

EINSTEIN
TELESCOPE

The Einstein Telescope e-Infrastructure

Patrice Verdier, IP2I Lyon – IN2P3
On behalf the e-Infrastructure Board of ET

March 17th, 2026

Introduction: EIB mandate

«...to design, create and operate an evolving, efficient and functional e-infrastructure environment at a reasonable cost for the collaboration. Initially the focus will be the development of a Computing Model for the ET »

- Prepare a plan of the studies and activities that need to be undertaken for the development of the ET computing.
- Propose a computing model and its updates to the collaboration.
- Provide a software framework allowing **traceability and reproducibility, efficient job submission and data access**

Out of scope: actual science code, physics and engineering tools

Organization

EIB Chairs: Stefano Bagnasco (INFN), Paul Laycock (Unige) ~~Patrice Verdier (IP2I Lyon - IN2P3)~~

ET-PP WP8 leaders: Achim Stahl (U. Aachen), Nadia Tonello (BSC)

Division 1: Software, frameworks, and data challenge support

Andres Tanasijczuk (Université Catholique de Louvain)

Division 2: Services and Collaboration Support

Antonella Bozzi (EGO)

Division 3: Computing and data model, Resource Estimation

Gonzalo Merino (PIC)

Division 4: Multimessenger alerts infrastructure

Steven Schramm (Université de Genève)

TTG: Technology Tracking working Group

Sara Vallero (INFN Torino)

Task 8.1: T0 data center

Leader: Patrice Verdier (IP2I-IN2P3)

Task 8.2: Computing and Data Model

Leader: Paul Laycock (U. Geneva)

