SETTING THE SCENE

THE STRAWMAN ET COMPUTING MODEL

Stefano Bagnasco | INFN Torino

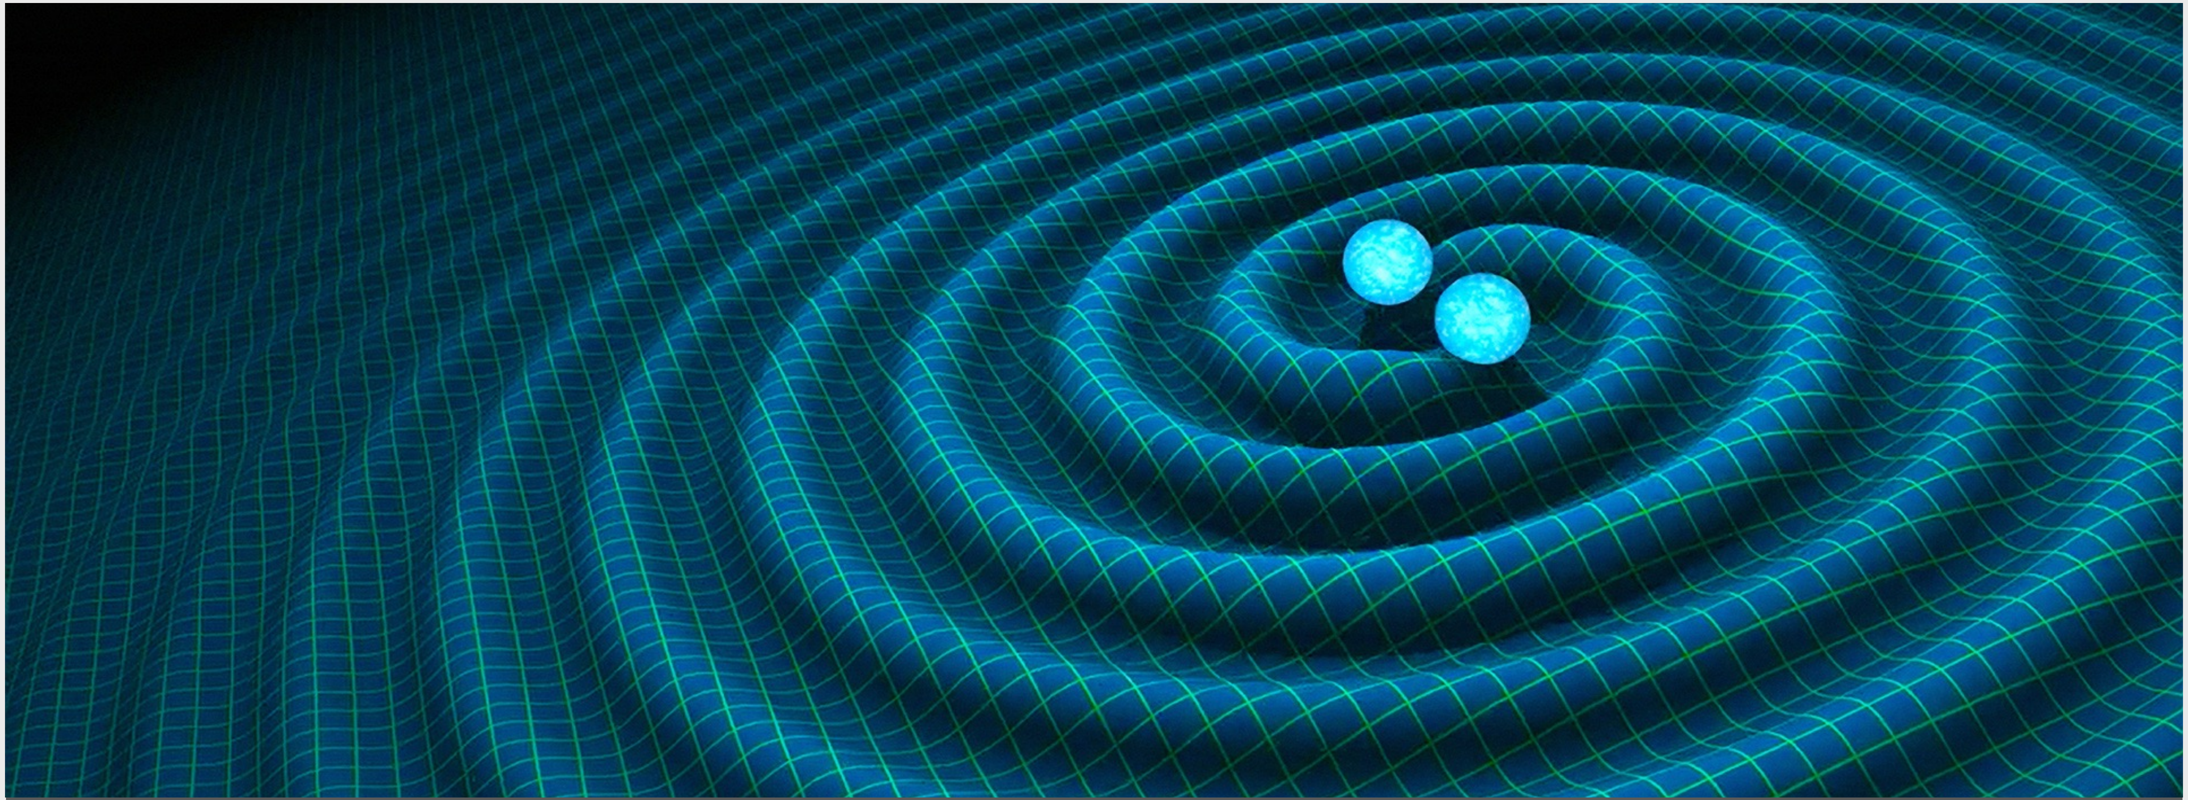

GRAVITATIONAL WAVES RESEARCH

LOOKING FOR...

Transient sources:

- CBC: Compact Binary Coalescence
  - Binary Black Holes (BBH), Binary Neutron Stars (BNS), NSBH
  - Strongest emitters, well modelled for much of the parameter space
  - Matched filtering very effective

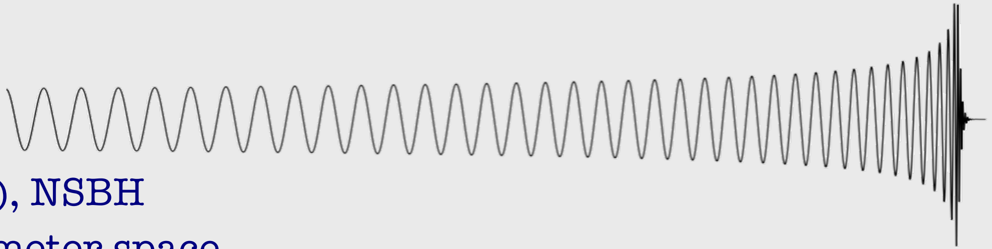

- Burst: Unmodeled transient bursts
  - E.g., Core Collapse Supernovae (CCS, and anything else)
  - Weaker and no (or poor) model, so coherence methods more effective

Continuous sources:

- CW: Continuous waves
  - E.g., Asymmetric spinning neutron stars
  - Usually well-modelled
  - All-sky and targeted searches

