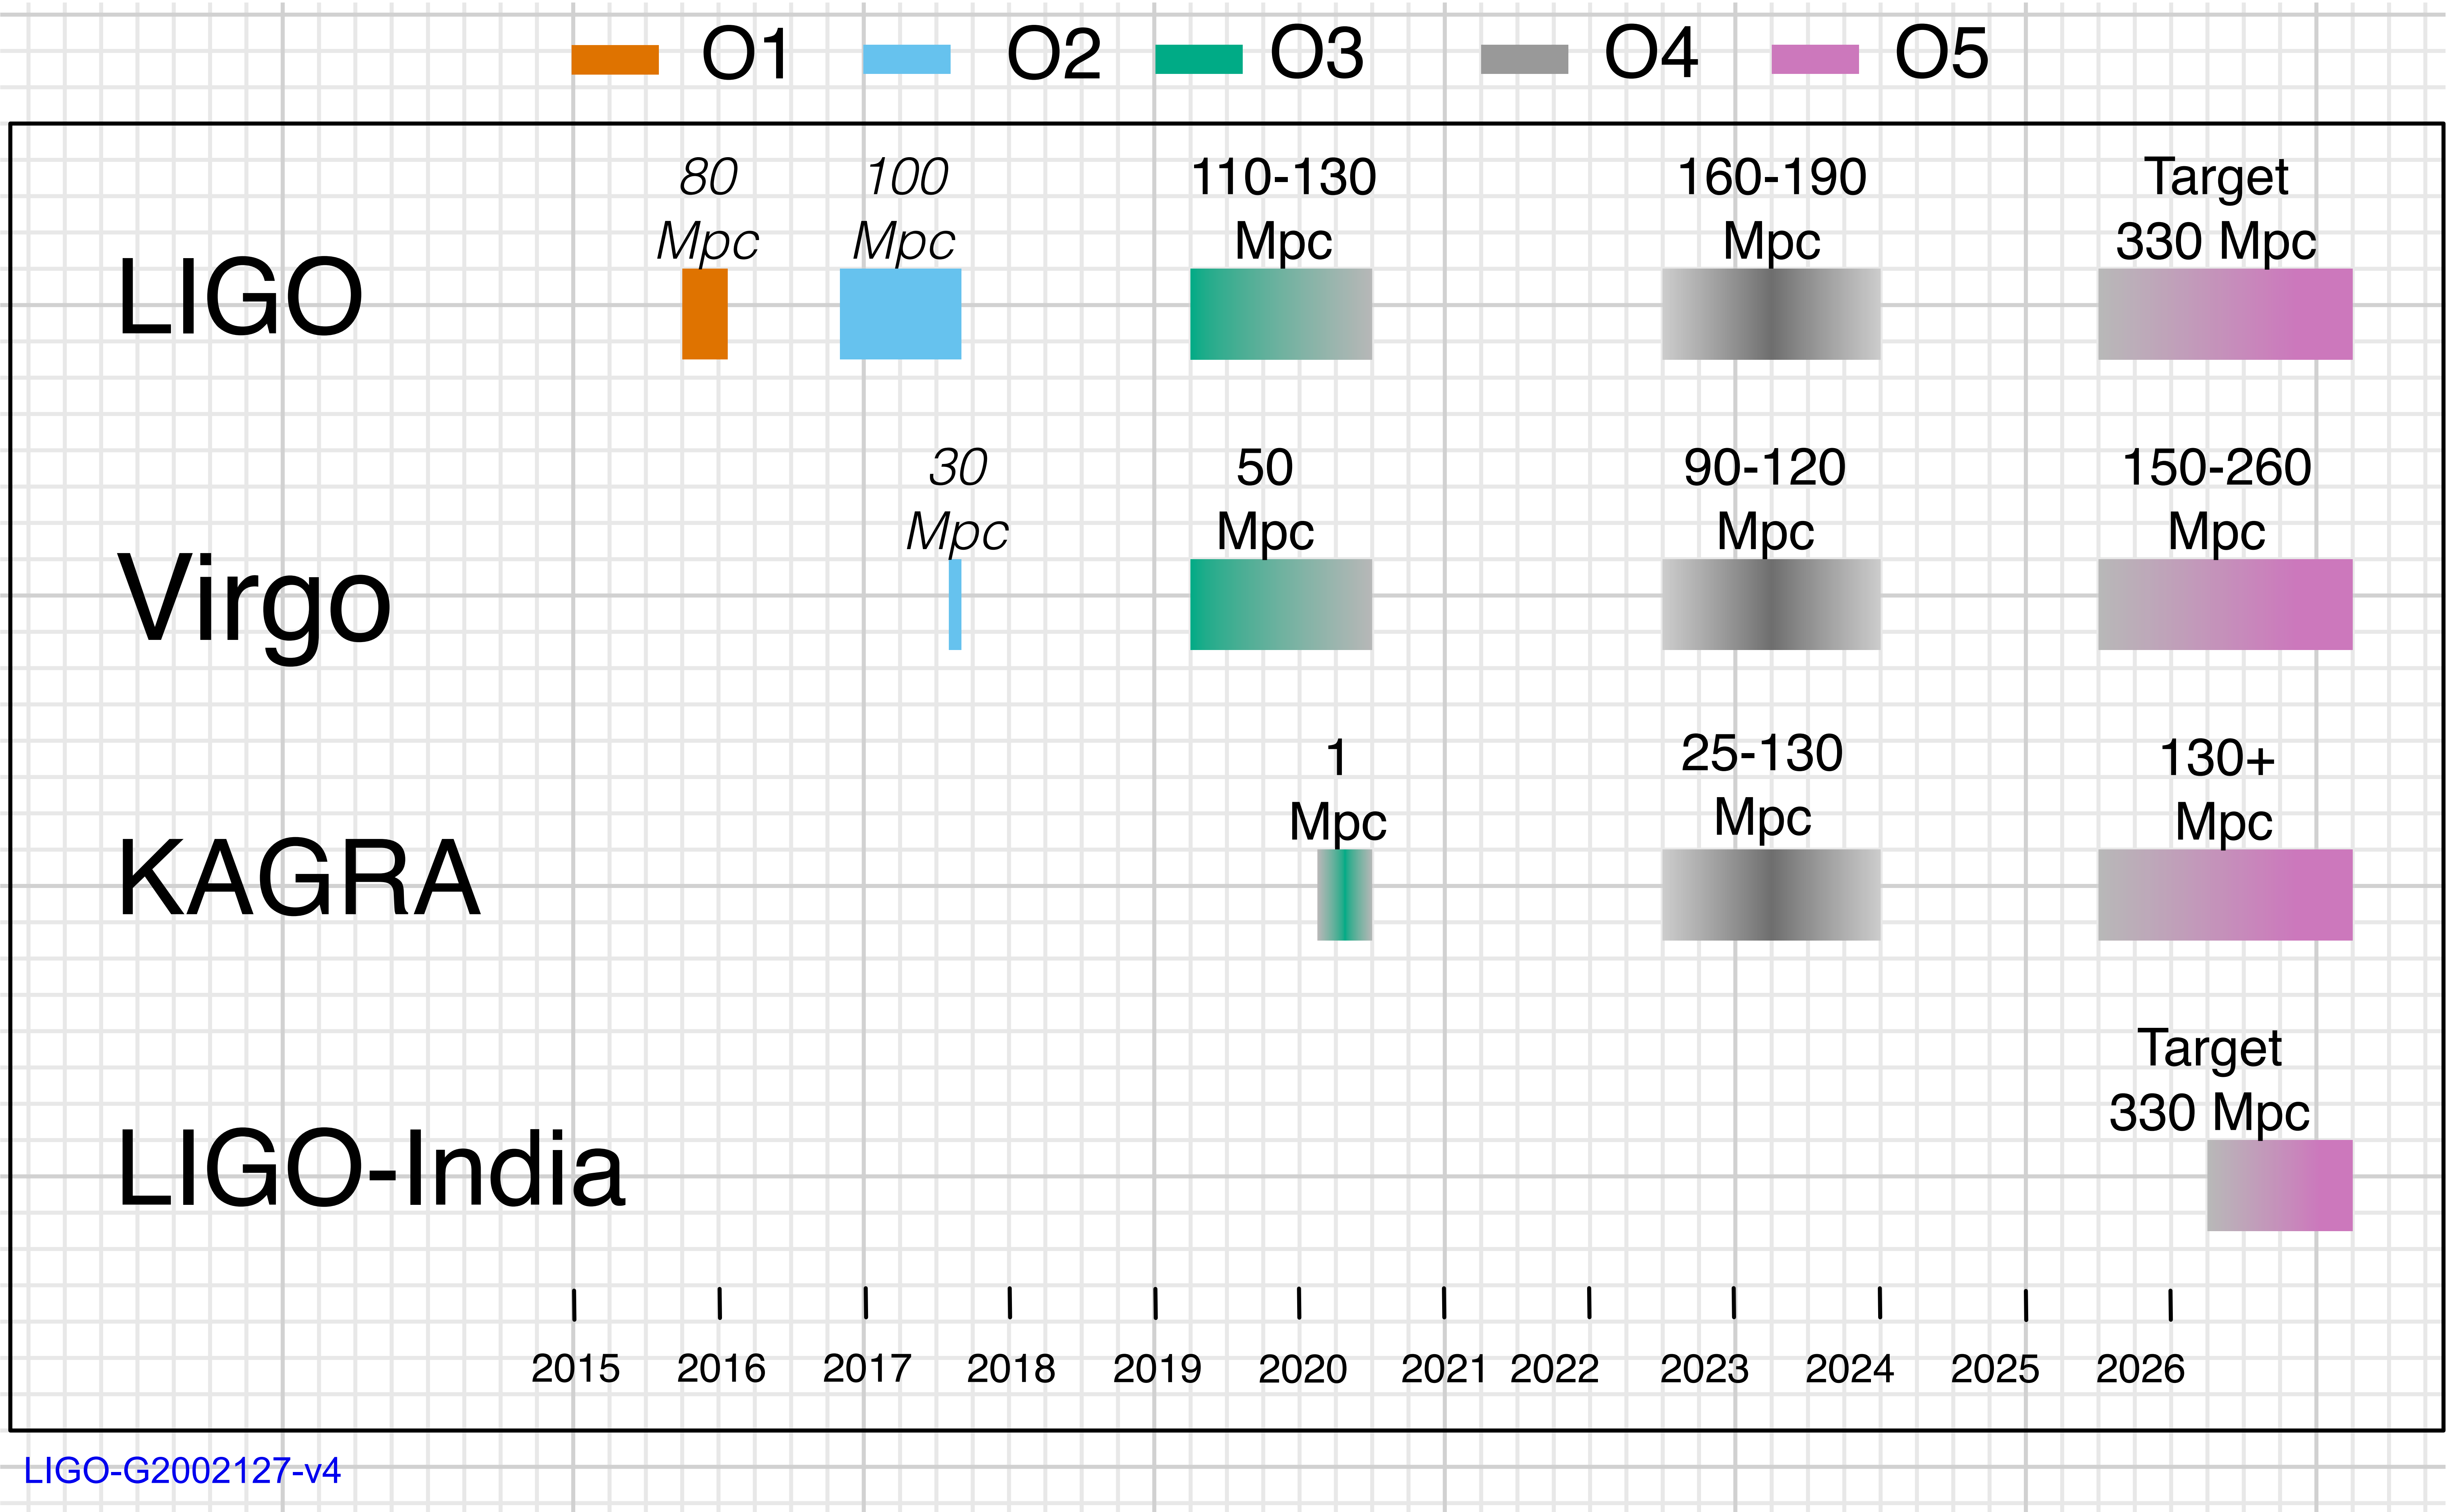

The number of events is proportional to the covered volume (the third power of the range).

O3 highlights

GWOSC: Gravitational Wave Open Science Center
GWTC: Gravitational Wave Transient Catalog

O3a (Apr 1 – Oct 1, 2019)

- GWTC-2 catalog paper, 39 candidate events
- data publicly available in GWOSC since May 1st

O3b (Nov 1, 2019 – Mar 27, 2020)

- One month earlier than scheduled due to pandemics
- GWTC-3 catalog paper (not yet published)

Fraction of CPU power used by different analyses (online+ offline)

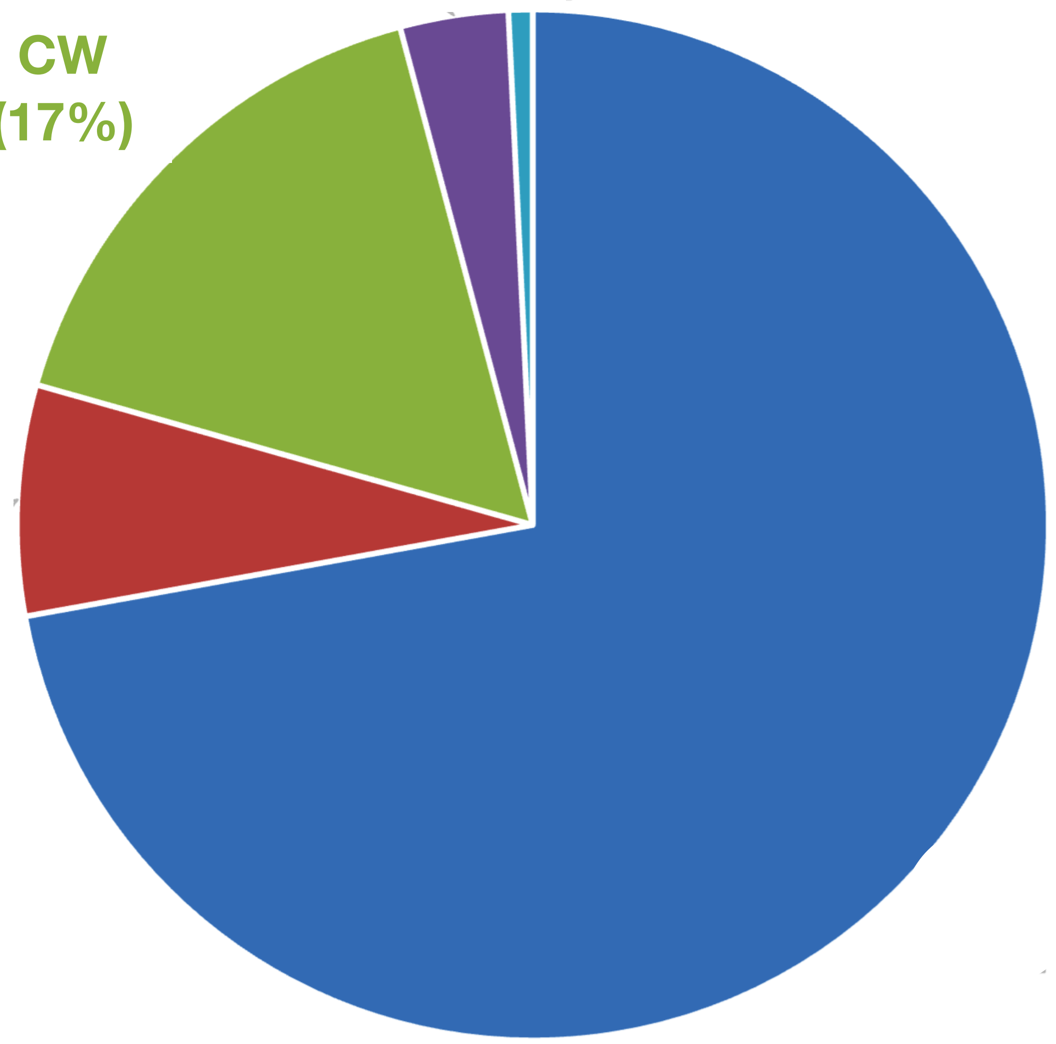

FEATURED EVENTS

- GW190425: large BNS coalescence

  Total mass exceeds that of known galactic neutron star binaries
- GW190412: asymmetric BBH coalescence