Task 8.3: Resources

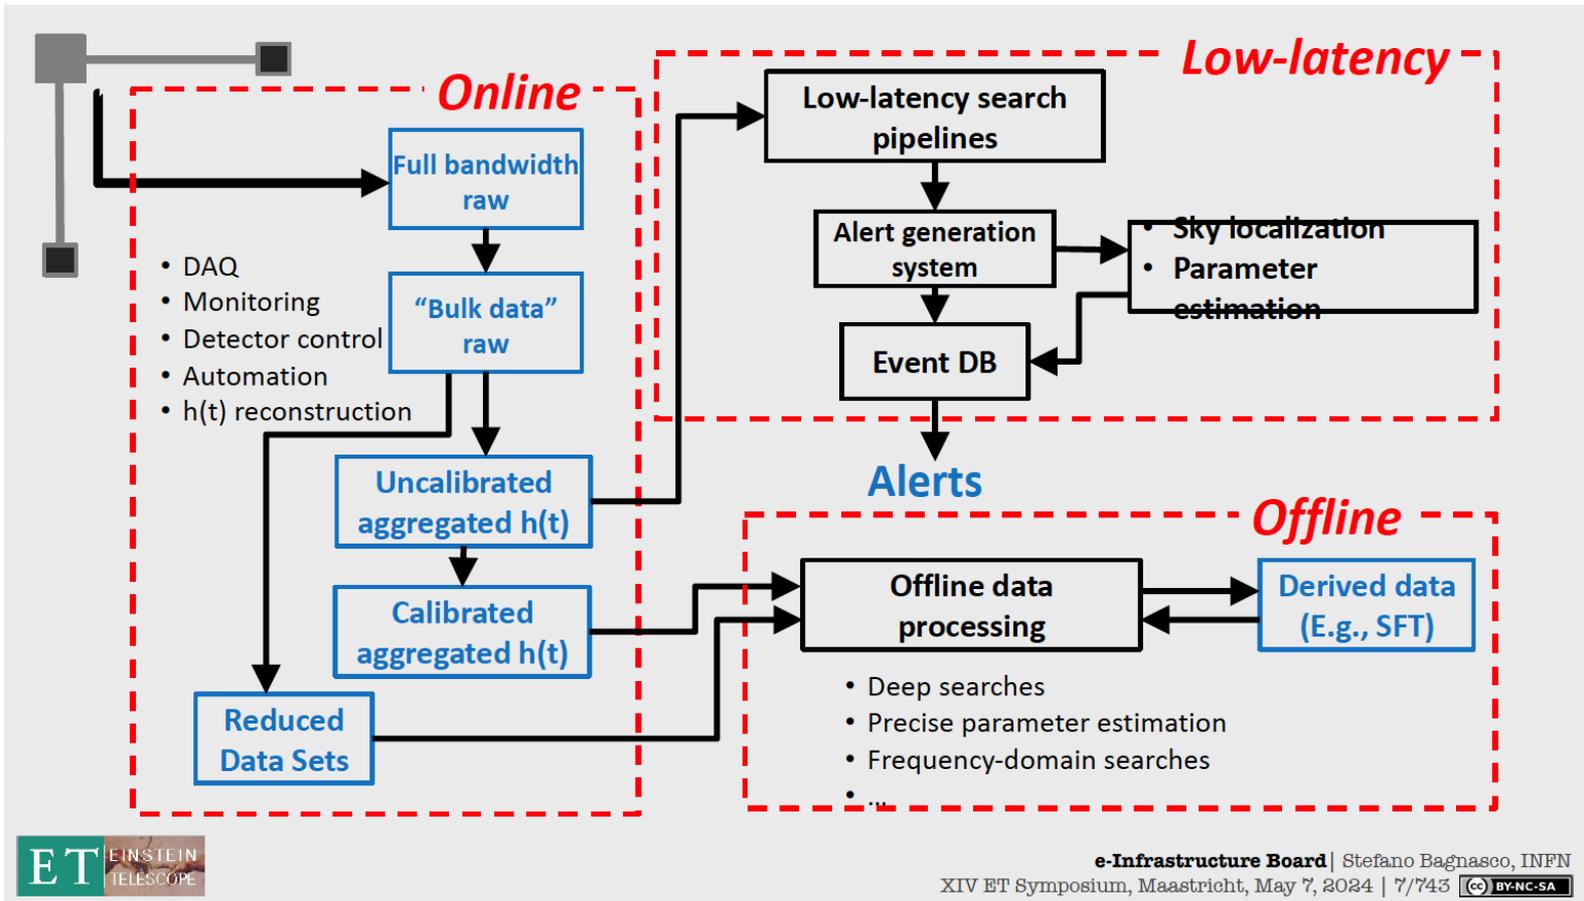
Leader: Silvio Pardi (INFN Napoli)

Task 8.4: Data Access

Implementation

Leader: Nadia Tonello (BSC)

Joint WP8+EIB weekly call for coordination





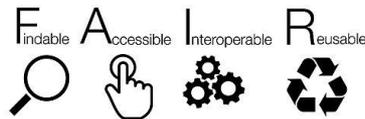
EU policy

The EU's policy framework fosters **a sustainable, sovereign, and innovative digital infrastructure** for science, ensuring that digital growth supports both **research excellence and climate neutrality**.

Embedding sustainability, interoperability, and FAIR principles into the design of new RIs, ensures scientific excellence, compliance with EU policies, and long-term viability.

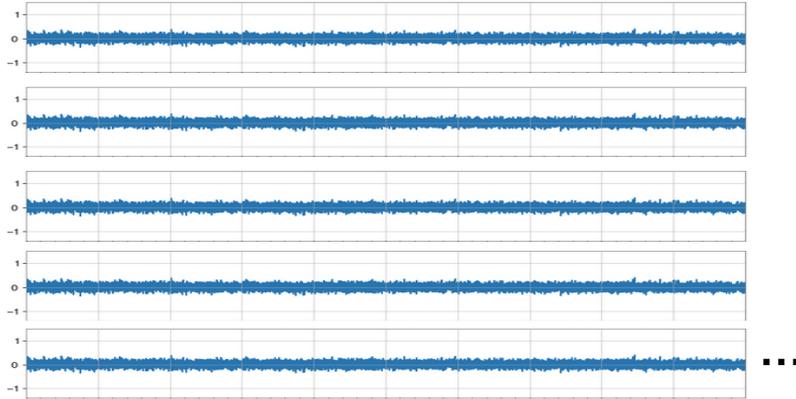
Climate-Neutral and Energy-Efficient Data Centres by 2030: The European Commission has set a clear target: all data centres in the EU must be climate-neutral, highly energy-efficient, and sustainable by 2030