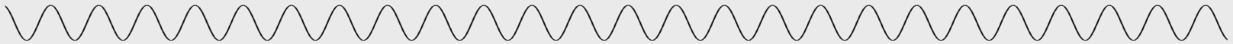

- SGWB: Continuous stochastic background
  - Astrophysical & cosmological

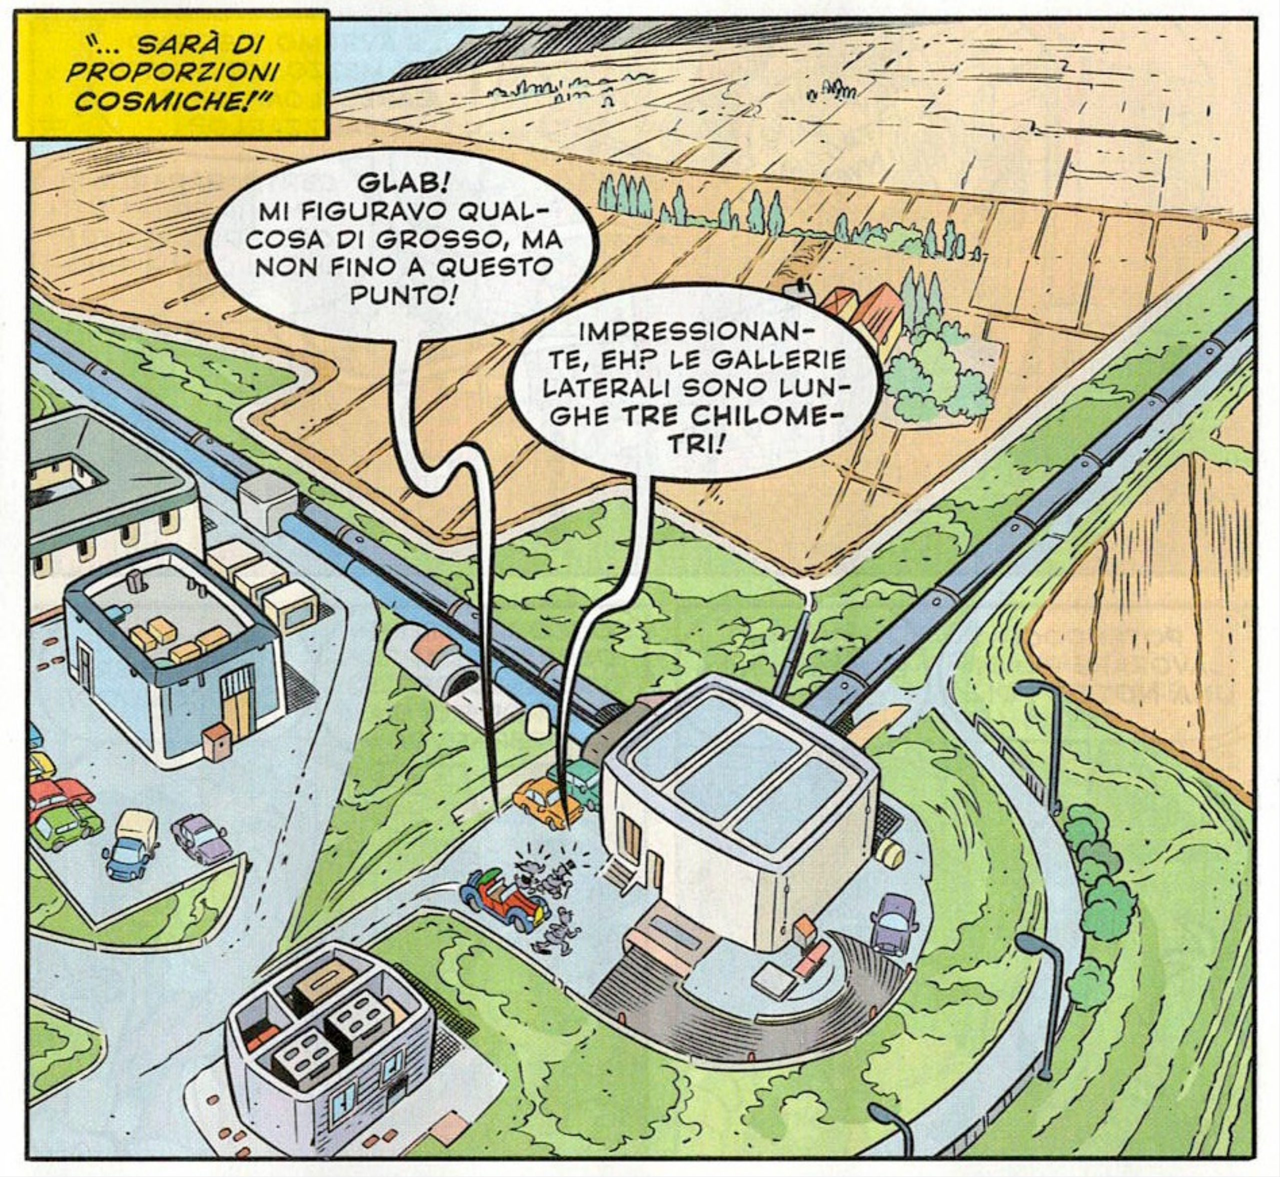

M. Bosco, G. Soldati, "Sogni d'oro zio Paperone"
Topolino **3538**:45-70 (2023)

A WORLDWIDE NETWORK TODAY

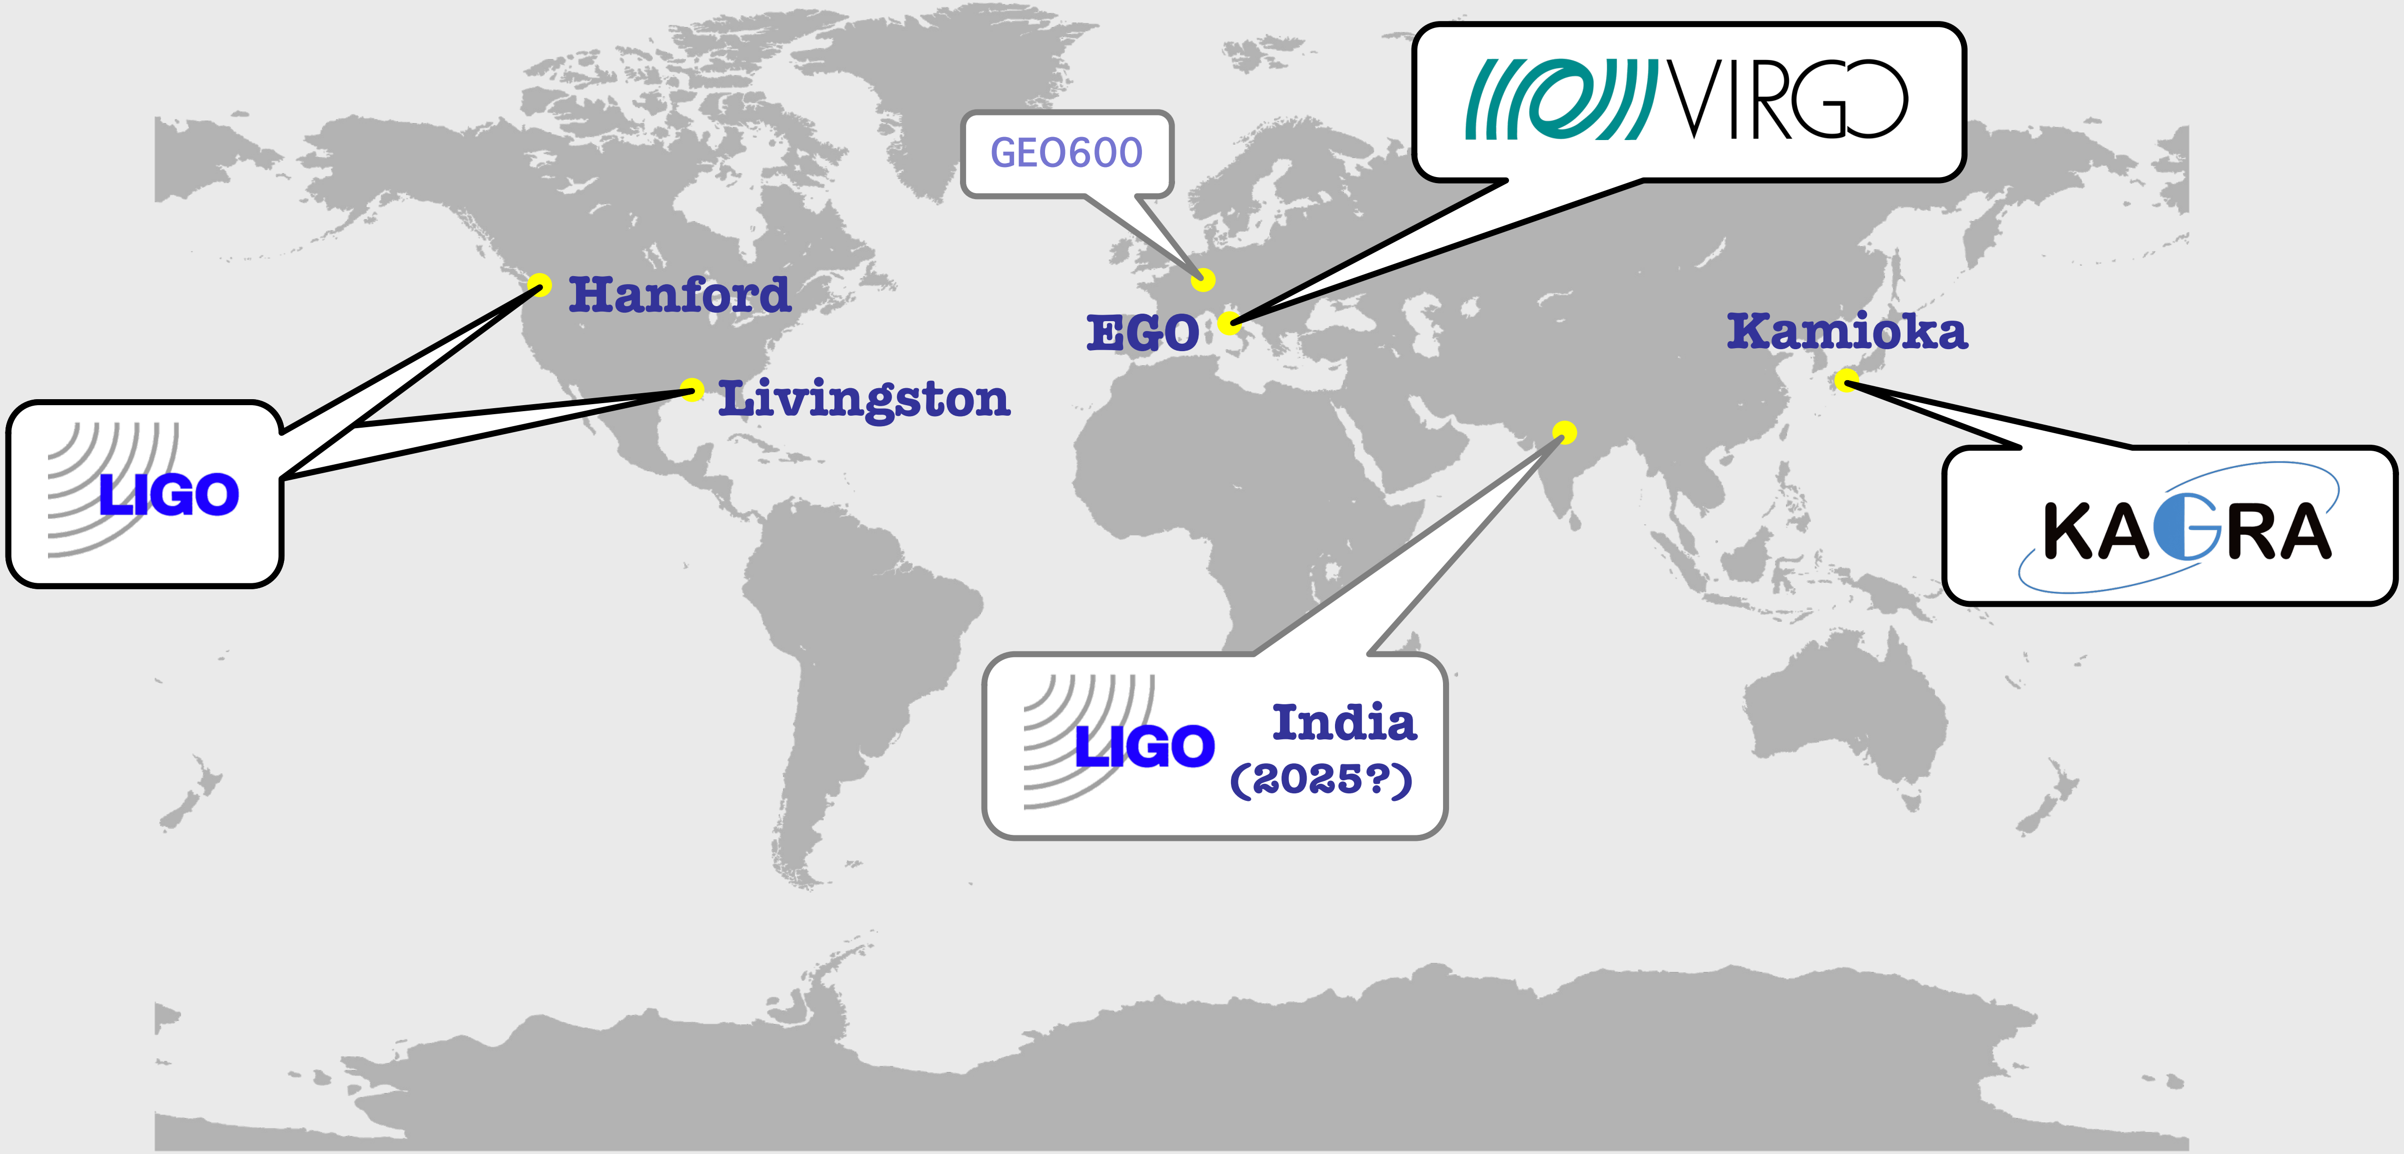

FROM DISCOVERY TO OBSERVATION

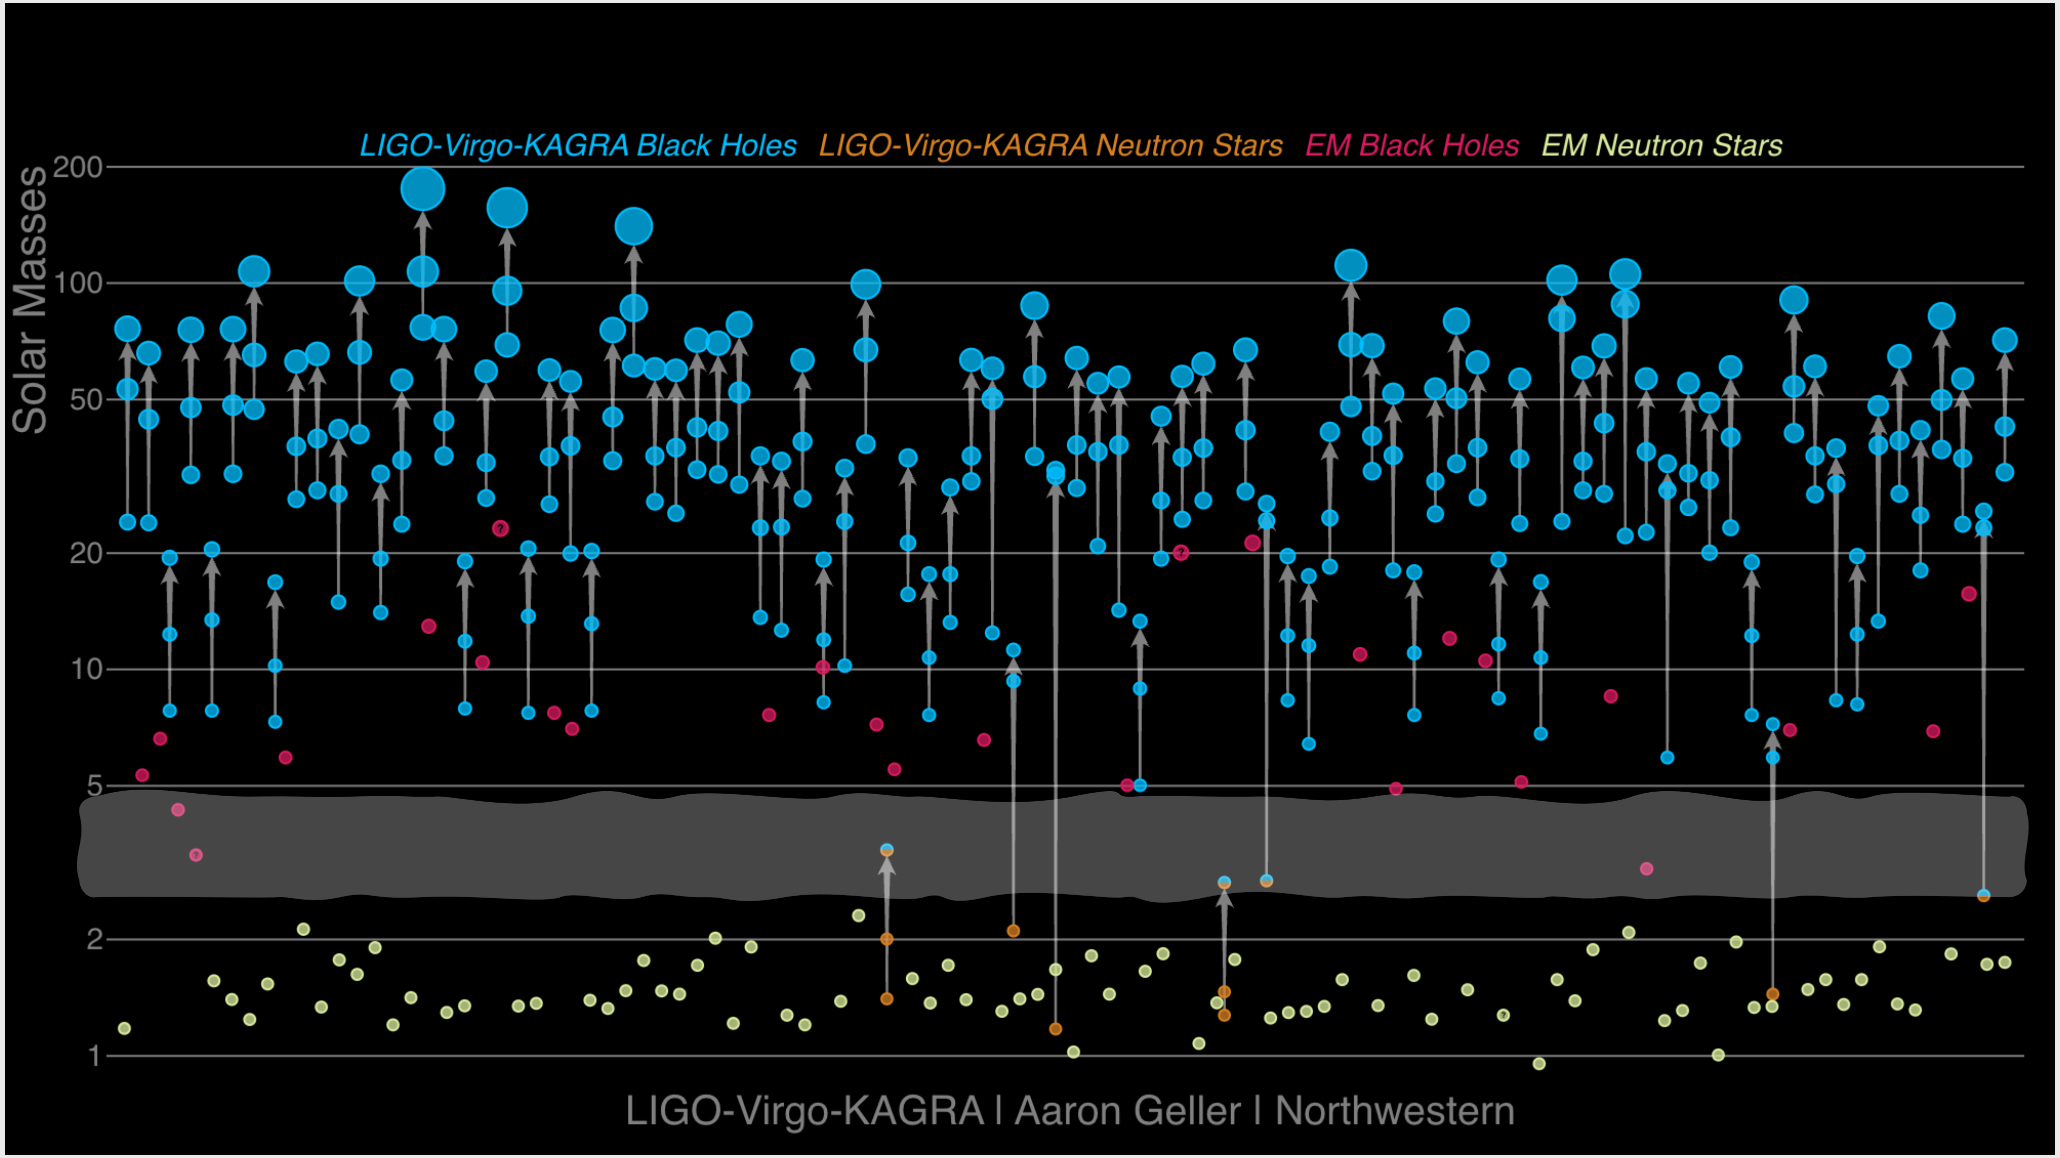

https://media.ligo.northwestern.edu/gallery/mass-plot

COMPUTING FOR GW: VIRGO

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

- 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 already possible
  - This will be the case for most events in ET

