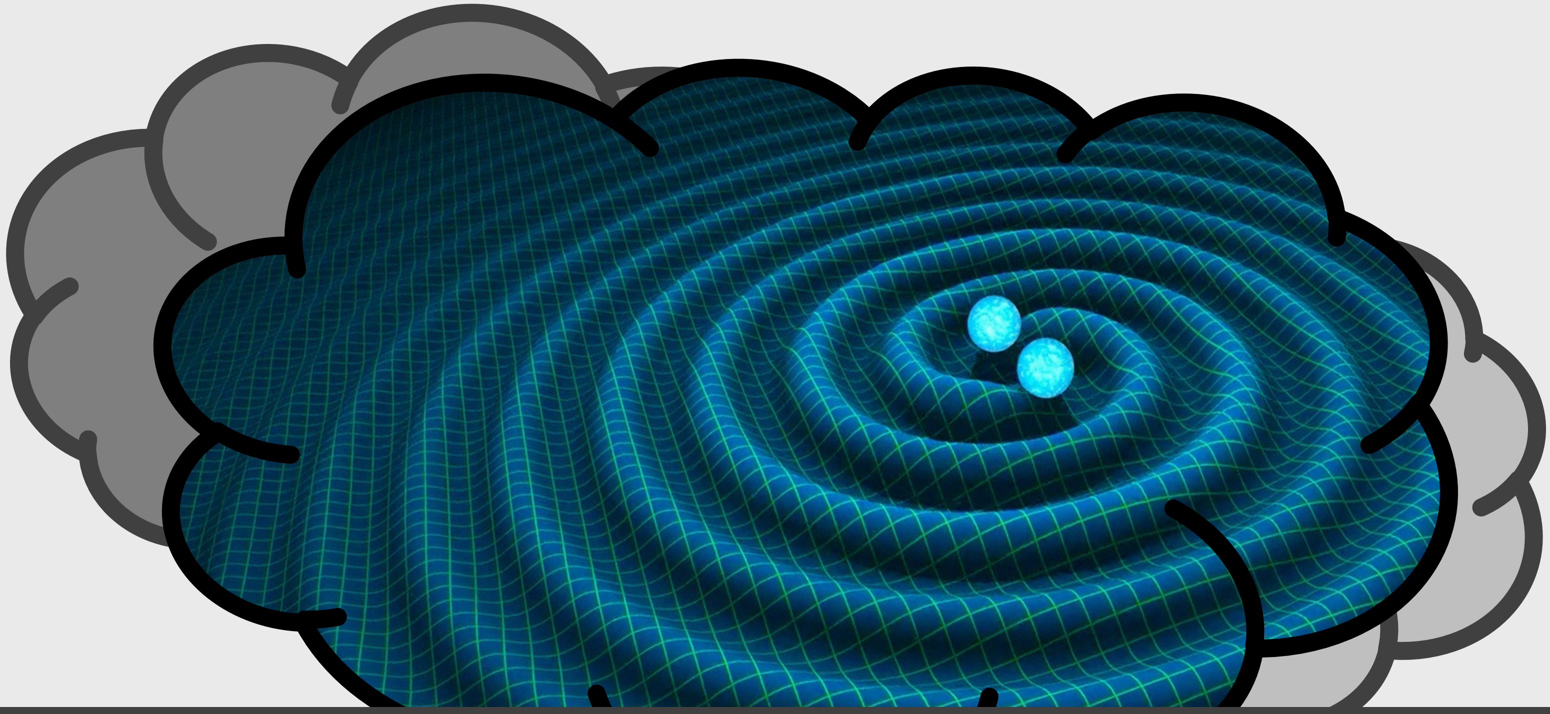

EINSTEIN TELESCOPE & VIRGO

Stefano Bagnasco, INFN

Riunione C3SN

VIRGO EINSTEIN TELESCOPE

Sep 4, 2025

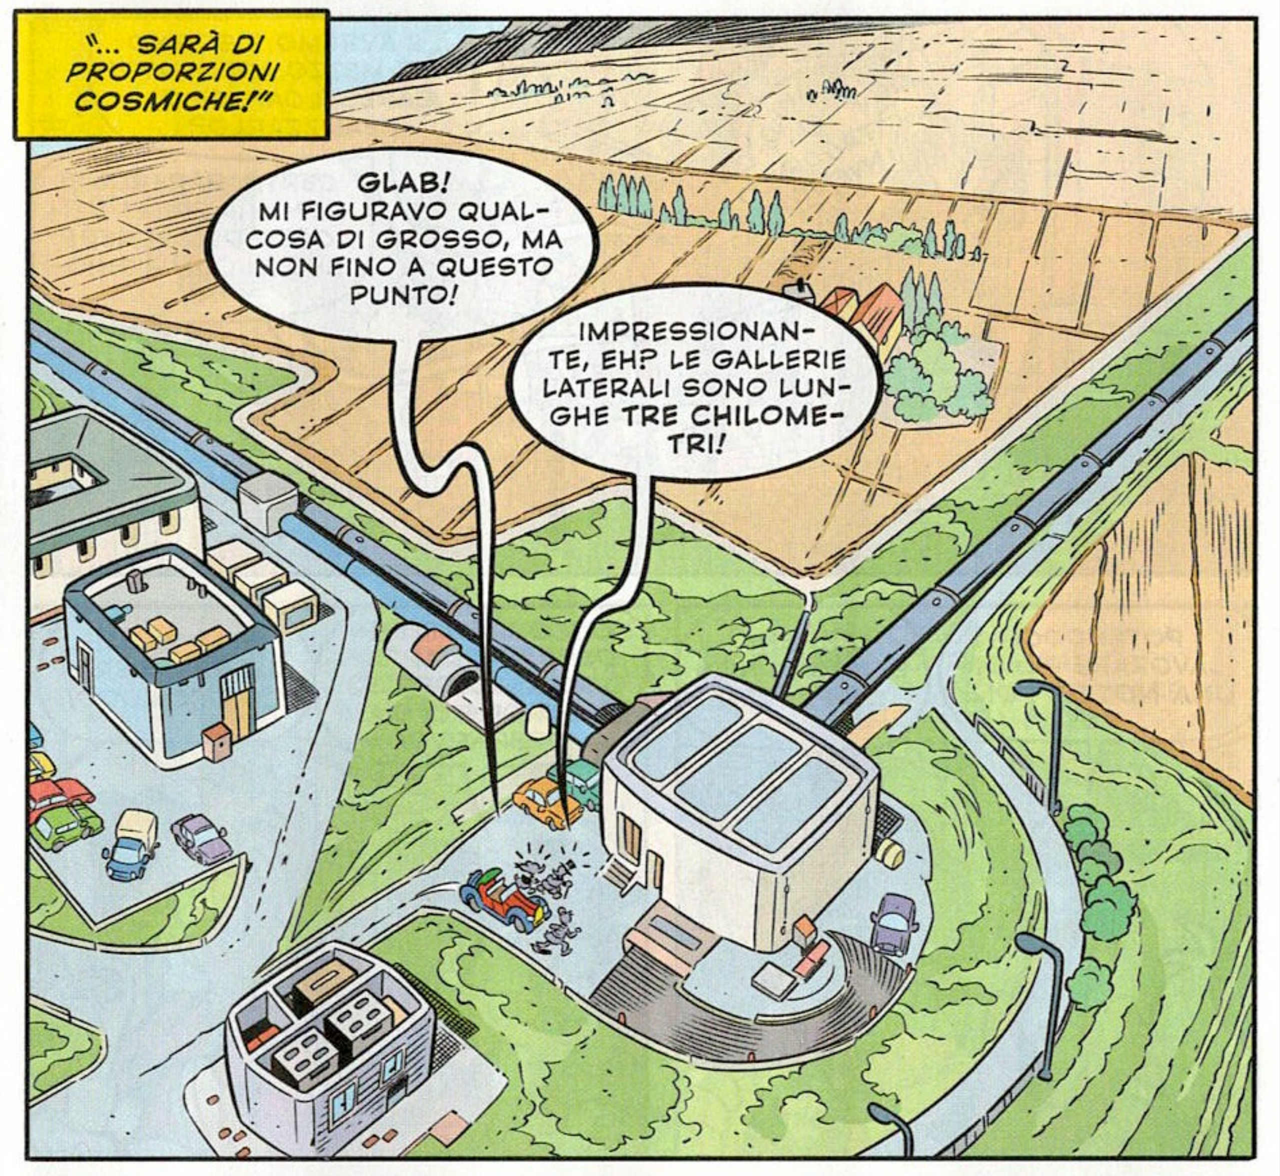

M. Bosco, G. Soldati, "Sogni d'oro zio Paperone"
Topolino **3538**:45-70 (2023)

WIP

GW231123: THE MOST MASSIVE BLACK HOLE BINARY DETECTED THROUGH GRAVITATIONAL WAVES

On November 23, 2023, at 13:54:30 UTC, the LIGO-Virgo-KAGRA (LVK) collaboration detected GW231123, a gravitational-wave signal likely caused by the merger of two black holes with the highest combined mass the LVK collaboration has ever observed. These black holes would have been spinning incredibly fast, and their individual masses appear to fall into a range that challenges existing theories about how massive stars evolve and end their lives.

DETECTING THE SIGNAL

This gravitational wave was observed by the two Advanced LIGO detectors in Hanford and Livingston during the first part of the fourth LVK observing run (O4a). The coherence between the two observatories was essential in making a confident detection. As shown in **Figure 1**, the signal lasted about a tenth of a second but stood out clearly, about **20 times louder** than the typical detector noise. To ensure this was not a random blip in the data, we performed careful statistical checks. Using techniques that simulate thousands of years' worth of fake data, we found that the probability of random noise mimicking GW231123 is less than once in 10,000 years! This gives us *extreme* confidence in the non-terrestrial origin of the signal, and thus in the reality of this gravitational-wave signal.

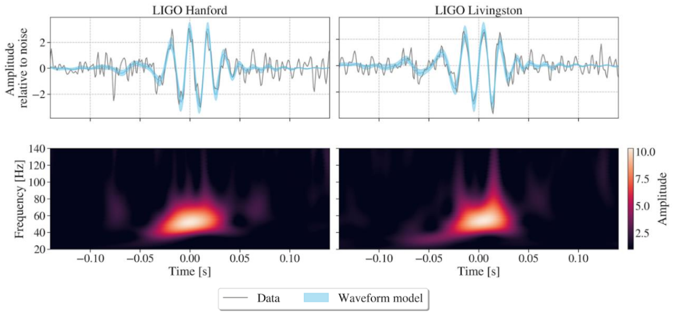

Figure 1: GW231123 signal in data from the LIGO Hanford (left) and Livingston (right) detectors. The top panels show the amplitude of the data over time (grey traces). The shaded blue band shows our estimate of the true signal. The bottom panels are spectrograms, also known as time-frequency maps, which show the signal amplitude over time (horizontal axis) and across frequencies (vertical axis). Brighter colors represent a stronger signal.

THE SOURCE BEHIND THE SIGNAL

The data strongly suggest that this signal came from the violent merger of two black holes. To learn more about these black holes—like how massive they were and how fast they were spinning—we used several models based on Einstein's theory of general relativity to simulate what such a signal would look like for different black hole pairs.

By comparing the data to these models, we found that these black holes weighed approximately 137 and 103 times the mass of the Sun, respectively. Taking all uncertainties into account, their total mass was likely between 190 and 265 solar masses, dethroning GW190521 as the most massive black hole binary observed so far.

FIND OUT MORE:

Visit our websites: www.ligo.org
www.virgo-gw.eu
gwcenter.icrr.u-tokyo.ac.jp/en/

GWTC-4.0: UPDATING THE CATALOG WITH OBSERVATIONS FROM THE FIRST PART OF THE FOURTH LIGO-VIRGO-KAGRA OBSERVING RUN

In August 2025, the LIGO-Virgo-KAGRA (LVK) Collaborations released the interferometric strain data from the first part of the fourth observing run (O4a), which ran from May 2023 to January 2024. In this data, we have discovered 128 new confident gravitational-wave (GW) signals originating from merging black holes and neutron stars. Alongside the strain data release, we publish version 4.0 of the Gravitational Wave Transient Catalog (GWTC-4.0), which contains lists of the candidate signals and measurements of their properties. We are also publishing a set of papers that accompany the catalog. These papers are submitted to the Astrophysical Journal Letters for publication as a Focus Issue. Here, we summarise the first three of these papers that focus on the production and results from GWTC-4.0 itself.

INTRODUCTION

This paper, "GWTC-4.0: An Introduction to Version 4.0 of the Gravitational-Wave Transient Catalog", provides an overview of the focus issue, including details of the other papers and important information to help a reader understand the nomenclature of the field.

What's in a name?

We name our gravitational waves using the date and time that we detect them, in Coordinated Universal Time (UTC). The (chronologically) first new event in our Catalog, denoted GW230518_125908, was detected on 18 May 2023 at 12:59:08 UTC. We add the prefix GW to all of the candidate signals we detect (this is slightly different to what we did in the past, as signals which we think were likely to be caused by terrestrial effects are now included in the catalog, but with additional data about why we think they might not be genuine gravitational waves).