This ambition is part of the broader **European Green Deal and the Green Deal Industrial Plan**, which aim to achieve climate neutrality by 2050 while boosting the competitiveness of Europe's net-zero industry

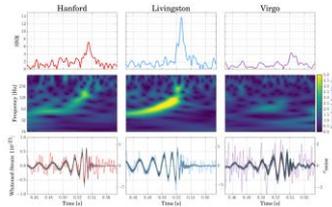


Challenges : from 2G to 3G & LISA

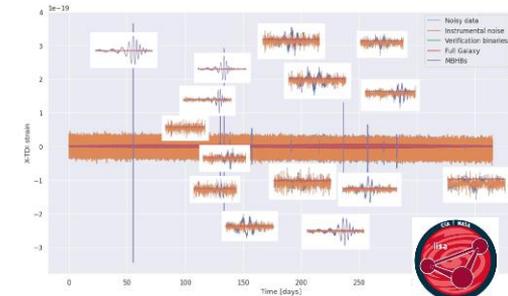
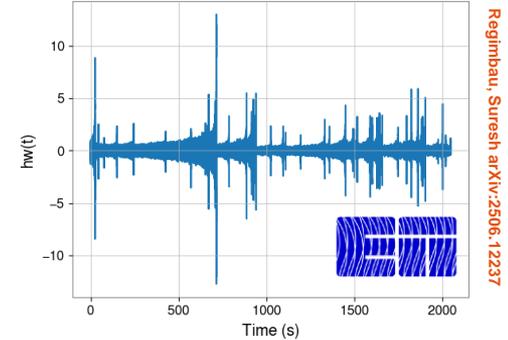
2G : LVK are dominated by backgrounds



And 2-3 times per



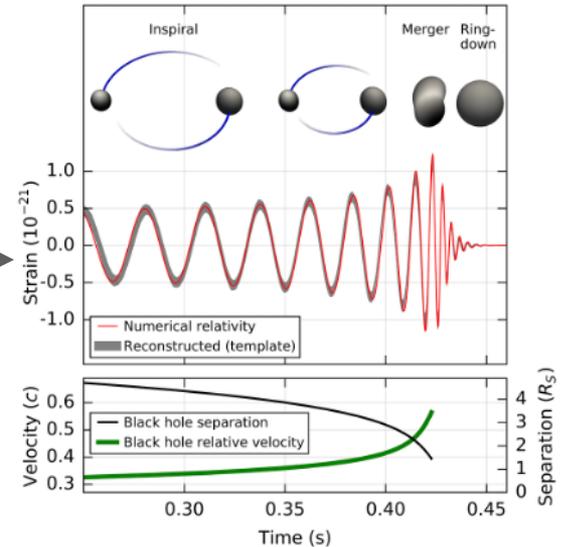
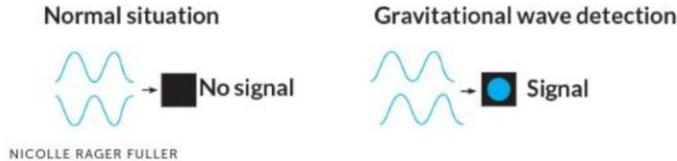
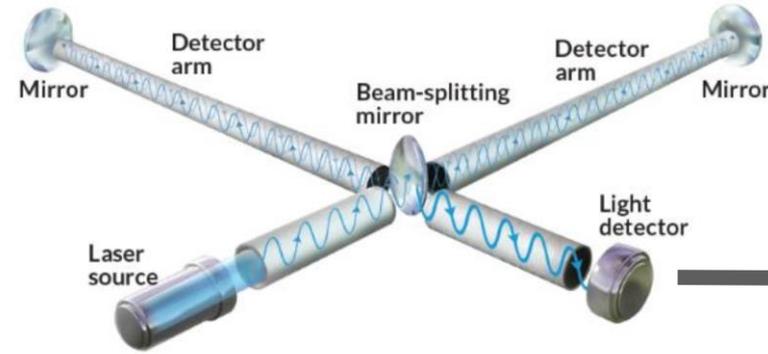
ET and LISA will be dominated by signals