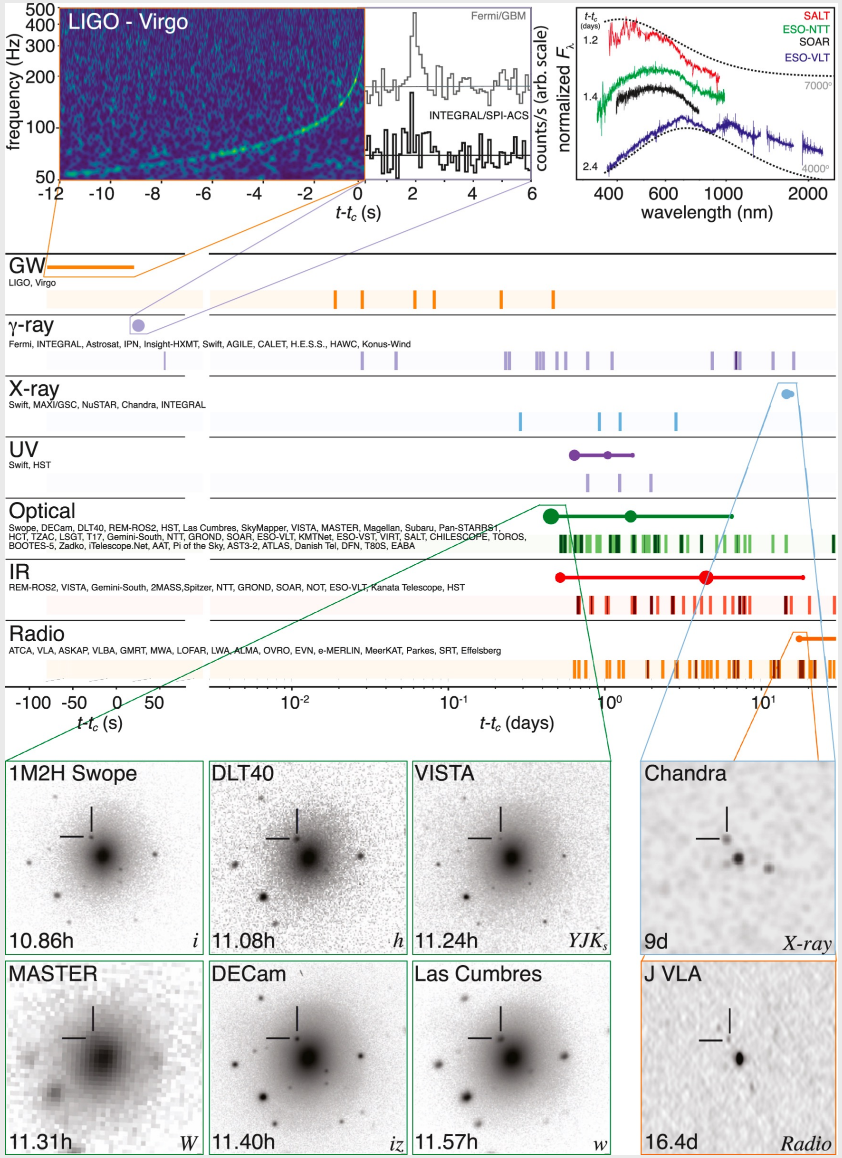

“Multi-messenger Observations of a Binary Neutron Star Merger”
B. P. Abbott *et al.* 2017 *ApJL* 848 L12
doi:10.3847/2041-8213/aa91c9

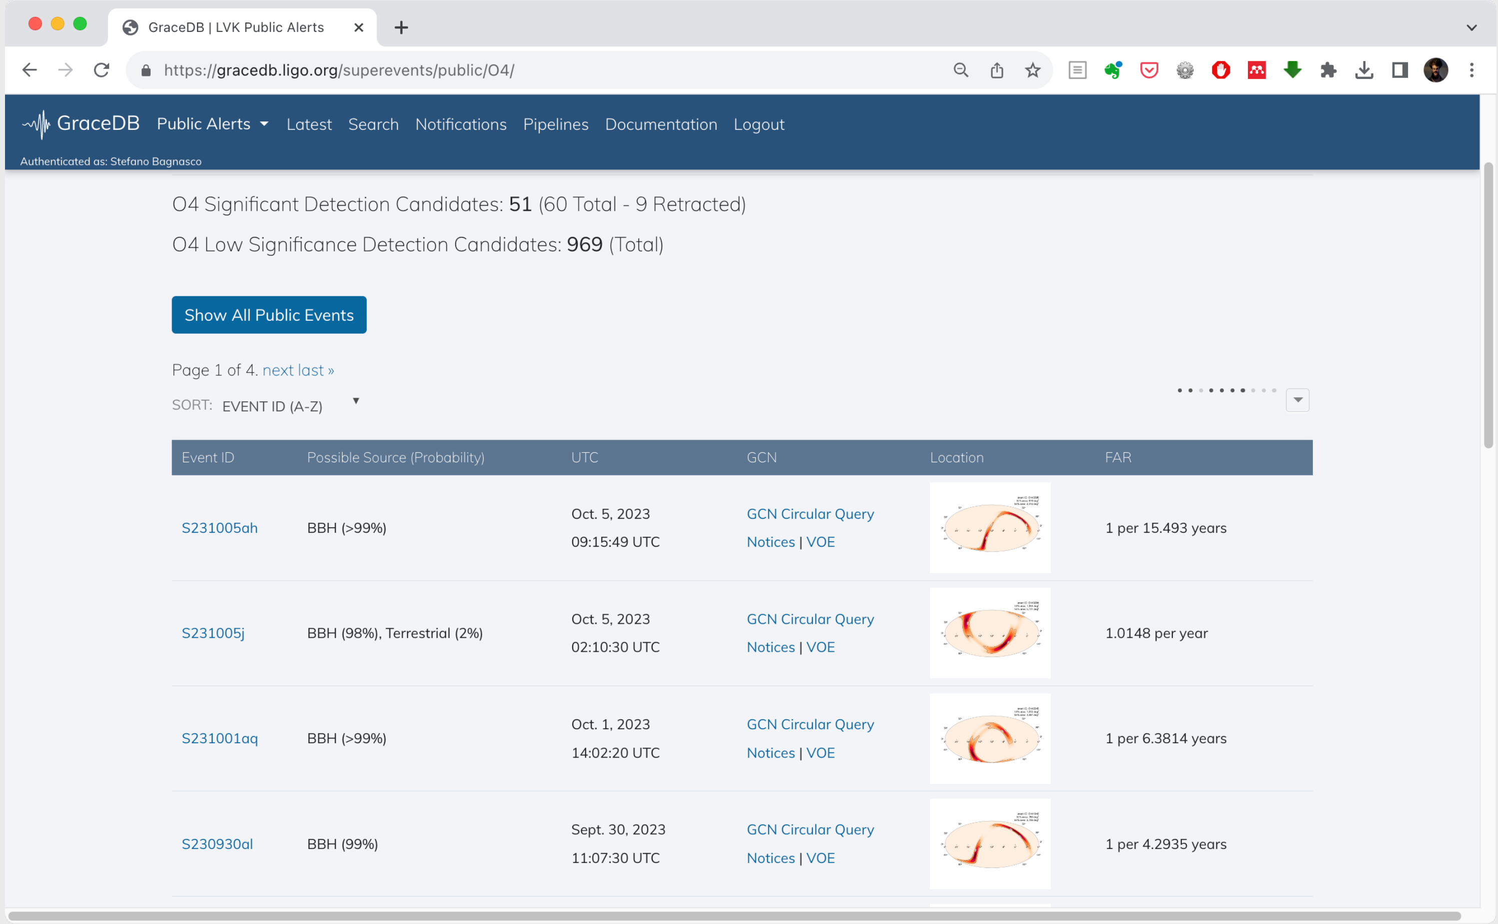

- Centralized aggregation point for information about candidate gravitational-wave events
- Provides web interface and programmatic API
- Works together with a notification service (HOPSKOTCH/Kafka-based igwn-alert)
- GWCelery: a package for annotating and orchestrating LIGO/Virgo alerts, built on the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