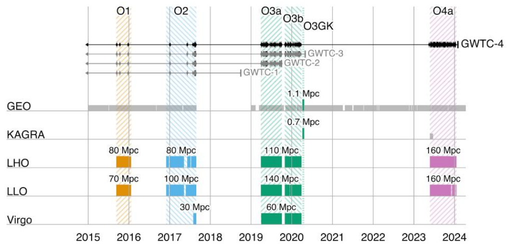

Figure 1: A timeline of observing runs showing the data-taking periods for the gravitational-wave observatories GEO, KAGRA, LIGO-Hanford (LHO), LIGO-Livingston (LLO), and Virgo. Numbers above the coloured blocks correspond to the approximate distance the detector can see a standard binary neutron star merger, providing a measure of the sensitivity. Along the top, we provide markers for when events have been added to the GWTC as well as horizontal bars showing the span of data in each catalog.

Cumulative catalogs

The GW transient catalog is cumulative - in other words, GWTC-4.0 contains not only the new results observed in O4a but also all past catalogs. To illustrate this, in **Figure 1** we show a timeline of all observing runs (colored bands) and the data span of each catalog. You'll notice that our catalogs follow a naming convention GWTC-<major>.<minor>, where the major number is incremented when the span of data increases and the minor number increments if there is a change in the methods/data describing the signals (e.g. a re-analysis of the data as happened with GWTC-2.1). In the past, we have routinely omitted the minor number when it was 0 (see, e.g. GWTC-3). However, we found that this can lead to some confusion - for example does GWTC-2 refer to GWTC-2.0 or GWTC-2.1? We will therefore now always include the minor number when referring to a specific catalog. However, we still omit the minor number when referring only to the data span (as shown in Figure 1).

From Figure 1, you will also see that the two LIGO detectors were online during O4a: all candidate signals and measurements were made on just data from these two detectors. KAGRA did briefly take data during the beginning of the run, but the sensitivity was insufficient to

FIND OUT MORE:

Visit our websites: www.ligo.org
www.virgo-gw.eu
gwcenter.icrr.u-tokyo.ac.jp/en/

Conclusions

- Virgo O5 upgrade is a big enterprise, involving also infrastructural modifications to the Central Building
  - Good prospects for increasing the sensitivity, reducing the gap with LIGO
  - Will keep us busy for the next years

- First version of the TDR delivered in May 2025
  - Consolidated resource loaded schedule will be available when money matrix (who pays what and when) will be finalized
  - Second version of the TDR will come after incorporating recommendations from External Review Committee (that gave a positive evaluation of the project)

- Project Management
  - Given the nature/size of the Project, the management structure has been deeply revised
  - All major processes are defined by appropriate documentation