Dataset size won't change so much

Complexity and computing power to reduce the data and reconstruct events will increase a lot

GW detection and GW data



- Strain data, $h(t)$, is the time-series needed for analysis, recorded at **16 kHz**
- This will be the same for **Einstein Telescope** as it is today for **LVK**
 - **Computing power needed scales ~number of events - a factor of 1000 using today's s/w**
 - **Assume a sister detector in the US - crucial to share data at low latency**

ET computing model

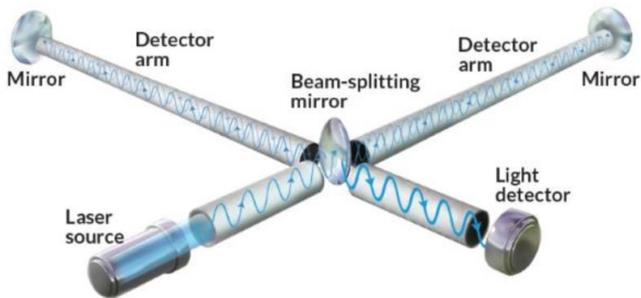
Deliverable for Feb. 2026

Contents

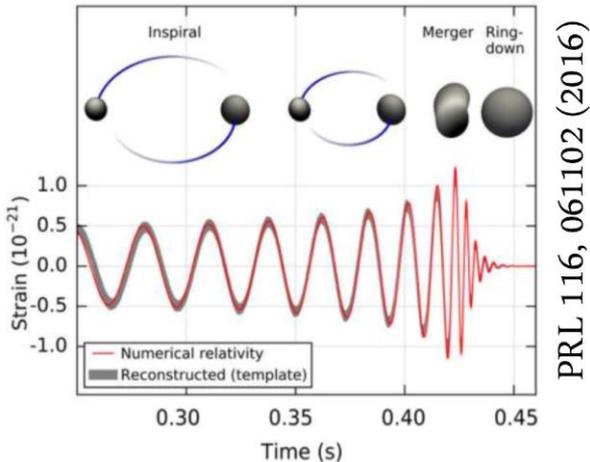
I	Introduction	9
1	Scope	10
2	ET Science Goals	13
3	Gravitational Wave Detection	18
3.1	Basic principle of GW detectors	18
3.1.1	Interferometers	18
3.1.2	Data Output	20
3.2	Analysis of gravitational wave signals	21
3.2.1	Gravitational wave events	21
3.2.2	Analysis methods	24
3.3	Gravitational Wave Alerts	27
4	Computing Model Design Principles	30
4.1	FAIR Principles	31
4.1.1	A FAIR implementation for the Einstein Telescope	31
4.1.2	Beyond FAIR	33
4.2	Computing Tasks and Resource Classification	34
4.2.1	On-site Computing	34
4.2.2	Off-site Computing	35
4.3	Software management	36
4.3.1	Software Best Practices and Metrics for Success	36
4.3.2	Organisational Principles	37
4.3.3	Critical Engineering Practices	37
4.3.4	Metrics to Track (DORA)	38
4.3.5	Release management	38
II	Core Mission	40
5	On-site Computing	41
5.1	DAQ	41
5.1.1	Main data streams	41
5.1.2	Supervisory control and data acquisition (SCADA) system	42

- **Main question:** Can we analyse data that contains 1000 times more events, including long-lasting and overlapping signals?
- **Main strategy:** Build on the experience of **LVK**, leverage expertise particularly from **CERN**
- Much of **LVK** computing relies on the “Open Science Grid” ([OSG](#)), used to provision resources for the CERN experiments ([WLCG](#)) - tools like [HTCondor](#) and [Rucio](#)
- For ET, there are also synergies with [EOSC](#) / [ESCAPE](#)
- **Main challenges:** There is no formal **LVK** Computing Model document