COMPLEX OVERALL DATA FLOWS

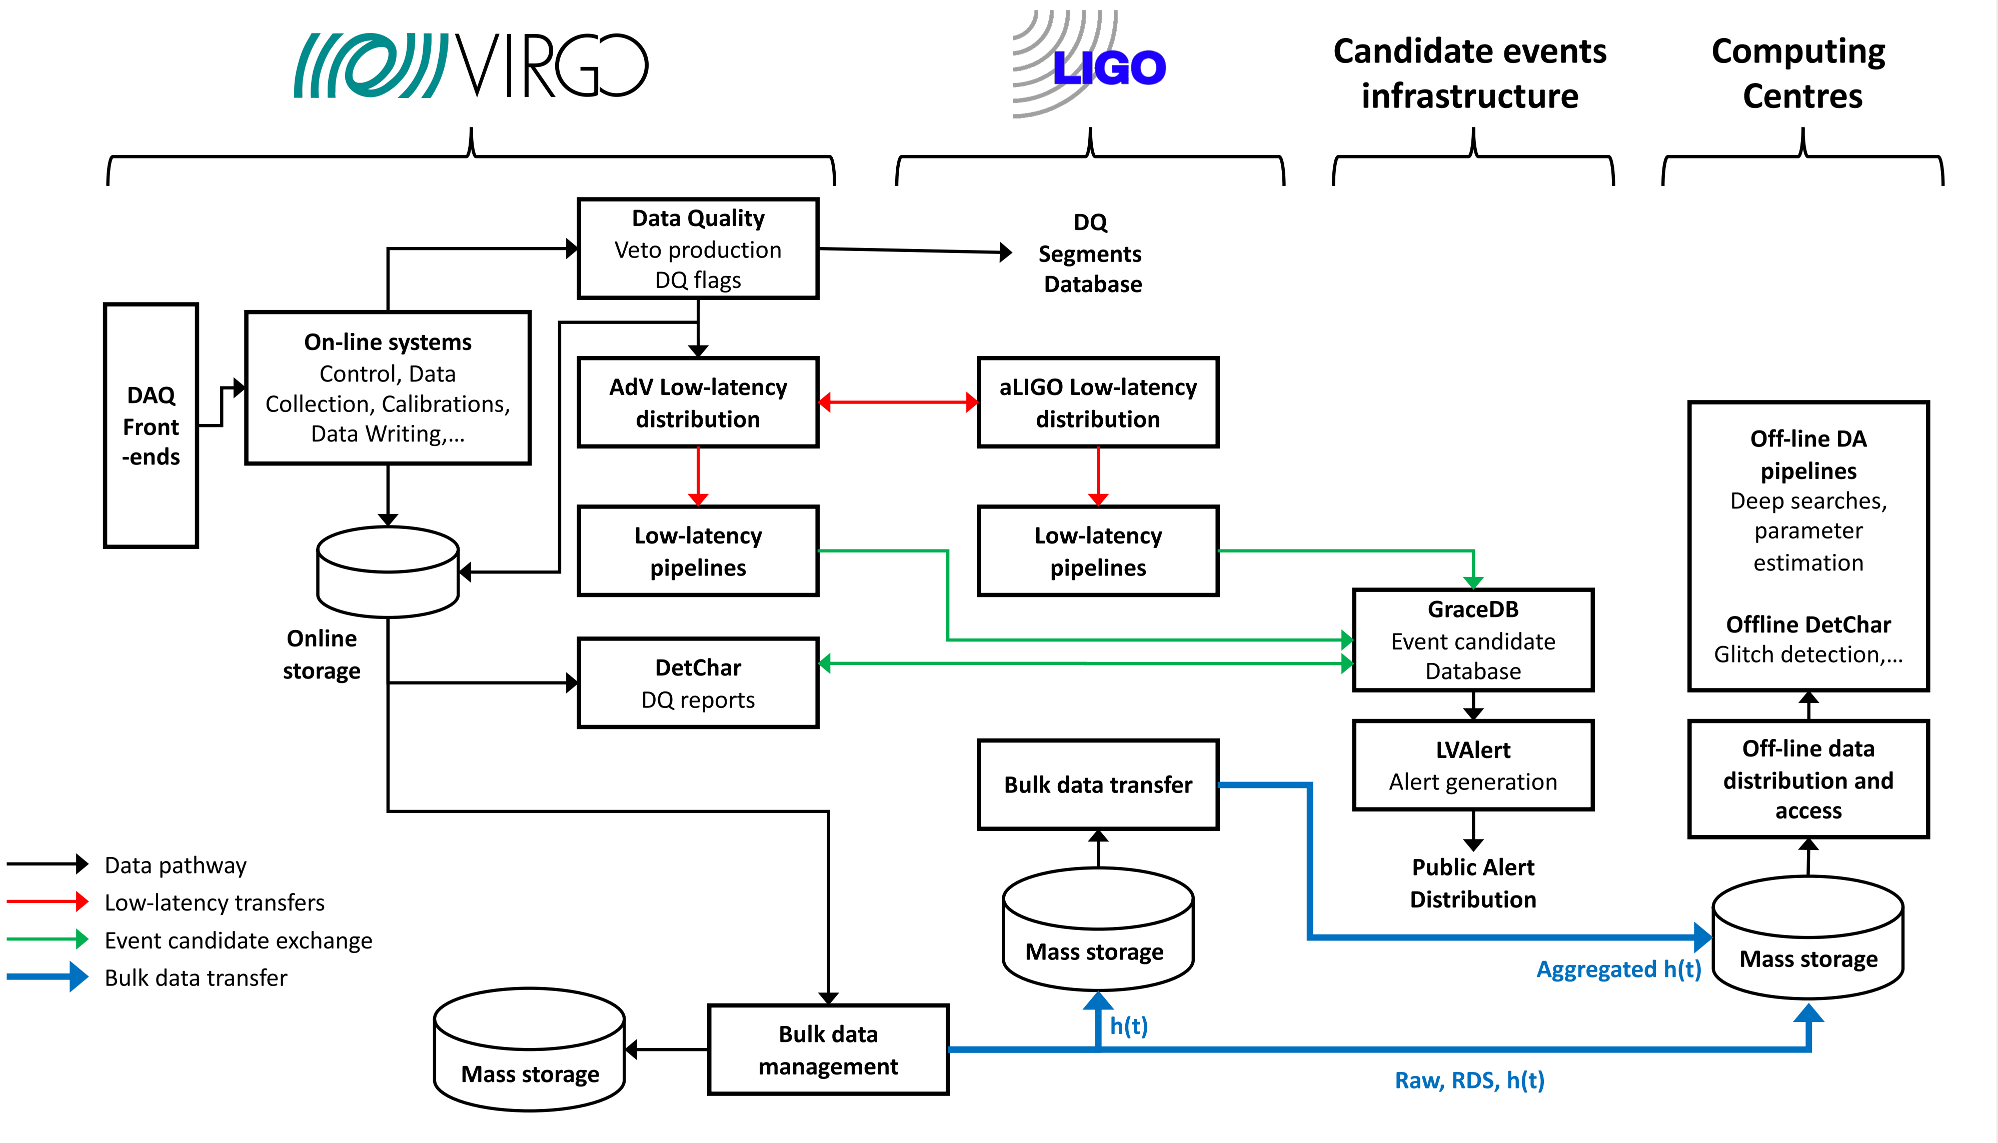

OVERALL COMPUTING RESOURCES USAGE

- 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 (and old data becomes public)
- Projected computing for O4 about 1.5÷2 × O3
- Overall: about 10% of an LHC experiment

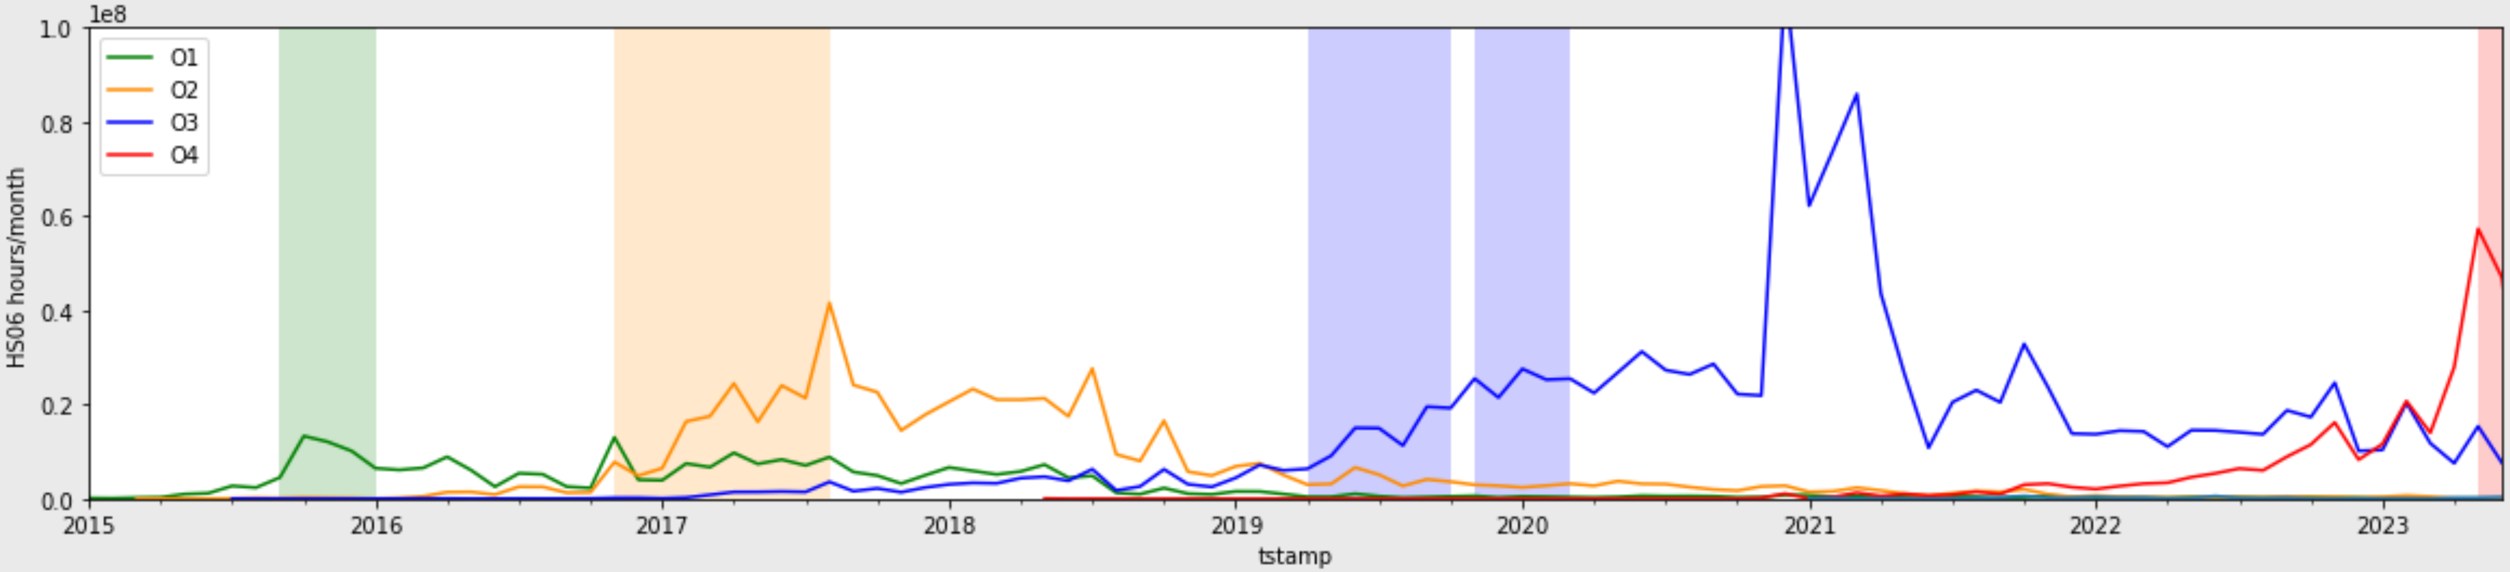

- **Raw data size:** about 120TB per month of observation per observatory
  - Includes all control channels from the instrument
  - Transferred to custodial storage for safekeeping
  - Raw data don't grow much with increasing sensitivity (they do grow with instrument complexity: 1.5 × from O3 to O4)
- **"Aggregated" data for analysis:** 10TB/year per observatory
  - Includes the single physics channel h(t) and summary "data quality" information
  - Distributed to computing centres for low-latency and offline analysis
  - Published to GWOSC after proprietary period
- **Computing:** nearly 10^9 CPU core hours
  - to process O3 data, both low-latency and offline
  - About 10% of one of the LHC experiments
  - And this does grow with sensitivity!

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.