  First BBH detection with clear evidence for unequal-mass components
- GW190814: most asymmetric CBC

  Either the lightest black hole or the heaviest neutron star ever discovered in a double compact-object system
- GW190521: first Intermediate Mass BH

  The most massive gravitational wave binary observed to date

European computing centers

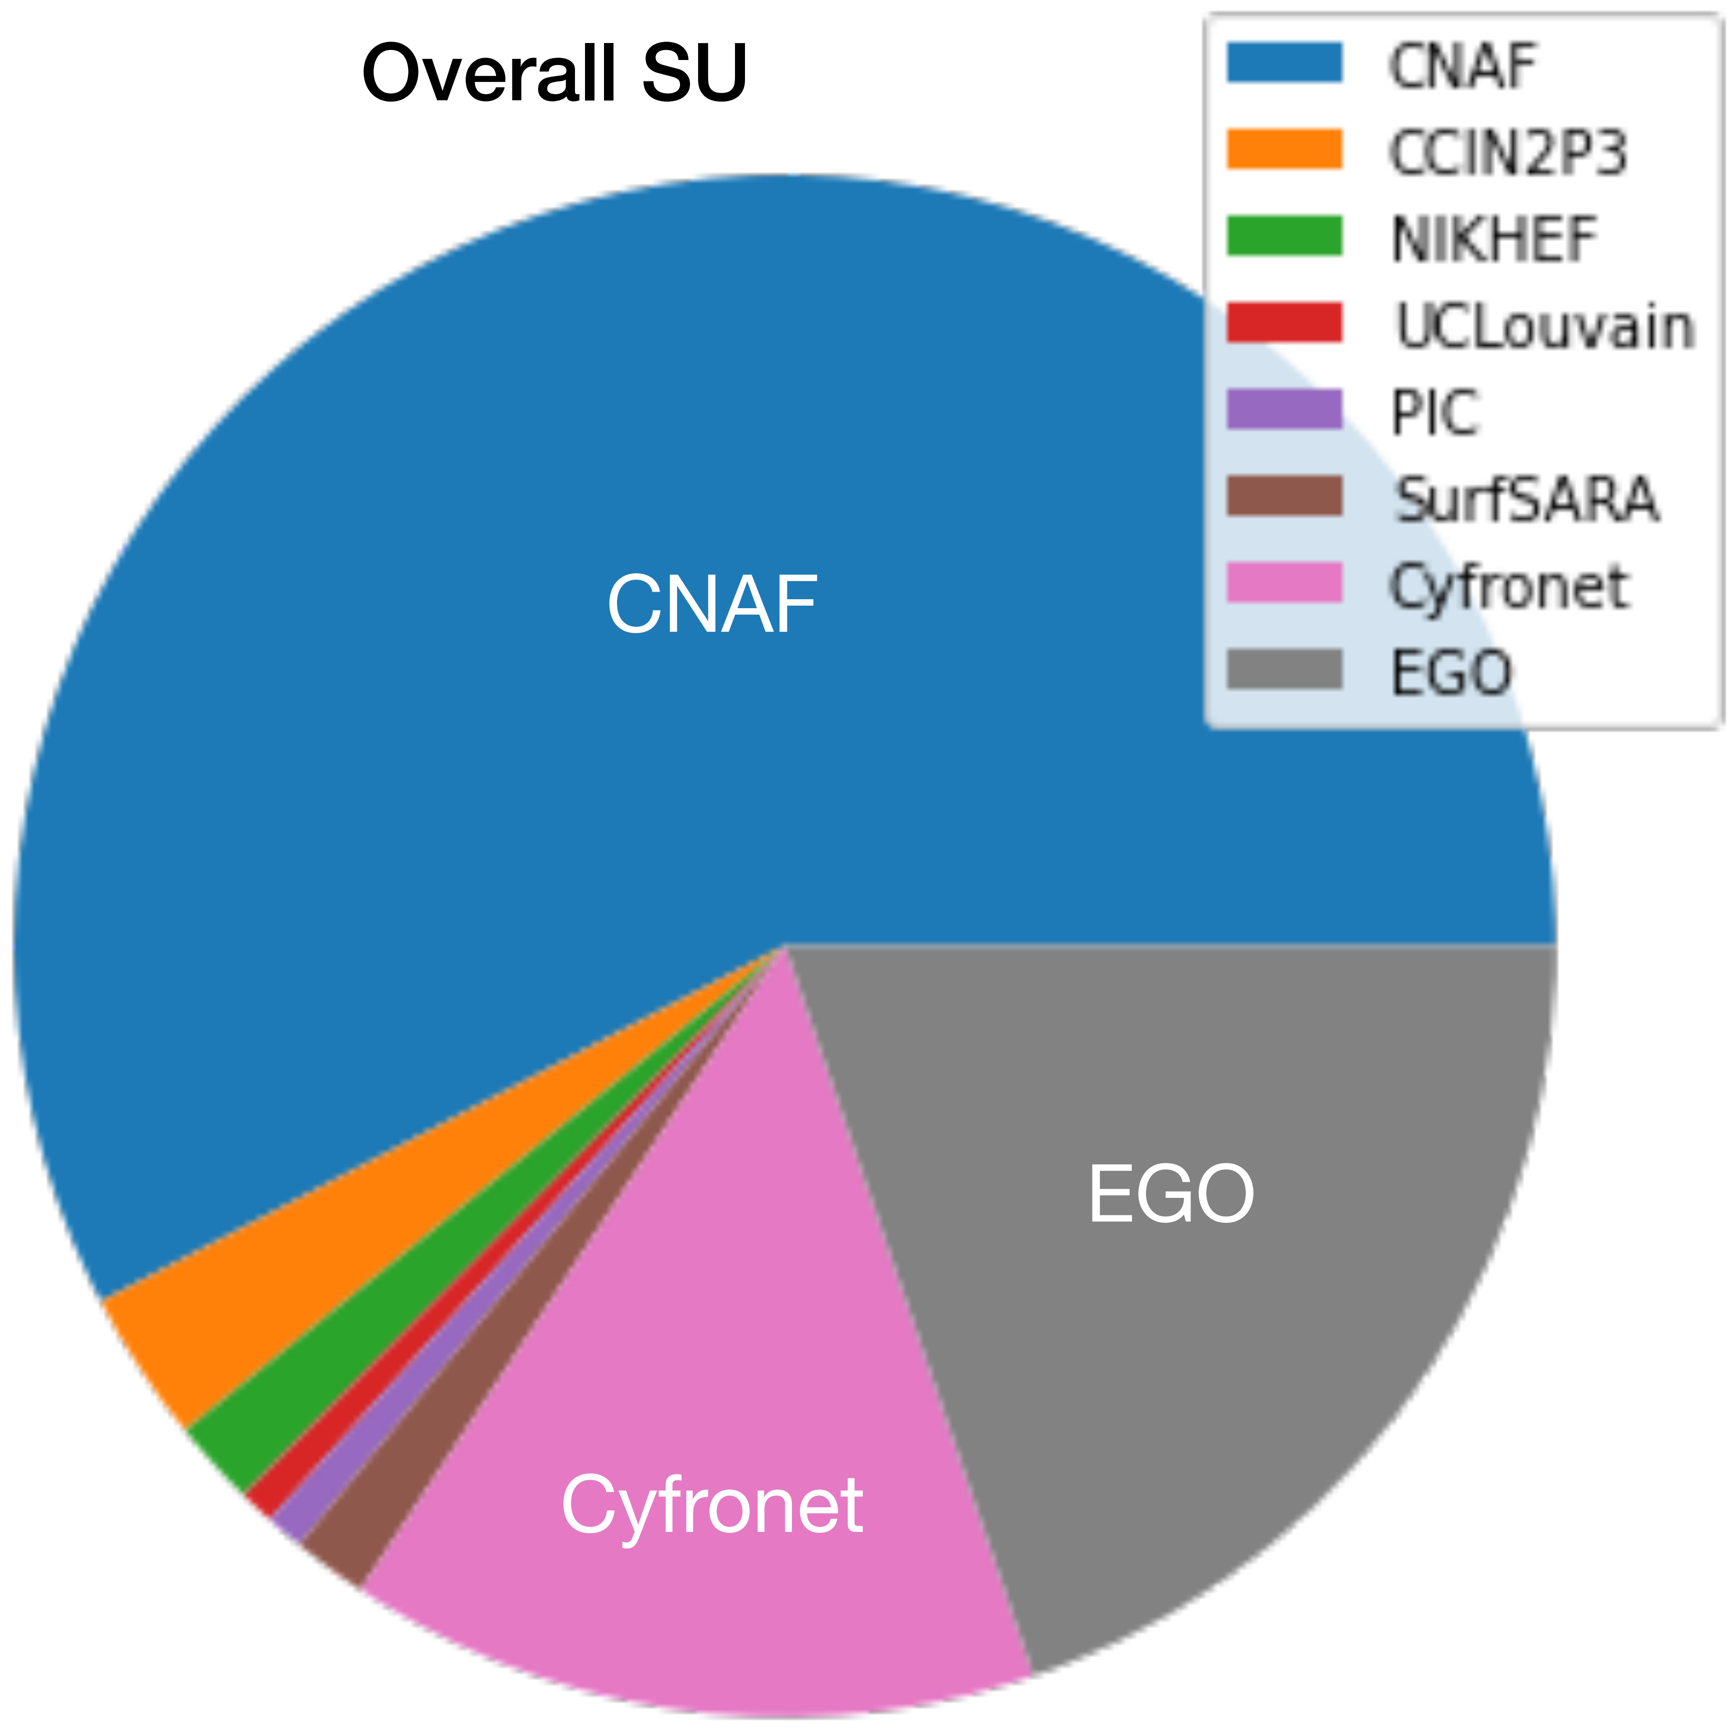

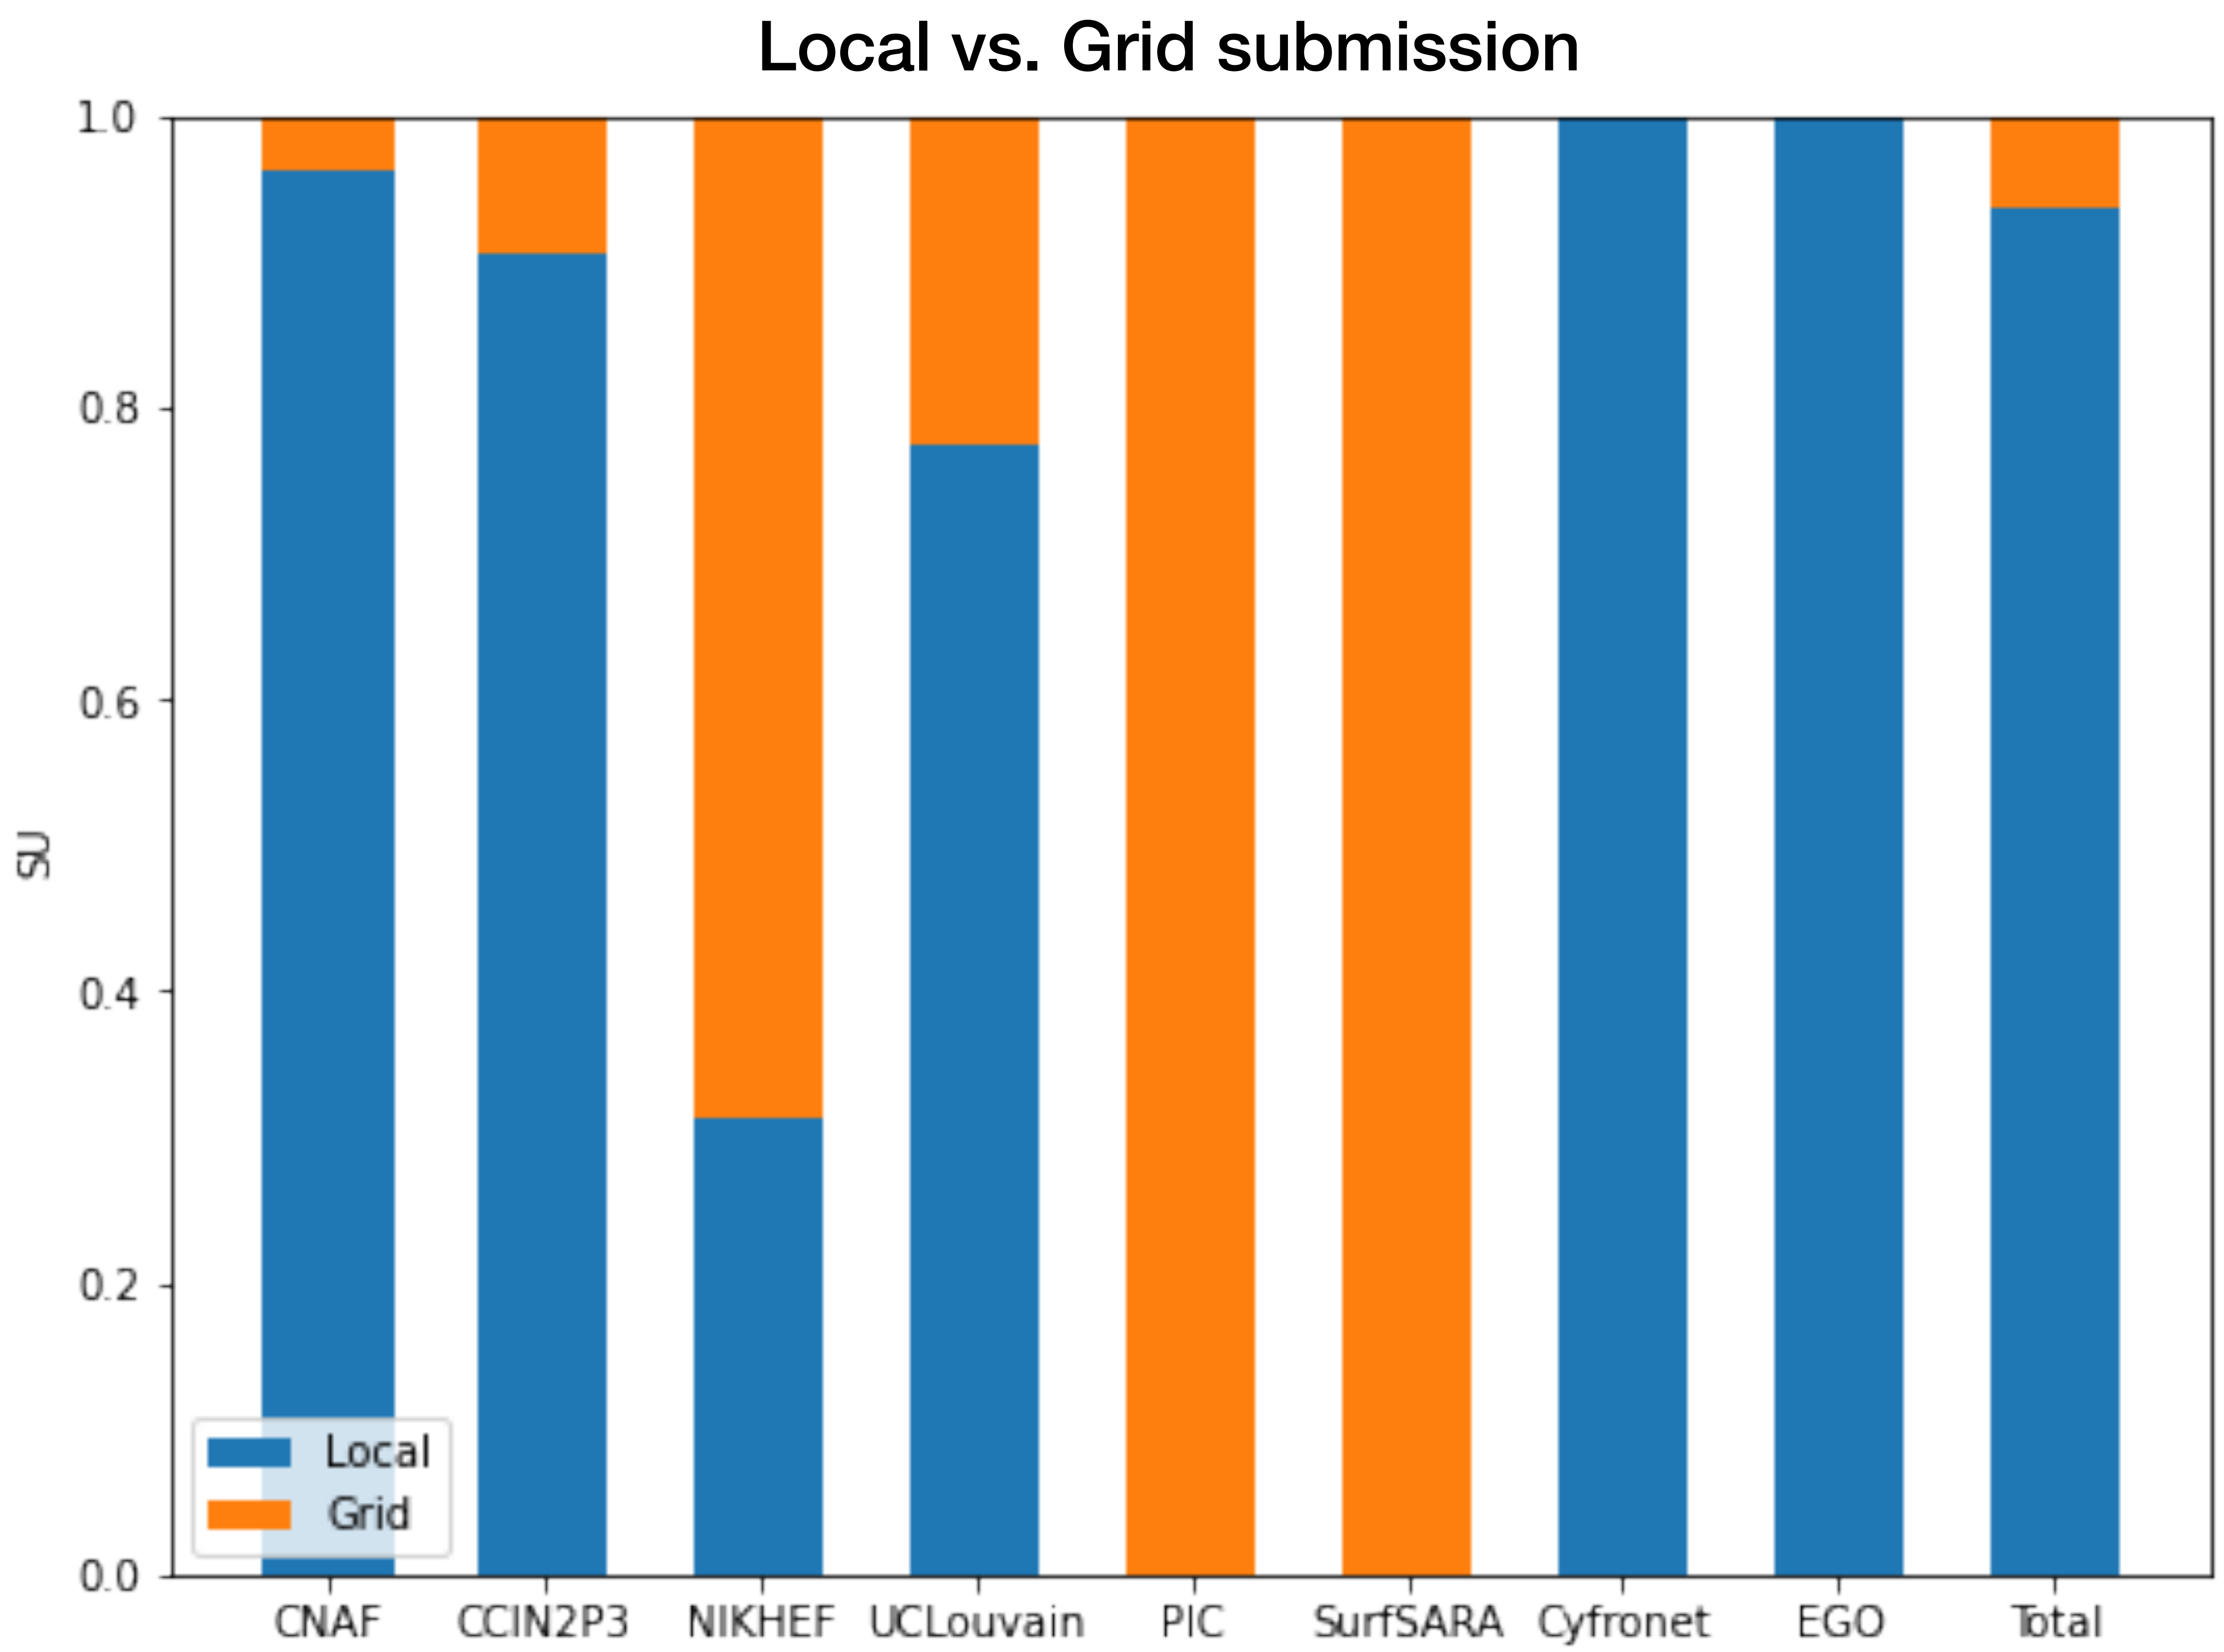

Local: direct submission
Grid: through GlideinWMS on WLCG sites

Towards O4

IGWN: International Gravitational Wave Observatory Network

THE MASTER PLAN: switch to a Distributed Computing Model

- LIGO, Virgo and KAGRA to join their computing resources in the common IGWN infrastructures
- Evolve online, offline and LL computing services to use (a small number of) common, modern, mainstream, widely used tools to reduce support burden
- Have analysis pipelines run on a uniform runtime environment, using common submission and data access interfaces and tools
- Scale up the exploitation of large-scale external shared computing resources
- Ramp up Virgo contribution (power, services, effort)

VIRGO contribution should be proportional to the collaboration size (30%), we are almost there considering only distributed computing