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

THE GENERAL MODEL

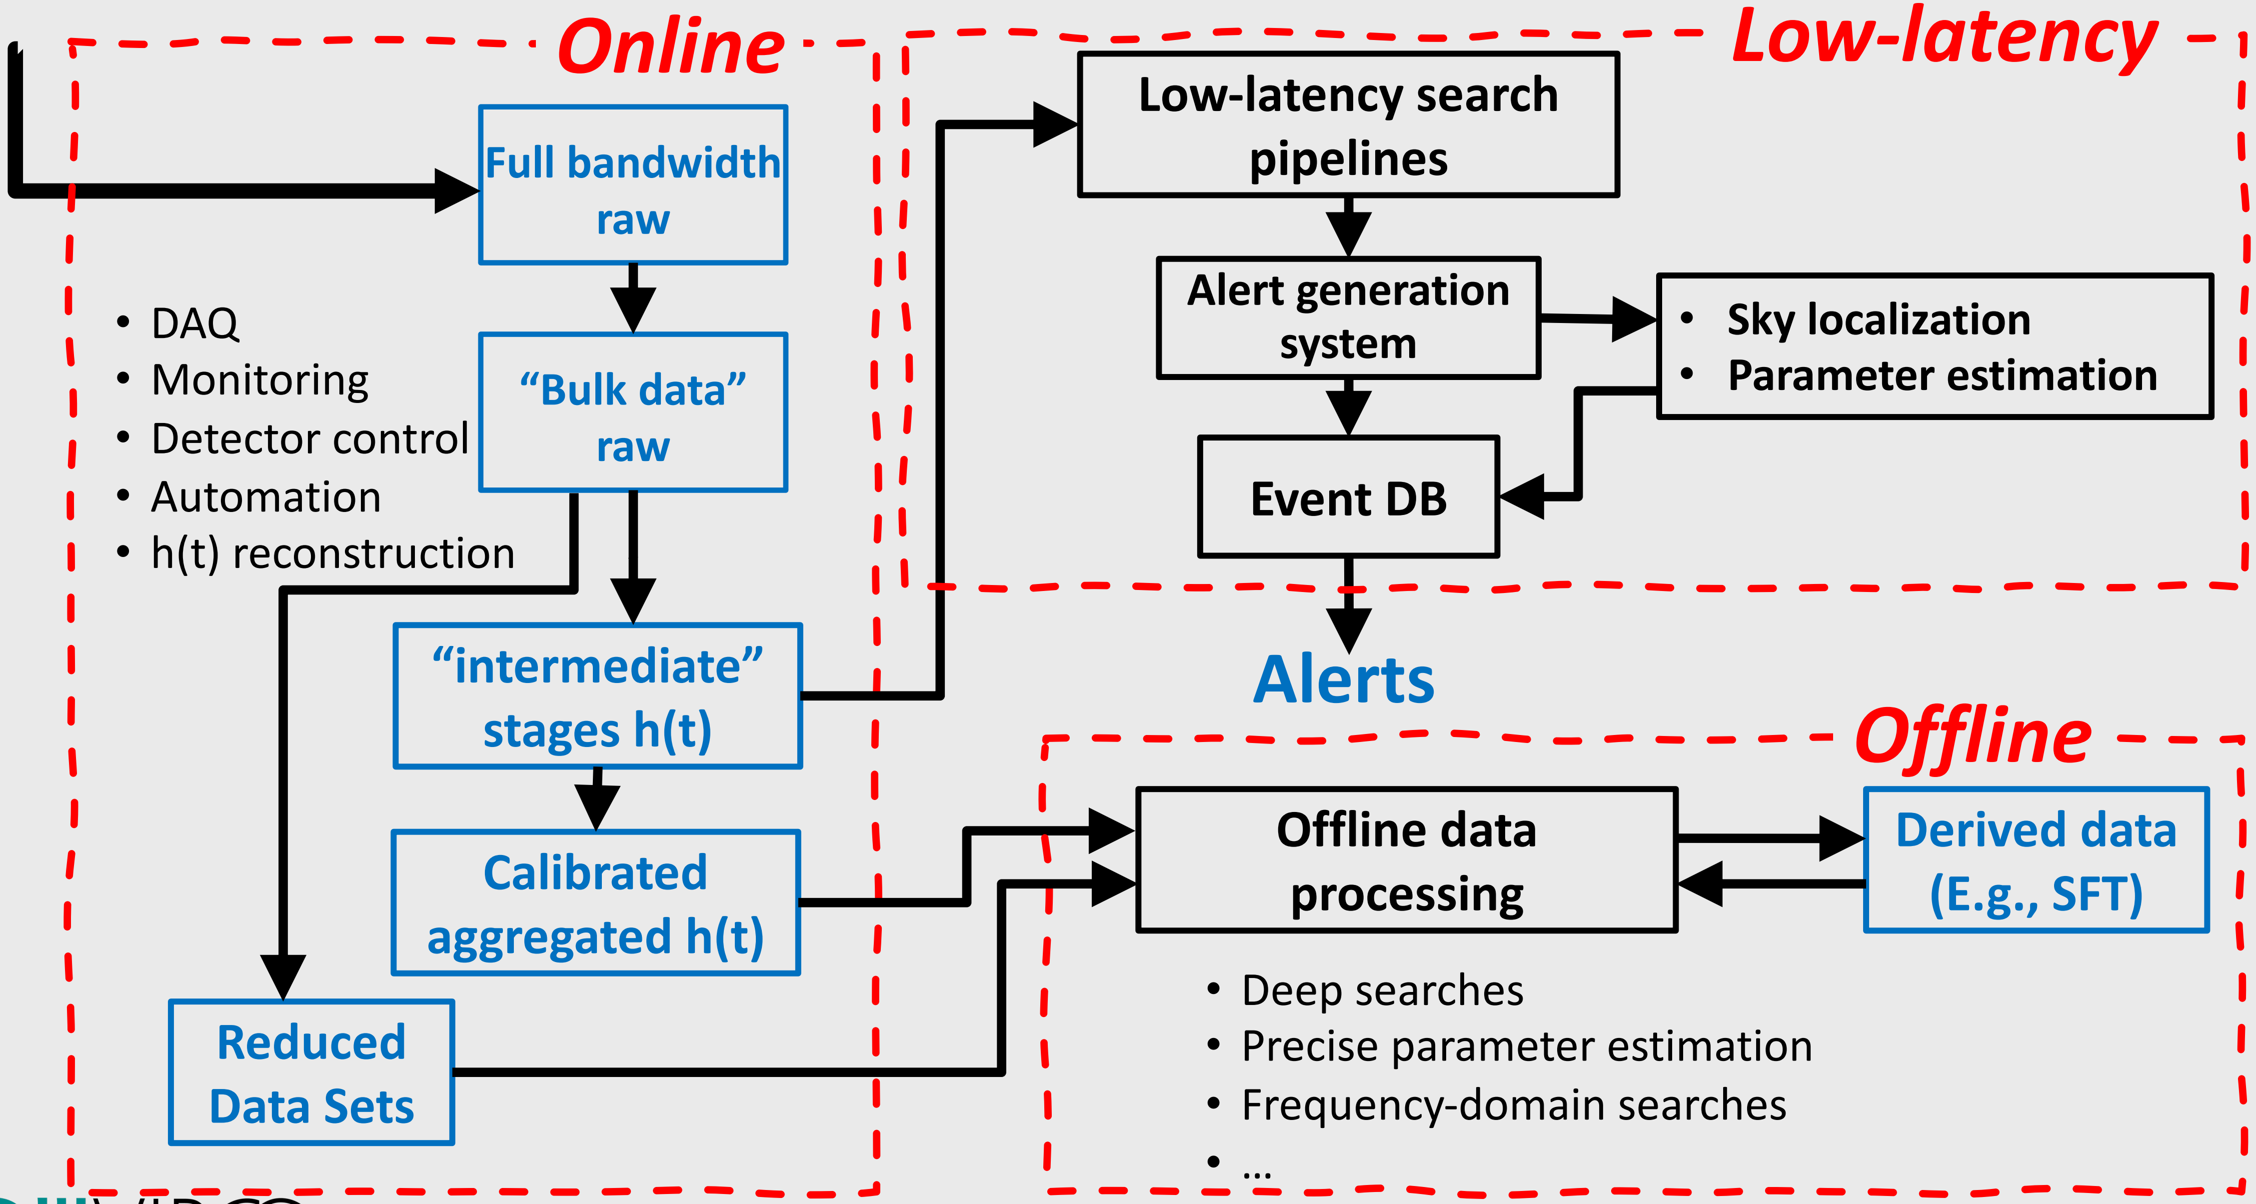

OFFLINE DATA DISTRIBUTION

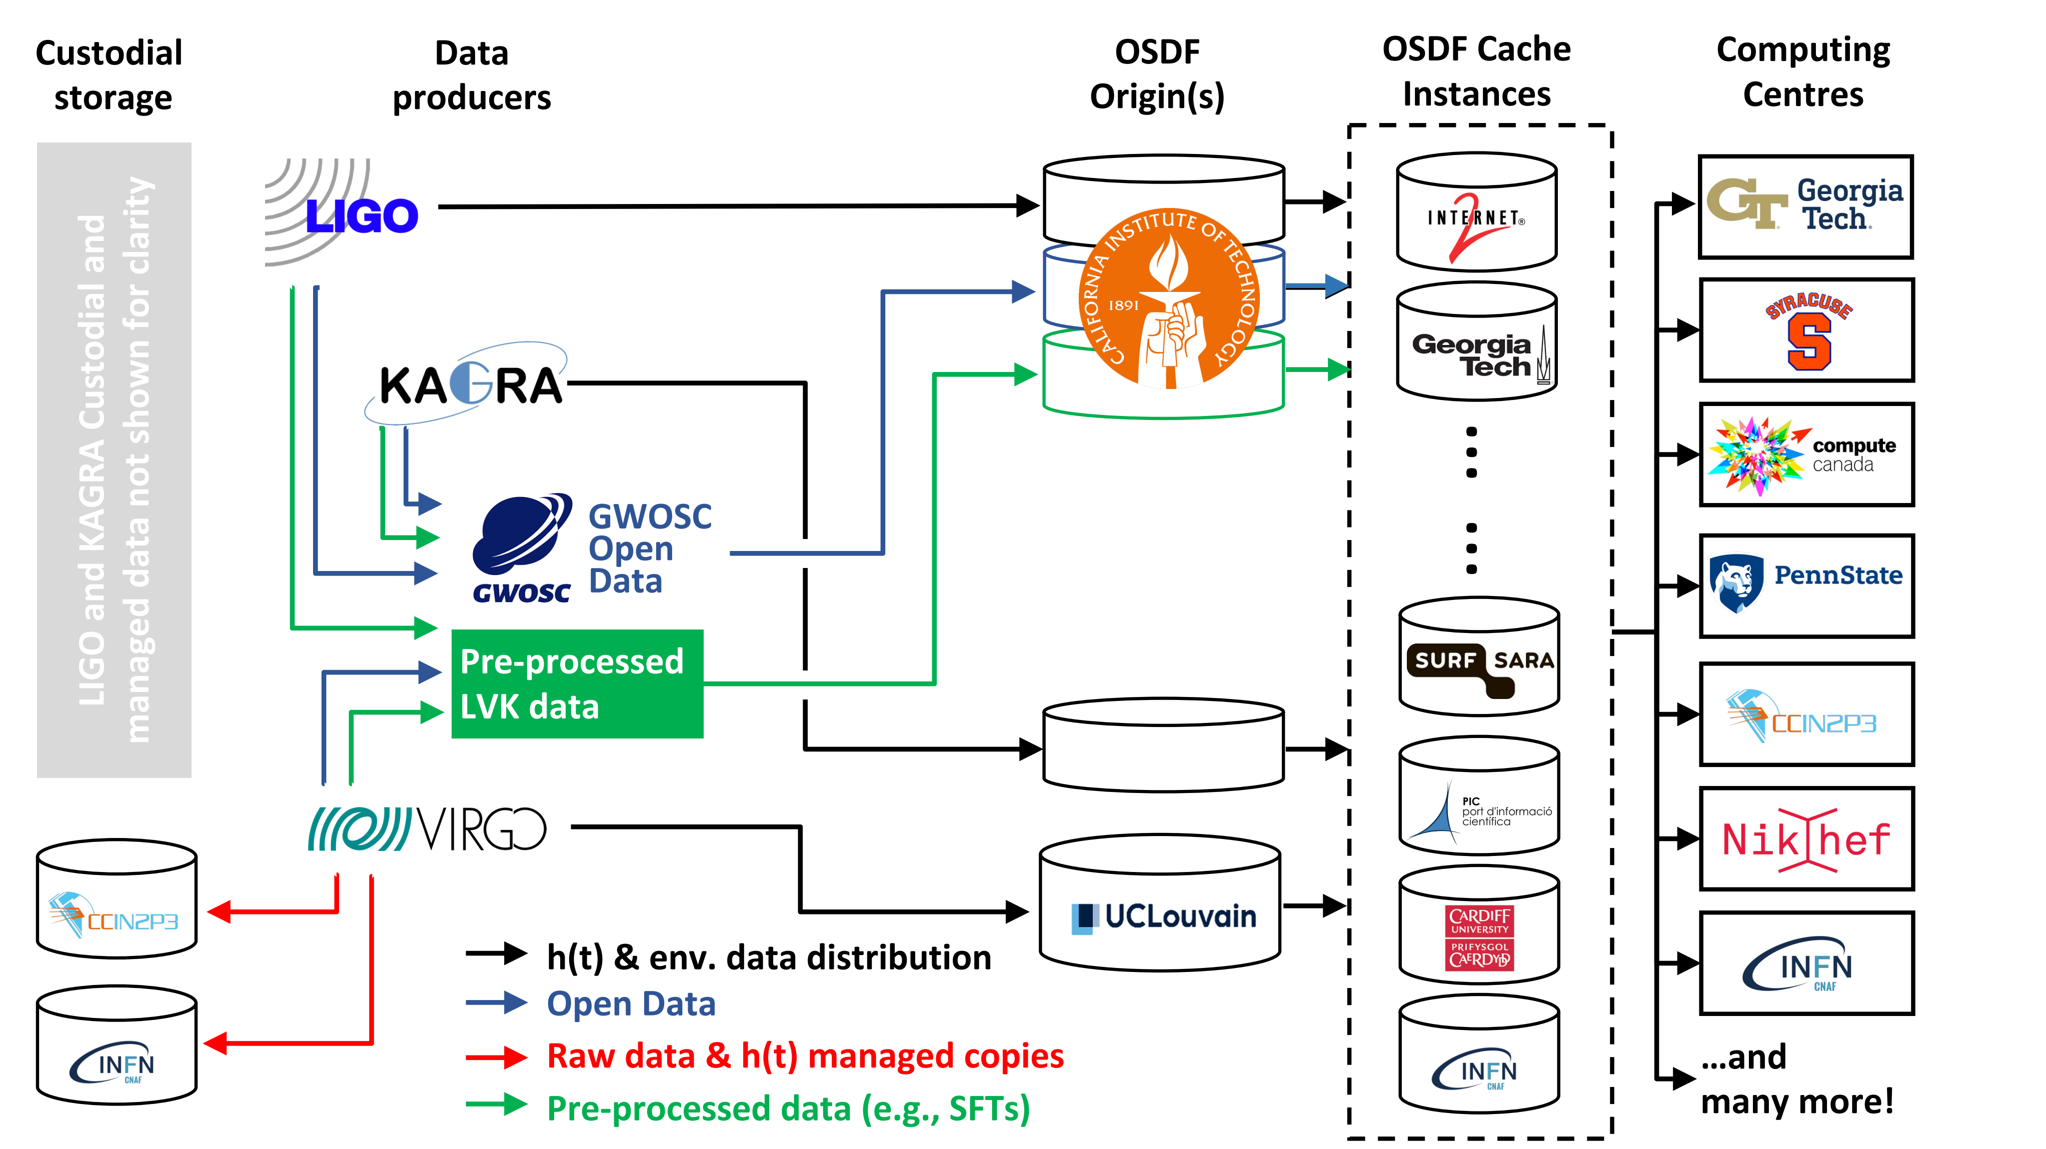

LVK OBSERVING SCENARIO TIMELINE

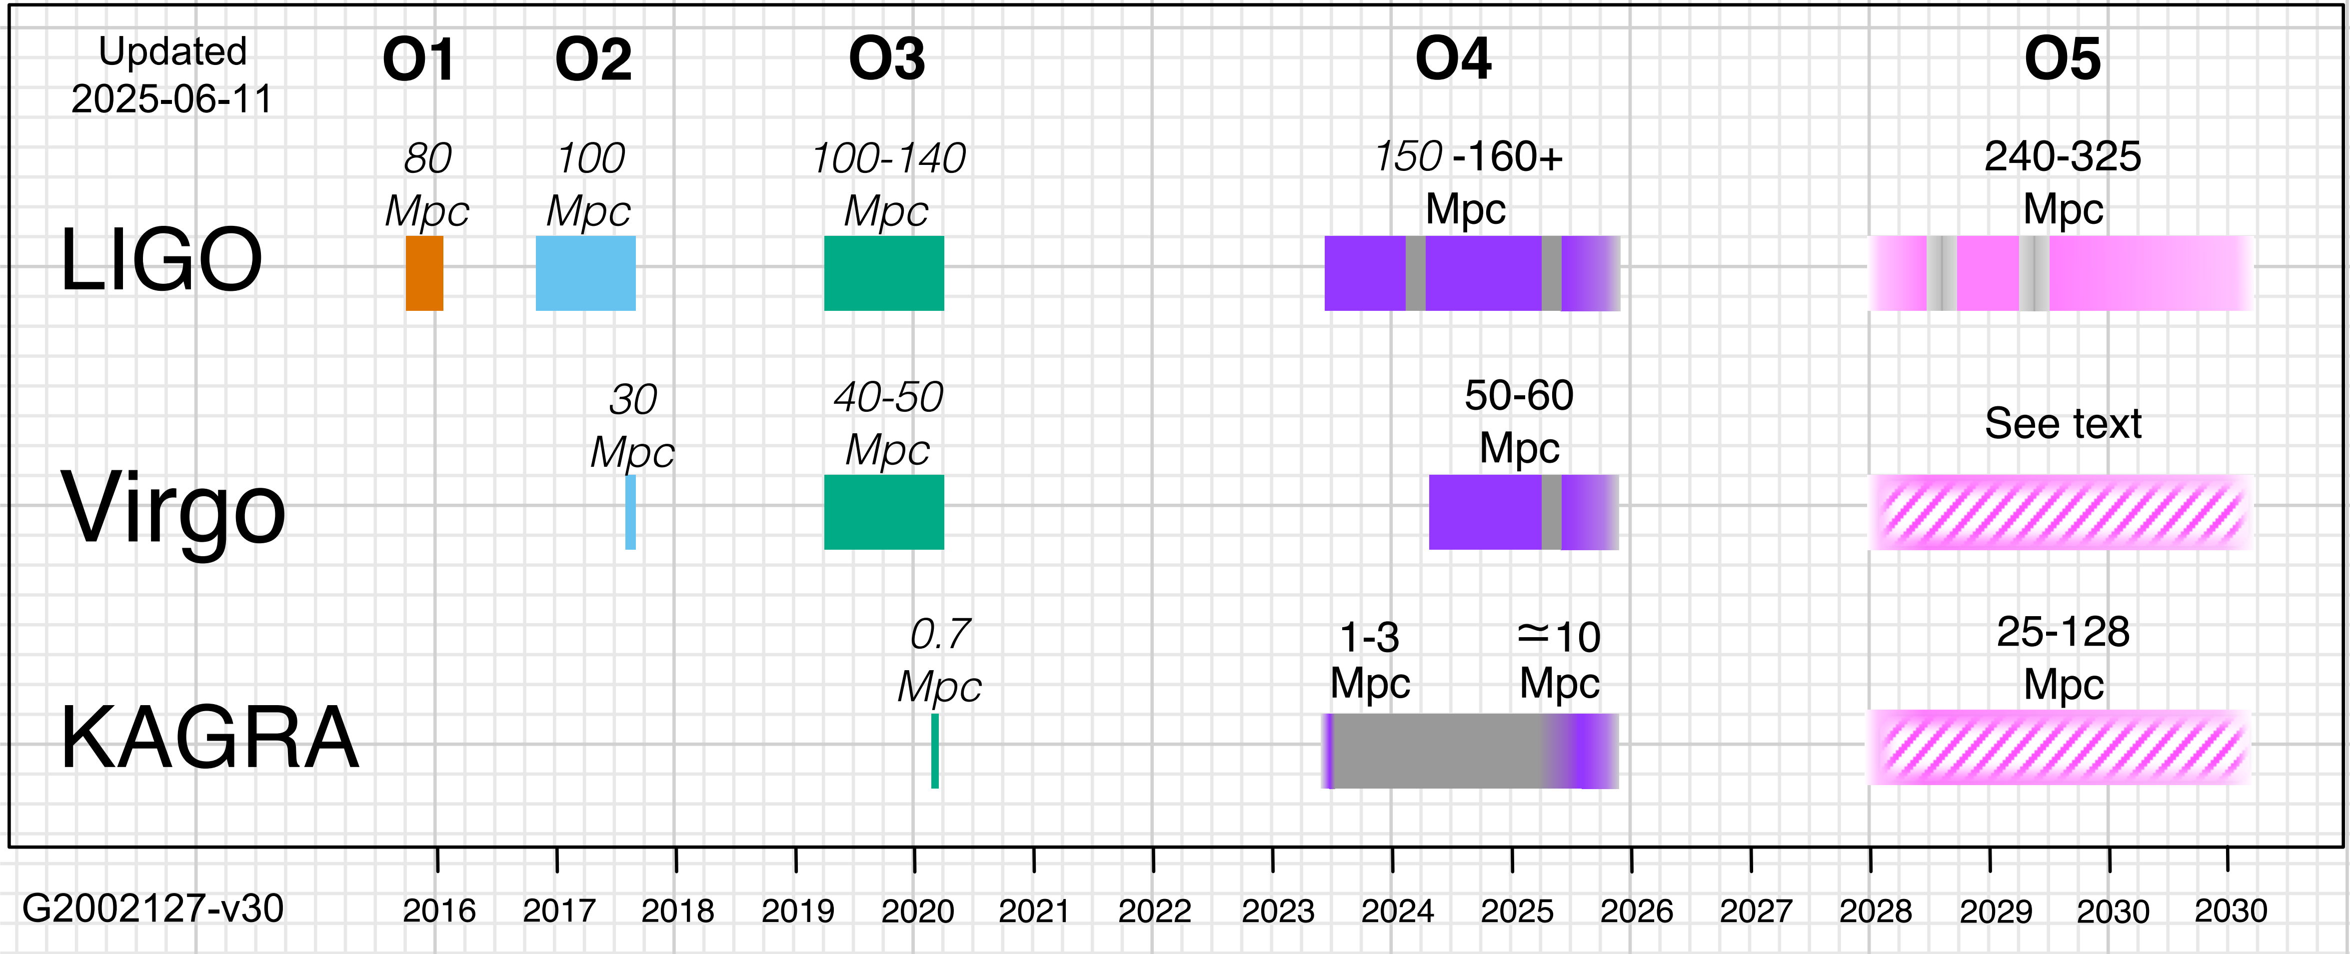

VIRGO

BOUNDARY CONDITIONS

- Virgo is supposed to provide ~30% of the LVK computing resources
  - CPU, storage, services, personpower, support
  - Nearly there for CPU and storage
  - Very far for everything else
- Strong push from the Council to reduce on-site computing
  - i.e., mostly LL analyses
- Long-standing "Virgo Computing Cloud" MoU ready to be signed
  - By CNAF and CC-IN2P3, initially
  - More and more sites will contribute CPU and storage (and possibly cloud)