IGWN LOWER-LATENCY DATA DISTRIBUTION

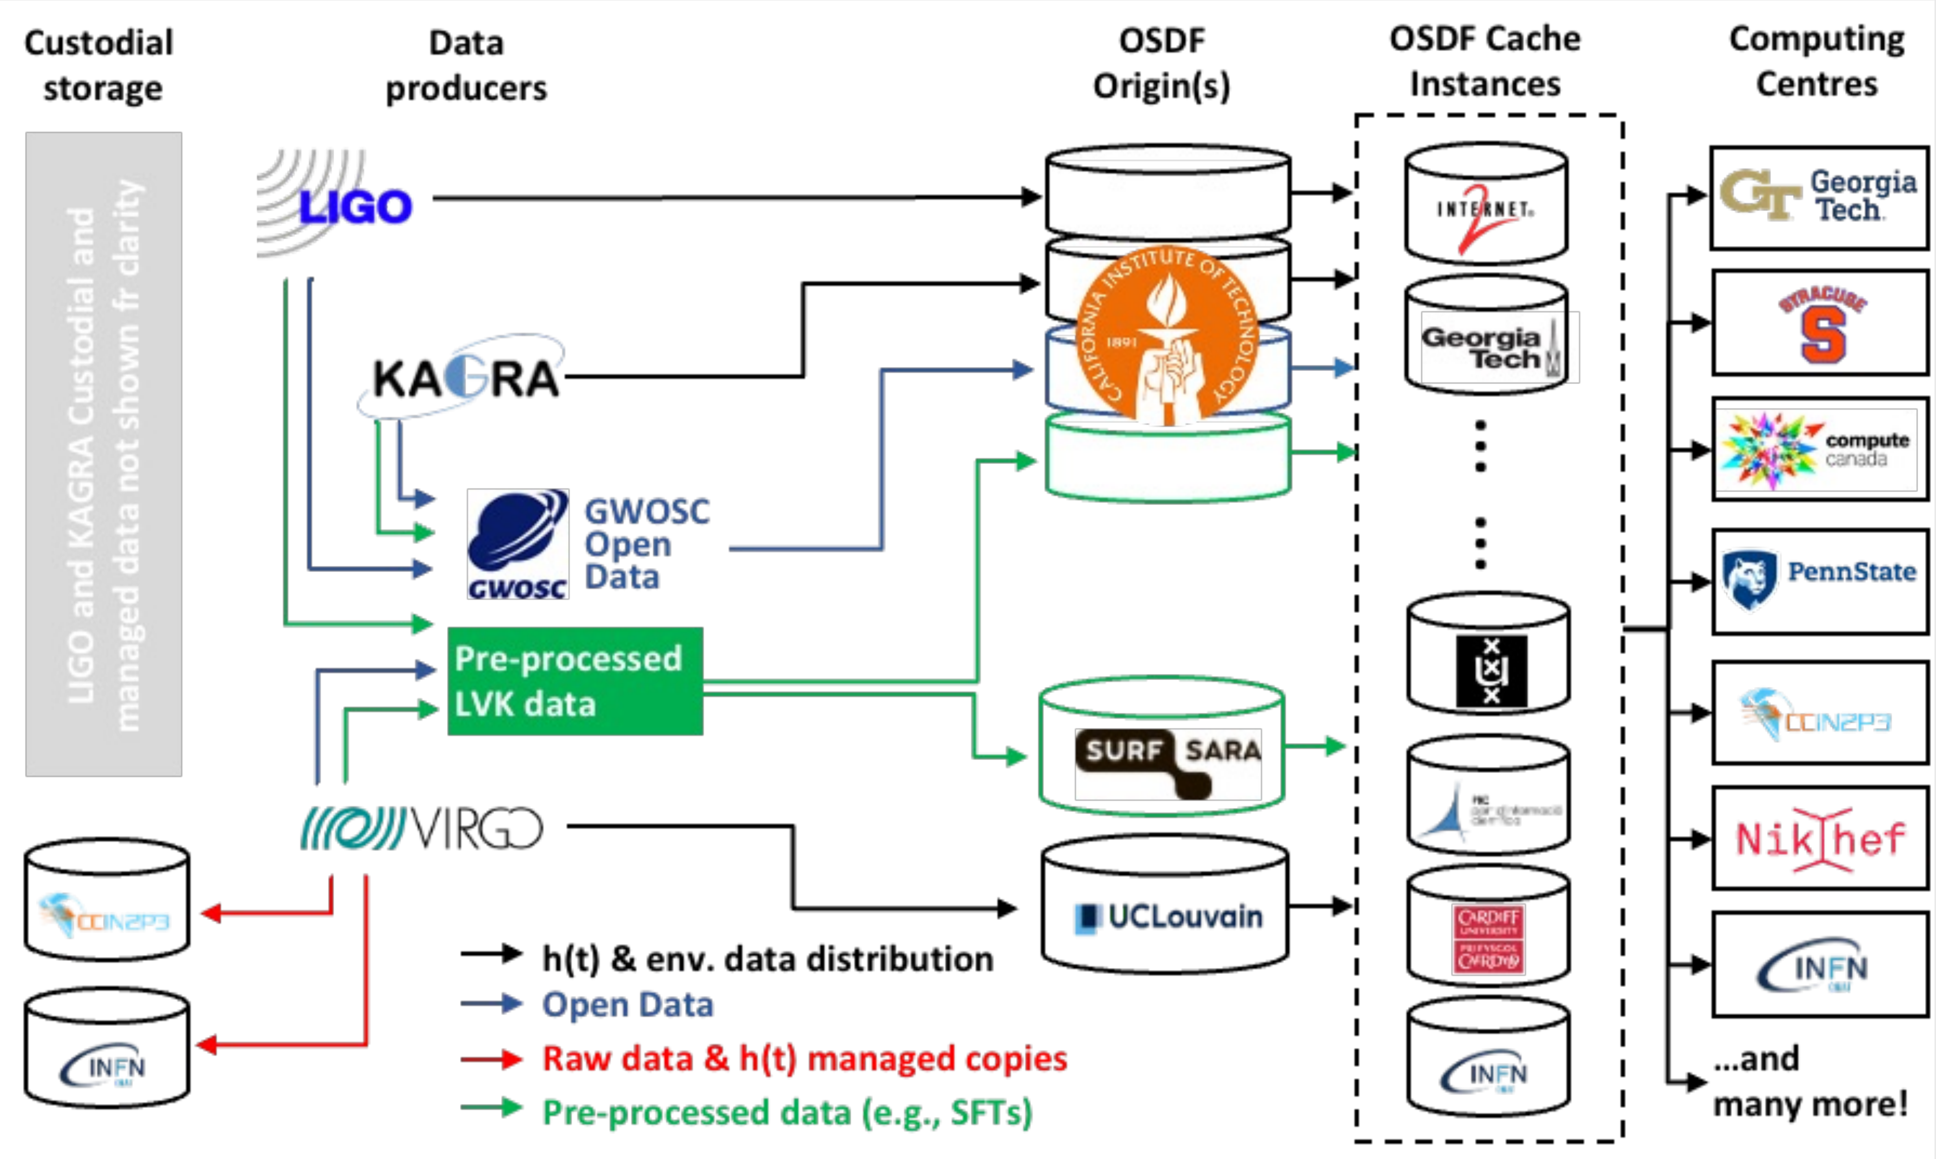

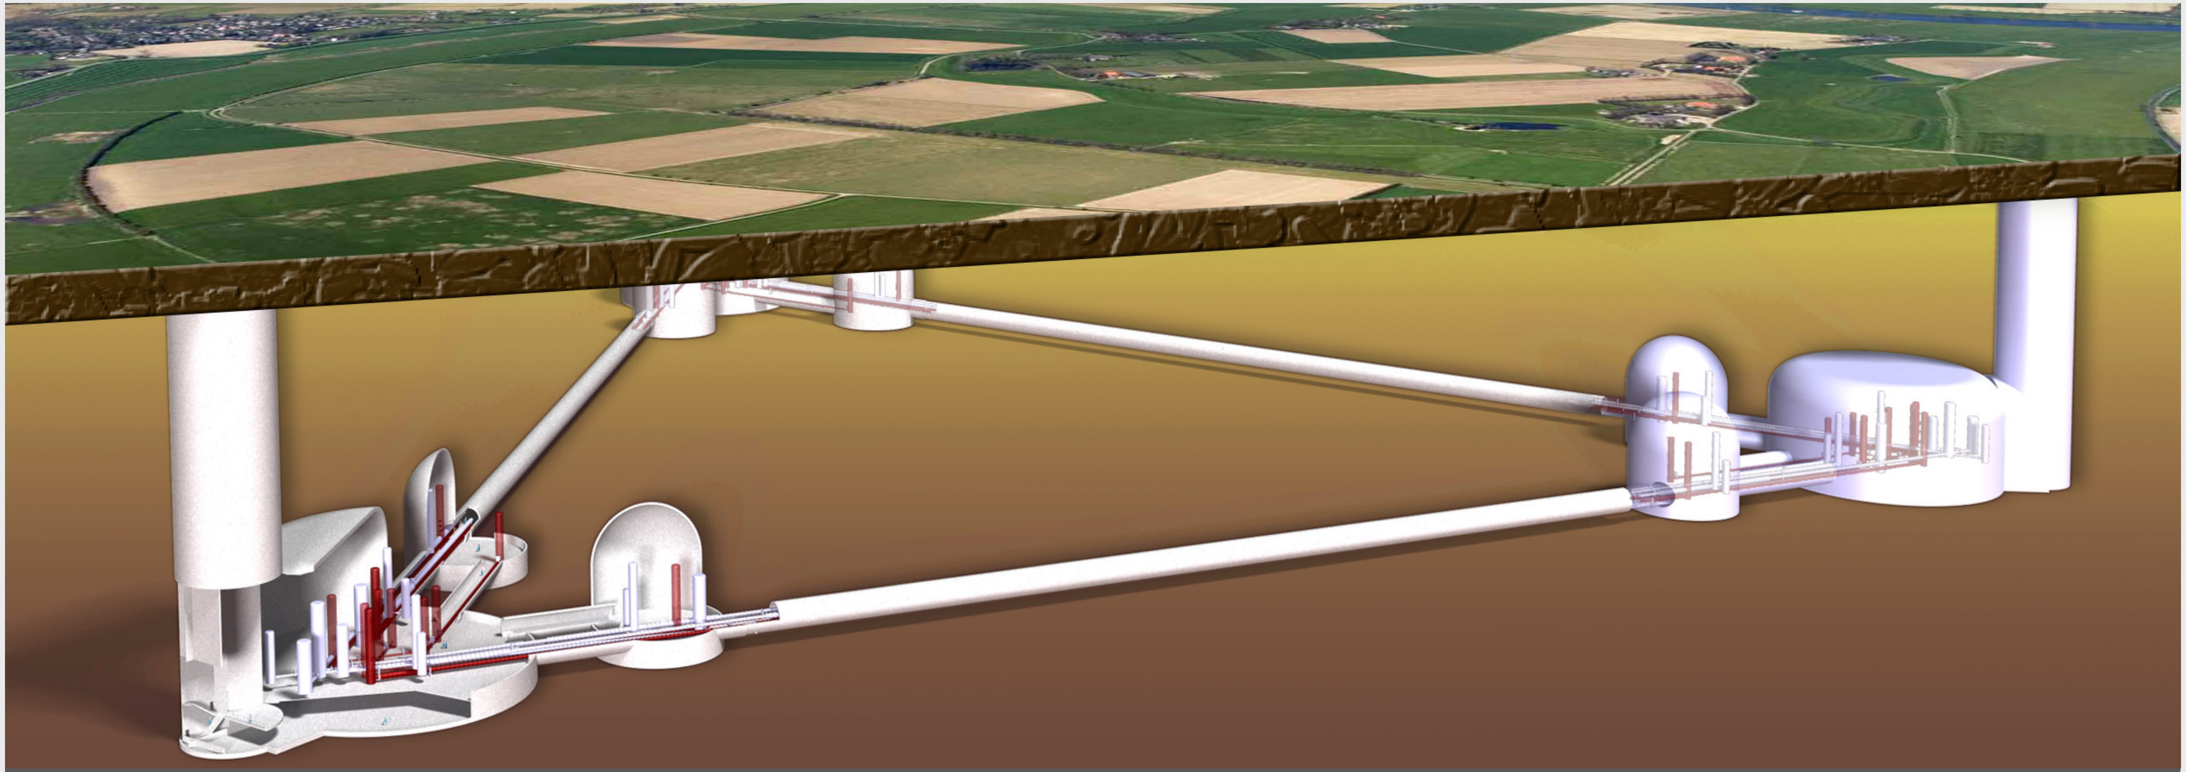

COMPUTING FOR ET

- ET is the project aiming to realise the **European 3rd Generation Gravitational Wave observatory**

- ET has been a pioneer idea that defined the concept of 3rd generation GW observatory:
  - A sensitivity at least 10 times better than the (nominal) advanced detectors on a large fraction of the detection frequency band
  - Wideband (possibly wider than the current detectors) accessing the frequency band below 10Hz
  - High reliability and improved observation capability

- ET has a long and important history that formed first the ET community and now the ET project

- ET is also a formal scientific collaboration since June 2022

- Currently there are two candidate sites being characterized to host ET:
  - The Sardinia site, close to the Sos Enattos mine
  - The Euregio Meuse-Rhine site, close to the NL-B-D border
  - A third option in Saxony (Germany) was recently proposed and is under discussion

ET: WHEN

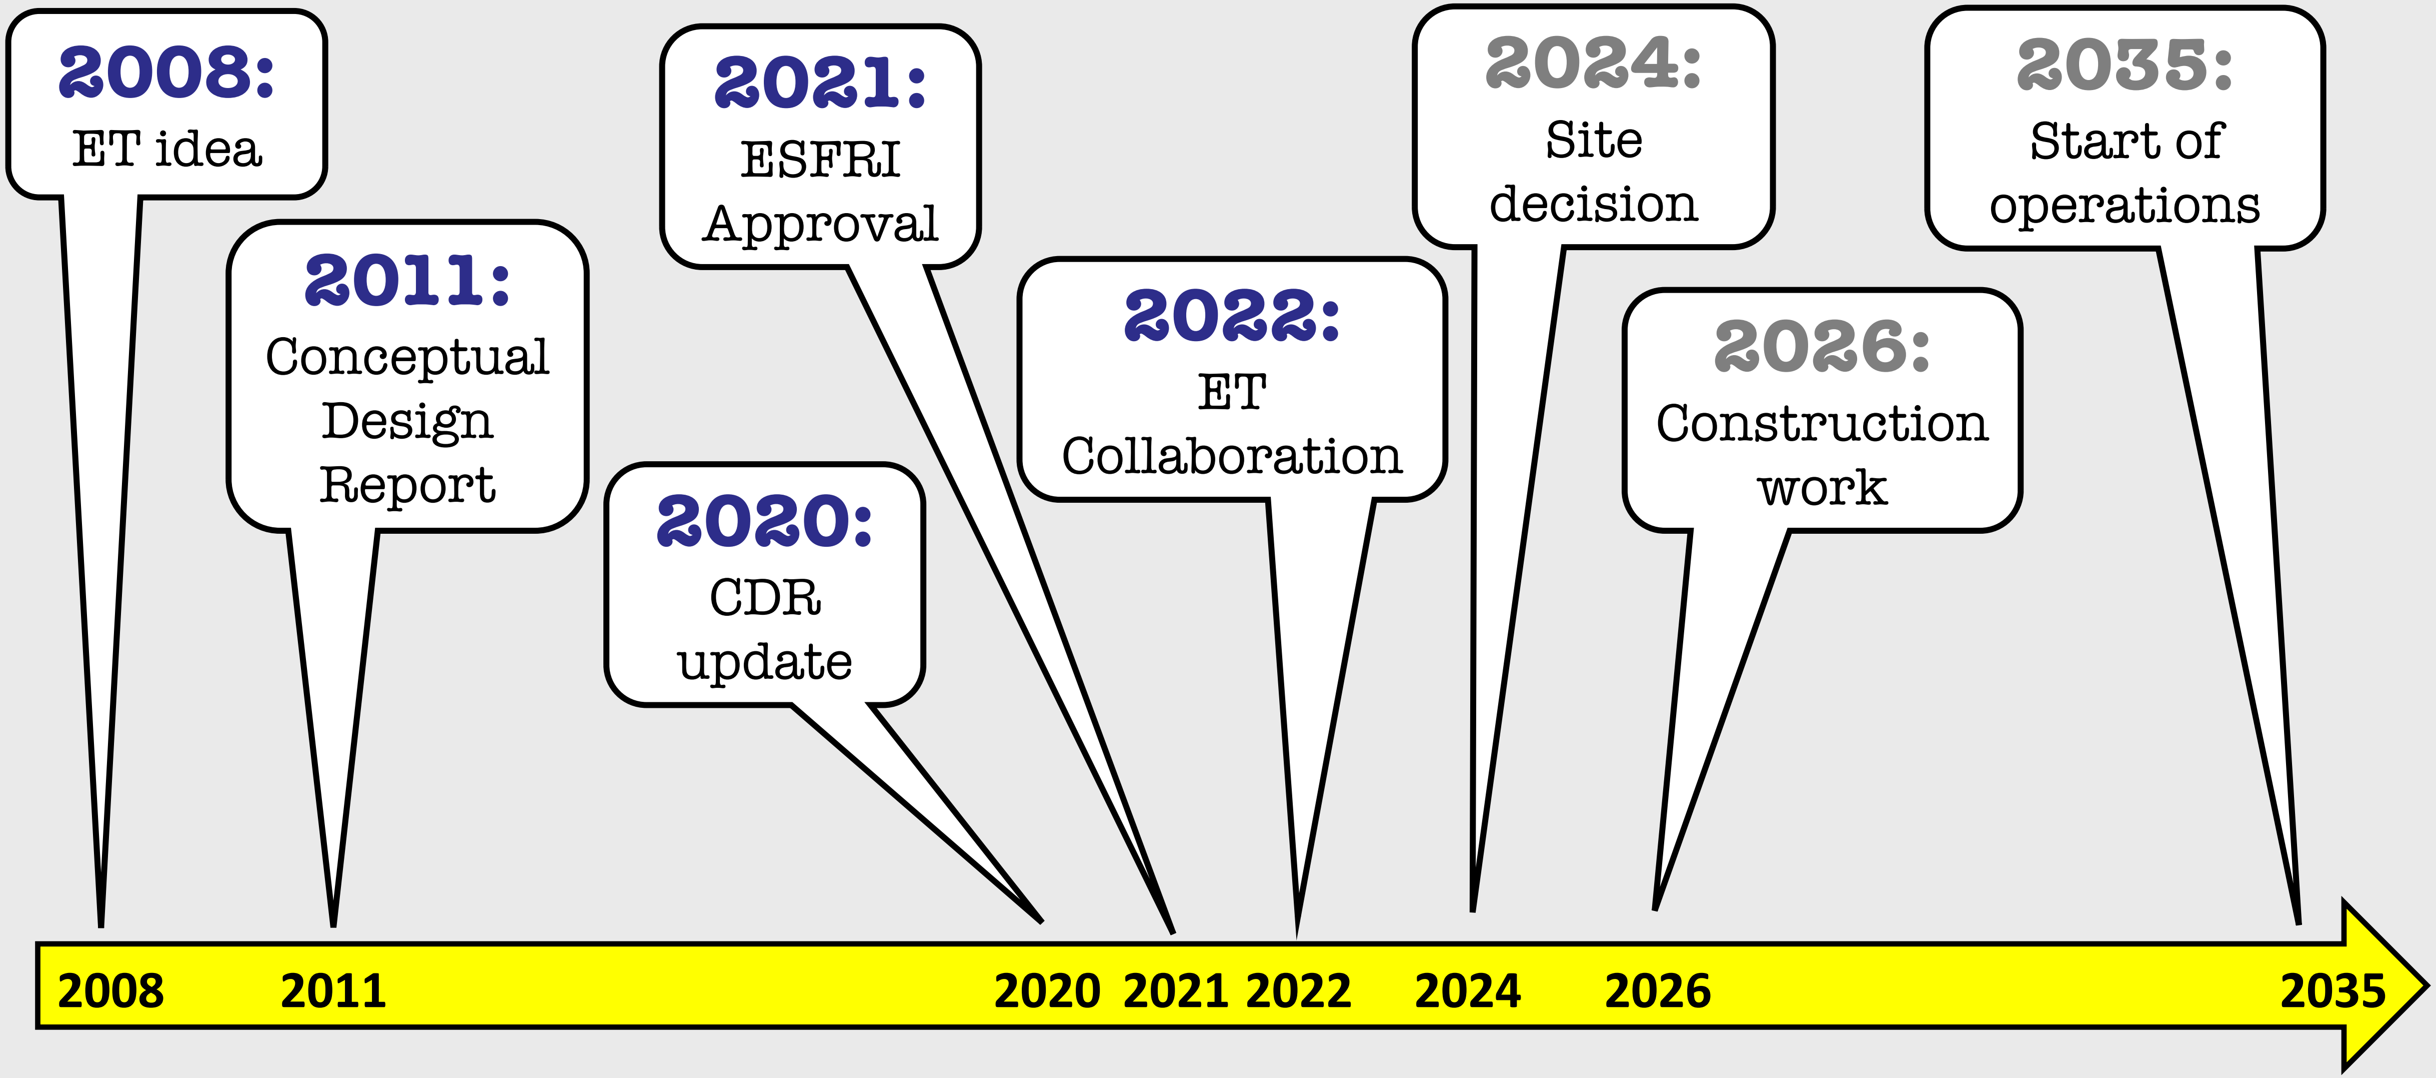
2008:
ET idea
2011:
Conceptual Design Report
2020:
CDR update
2021:
ESFRI Approval
2022:
ET Collaboration
2024:
Site decision
2026:
Construction work
2035:
Start of operations
2008
2011
2020
2021
2022
2024
2026
2035