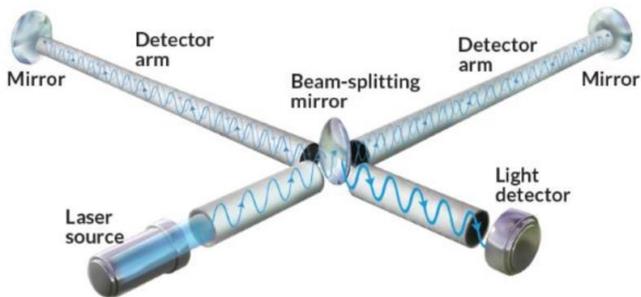
Data Challenges for Einstein Telescope



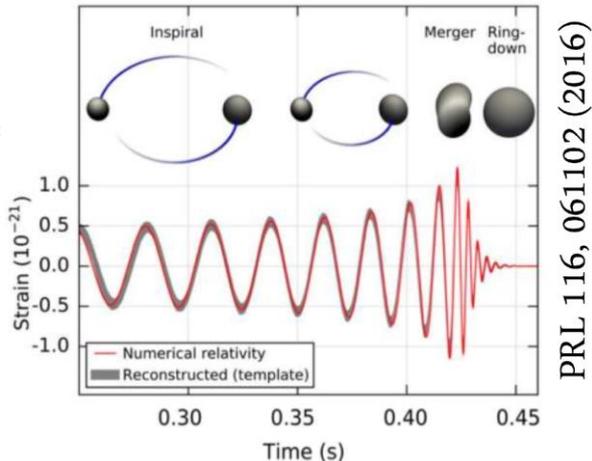
- Strain data, $h(t)$, is one ***time series sampled at 16 kHz***, but ET will have more auxiliary channels than Virgo (cryogenics, advanced detectors)
 - Over-estimate **~100 PB/year** of raw detector data (very many auxiliary channels)
 - This is significant, but largely for archive
- To first order, most analysis only needs $h(t)$
 - about **10-100 TB/year**
 - Plus data quality and some metadata



Compute Challenges for Einstein Telescope



- **Low latency computing** will deal with 1000 times more events, some of which will overlap
 - It currently takes ~ 30 s to analyse data promptly and generate alerts using **~ 20 k cores + ~ 500 GPUs**
- Increased signal sensitivity for **ET** means signals are also in-band for much longer (minutes, hours, days)
 - A huge deal for multi-messenger science !
 - *Using today's s/w, ET could need 40M cores just for low latency - this will NOT be what ET will use!*



AI/ML and GPUs are being aggressively pursued, see the [Science Blue Book](#)

Latency and computing resources

Latency

On-site dedicated resources

- Lowest latency
- Full control of hardware specification
- Operation is our responsibility

We require some resources that are completely dedicated to ET, e.g. DAQ

Off-site dedicated resources

- <1s latency of network to centre(s)
- 1s - 1min latency to start ET jobs
- ET defines hardware requirements
- Operation is centre(s) responsibility

LVK strain data is distributed (< 1s) to different sites to look for GW event candidates, we will do the same for ET Jobs must start 1s < 1 min of receiving strain data

Off-site fixed allocations

- <1s latency of network to centre(s)
- 1min - 10min latency to start ET jobs
- ET defines hardware requirements
- Operation is centre(s) responsibility

Some post-processing (enrichment) for GW event candidates takes hours, added value but not urgent Jobs should start 1min-10min of receiving strain data