→ gradually move away from local submission

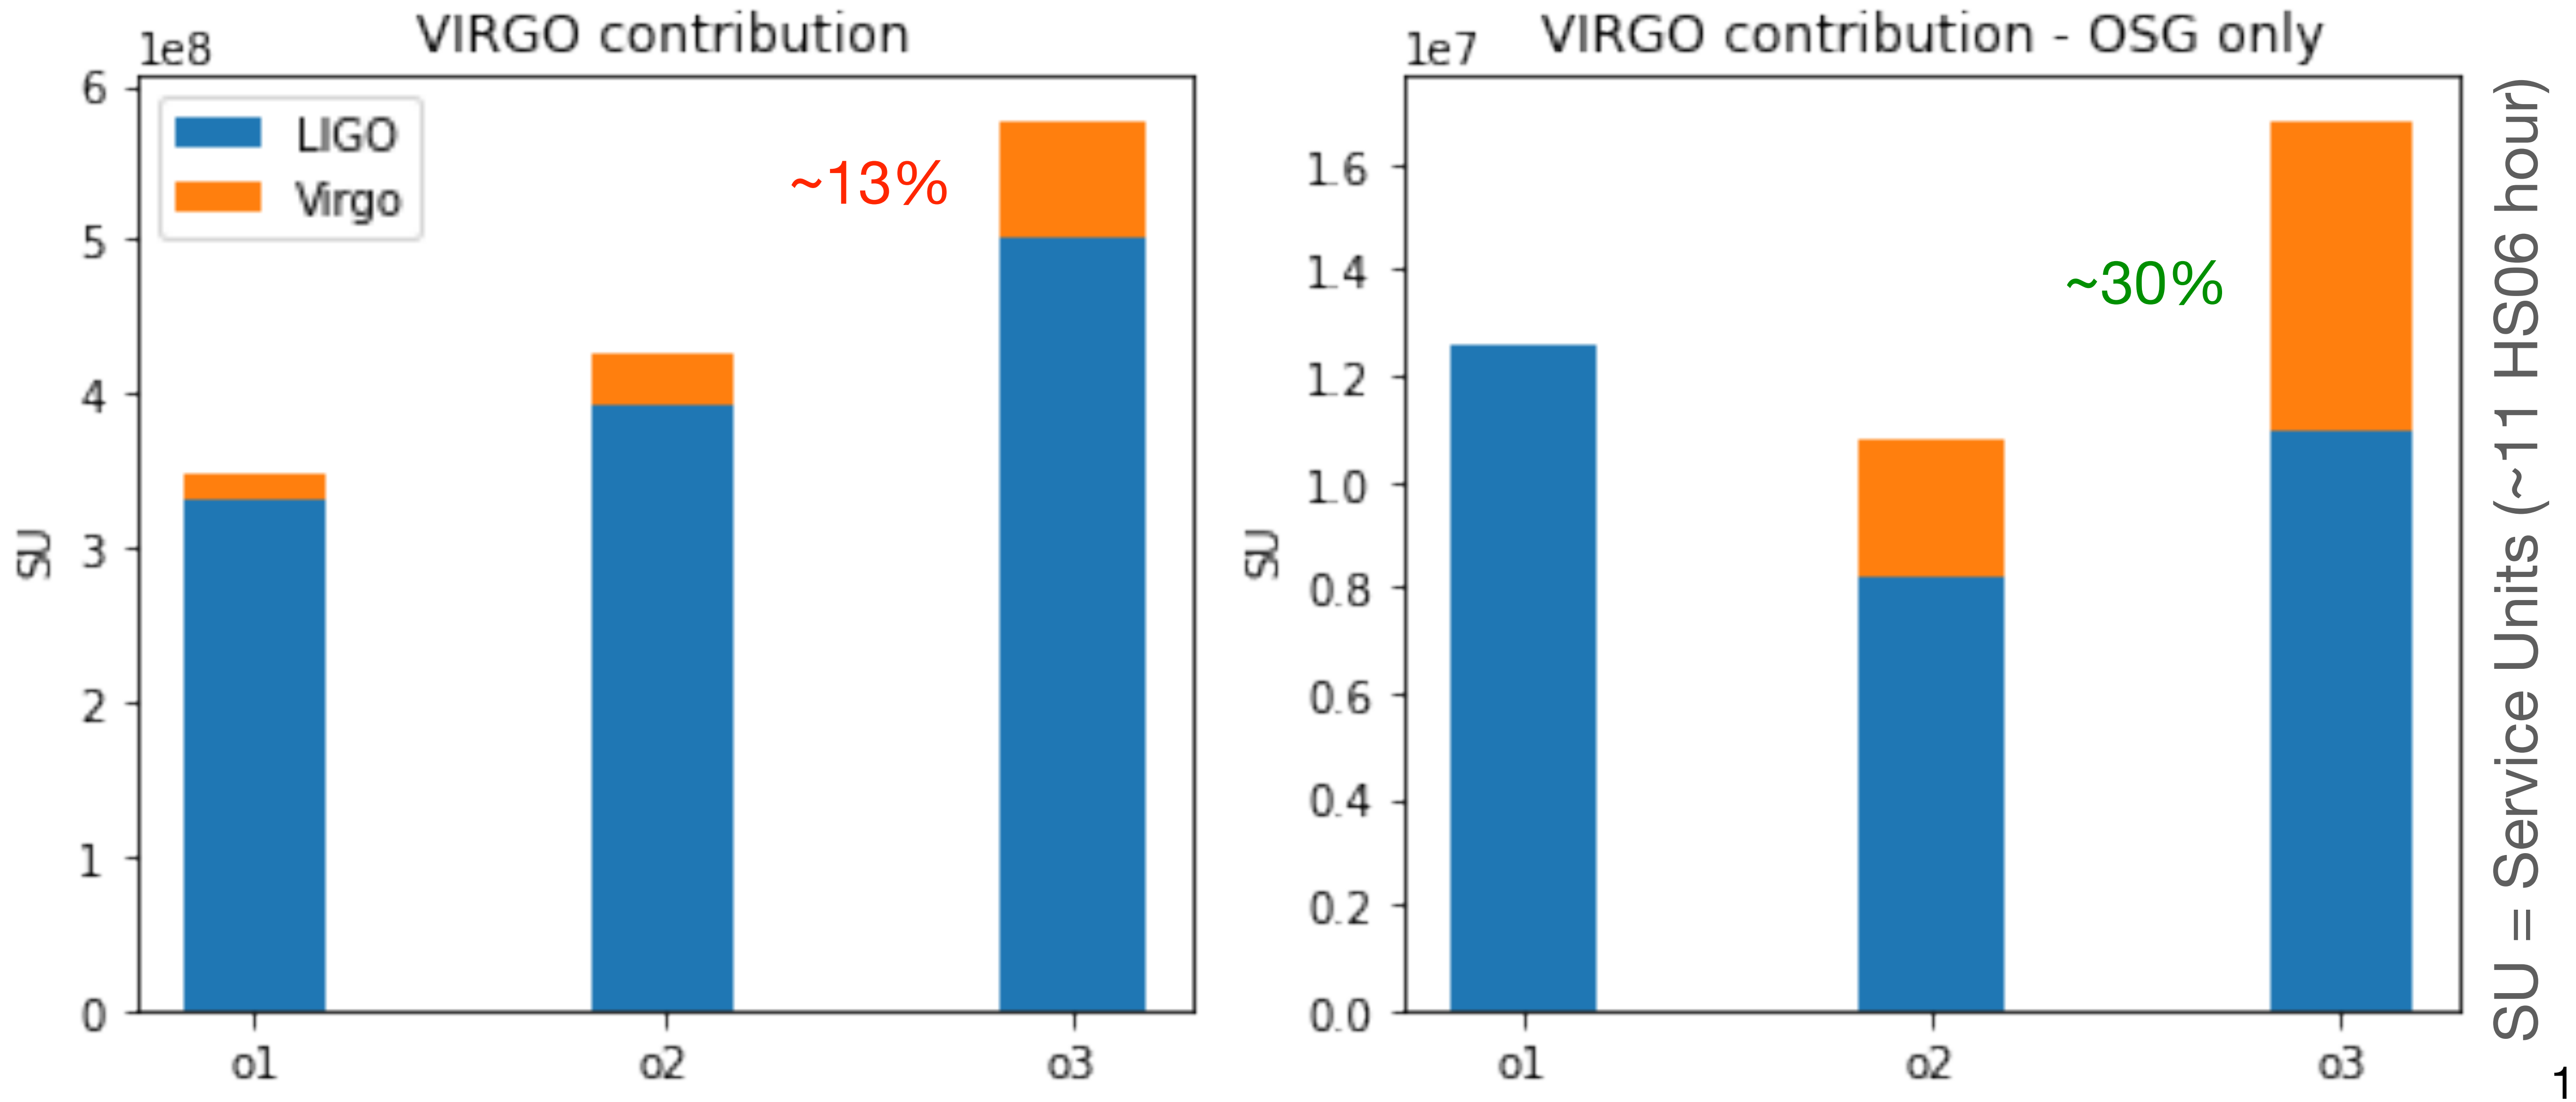

Overall data flows

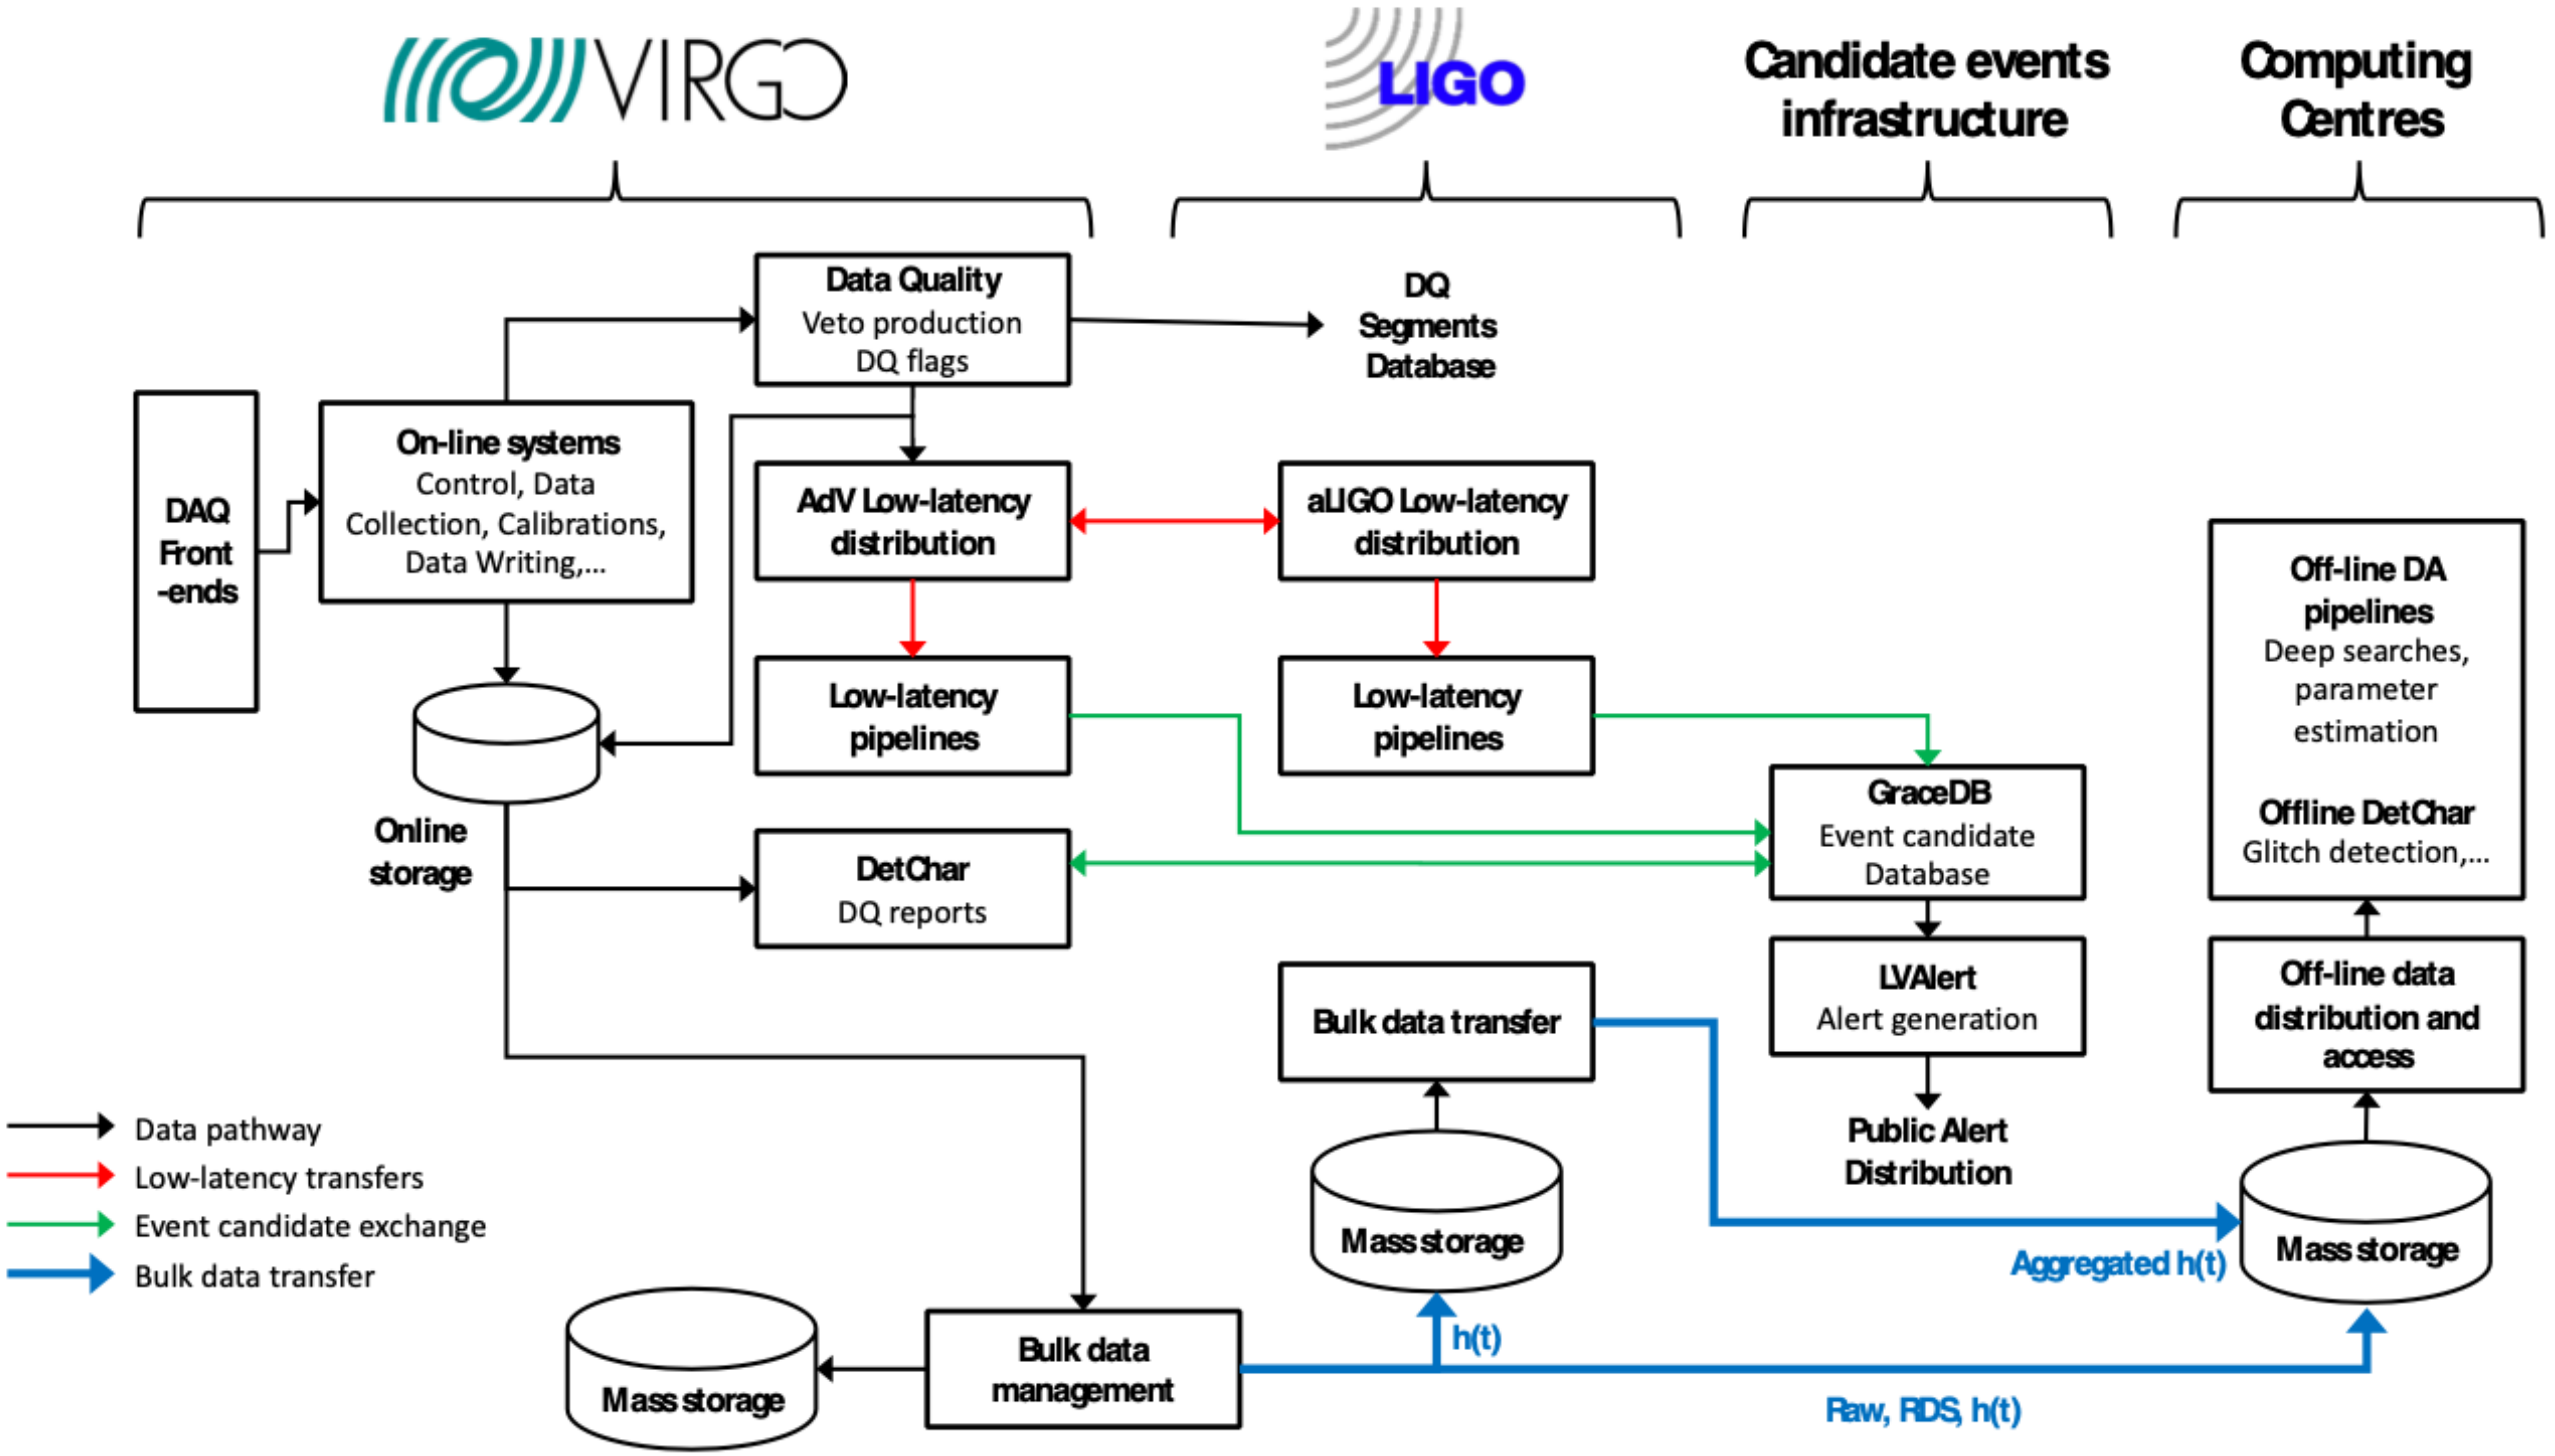

Offline data distribution: now

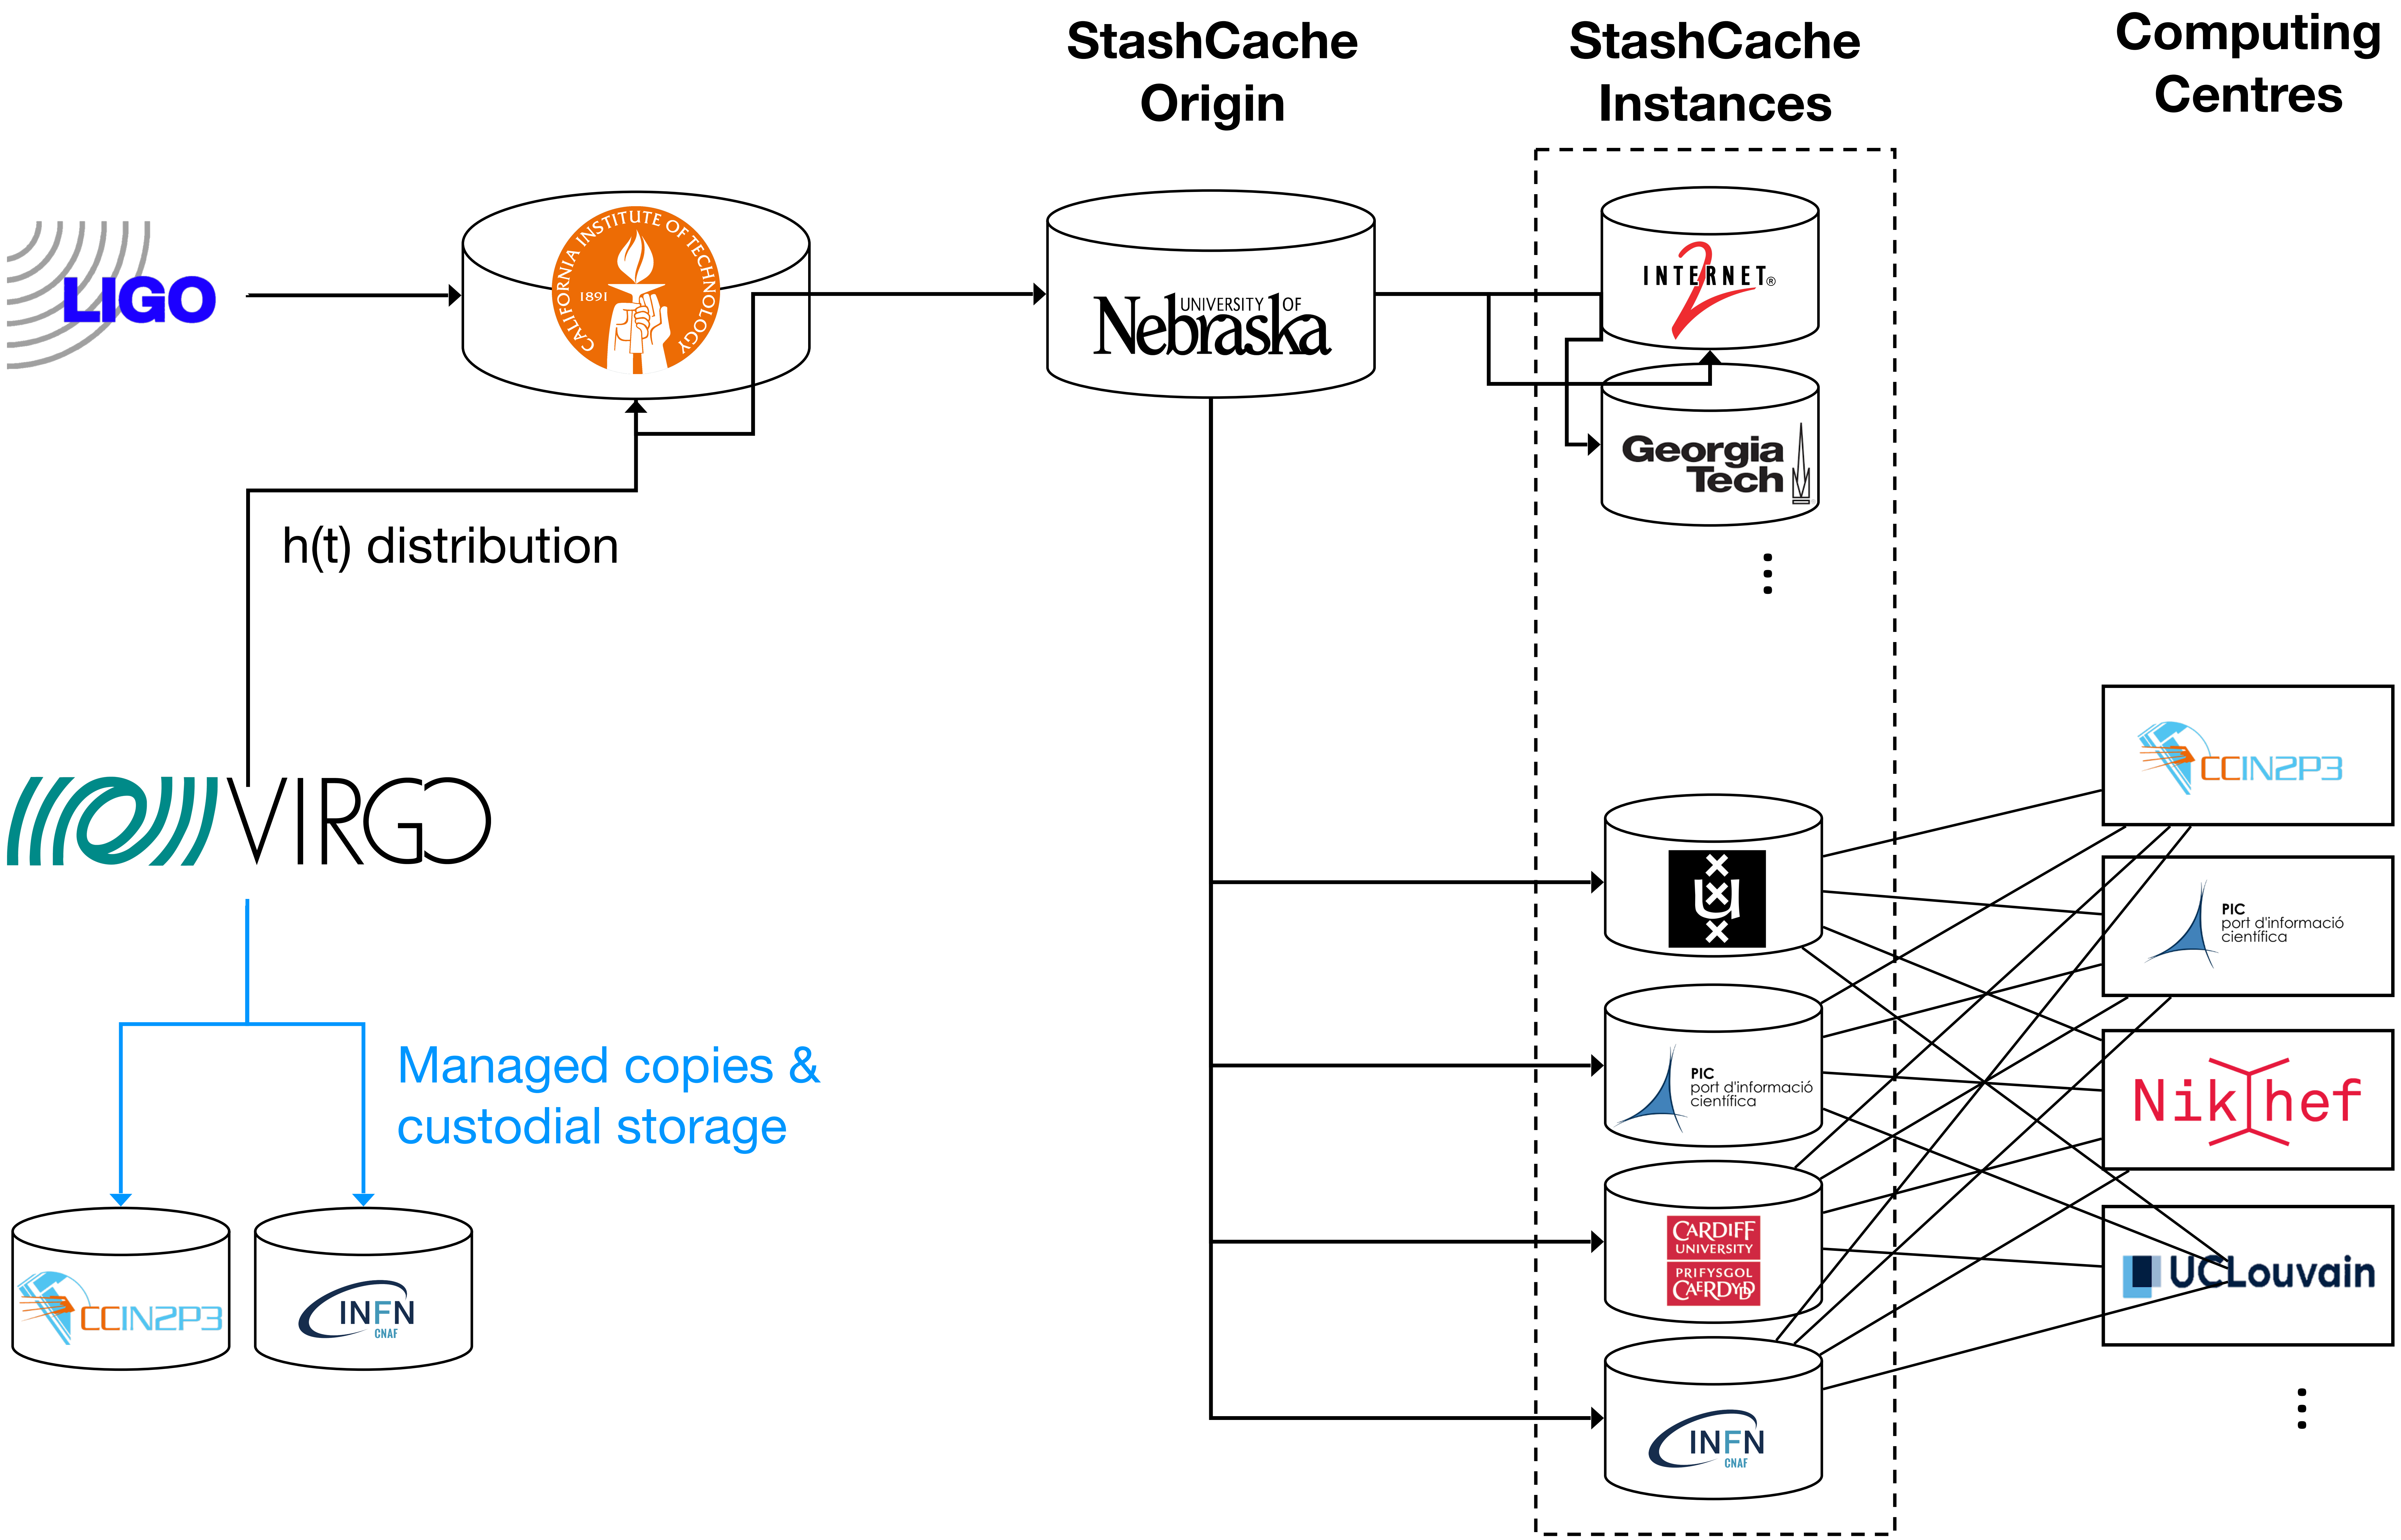

Offline data distribution: then

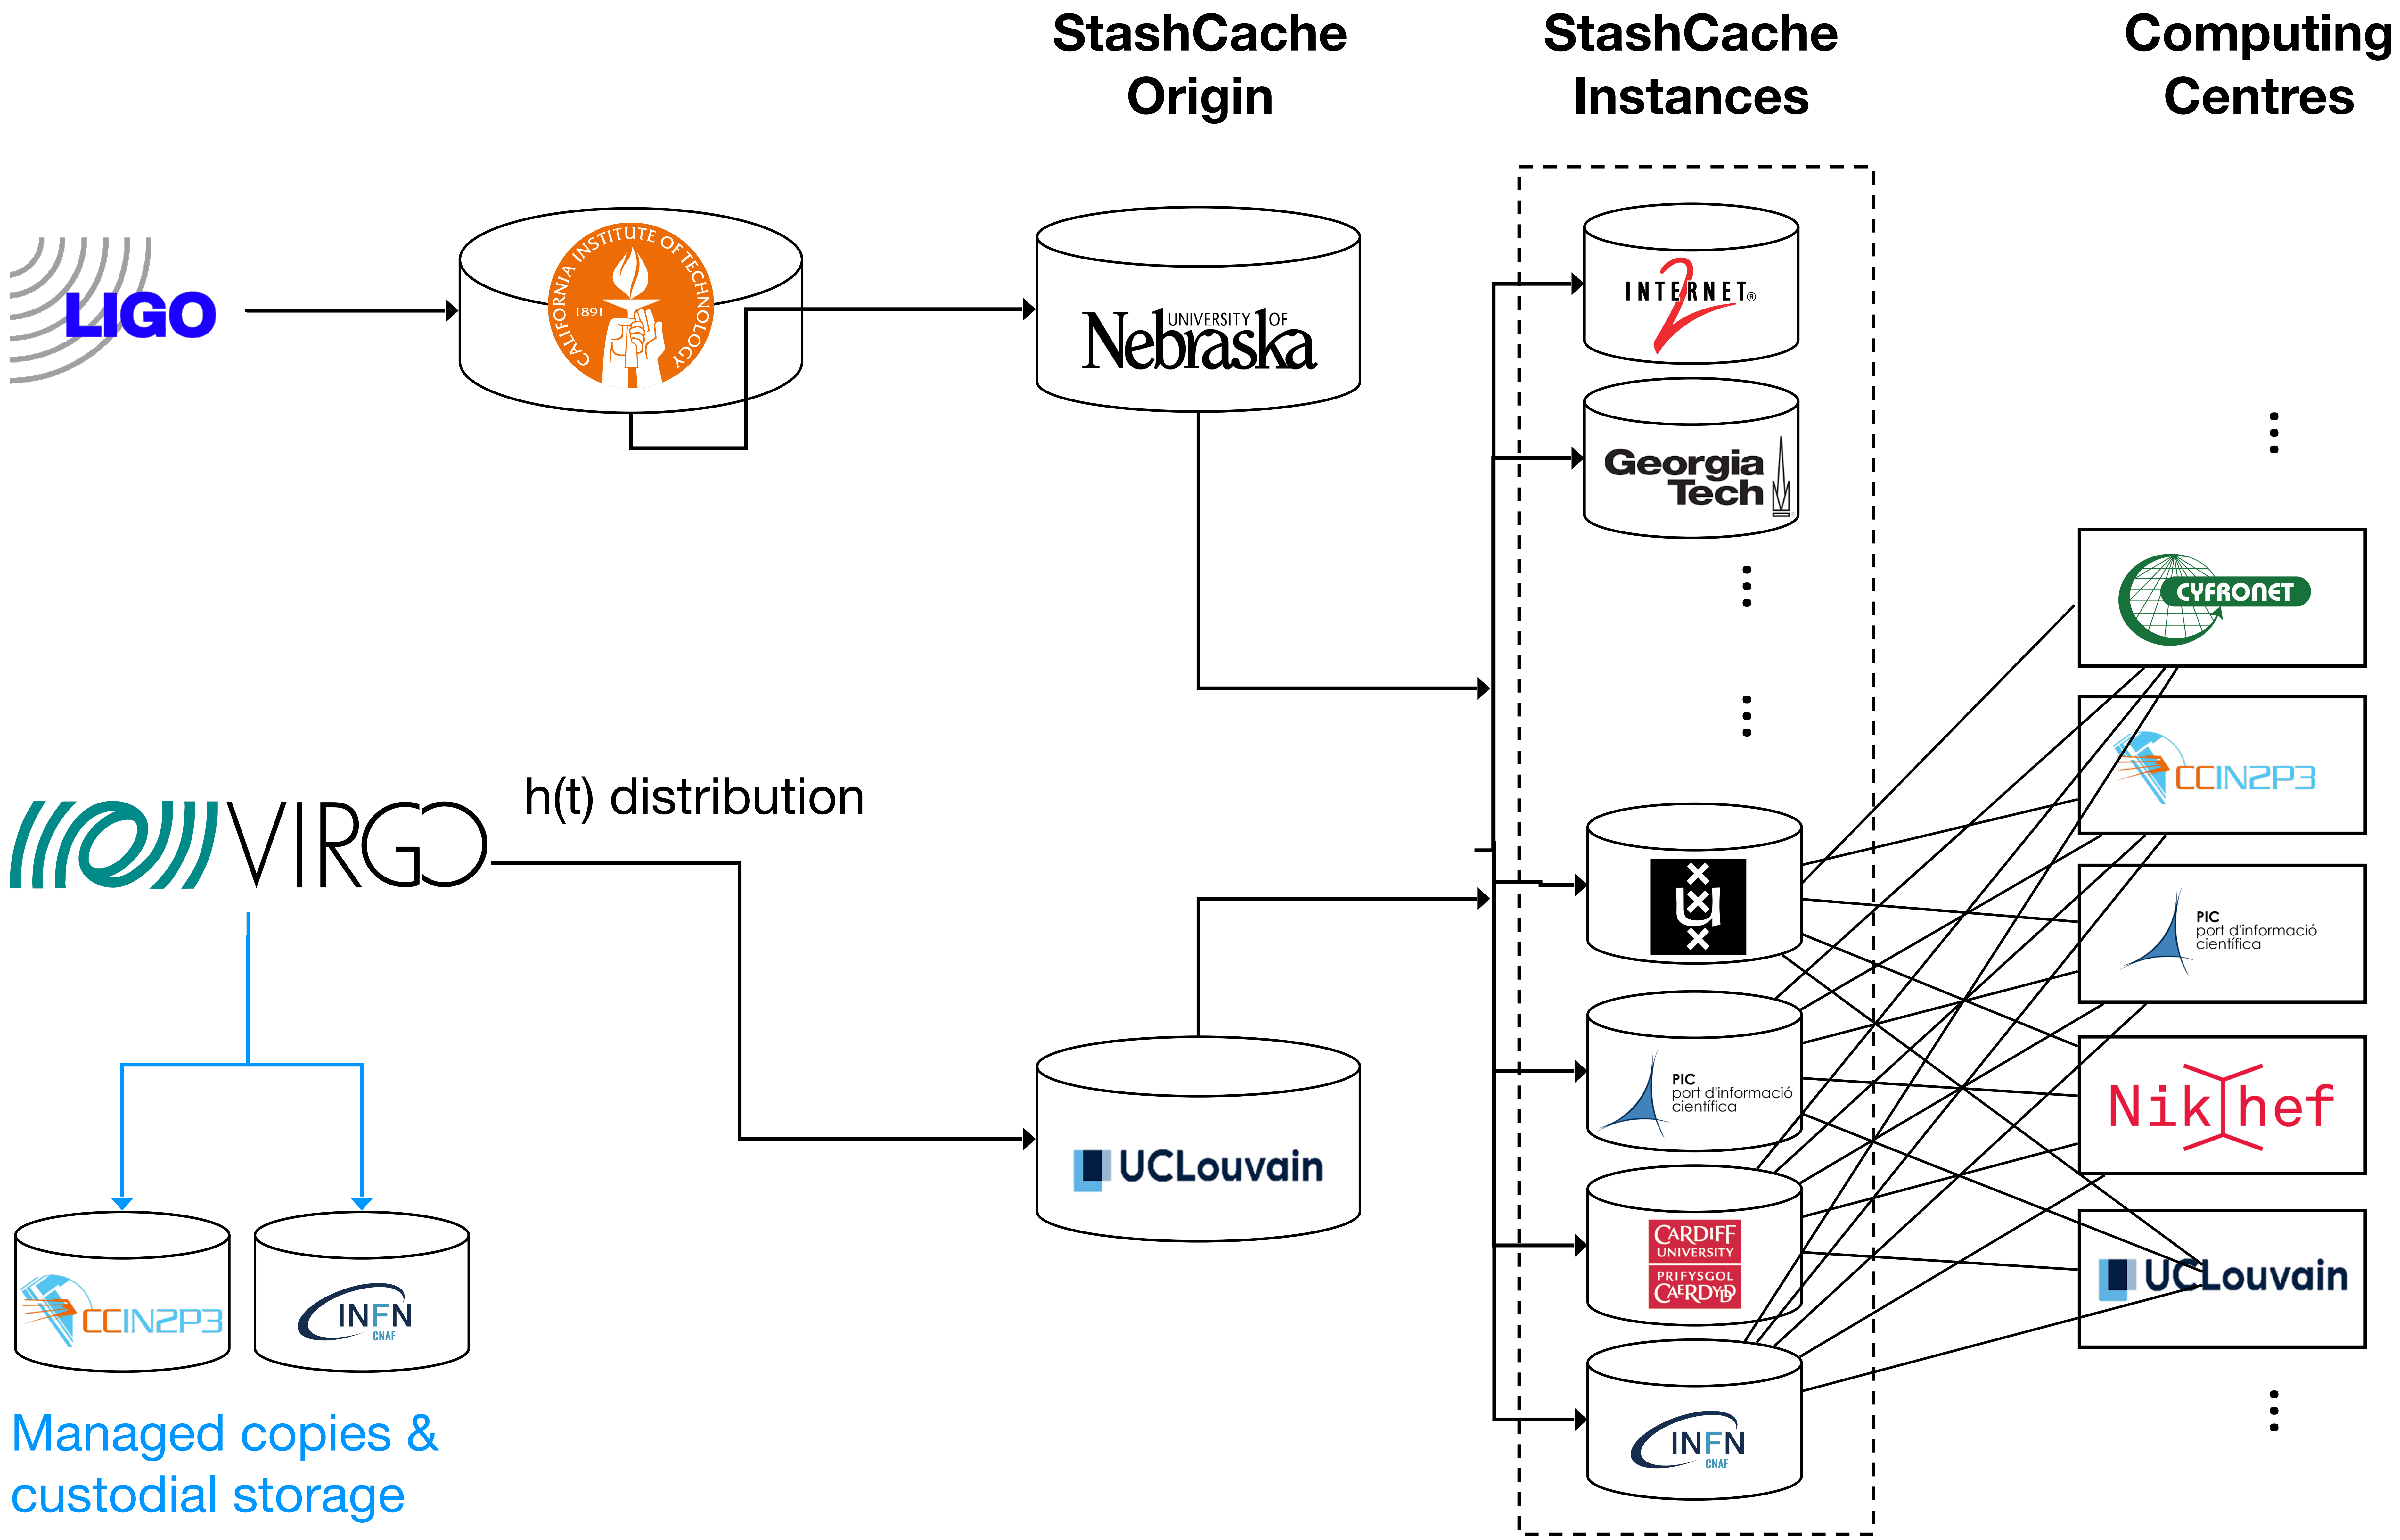

Low Latency in a nutshell

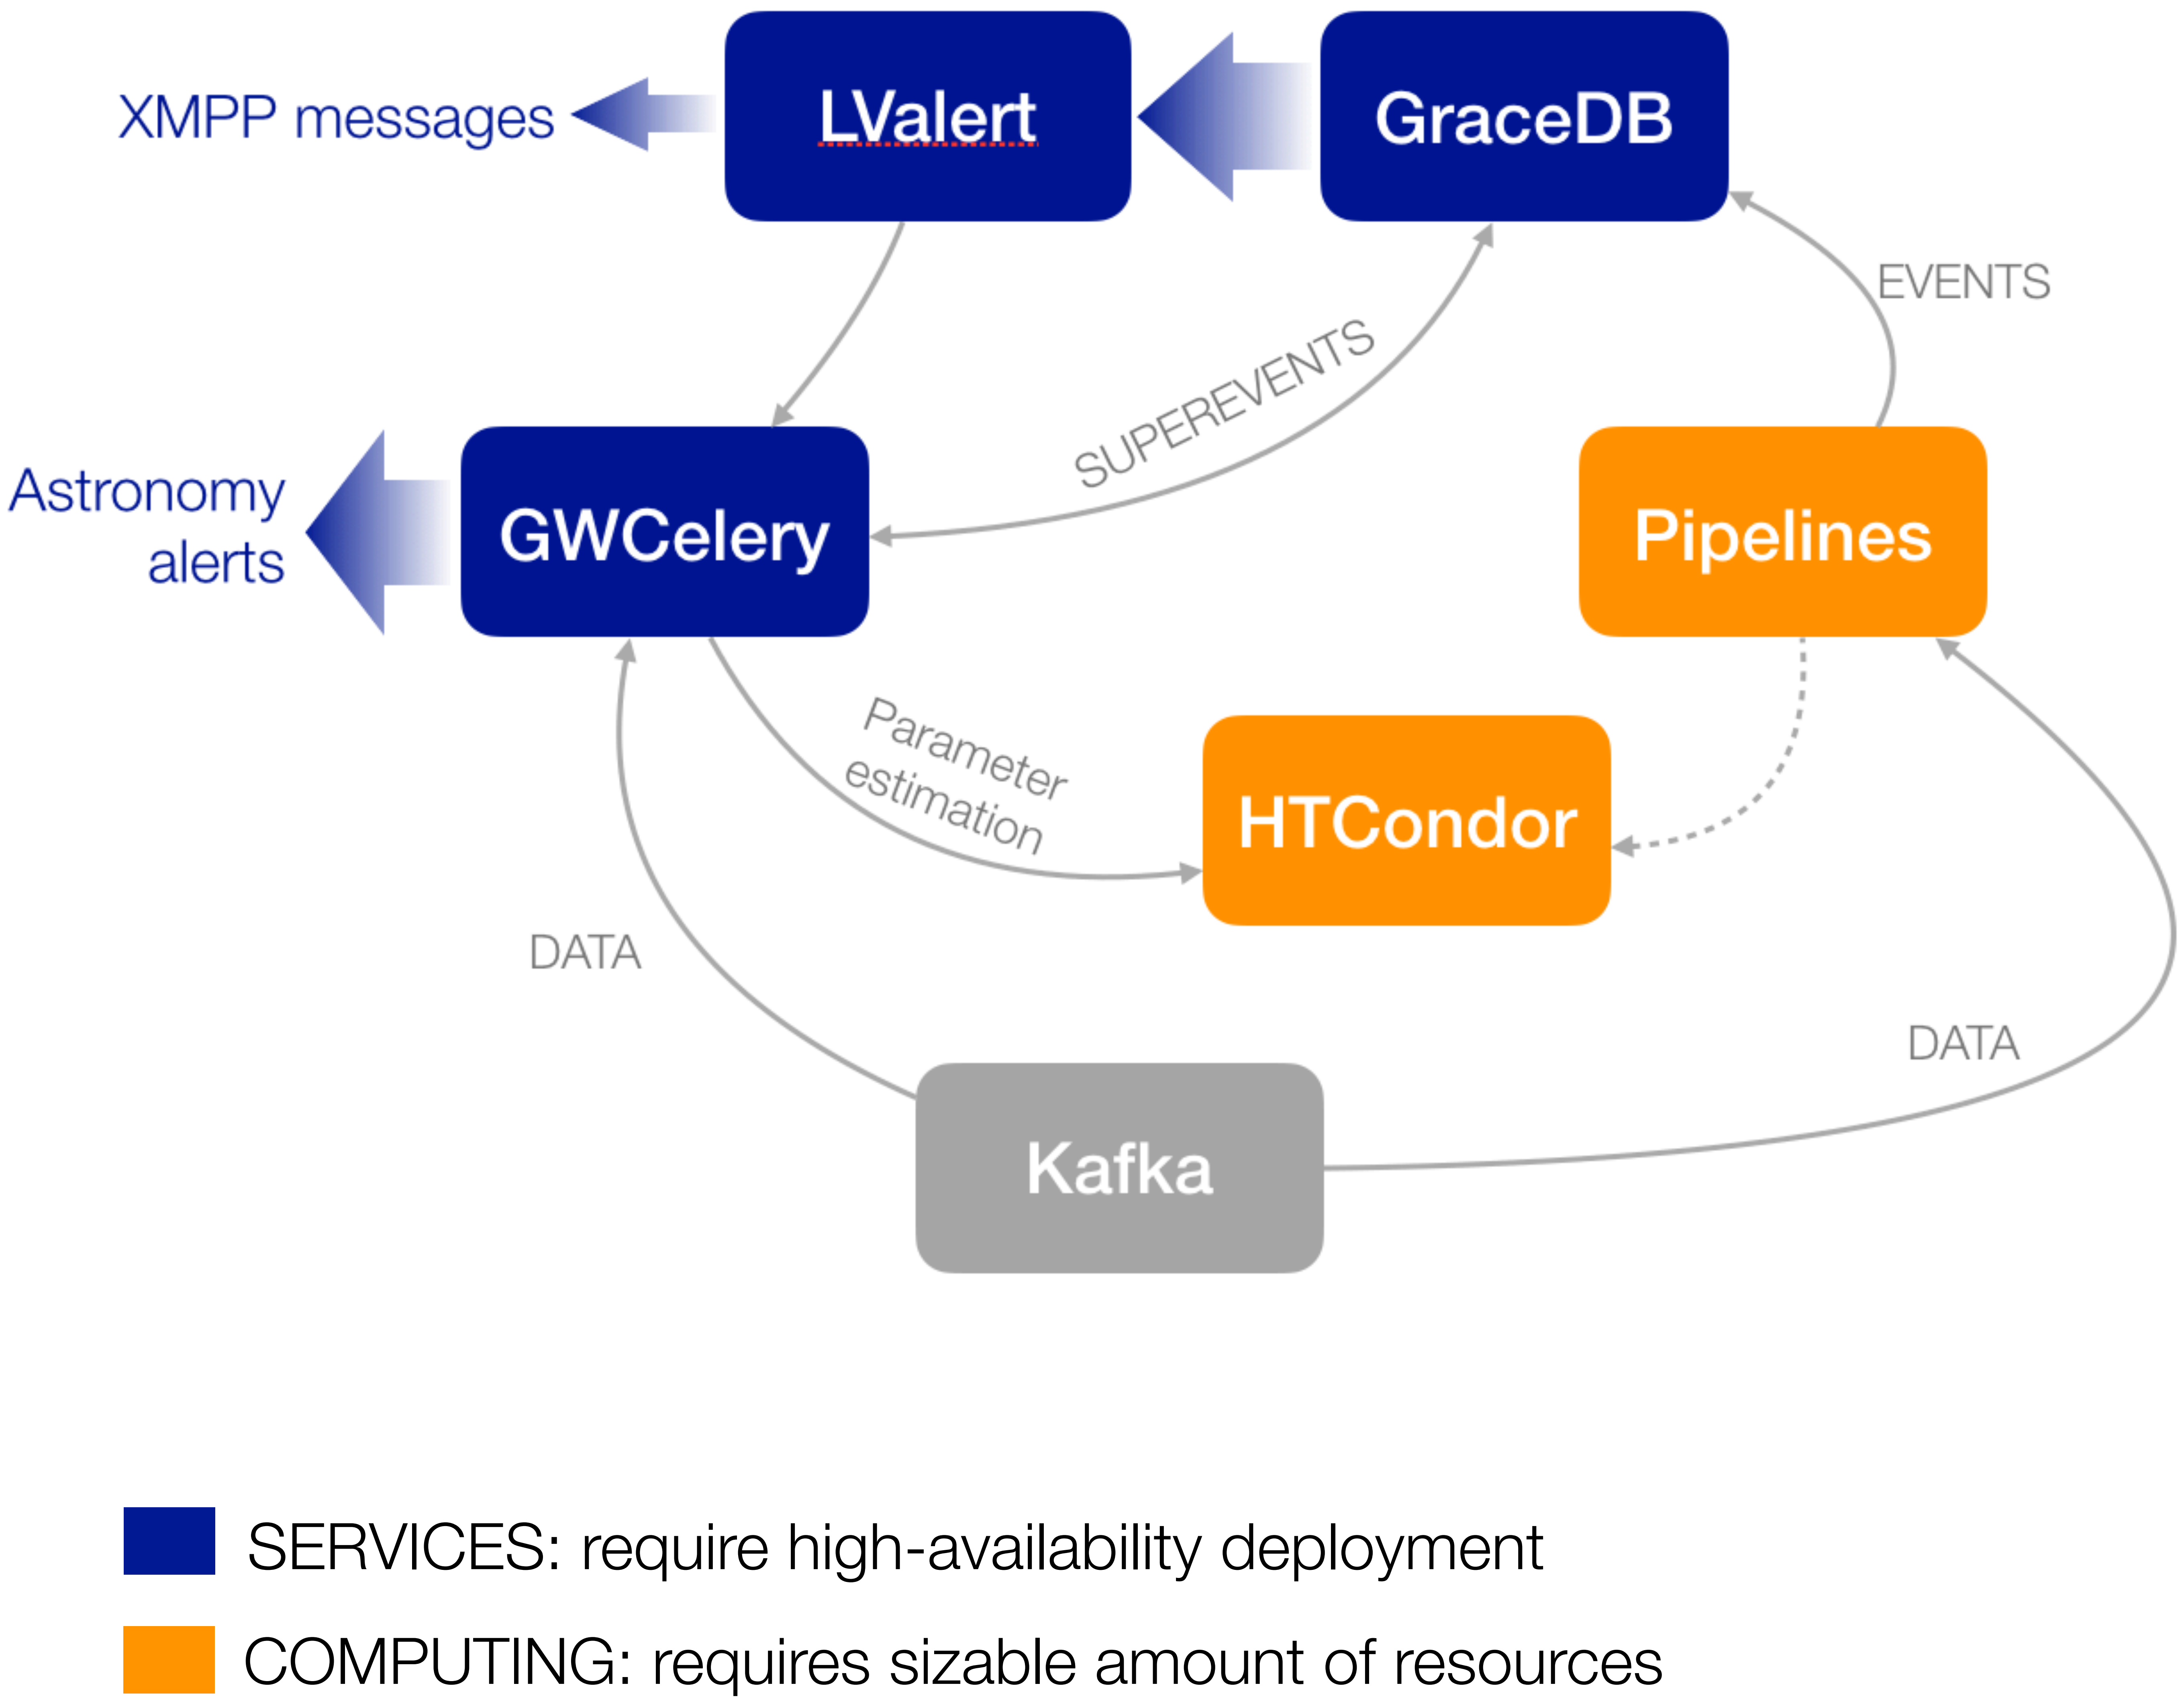

- Different Tiers of **GracedB and LValert** instances (Playground, Dev, Test, Production) managed by Ligo and currently deployed on AWS
- Bring some key services to Europe → IGWN Test instances of GraceDB and LValert deployed on Kubernetes cluster at CNAF (consolidation work ongoing)
- **GWCelery** deployed on virtual machine at CNAF
- Kafka data producers running at CIT and EGO, Kafka data consumers running at EGO and CNAF
- **Pipelines and HTCondor** running at EGO → feasibility study to run also off-site

Low Latency in O4 and beyond

- Scale up to cope with increasing event rate
- Automatize and harden
- Speed up to reduce latency ("negative latency" alerts!)

EGI-ACE: INFRAEOSC-07 project for "capacity building". Its main goal is to **implement the compute platform of the European Open Science Cloud and contribute to the EOSC Data Commons** by delivering integrated computing platforms, data spaces and tools as an integrated solution that is aligned with major European cloud federation projects and HPC initiatives.