- Three detectors in a **triangular** structure
  - Closed geometry allows the use of the null data stream
  - Extra detector adds redundancy and makes up for the non-right 60° angle
- Each detector ("red", "green" and "blue") consists of **two** Michelson interferometers
  - High-frequency and Low-frequency

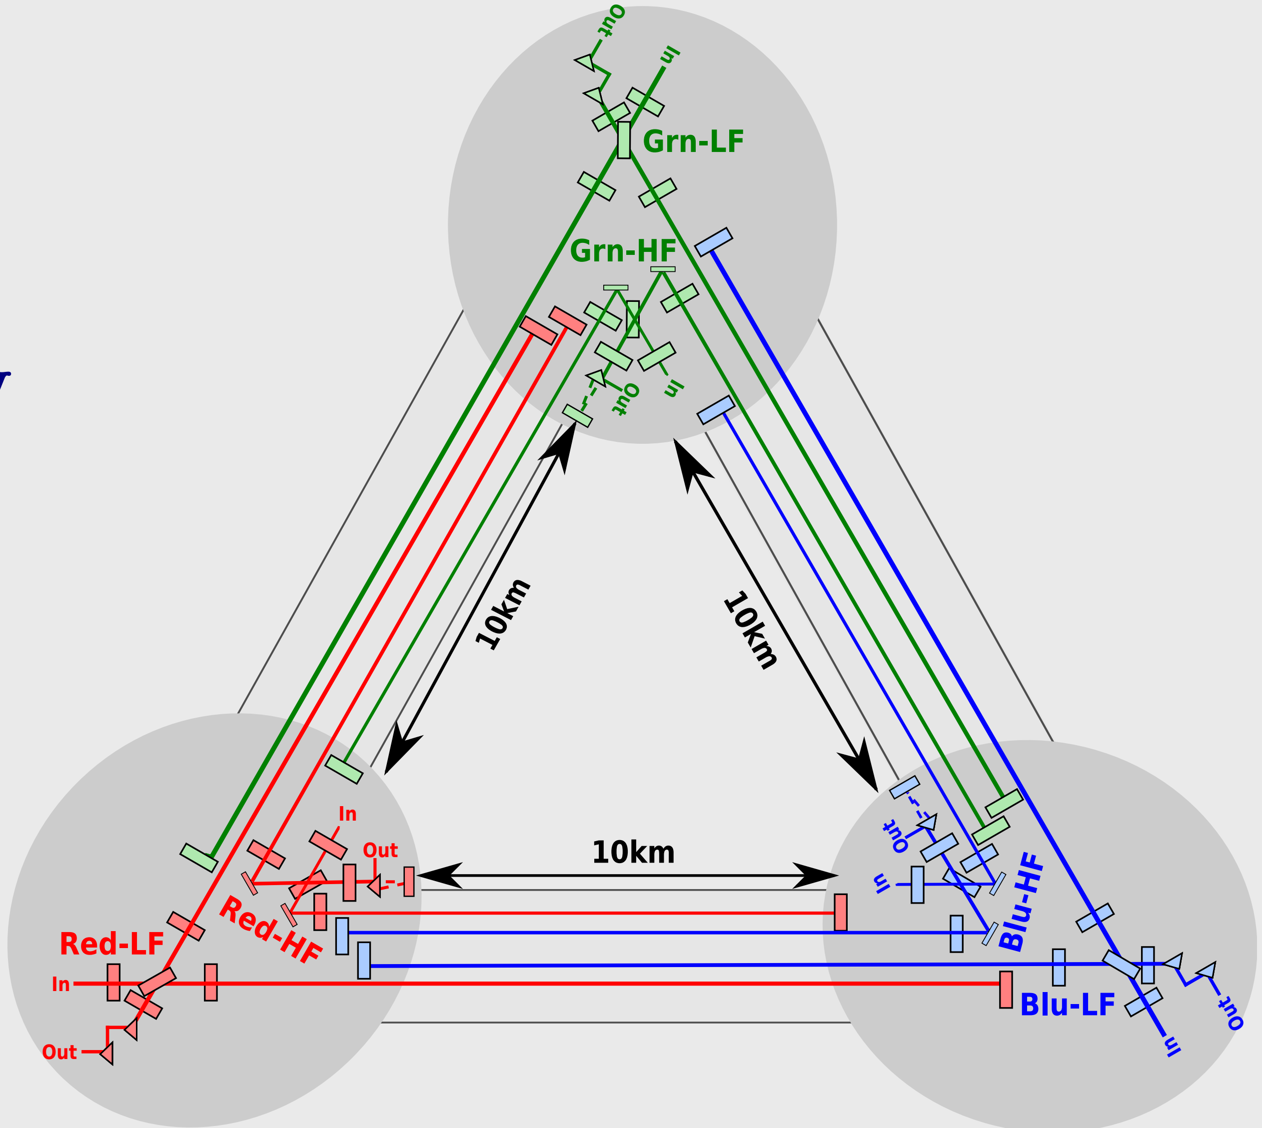

ENHANCED SENSITIVITY

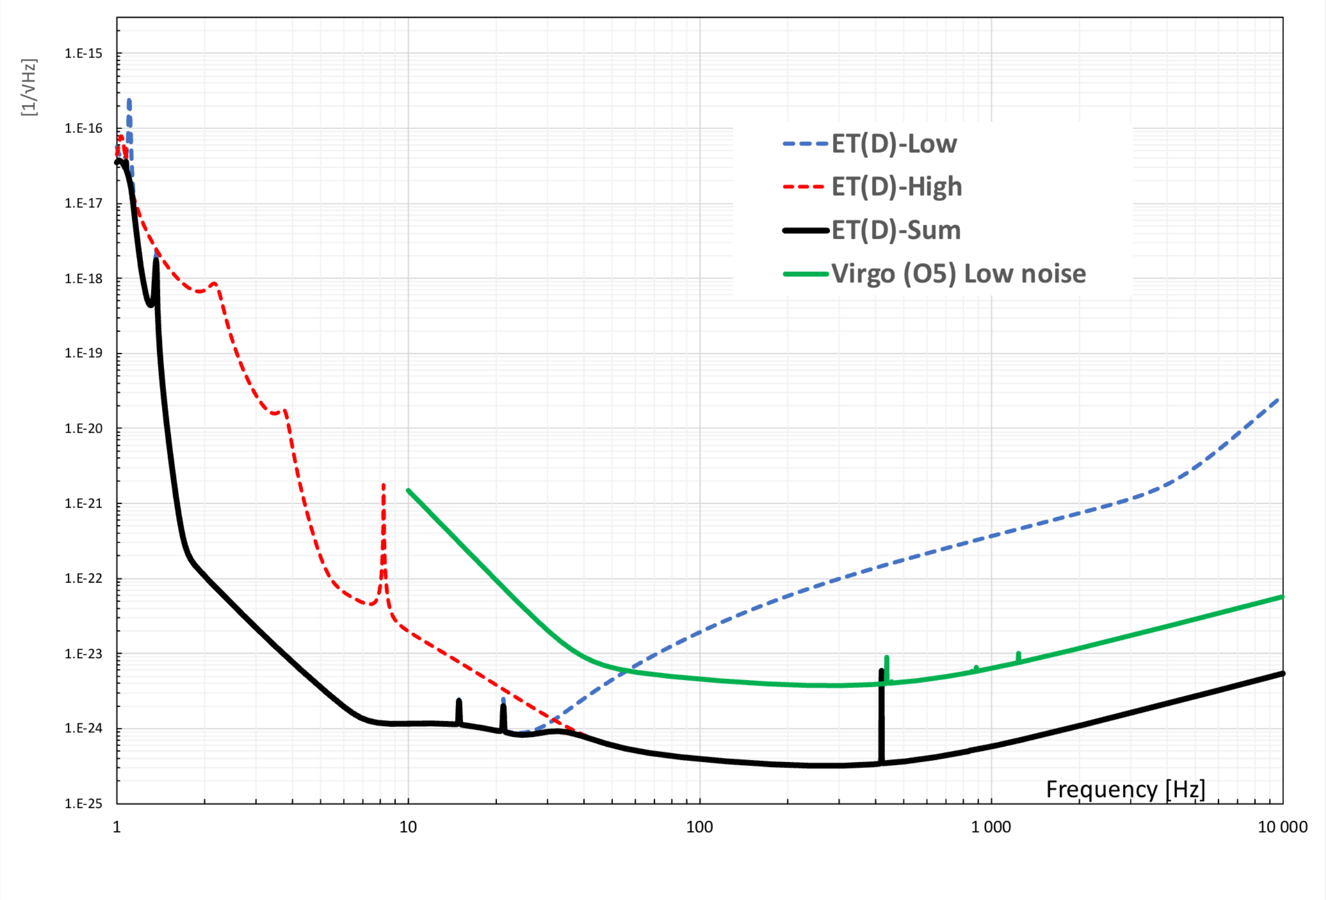

Binary Neutron Star "range"

[Mpc ≈ 3×10^{13} km ≈ 3×10^{6} ly]

	LIGO	Virgo	ET
O1	80		
O2	100	30	
O3	100-140	40-50	
O4	160-190	(60-100)	
O5	230-325	150-260	
			>2000