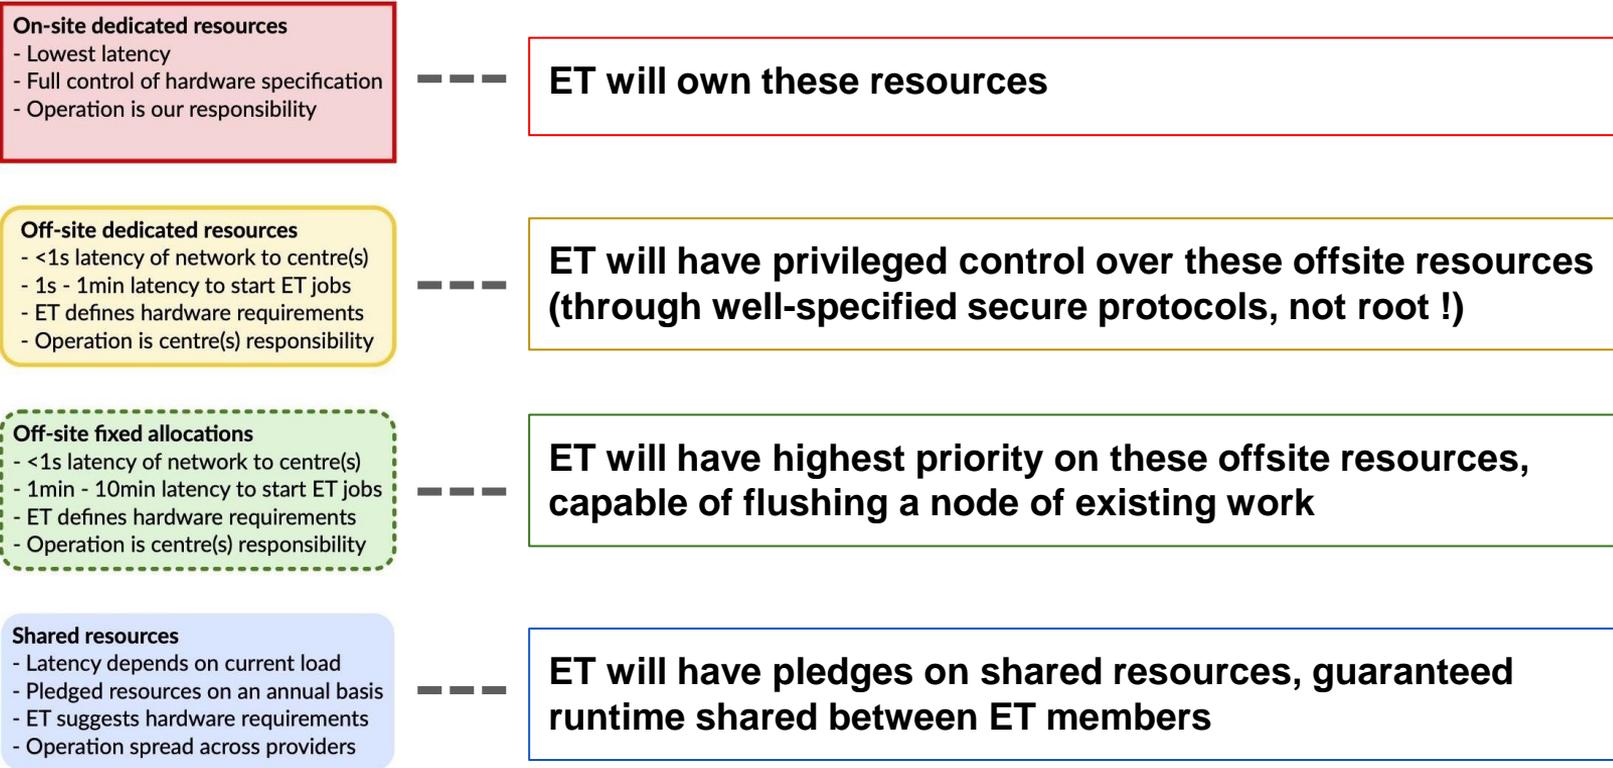
Shared resources

- Latency depends on current load
- Pledged resources on an annual basis
- ET suggests hardware requirements
- Operation spread across providers