CLOUD RESOURCES

- Existing: LLAI ("Low-latency Alert Infrastructure") r&d deployment on K8S
  - Currently also used for Mock Data Challenges
- Ongoing: moving DQSegDB service away from CIT
  - Possibly other "small footprint" service in future
  - Then, move to K8S deployment
- Up next: LL analysis on virtualized HTCondor cluster
  - Critical service, uptime is crucial
- The future: high-availability LLAI for production instance?
  - The most critical service for LVK

GRACEDB (LLAI)

GraceDB | LVK Public Alerts

gracedb.ligo.org/superevents/public/O4/

GraceDB Public Alerts Latest Search Notifications Pipelines Documentation Logout

Authenticated as: Stefano Bagnasco

Event ID	Possible Source (Probability)	Significant	UTC	GCN	Location	FAR	Comments	Ω Scan
S250904ae	BBH (>99%)	Yes	Sept. 4, 2025 03:33:07 UTC	GCN Circular Query Notices | VOE	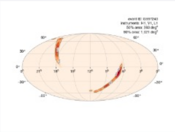	1 per 100.04 years		Ω H1 Ω L1 Ω V1
S250901cb	BBH (>99%)	Yes	Sept. 1, 2025 18:59:41 UTC	GCN Circular Query Notices | VOE	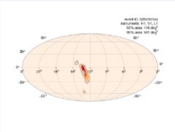	1 per 9.9253e+05 years		Ω H1 Ω L1 Ω V1
S250830bp	BBH (>99%)	Yes	Aug. 30, 2025 10:24:18 UTC	GCN Circular Query Notices | VOE	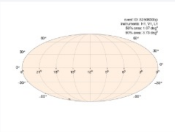	1 per 100.04 years		Ω H1 Ω L1 Ω V1
S250830m	BBH (98%), Terrestrial (2%)	Yes	Aug. 30, 2025 03:27:09 UTC	GCN Circular Query Notices | VOE	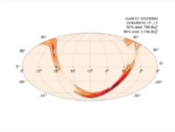	1 per 2.6632 years		Ω H1 Ω L1 Ω V1
S250827fo	BBH (>99%)	Yes	Aug. 27, 2025 22:49:40 UTC	GCN Circular Query Notices | VOE	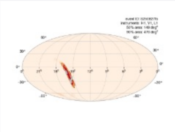	1 per 1405.1 years		Ω H1 Ω L1 Ω V1

ACCOUNTING INFO

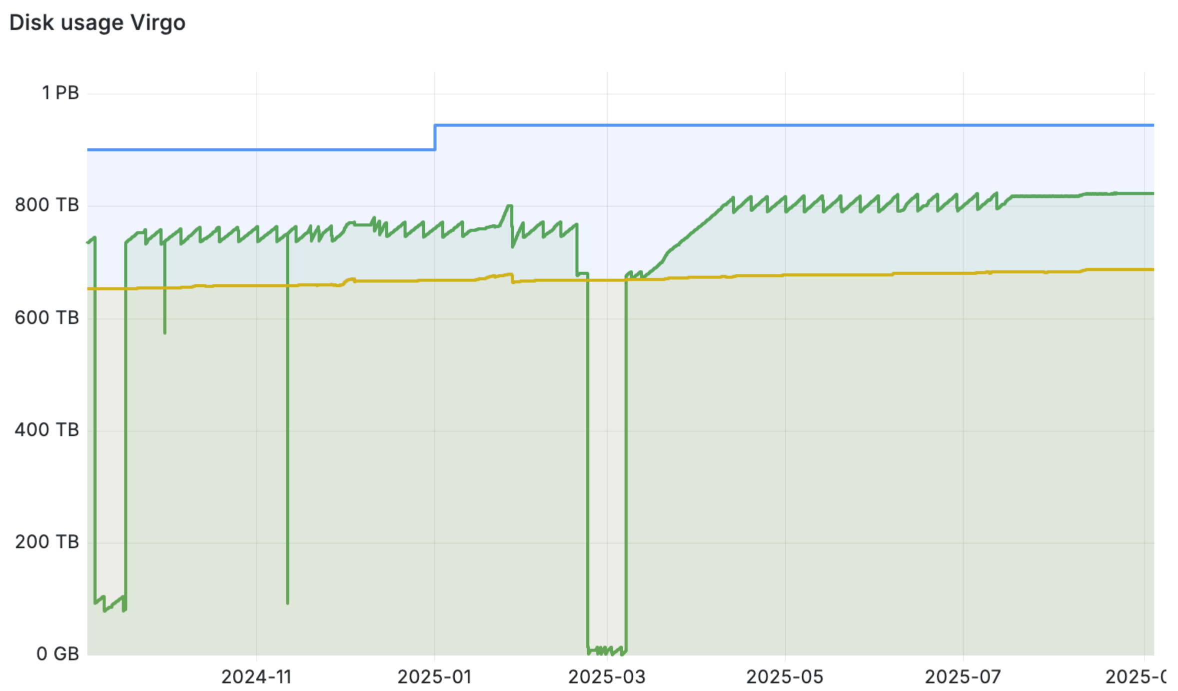

Name	Last *
disk+buffer used	822 TB
disk used	687 TB
Pledge	944 TB

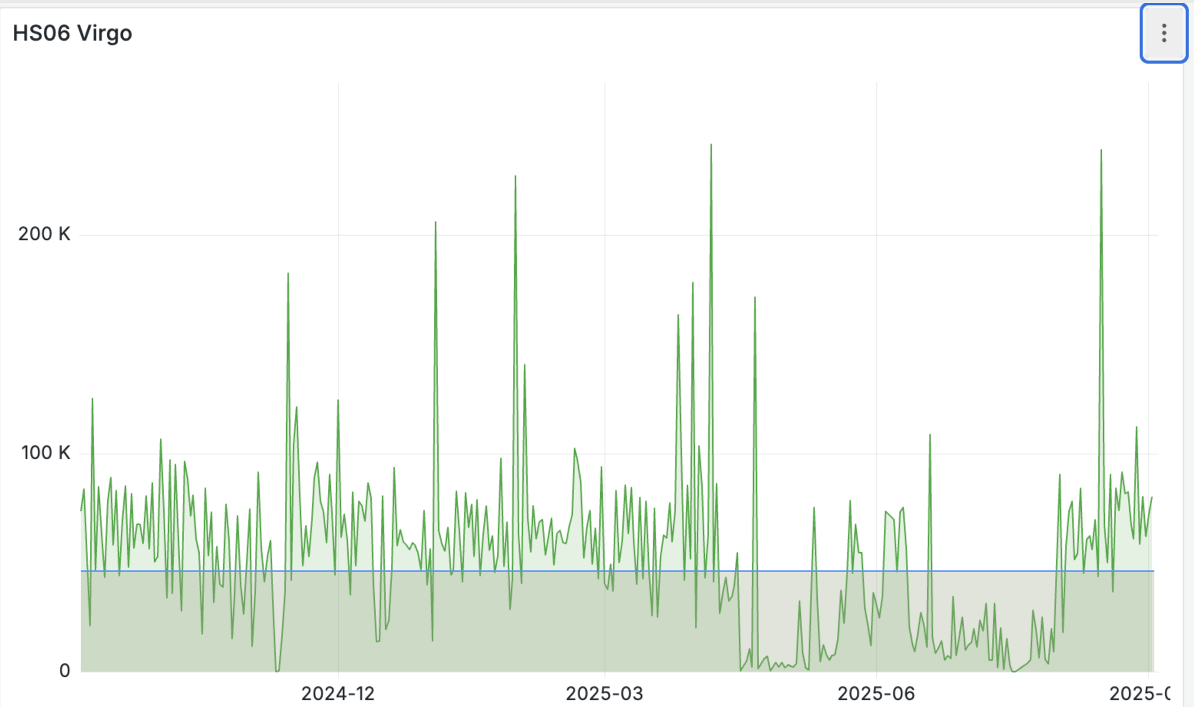

Name	Mean	Min	Max
used	53.3 K	0.00192	241 K
pledge	46 K	46 K	46 K
assigned	46 K	46 K	46 K