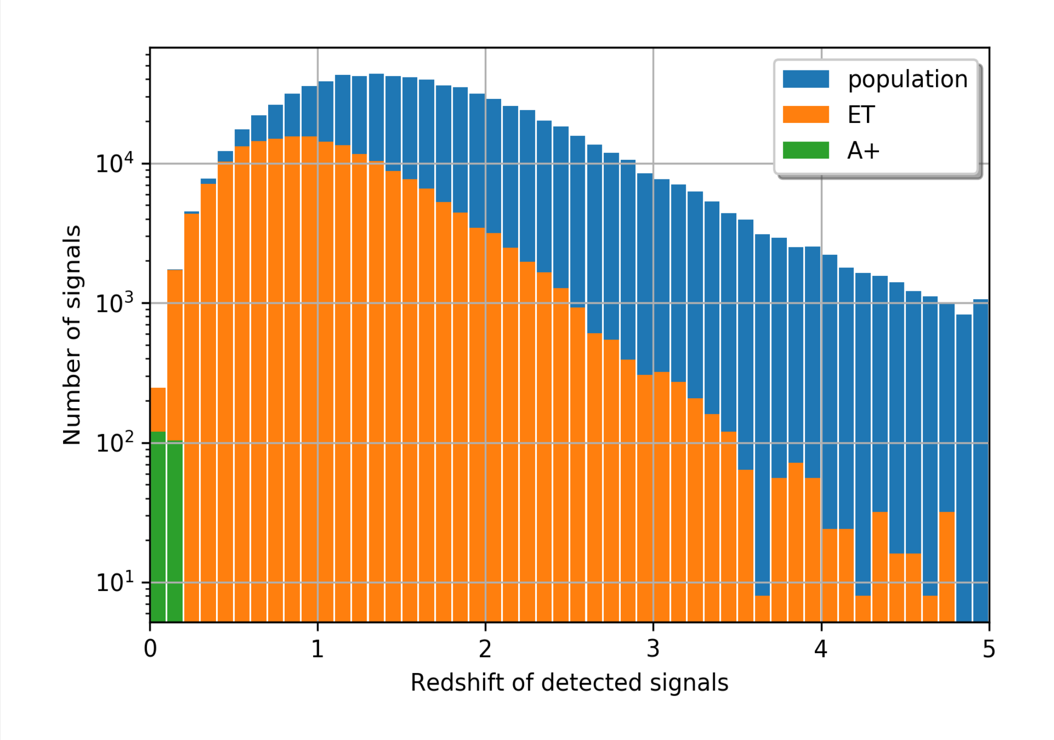

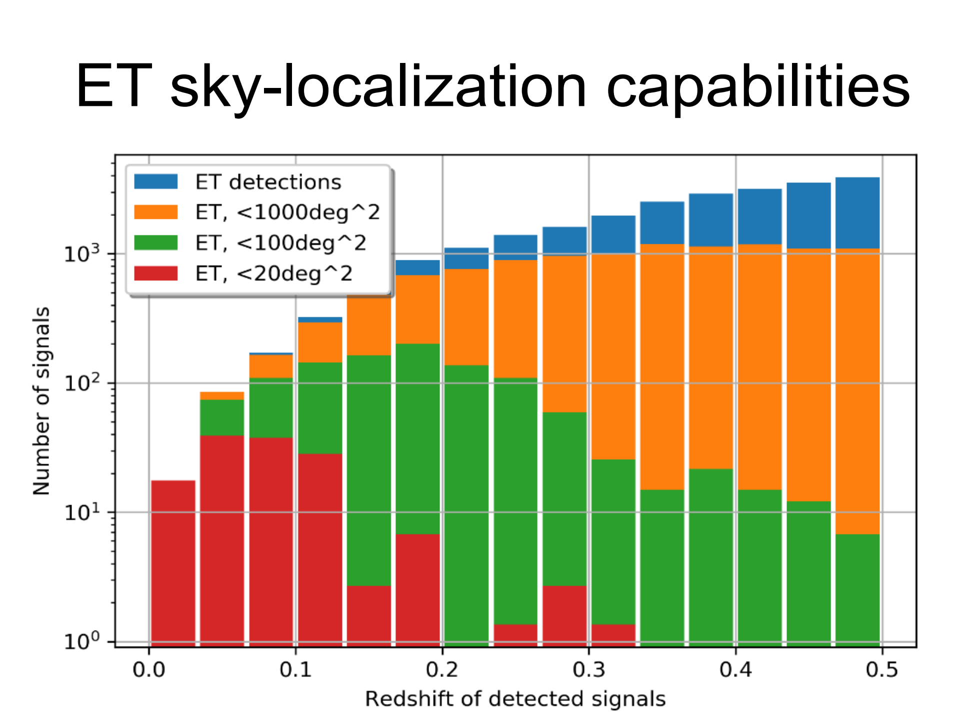

marica.branchesi@gssi.it

- 10^5 BBH detections per year
- 10^5 BNS detections per year
- o(100) detections per year with <20 deg^2
- Early warning by minutes (hours)

EARLY WARNING

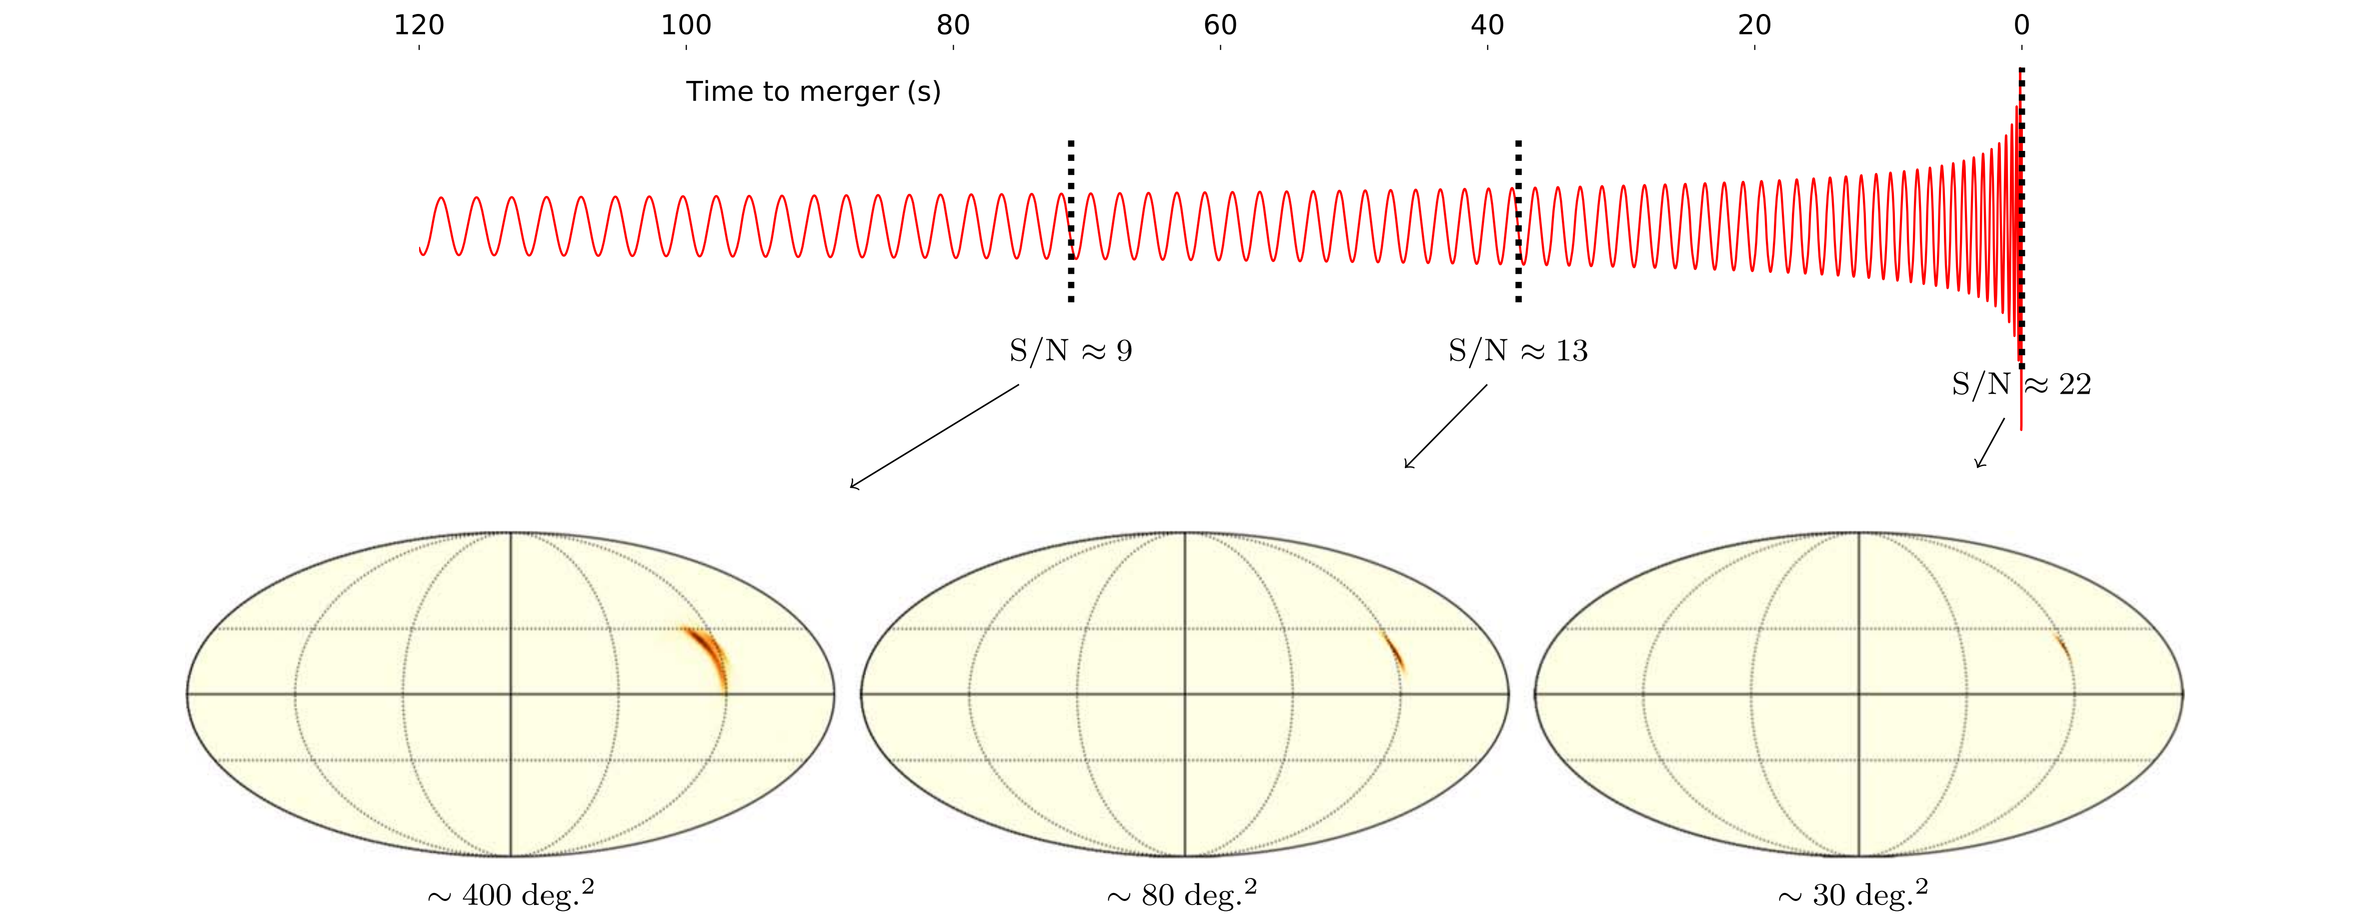

DATA AND MORE DATA AGAIN

- Again: luckily, raw interferometer data don't grow much with increasing instrument sensitivity
  - We're not exploding like HL-LHC!
  - In ET we expect about few tens of PB of raw data per year (baseline 6-interferometer design, more control channels,...)
  - No big deal today, piece of cake by 2035
- However, the amount of useful scientific information encoded in the data does grow (a lot)
  - And the computing power needed to wring it out
  - It's already a task to precisely estimate the computing power needs!