O3 data replay exercise

A truly end-to-end exercise:

- Provide a reasonable representation of low-latency detector data in O4
- Exercise the search pipeline's ability for finding GWs as well as the mitigation of generating event candidates from terrestrial noise; for this, a stretch of data that generated a variety of Bursts, BNS, NSBH and BBH events as well as retractions is desired
- Exercise the calibration pipeline's ability to provide calibrated h(t) in low latency
- Exercise low-latency detchar services and their ability to detect and/or mitigate nonstationary/non-Gaussian noise in h(t) using h(t) and/or auxiliary channel information
- Exercise the alert infrastructure's capability in providing timely alerts and evaluate its full latency
- Currently focusing on the January 6, 2020 to February 14, 2020 O3 period replay (Stage 1 + 2)

O3 data replay setup

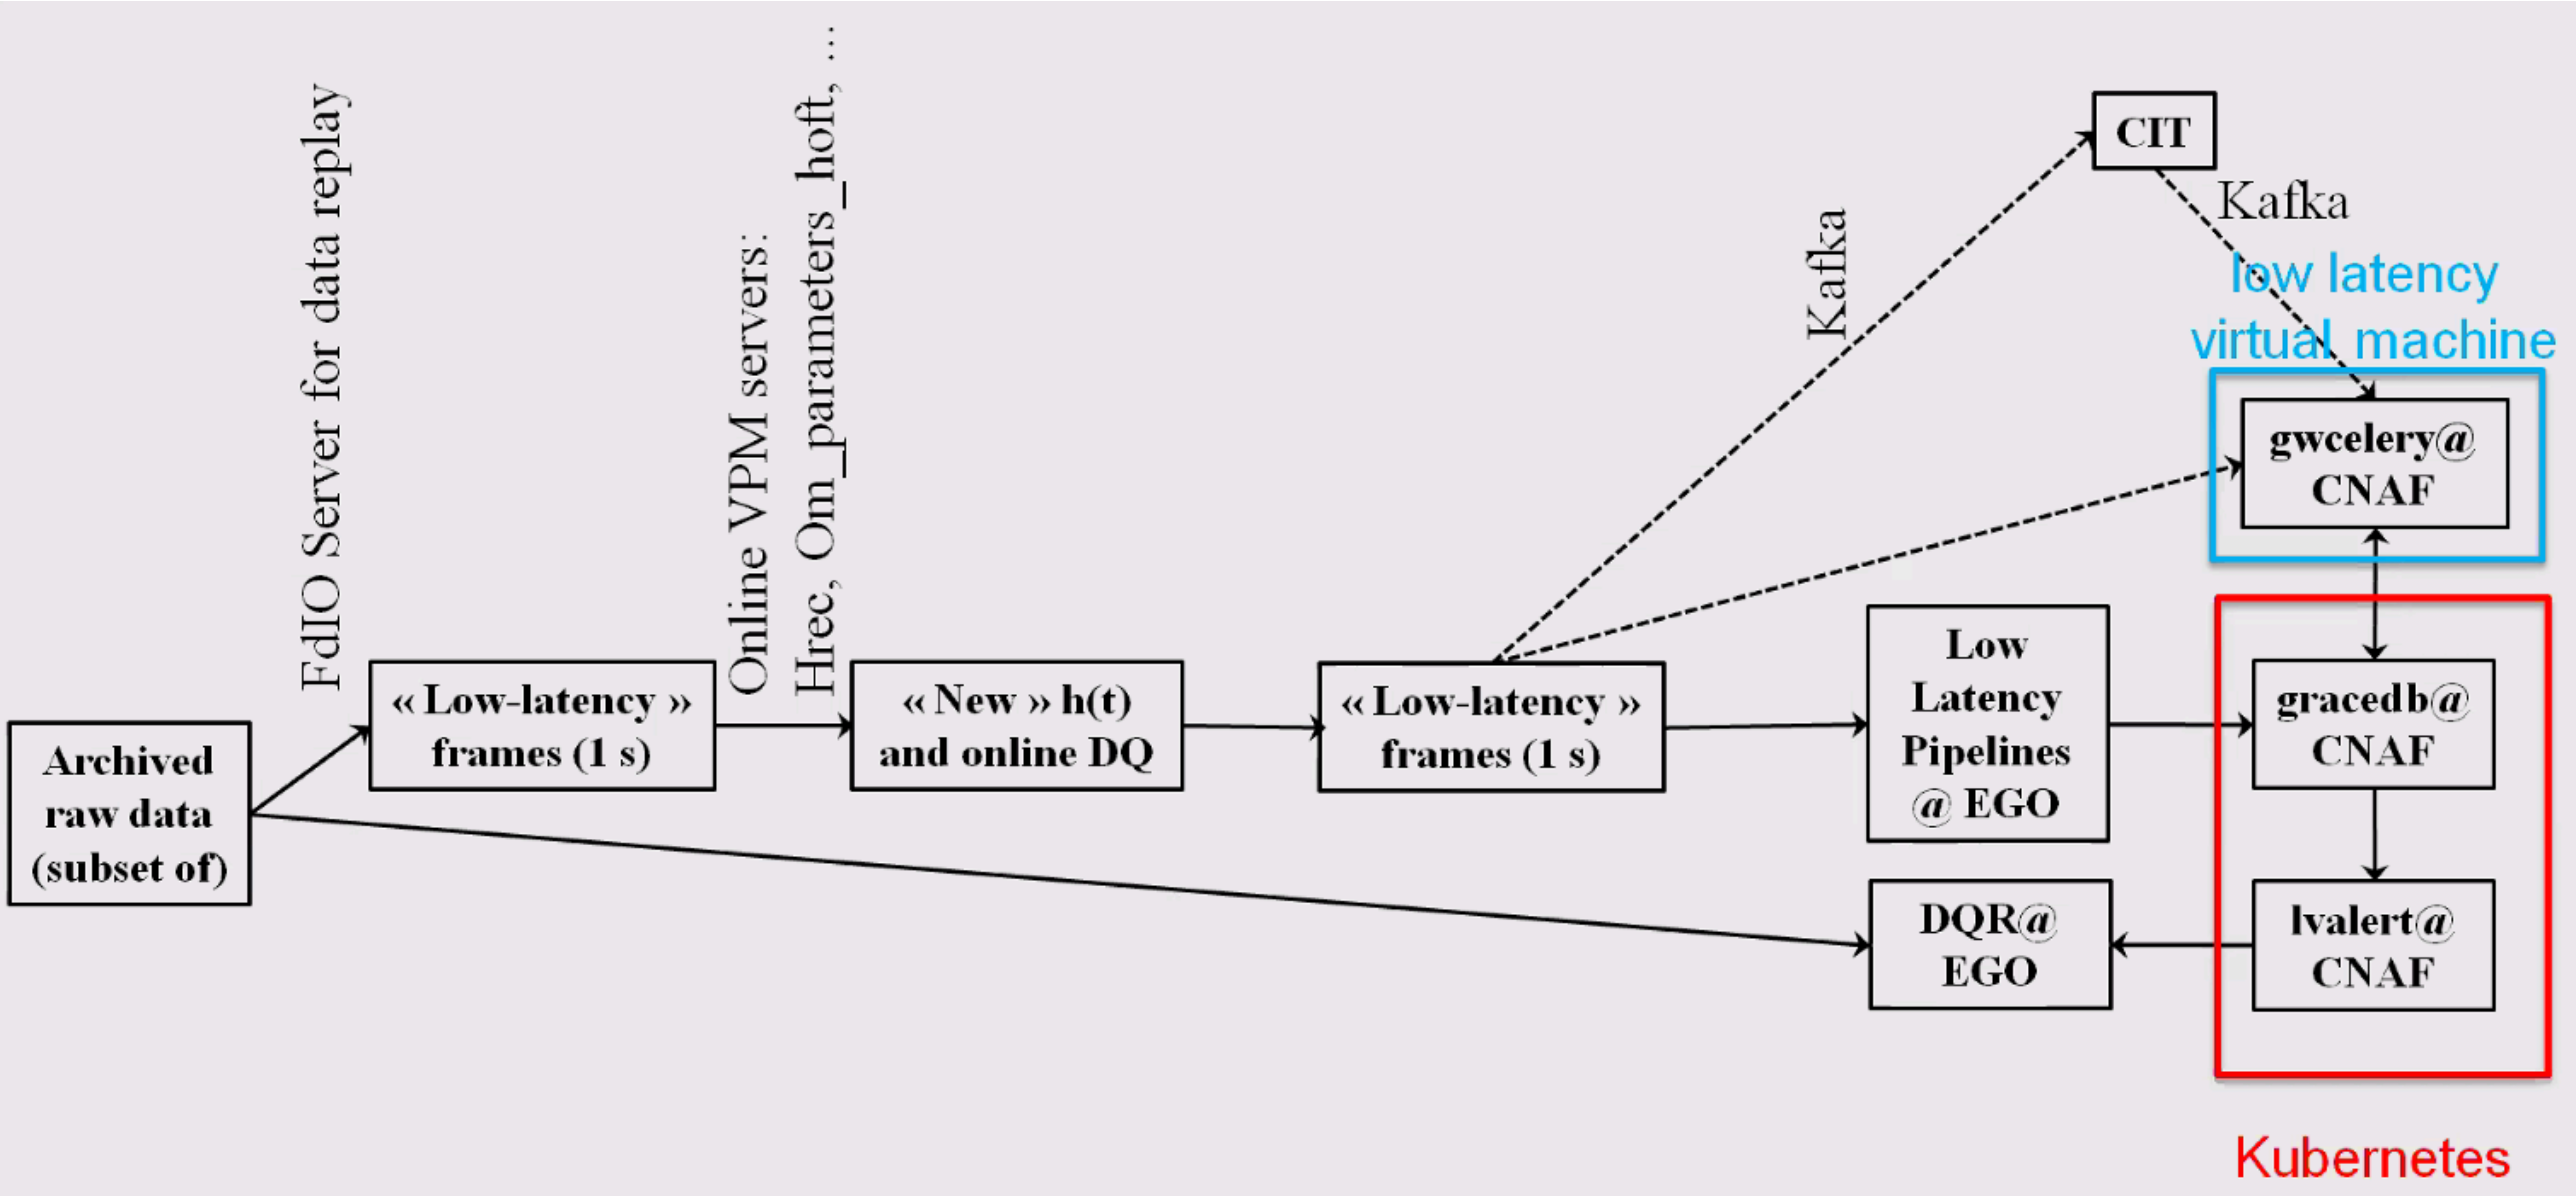

Slide from F. Carbognani

3G detectors: Einstein Telescope

- Einstein Telescope project proposed to ESFRI committee

- Triangular design with six 10km interferometers (3 low- and 3 high-frequency)