ACCOUNTING INFO

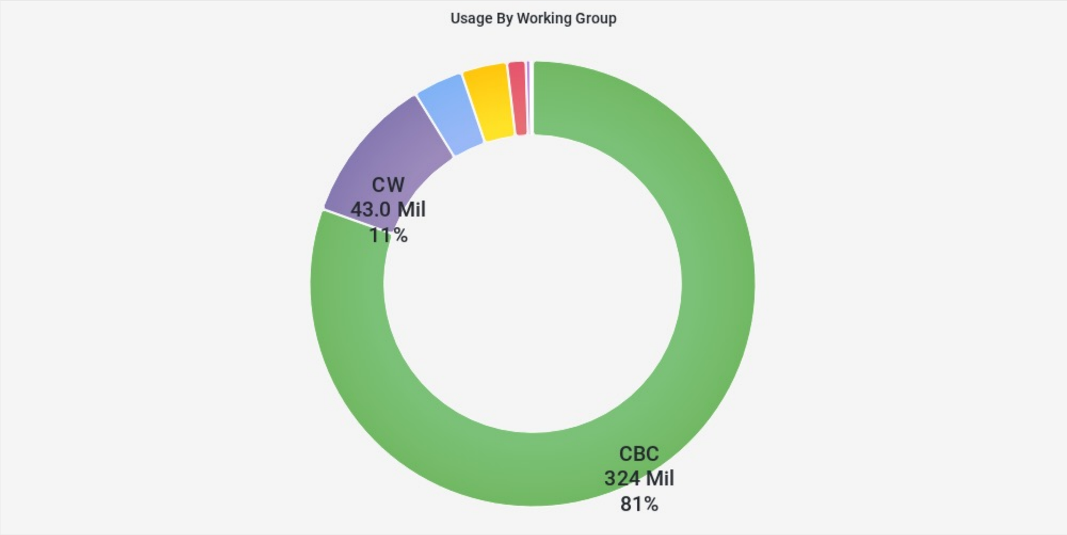

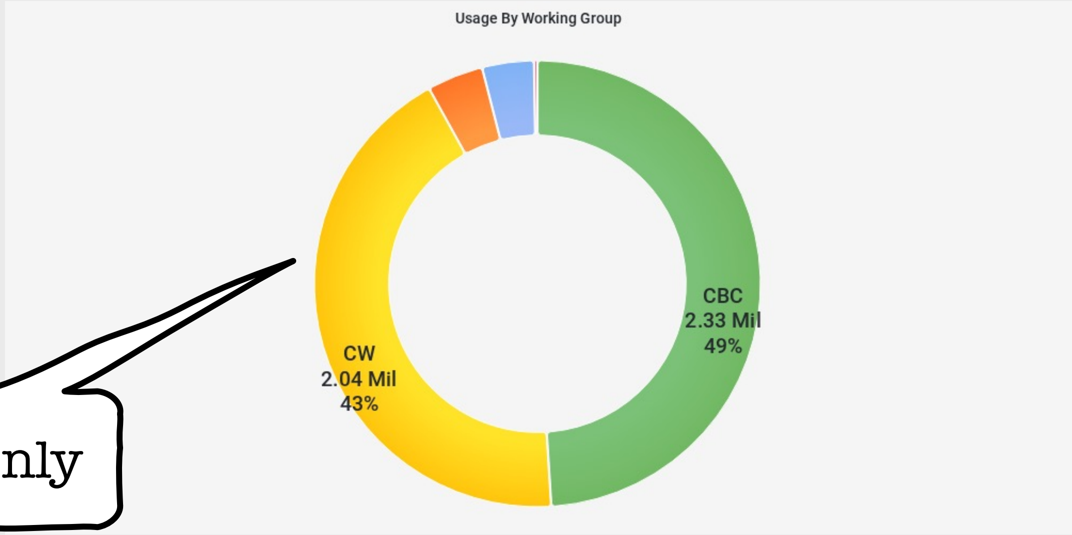

CNAF only

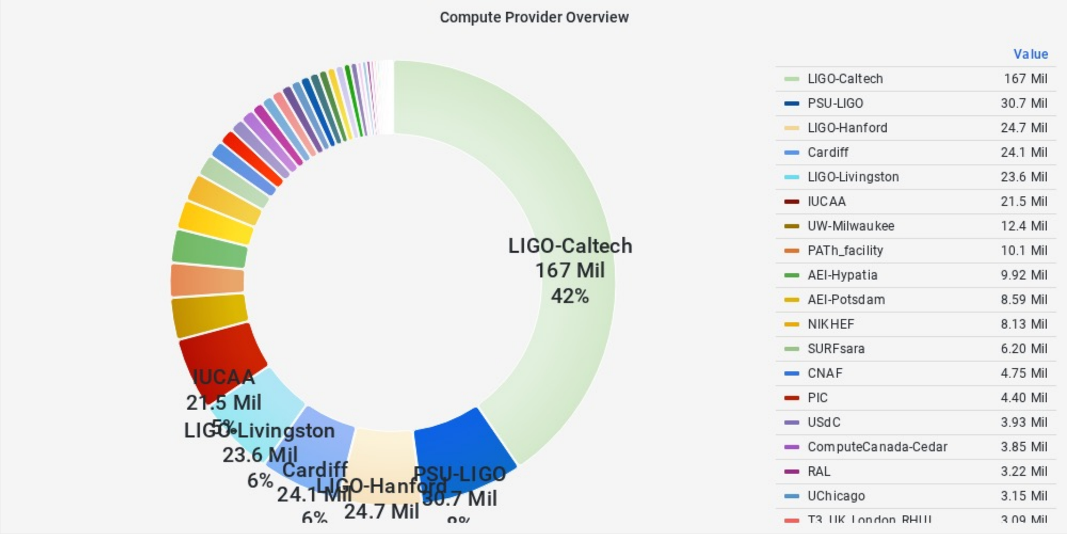

Virgo only

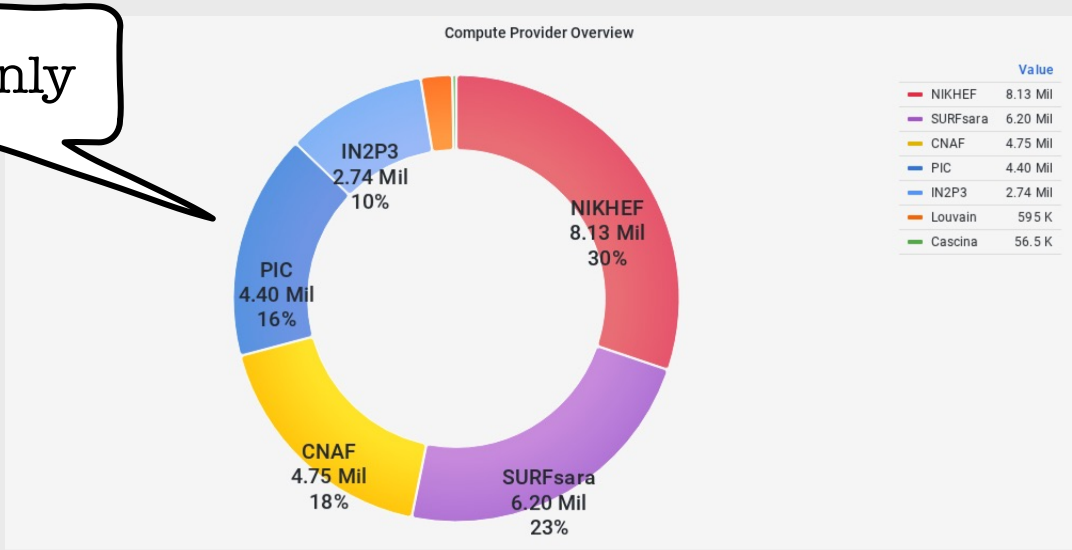

ACCOUNTING INFO

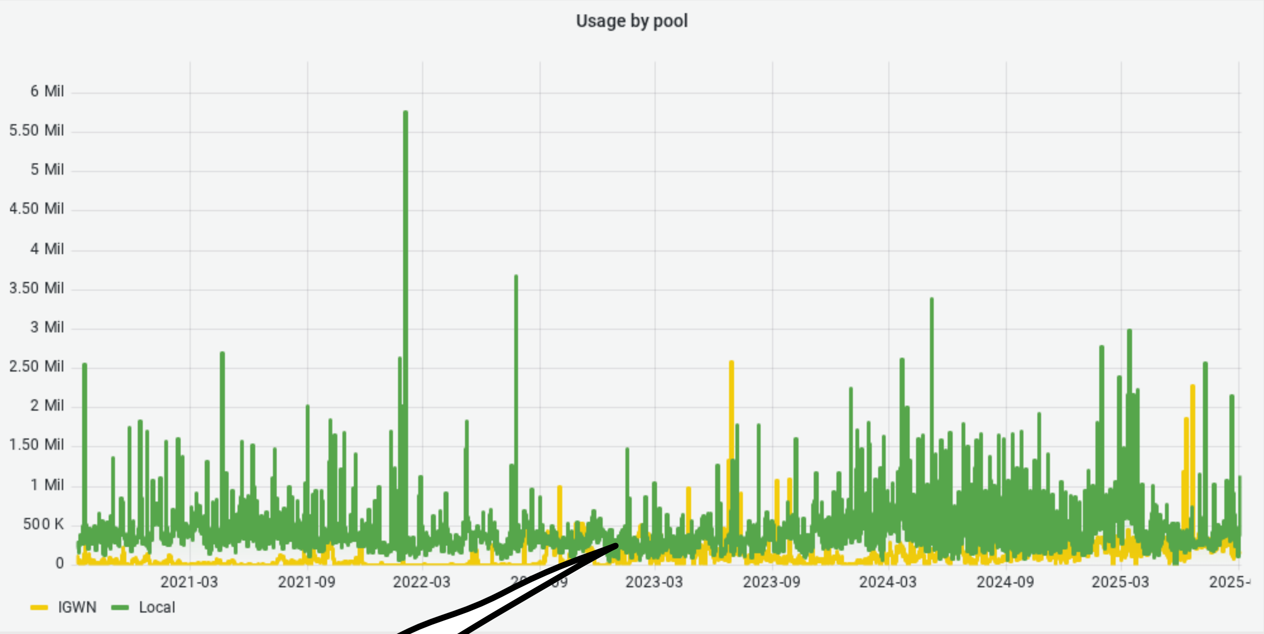

5 years

CNAF only

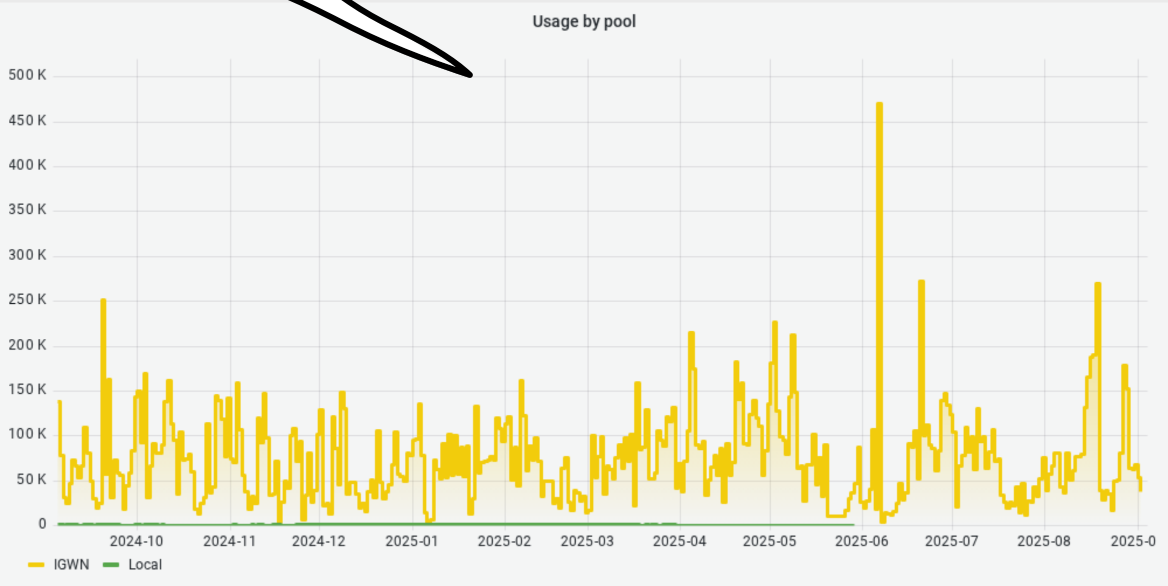

VIRGO

CNAF RESOURCES EVOLUTION

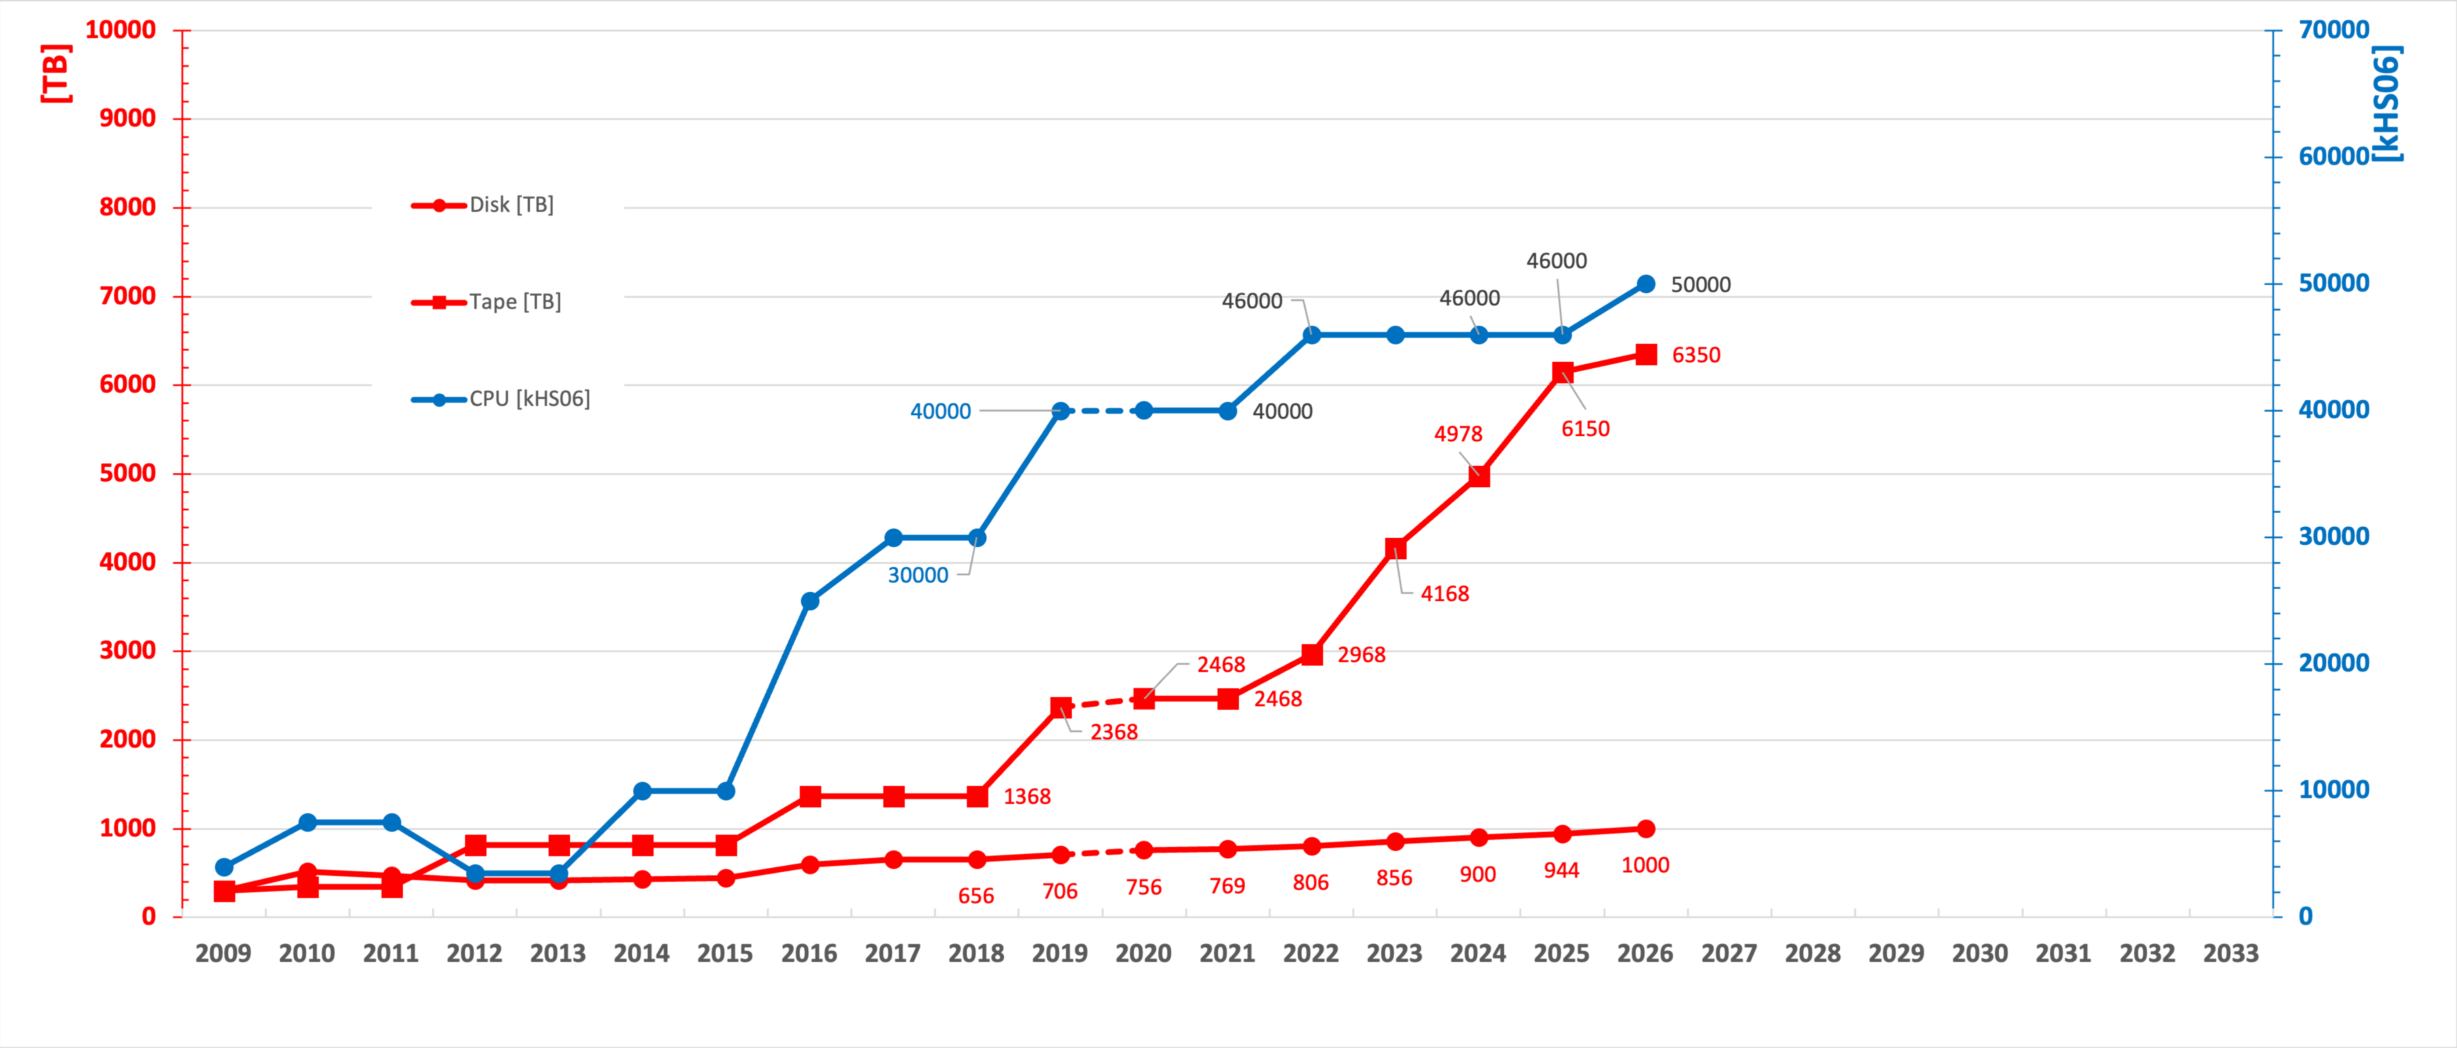

VIRGO

REQUESTS FOR 2026

	CPU HS06	DISCO TB	TAPE TB	Infrastruttura
(Incremento)	5000	50	200	Tier-1
(Pledge)	50000	1000		
	5M core hours			CINECA
	1500 cores (150 VMs)	5		INFN-Cloud

VIRGO

REQUESTS FOR 2026

I 50TB devono essere disco nVME ad alte prestazioni per un upgrade della cache OSDF/Pelican (AKA StashCache

	CPU HS06	DISCO TB	TAPE TB	Infrastruttura
(Incremento)	5000	50	200	Tier-1
(Pledge)	50000	1000		
	~~5M core hours~~			CINECA
	1500 cores (150 VMs)	5		INFN-Cloud

Richiesta fortemente ridotta (50K core hours)

ET


ET ITALY
Einstein Telescope
English
Italiano
Español
Deutsch
Français
Sardu
You Tube
ET in Itàlia
Sa traditzione Italiana
Sos Enattos, su giassu ideale
Su comitadu pro sa candidadura Italiana
Is progetos PNRR in agiudu de ET

- 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 newcomer in Saxony (Germany)

Journal of Cosmology and Astroparticle Physics
An IOP and SISSA journal

Science with the Einstein Telescope: a comparison of different designs

Marica Branchesi,[1,2,*] Michele Maggiore,[3,4,*] David Alonso,[5] Charles Badger,[6] Biswajit Banerjee,[1,2] Freija Beirnaert,[7] Enis Belgacem,[3,4] Swetha Bhagwat,[8,9] Guillaume Boileau,[10,11] Ssohrab Borhanian,[12] Daniel David Brown,[13] Man Leong Chan,[14] Giulia Cusin,[15,3,4] Stefan L. Danilishin,[16,17] Jerome Degallaix,[18] Valerio De Luca,[19] Arnab Dhani,[20] Tim Dietrich,[21,22] Ulyana Dupletsa,[1,2] Stefano Foffa,[3,4] Gabriele Franciolini,[8] Andreas Freise,[23,16] Gianluca Gemme,[24] Boris Goncharov,[1,2] Archisman Ghosh,[7] Francesca Gulminelli,[25] Ish Gupta,[20] Pawan Kumar Gupta,[16,26] Jan Harms,[1,2] Nandini Hazra,[1,2,27] Stefan Hild,[16,17] Tanja Hinderer,[28] Ik Siong Heng,[29] Francesco Iacovelli,[3,4] Justin Janquart,[16,26] Kamiel Janssens,[10,11] Alexander C. Jenkins,[30] Chinmay Kalaghatgi,[16,26,31] Xhesika Koroveshi,[32,33] Tjonnie G.F. Li,[34,35] Yufeng Li,[36] Eleonora Loffredo,[1,2] Elisa Maggio,[22] Michele Mancarella,[3,4,37,38] Michela Mapelli,[39,40,41] Katarina Martinovic,[6] Andrea Maselli,[1,2] Patrick Meyers,[42] Andrew L. Miller,[43,16,26] Chiranjib Mondal,[25] Niccolò Muttoni,[3,4] Harsh Narola,[16,26] Micaela Oertel,[44] Gor Oganesyan,[1,2] Costantino Pacilio,[8,37,38] Cristiano Palomba,[45] Paolo Pani,[8] Antonio Pasqualetti,[46] Albino Perego,[47,48] Carole Périgois,[39,40,41] Mauro Pieroni,[49,50] Ornella Juliana Piccinni,[51] Anna Puecher,[16,26] Paola Puppo,[45] Angelo Ricciardone,[52,39,40] Antonio Riotto,[3,4] Samuele Ronchini,[1,2] Mairi Sakellariadou,[6] Anuradha Samajdar,[21] Filippo Santoliquido,[39,40,41] B.S. Sathyaprakash,[20,53,54] Jessica Steinlechner,[16,17] Sebastian Steinlechner,[16,17] Andrei Utina,[16,17] Chris Van Den Broeck[16,26] and Teng Zhang[9,17]

[1]Gran Sasso Science Institute (GSSI), I-67100 L'Aquila, Italy
[2]INFN, Laboratori Nazionali del Gran Sasso, I-67100 Assergi, Italy

*Corresponding author.

https://doi.org/10.1088/1475-7516/2023/07/068

JCAP07(2023)068

ET-0036E-25

The Science of the Einstein Telescope

Einstein Telescope collaboration