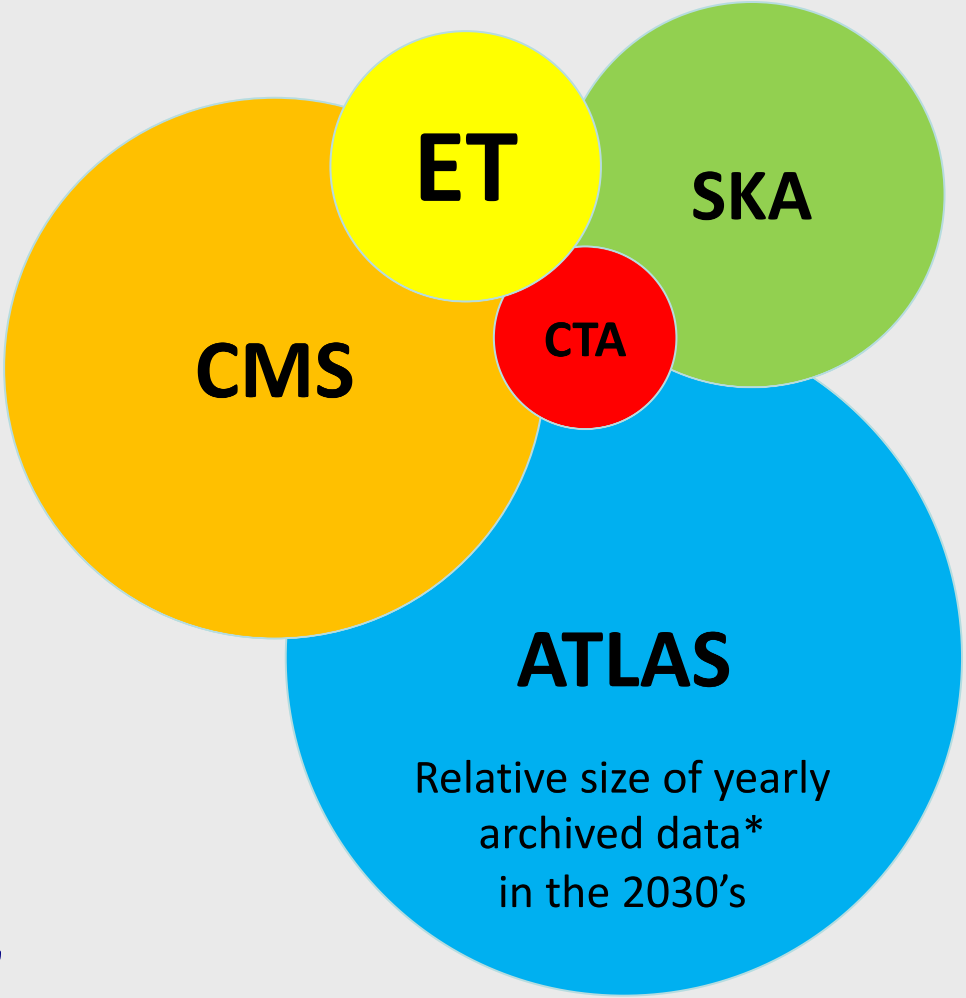

*to say nothing of weather forecast, genomics, Earth observation, oil industry, GAFAM and everybody else

THE MANDATORY SLIDE WITH BOXES AND ARROWS

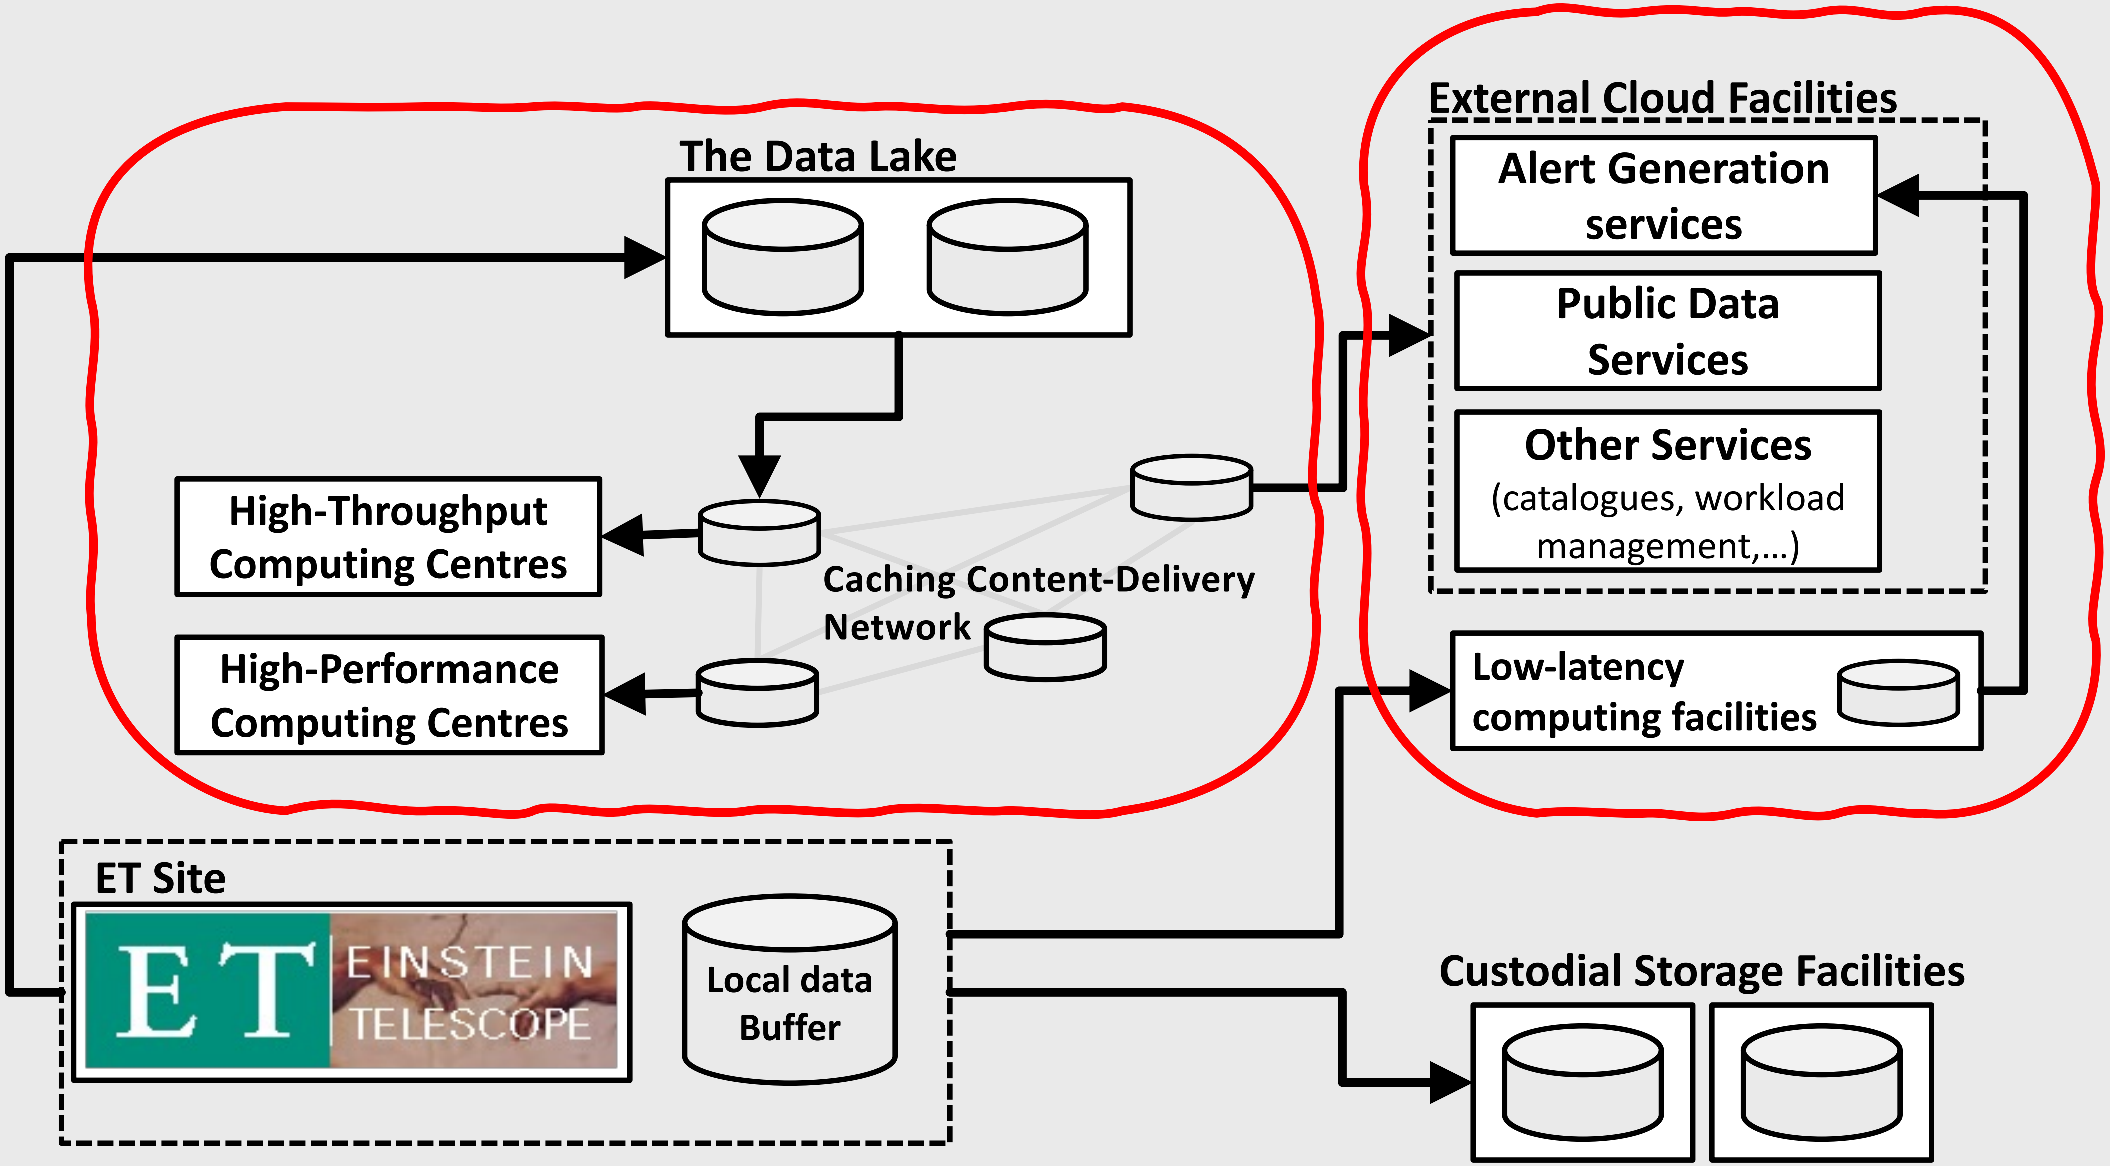

Figure from the ET ESFRI proposal

THE E-INFRASTRUCTURE BOARD

Chairs: SB, Patrice Verdier (IP2I-IN2P3)

Division 1: Software, frameworks, and data challenge support
- Chair: Andres Tanasijczuk (U. Catholique de Louvain)
- ISB Div10 liaison: John Veitch (U. Glasgow)

Division 2: Services and collaboration support
- Chair: Antonella Bozzi (EGO)

Division 3: Computing and data model, Resource Estimation
- Chair: Gonzalo Merino (PIC)

Division 4: Multimessenger alert infrastructure
- Chair: Steven Schramm (U. Geneva)

TTG: Technology Tracking Working Group
- Chair: Sara Vallero (INFN Torino)

ET-PP WP8: COMPUTING AND DATA ACCESS

Leaders: Nadia Tonello (BSC), Achim Stahl (U. Aachen)

Task 8.1: T0 data center

- Leader: Patrice Verdier (IP2I-IN2P3)

Task 8.2: Computing and Data Model

- Leader: Anastasios Fragkos (U. Geneva)

Task 8.3: Resources

- Leader: Silvio Pardi (INFN Napoli)

Task 8.4: Data Access Implementation

- Leader: Nadia Tonello (BSC)

MOCK DATA CHALLENGES

- MDC as multipurpose tools
  - Develop and exercise analysis code and strategies
  - Build the data analysis community and bootstrap new groups
  - Educate the community in the use of common distributed computing tools and best practices
  - Iteratively test the distributed computing infrastructure

- Mock Data Challenge support
  - MDC1: provide data distribution layer (OSDF: CVMFS + cache) and survey the activities
  - MDC2: provide (possibly a set of) prototype tools for workload management etc.
  - MDC3-n: iterate

Bottom line: computing needs grow more or less as a continuum from O4 to O5 to post-O5 to ET, and technologies keep evolving; not something specific to post-O5, but it extends the scope for synergies.

- Distributed computing infrastructure
  - CPU power needs grow continuously with sensitivity (CBC PE)
  - ET already needs a working and evolving computing infrastructure (for MDCs, simulations,…)
- Low-latency alert distribution network
  - High rates imply high automation, long signals imply new features (e.g., continuous alert updates)
  - In the coming years the developments may be driven by running experiments, the GW community already needs to be present
- Carbon-efficient computing
  - And, in general, technology tracking: heterogeneous computing, efficient algorithms, ML,…
  - Same message: development is a continuum

STRATEGY 1.0

- Use IGWN infrastructure as baseline
  - IGWN uses the European computing centres as an extension of the OSG (which is nonoptimal…)
  - However, the functionality is there (OSDF + HTCondor)
- Use ESCAPE as the first toolbox
  - First the "Data Lake" (DIOS), then the Virtual Research Environment
  - Also, Virtual Observatory (INFN Perugia)
- Develop a common initial R&D program
  - Data Lake (Rucio) for data distribution
  - VRE for data access and job management
  - …

1/10TH OF AN LHC EXPERIMENT

- Current computing needs of the entire GW network are roughly o(10%) of an LHC experiment of today

- In ET the event rate will be 10^3 -10^4 times the current one
  - Analysis of the “golden” events (EM counterparts, high SNR or “special” events) would already be within reach using current technologies
    - O(500) events per year = 12.5MHS06-y per year, the same order of magnitude of a LHC experiment in Run 4
    - Target: 1/10th of an LHC experiment in Run 4

- But: low-latency!

CONCLUSIONS

- The Einstein Telescope instrument itself will generate similar quantities of data to Virgo/LIGO
  - BUT it will contain many more, longer, stronger, overlapping signals
- IGWN is the de facto starting point for Einstein Telescope, and it's a great start
  - Distributed computing model, borrowing infrastructure particularly in the US (OSG)
  - Low latency workflows require experiments to share data!
  - ESCAPE is an obvious candidate as a toolkit we will use to start building infrastructure prototypes
- Offline analysis is an active topic of research
  - AI/ML will likely play an important role
  - Expectation that computing demands stays comparable with 1/10th of a Run5 LHC experiment
- ET will open new possibilities, and create new challenges, for multi-messenger astronomy in the time domain
  - Expect our understanding to evolve significantly in the next years, and meanwhile other observatories likely to lead the way in infrastructure (particularly Vera Rubin observatory)