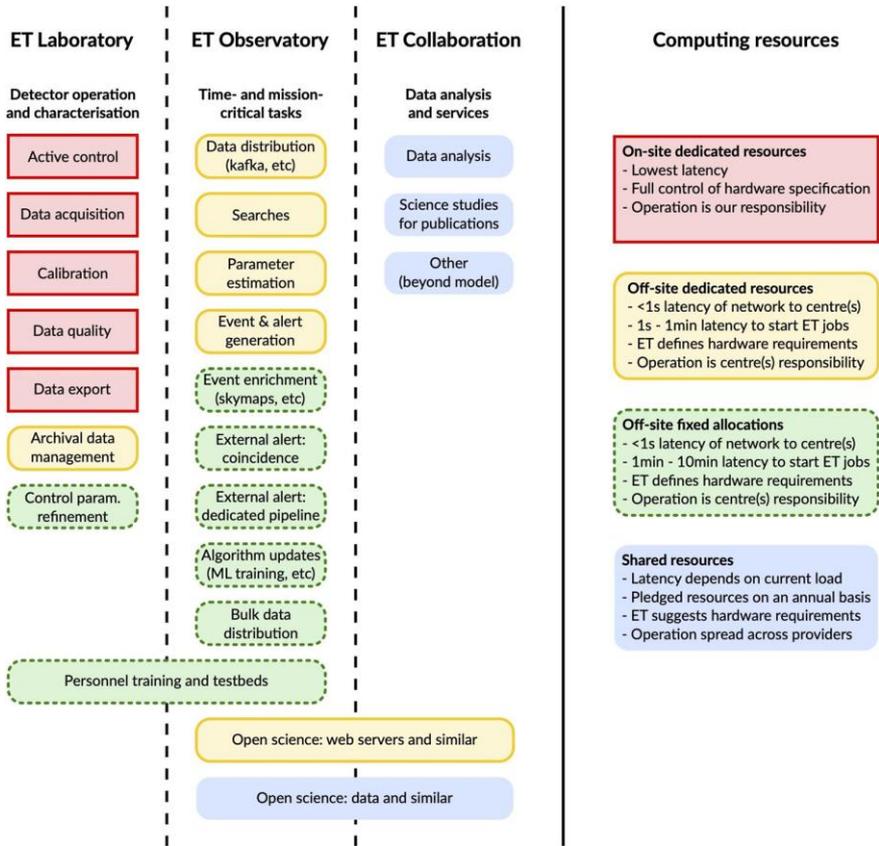
Data analysis of strain data, by collaboration members and later by the whole astro community ~asynchronous wrt availability of strain data