Adrian Abac,[1] Raul Abramo,[2] Simone Albanesi,[3,4] Angelica Albertini,[5,6] Alessandro Agapito,[7,8,9] Michalis Agathos,[10,11] Conrado Albertus,[12] Nils Andersson,[13] Tomás Andrade,[14] Igor Andreoni,[15,16] Federico Angeloni,[7,17,8,18] Marco Antonelli,[19] John Antoniadis,[20,21] Fabio Antonini,[22] Manuel Arca Sedda,[23,24,25,26] M. Celeste Artale,[27,28] Stefano Ascenzi,[23] Pierre Auclair,[29] Matteo Bachetti,[30] Charles Badger,[31] Biswajit Banerjee,[23] David Barba-González,[12] Dániel Barta,[32] Nicola Bartolo,[26,33,34] Andreas Bauswein,[35] Andrea Begnoni,[26,33] Freija Beirnaert,[36] Michał Bejger,[37,38] Enis Belgacem,[39,40] Nicola Bellomo,[26,33,34] Laura Bernard,[41] Maria Grazia Bernardini,[42] Sebastiano Bernuzzi,[3] Christopher P. L. Berry,[43] Emanuele Berti,[44] Gianfranco Bertone,[45] Dario Bettoni,[46,12] Miguel Bezares,[47,48] Swetha Bhagwat,[49] Sofia Bisero,[50] Marie Anne Bizouard,[51] Jose J. Blanco-Pillado,[52,53,54] Simone Blasi,[55,56] Alice Bonino,[49] Alice Borghese,[18] Nicola Borghi,[57,58] Ssohrab Borhanian,[3,59,60] Elisa Bortolas,[34,59,60] Maria Teresa Botticella,[61] Marica Branchesi,[23,24] Matteo Breschi,[3] Richard Brito,[62] Enzo Brocato,[25,18] Floor S. Broekgaarden,[63] Tomasz Bulik,[64] Alessandra Buonanno,[1,65] Fiorella Burgio,[66] Adam Burrows,[67] Gianluca Calcagni,[68] Sofia Canevarolo,[69] Enrico Cappellaro,[34] Giulia Capurri,[70,71] Carmelita Carbone,[72] Roberto Casadio,[57,73] Ramiro Cayuso,[74,75] Pablo Cerdá-Durán,[76,77] Prasanta Char,[12] Sylvain Chaty,[78] Tommaso Chiarusi,[73] Martyna Chruslinska,[79,80] Francesco Cireddu,[81,82,70] Philippa Cole,[59,60] Alberto Colombo,[42,60] Monica Colpi,[59] Geoffrey Compère,[83] Carlo Contaldi,[84] Maxence Corman,[1] Francesco Crescimbeni,[7,8] Sergio Cristallo,[25] Elena Cuoco,[57,73] Giulia Cusin,[85,39] Tito Dal Canton,[86] Gergely Dálya,[87] Paolo D'Avanzo,[42] Nazanin Davari,[18] Valerio De Luca,[88] Viola De Renzis,[9] Massimo Della Valle,[61] Walter Del Pozzo,[70,71] Federico De Santi,[70,71] Alessio Ludovico De Santis,[23,24] Tim

arXiv:2503.12263v2 [gr-qc] 29 Aug 2025

EINSTEIN TELESCOPE

- The unknown we know best* is the event rate
- Then there are less-known unknowns

Key Challenges

- Long-Duration Signals: ET's low-frequency sensitivity extends the duration of signals in its band. For example, binary neutron star signals may persist for hours, increasing computational demands for matched filtering and waveform modeling. Additionally, Earth's motion modulates the signal, complicating sky localisation and template bank design.
- Overlapping Signals: The high detection rate in ET will lead to frequent overlapping signals in the data. Traditional single-signal models may introduce biases, particularly for closely timed or comparable signal amplitudes. New methods to separate and analyze overlapping signals simultaneously are crucial.
- Noise Background Estimation: The dominance of GW signals complicates noise characterisation, as there will be minimal signal-free data segments. Traditional noise estimation techniques may overestimate the background. For a triangular configuration for ET, the signal-free null stream can be leveraged to produce correct background estimation.

arXiv:2503.12263

*Thanks to past US Secretary of Defense D. Rumsfeld for the concept

The scale of ET computing

EINSTEIN TELESCOPE

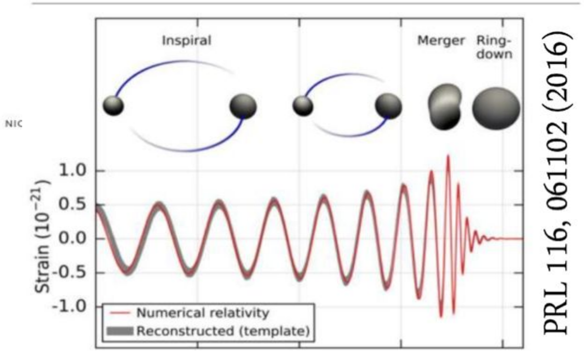

PRL 116, 061102 (2016)

- Increased signal sensitivity for ET means signals are also in-band for much longer (minutes, hours, days)
  - Signal "pile-up" complicates things
  - ***Naively, ET would need 40M cores just for low latency***
  - Naive here means assuming we only need to double the compute to handle pre-merger analysis

40M cores is 100 times what HL-ATLAS needs for a similar job; 400k cores is a lot of computing power!