- Underground to reduce seismic noise, cryogenic to reduce thermal noise

- Two candidate sites: *Sos Enattos* and *Euregio Meuse-Rhine*

- Current planning has first data in the 2030's

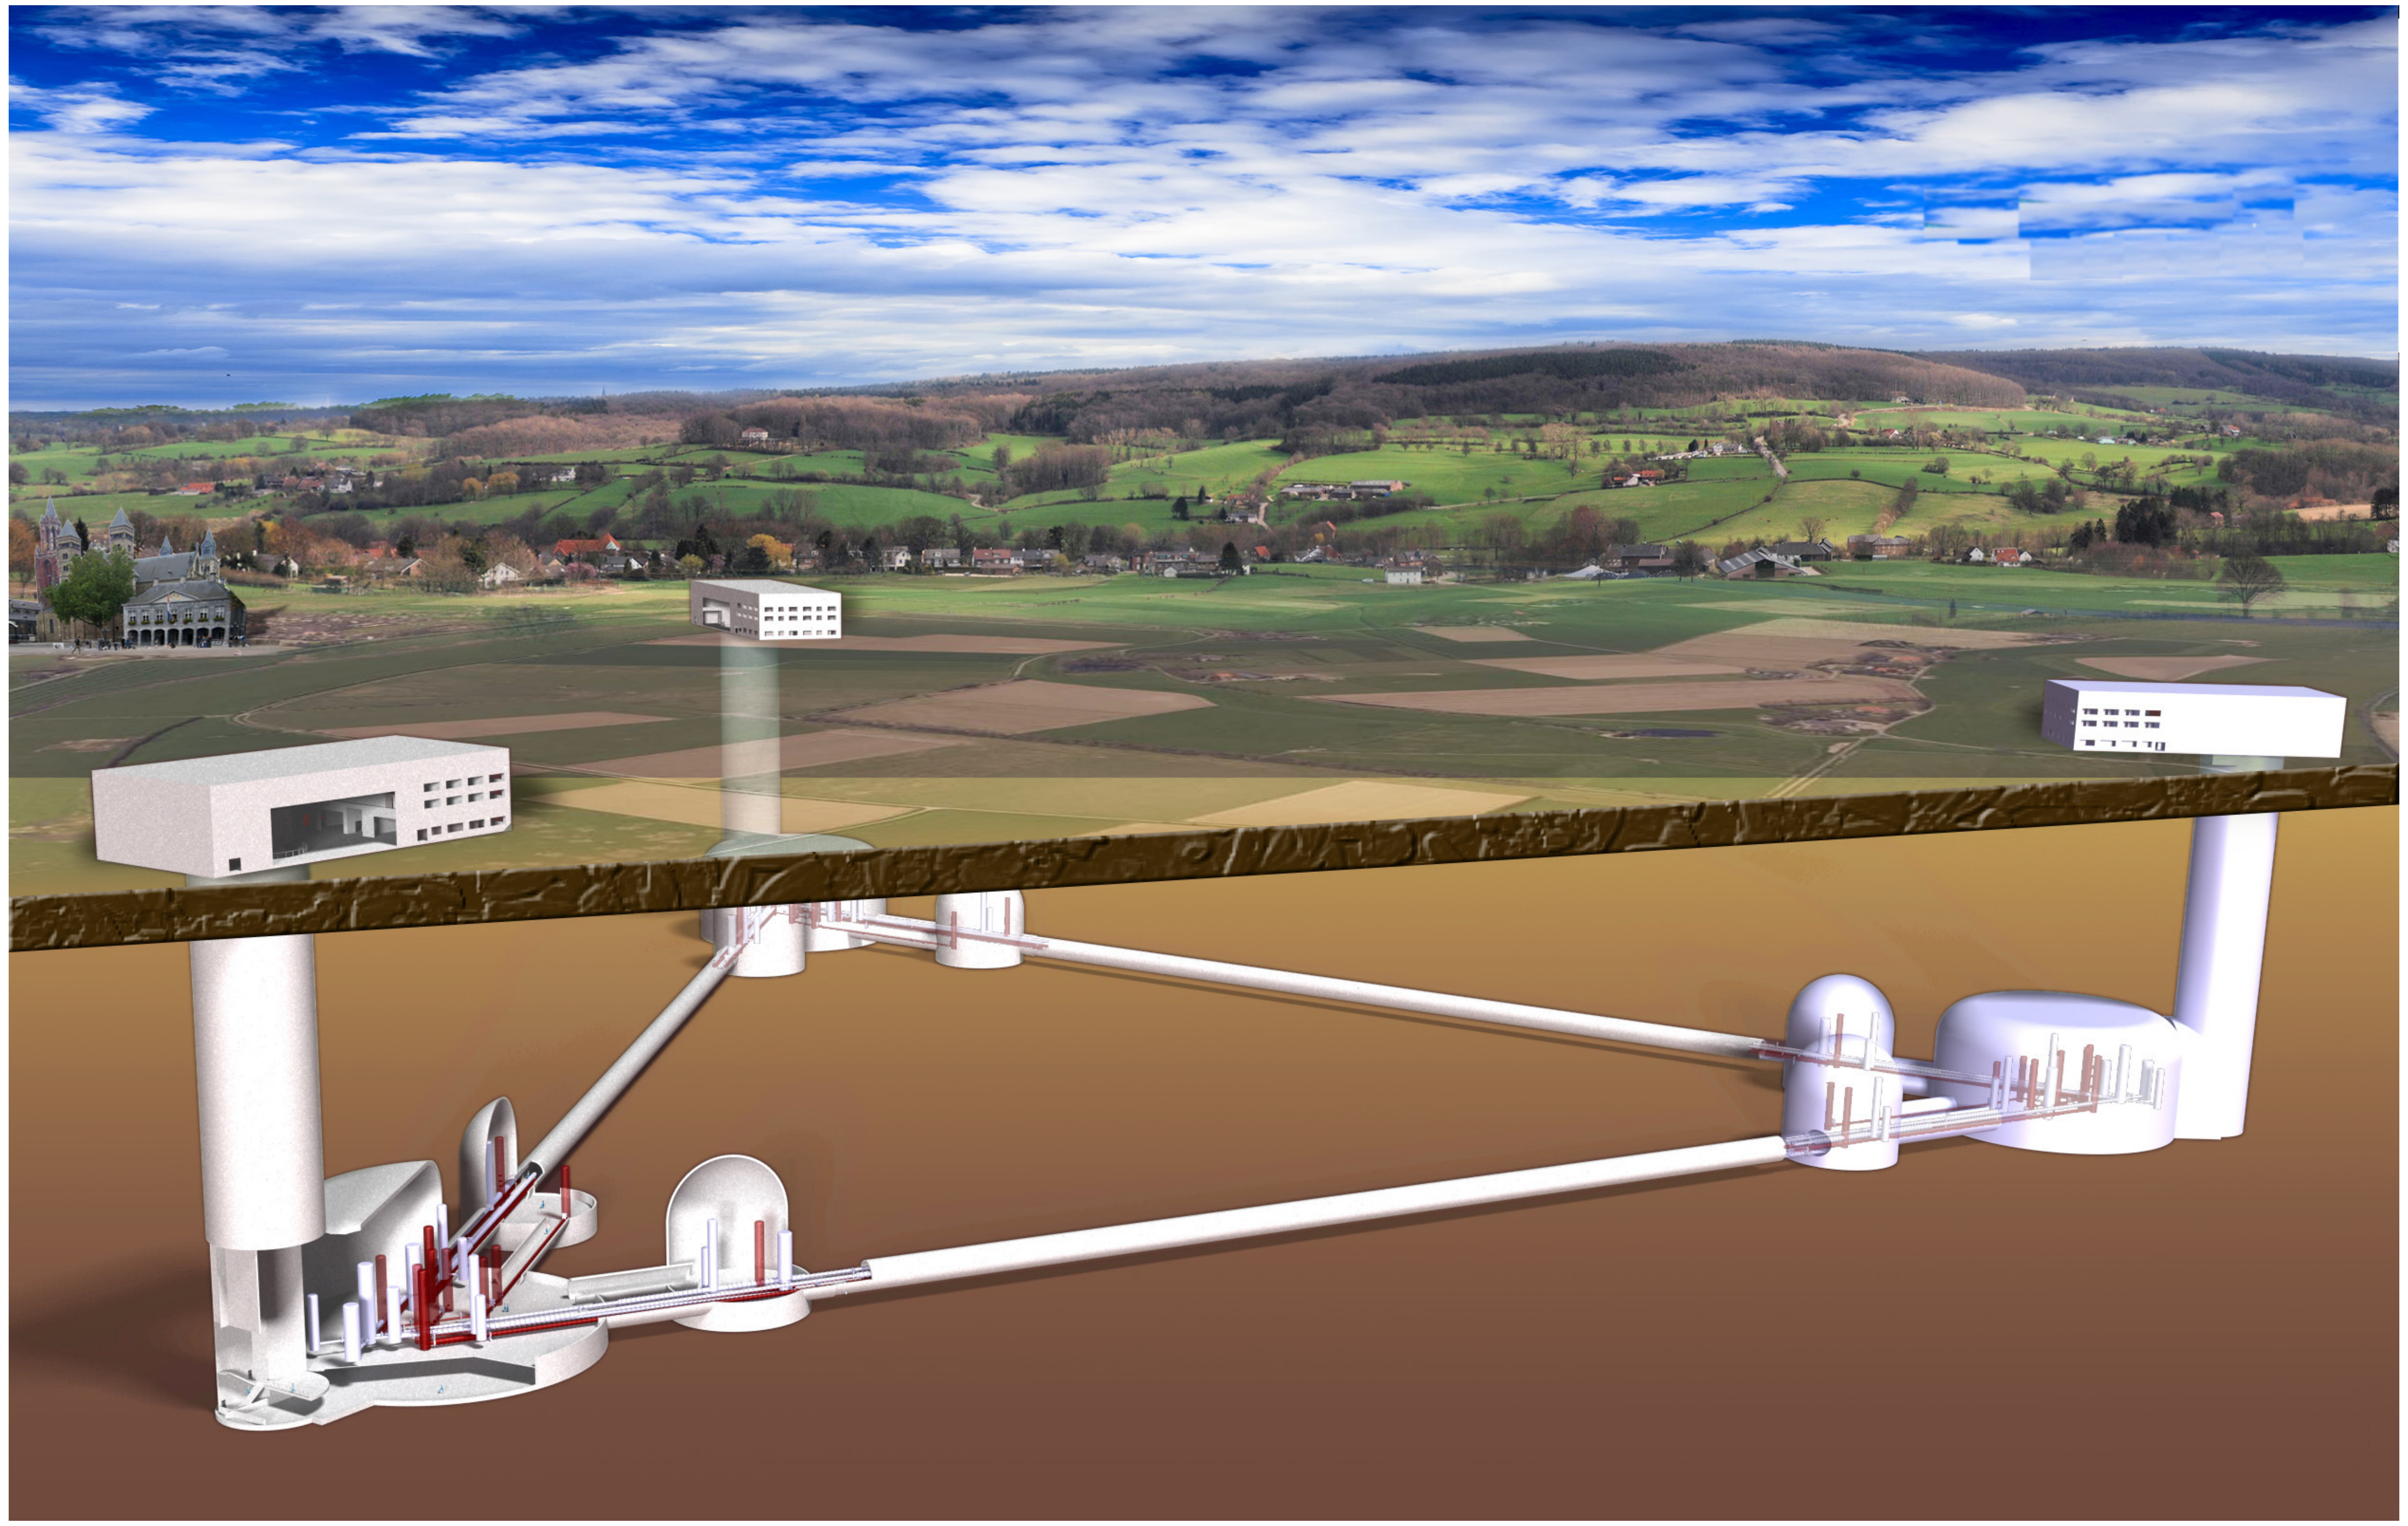

If you are willing to **join the ET EIB (E-Infrastructure Board)** please contact us (sara.vallero@to.infn.it). We will have a meeting to welcome newcomers and discuss possibilities to contribute on **Wednesday, June 16th, 11:00 – 12:00**.

3G computing challenges

- GW data size does not explode with instrument sensitivity like LHC…
- …but there are much more events embedded in the data: 10^3 - 10^4 what we have now.
- And will need to process them fast (“negative” latency): MM astronomy is the priority
  - Design the infrastructure for low-latency searches
- 2G interferometers generate o(1PB/yr) of data, we expect 10-100 times that
- Issue: it is difficult to estimate the amount of CPU power needed to extract science from the data
- We are currently ~10% of an LHC experiment, naïve scaling gives out-of-reach numbers:
  - Overlapping signals
  - Very long signals
  - Will matched filtering still work?
- Will need a lot of optimization work in the next years!

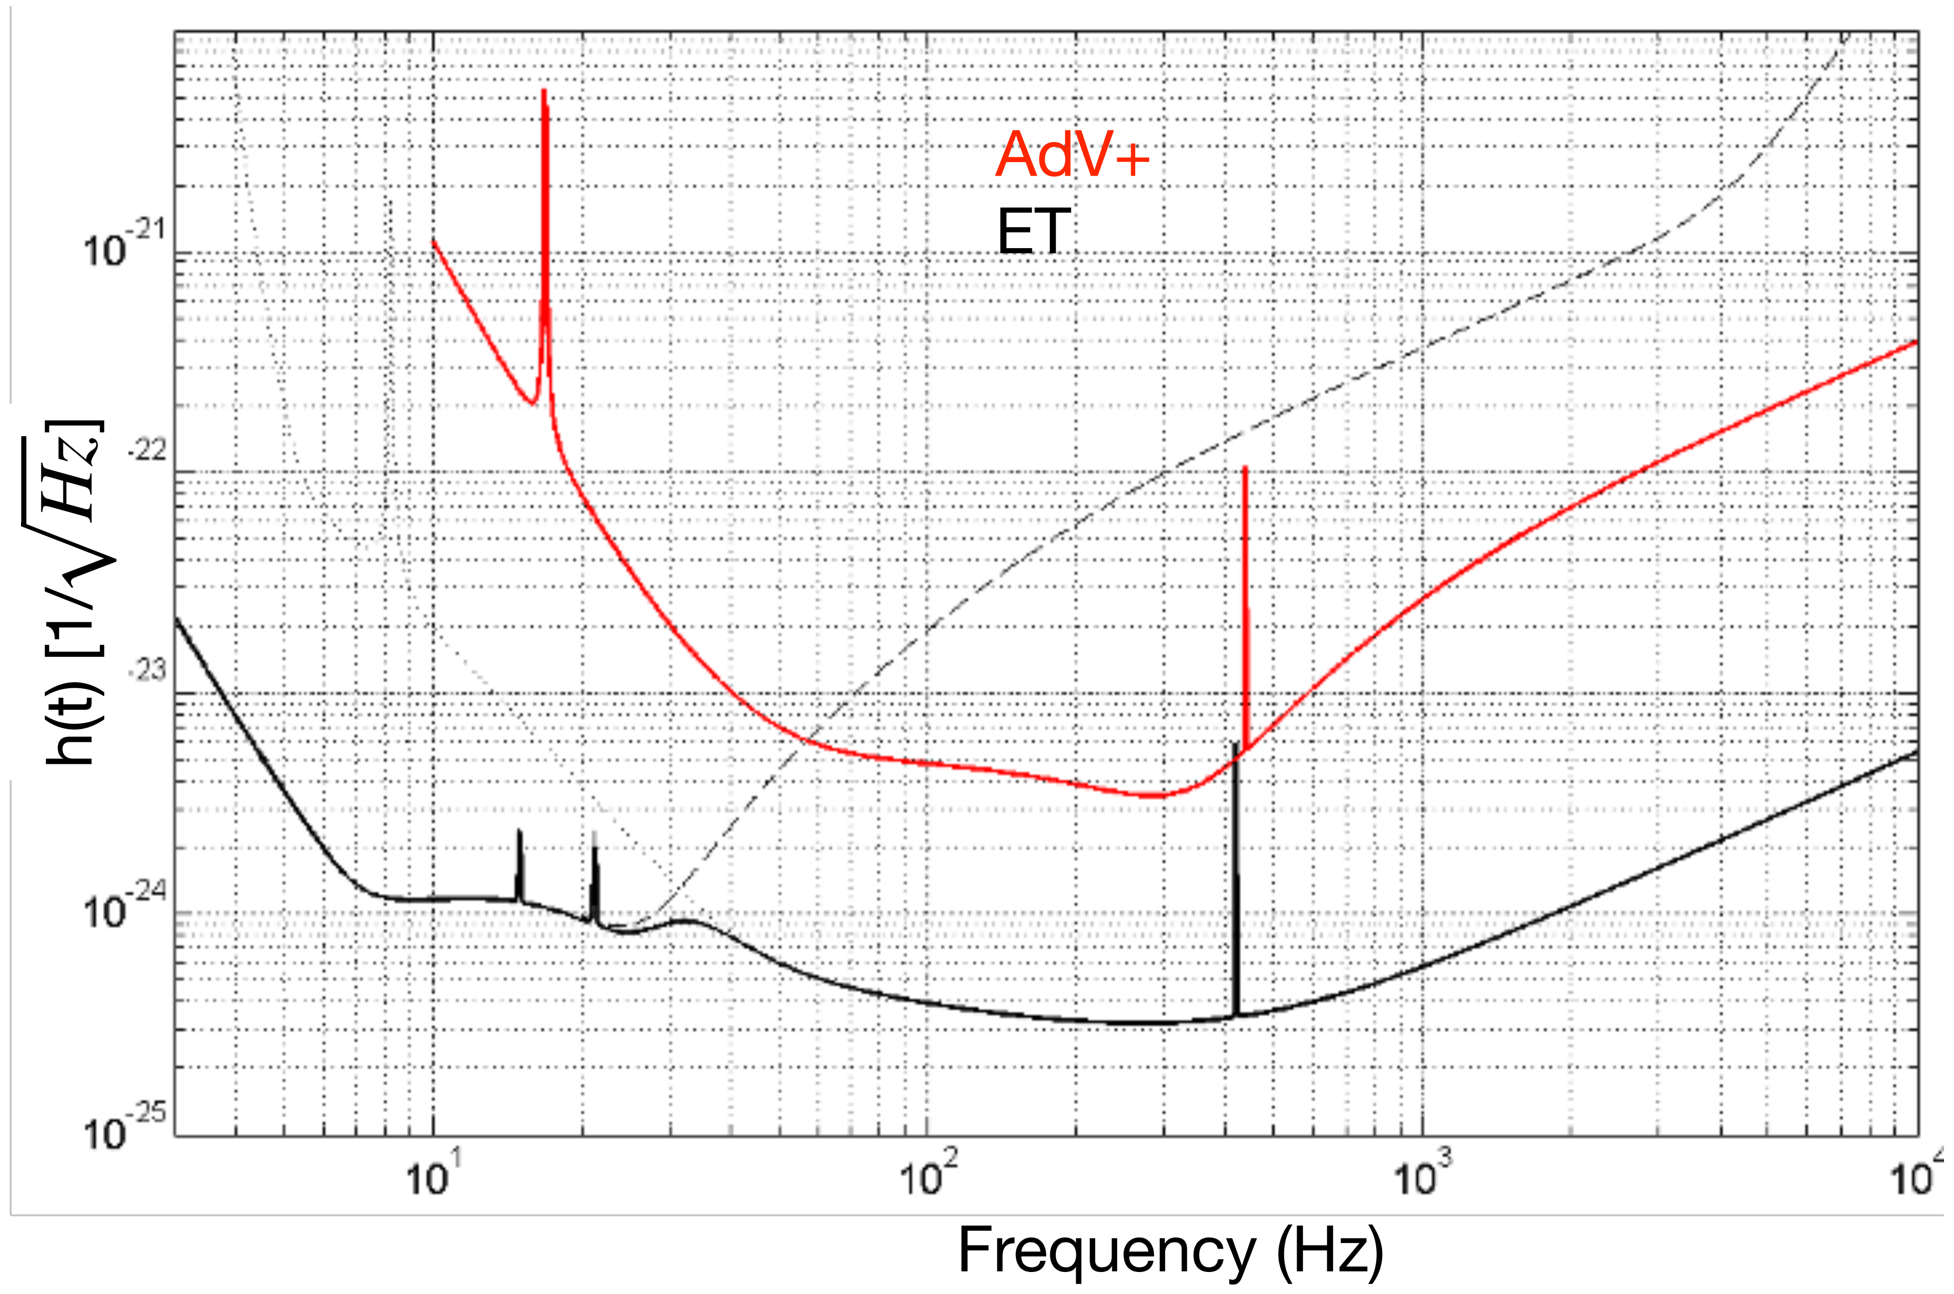

Wrap up

- Virgo computing is steadily evolving along the path outlined two years ago (use mainstream tools and adopt a distributed computing model)
- Virgo is no longer a "small experiment": already now it's 10% of one LHC experiment and it will continue to grow
- *Gruppo II* experiments (i.e. CTA) need tight data exchange in the low latency domain → we are working on a solid and scalable LL implementation leveraging INFN-Cloud
- In 10 years from now 3G experiments will pose a serious challenge to our computing capabilities (especially concerning cpu power) → we should start now to conceive sound infrastructures and innovative algorithms (help is needed)