Our target is ~10% of HL-ATLAS, significant but the wider community has experience and tools
- ***That implies a speed-up of 1000***

The BlueBook studies show very promising results at the level of speeding up algorithms
- Latency is determined by the slowest step, e.g. reading/writing to/from files or databases, CPU time…

We need end-to-end tests to prove we can achieve the required latency for multi-messenger science

UNIVERSITÉ DE GENÈVE
FACULTÉ DES SCIENCES
Département d'astronomie
Paul Laycock

GWSC GRAVITATIONAL WAVE SCIENCE CENTER

Paul.Laycock@unige.ch

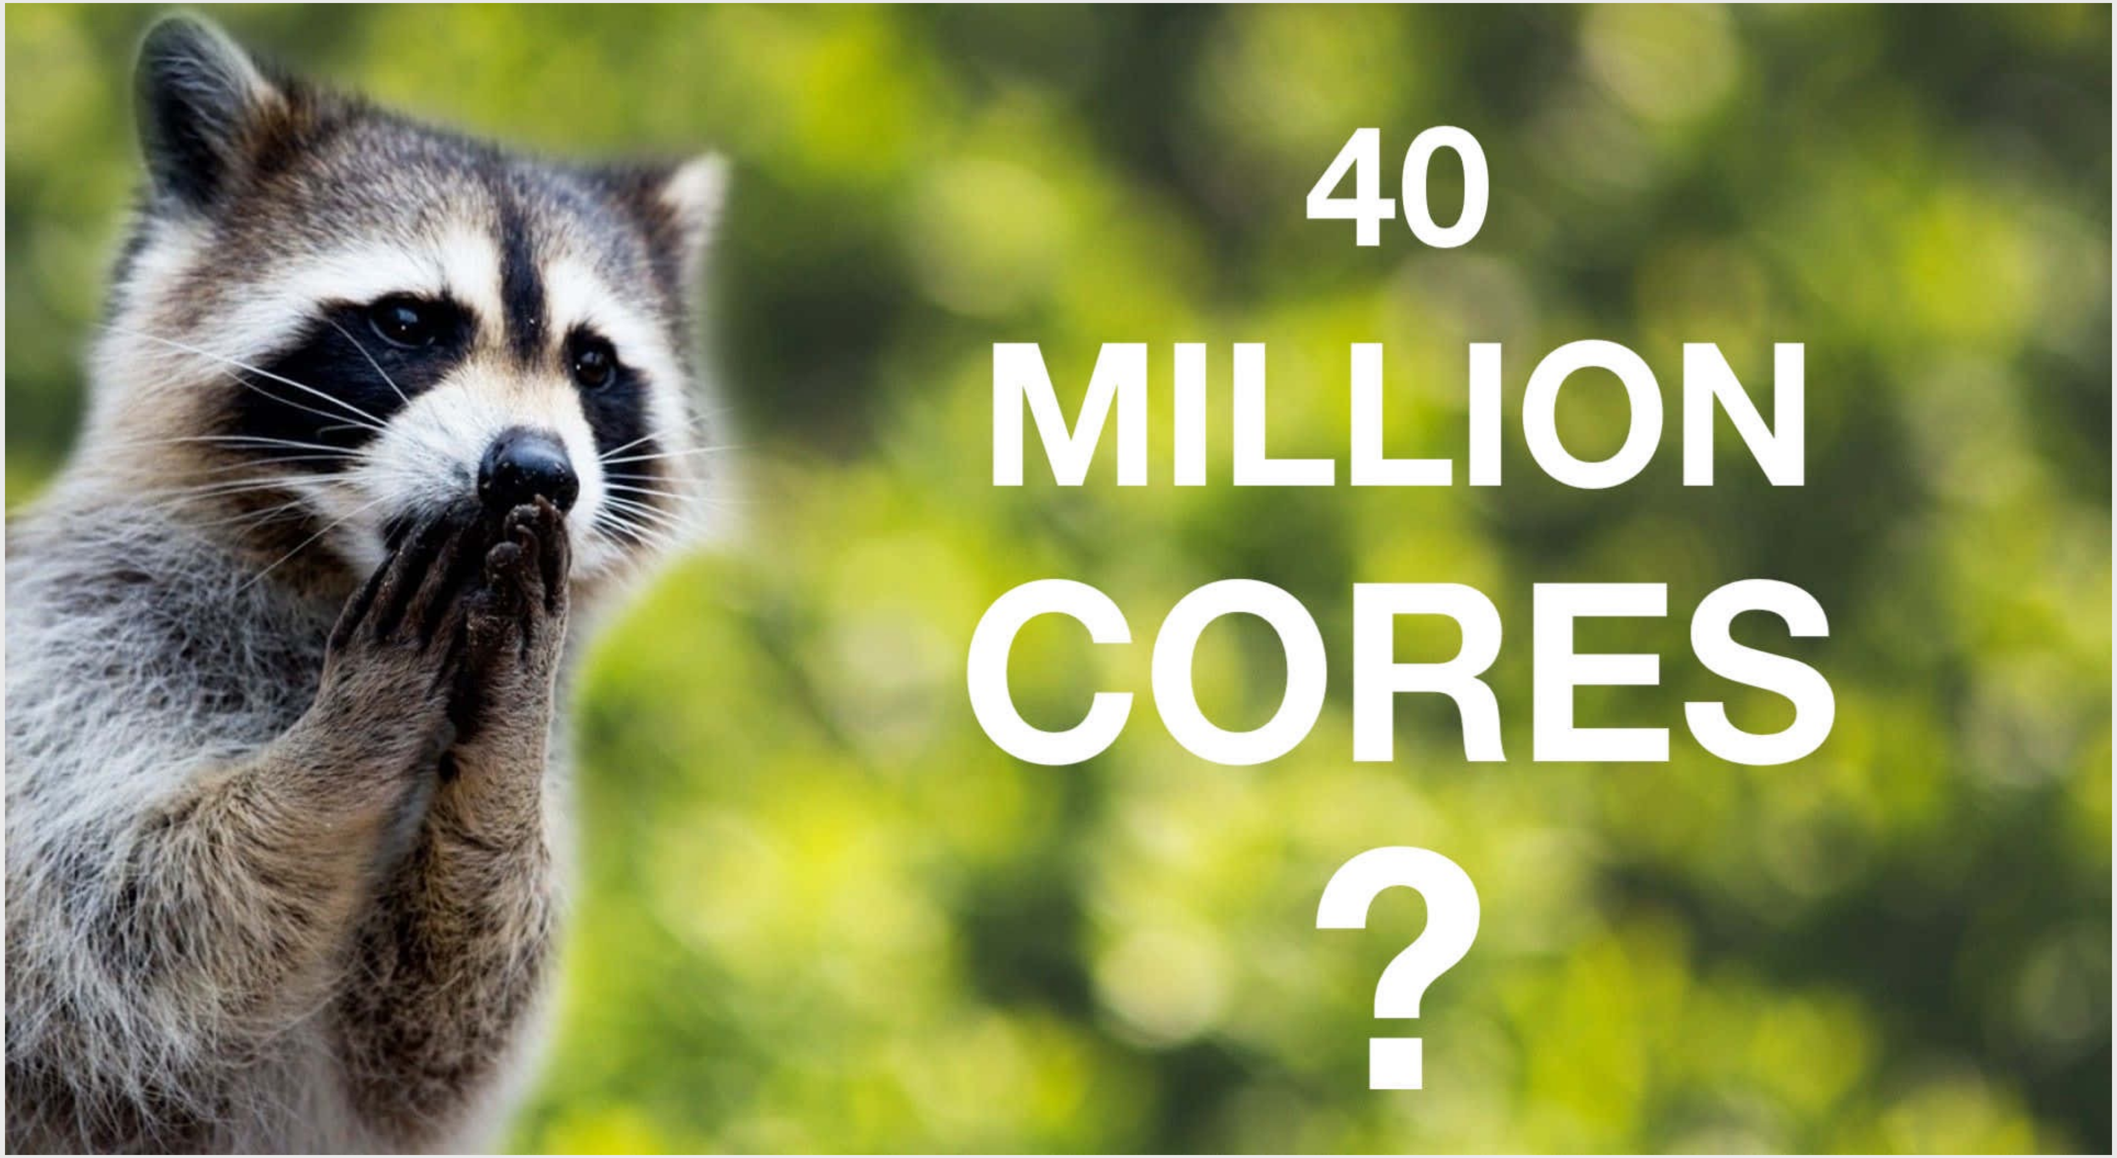

Sara.Vallero@to.infn.it

PE computational cost (J. Veitch)

EINSTEIN TELESCOPE

PIC port d'informació científica

Likelihood acceleration methods including relative binning, multibanding, and reduced order quadrature can reduce the amount of CPU required for PE by over three orders of magnitude.

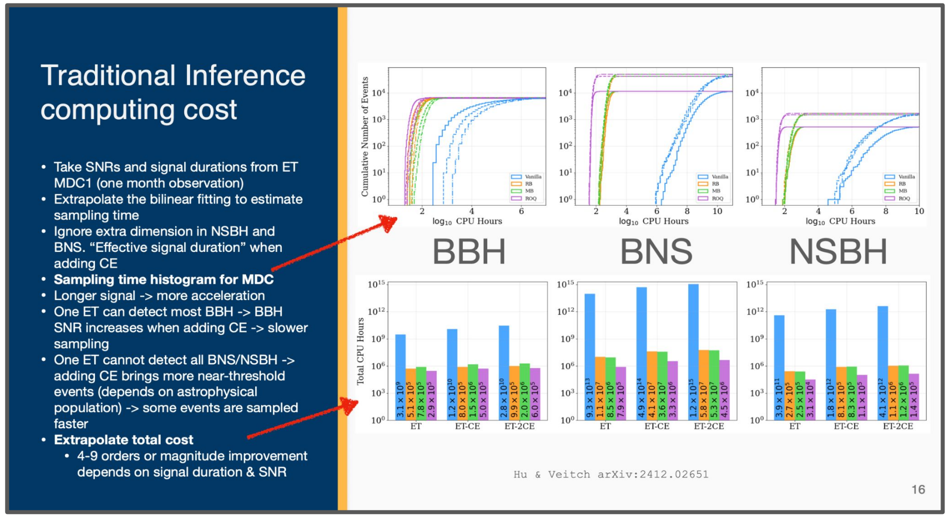

arXiv > gr-qc > arXiv:2412.02651

Search...

Help | Adv

General Relativity and Quantum Cosmology

[Submitted on 3 Dec 2024 (v1), last revised 5 Mar 2025 (this version, v2)]

Costs of Bayesian Parameter Estimation in Third-Generation Gravitational Wave Detectors: a Review of Acceleration Methods

Qian Hu, John Veitch

Bayesian inference with stochastic sampling has been widely used to obtain the properties of gravitational wave (GW) sources. Although computationally intensive, its cost remains manageable for current second-generation GW detectors because of the relatively low event rate and signal-to-noise ratio (SNR). The third-generation (3G) GW detectors are expected to detect hundreds of thousands of compact binary coalescence events every year with substantially higher SNR and longer signal duration, presenting significant computational challenges. In this study, we systematically evaluate the computational costs of source parameter estimation (PE) in the 3G era by modeling the PE time cost as a function of SNR and signal duration. We examine the standard PE method alongside acceleration methods including relative binning, multibanding, and reduced order quadrature. We predict that PE for a one-month-observation catalog with 3G detectors could require billions to quadrillions of CPU core hours with the standard PE method, whereas acceleration techniques can reduce this demand to millions of core hours. These findings highlight the necessity for more efficient PE methods to enable cost-effective and environmentally sustainable data analysis for 3G detectors. In addition, we assess the accuracy of accelerated PE methods, emphasizing the need for careful treatment in high-SNR scenarios.

6

Gonzalo.Merino@pic.es

THE CTLAB4ET IN TORINO

Inauguration on April 14th, 2025

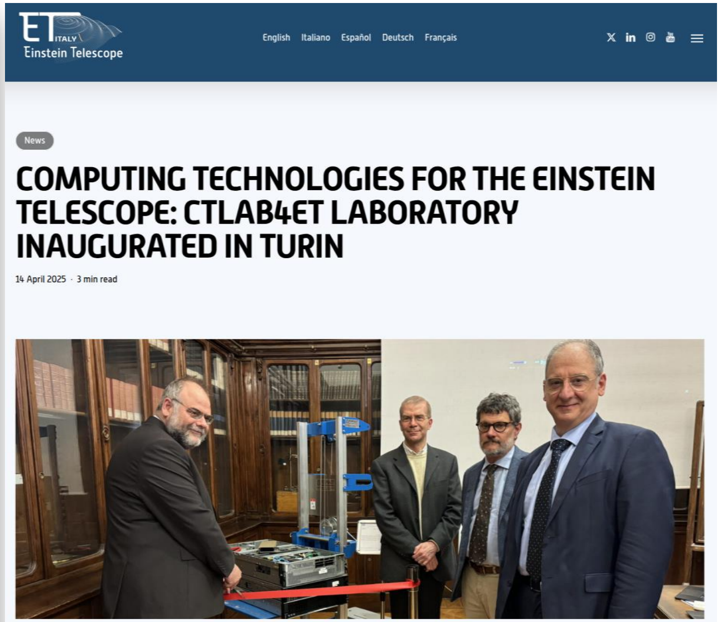

from www.einstein-telescope.it

from www.infn.it

Lia.Lavezzi@to.infn.it

MOCK DATA CHALLENGES

- What we need to write in the Computing Model by Feb 2026 is not a set of numbers but a **plan to get those numbers**
  - Also, we need to make sure they are realistic achievable numbers!
  - Both for online (DAQ), time-critical (LL) and asynchronous (offline)
- So Mock Data Challenges become a tool also to **demonstrate** we will be able to do all the science we want to do, both offline and for time-domain multimessenger (i.e., “low latency”)
  - We will have to run “low-latency” MDCs, not tomorrow but not in 2030 either.
  - Obviously this does not mean really running anything in LL now, but making sure that LL-related workflows (e.g., PE) are being progressively developed, tested and optimised – and refine the CM estimates accordingly
  - We have (many) years to do that, but we need to spell out the plan **now** for all computing domains to make sure we don’t miss pieces

GENERAL STRATEGY

- Use Mock Data Challenges as multi-purpose tools
  - More about this later
- Provide and deploy "workflow evaluation kits"...
  - Partial functionalities to evaluate tools and architectures
  - And quickly evolve towards a common and uniform environment
- ...using ESCAPE as the first toolbox
  - But not the only one
- Exploit synergies with Virgo as much as possible
  - IGWN computing infrastructure will be evolving, we cannot ignore it
- As usual, (skilled) personpower is the issue
  - Keep this in mind, I will not repeat it every other slide!

MADDEN and ETAP: A Global View

- We want to implement a **test setup** that **mimics** a possible future research **environment for GW experiments**.
- It will be used as a test bench for ET Mock Data Challenges, analysis algorithm development and technology tracking.

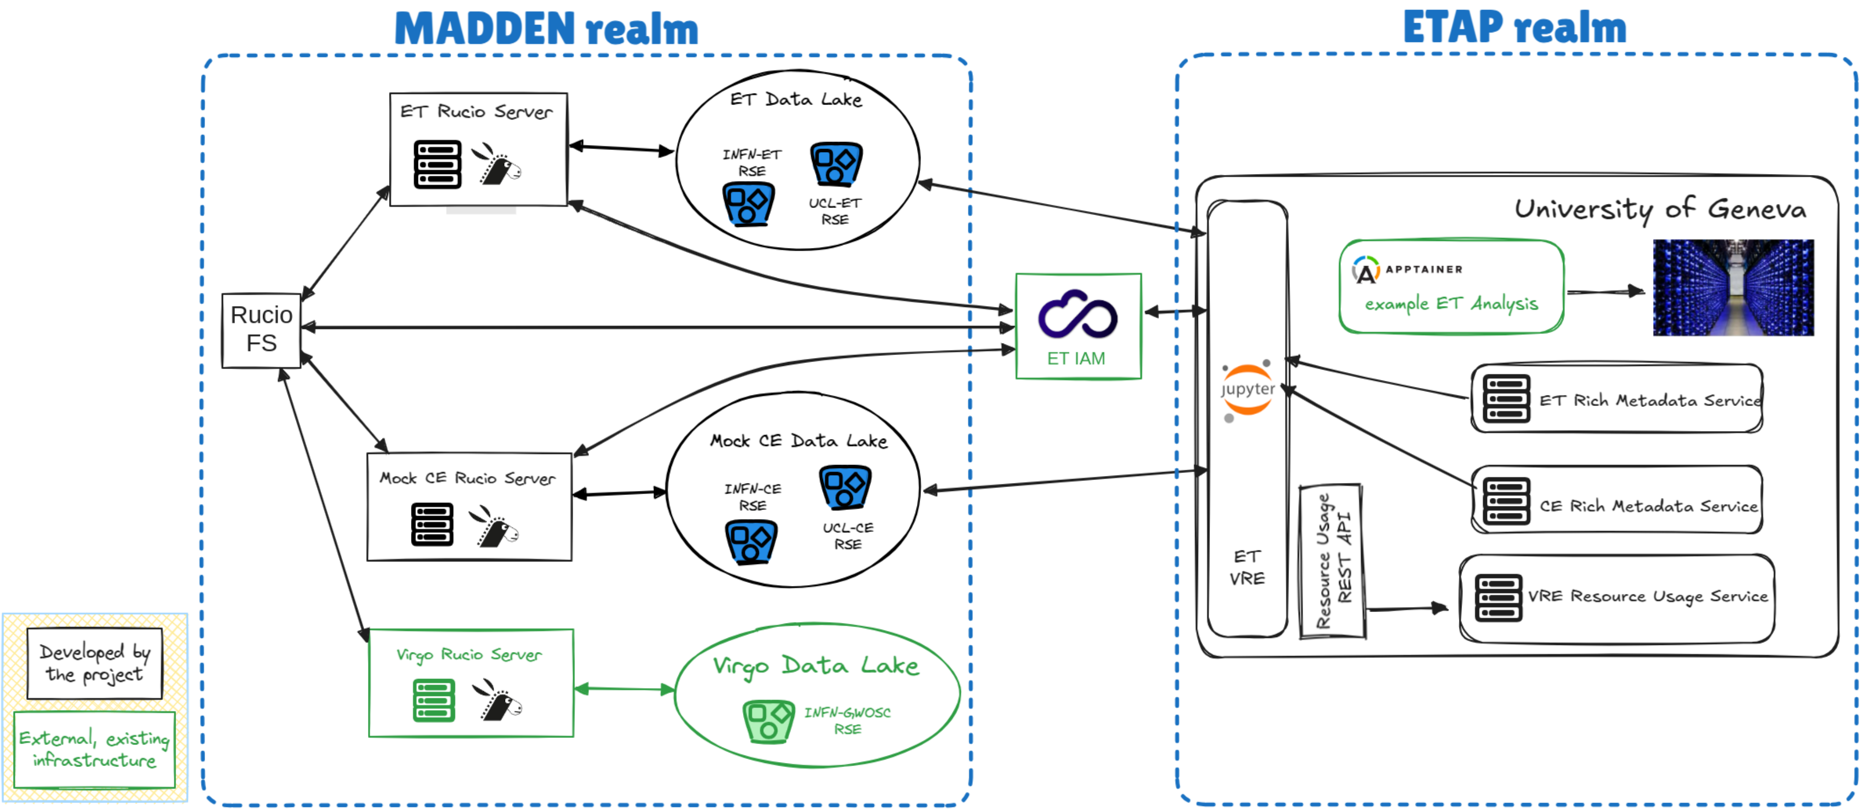

Lia.Lavezzi@to.infn.it

ACOUNTING INFO

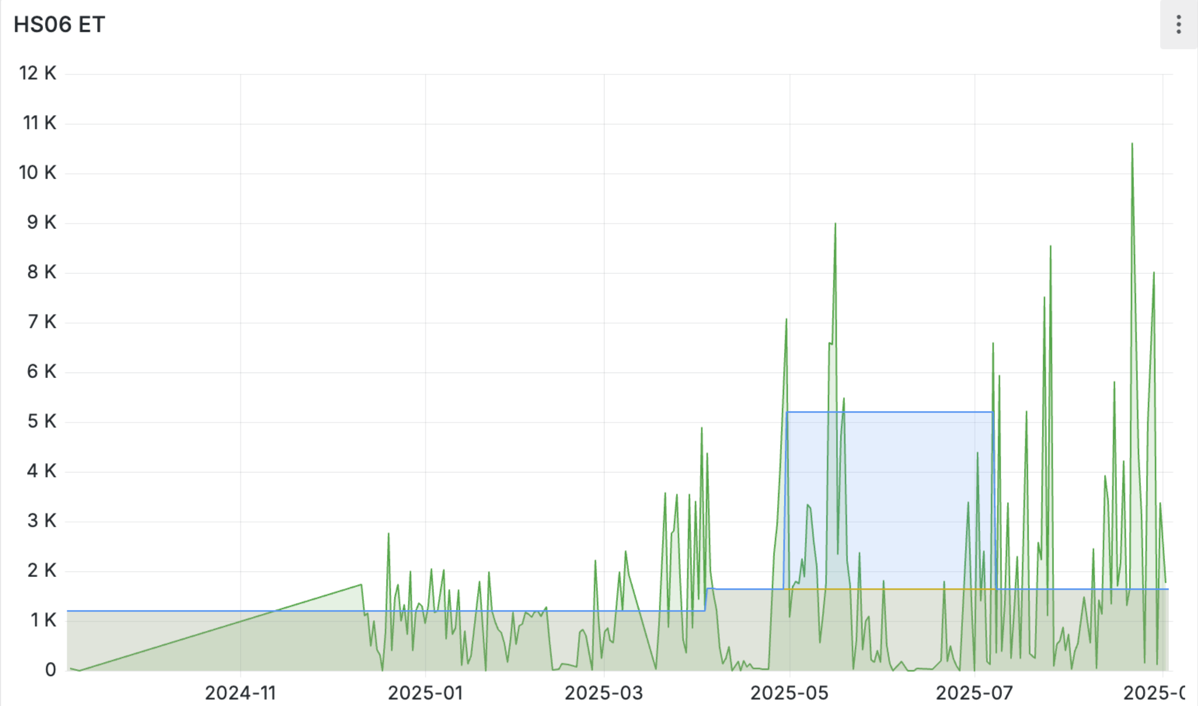

Name	Mean	Min	Max
used	1.54 K	0.381	10.6 K
pledge	1.38 K	1.20 K	1.66 K
assigned	1.96 K	1.20 K	5.20 K

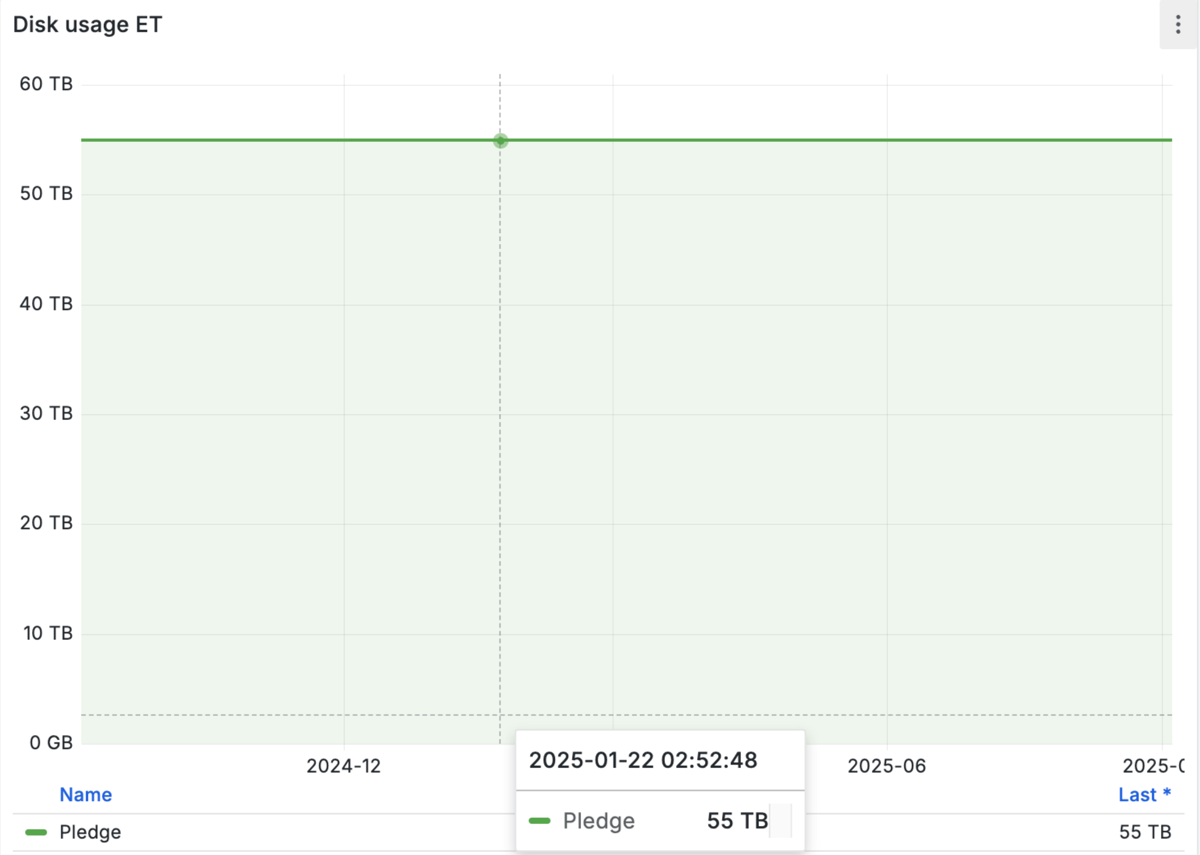

	CPU HS06	DISCO TB	TAPE TB	Infrastruttura
(Incremento)	3500	55	0	Tier-1
(Pledge)	4000	50?	0	
	10M core hours			CINECA

- Newtonian noise simulations for the Sos Enattos site
- ETO Task Force activities
- Contribution to generic MDC activities

EINSTEIN TELESCOPE

- Questions?