Latency and computing resources

Cost ↑



The ET Computing tasks

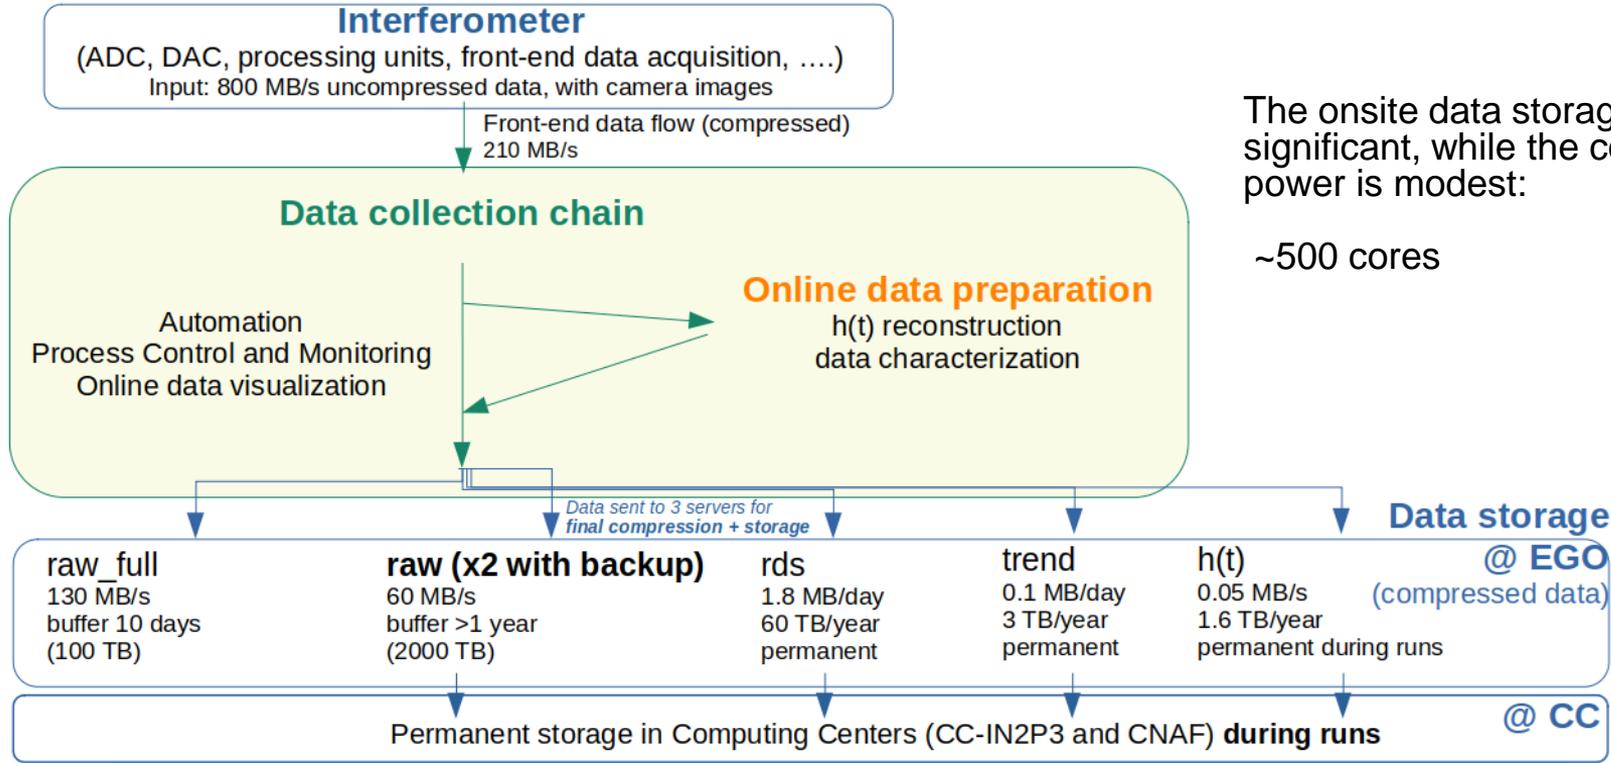


An important caveat, the ET Computing Model implementation will depend on the site decision

- The ET Computing Model concept defines the requirements
- There can be good reasons to minimise the number of computing sites, particularly for high priority tasks

In this talk, and the document, “Onsite computing” refers to things that are strictly *required to be at detector site(s)*

Onsite computing for Virgo



The onsite data storage needs are significant, while the computing power is modest:

~500 cores

Scaling data rates to ET

From ET-PP Deliverable 8.1 Computing and Data Requirements

Data rate	DAQ	Raw_full	Raw	RDS	Trend	Strain	DetChar
MB / s	250	130	2x60	1.8	0.2	0.05	0.2
TB / day	20.5	10.7	2x4.9	0.15	0.016	0.004	0.016
PB / year	7.4	3.8	2x1.8	0.05	0.006	0.001	0.006

Table 2: Data rates for the **Virgo** interferometer before the O4 run, “DAQ” is the bandwidth used by the entire DAQ system, while the subsequent columns correspond to data formats described in the text. Note that two full copies of the “Raw” data have been accounted for, while for Virgo the second copy currently only covers 6 months.

Data rate	DAQ	Raw_full	Raw	RDS	Trend	Strain	DetChar
MB / s	1250	650	2x300	9	1	0.25	1
TB / day	104	54	2x25	0.75	0.08	0.02	0.08
PB / year	37	19	2x9	0.26	0.029	0.007	0.029

Table 3: Data rates per **ET** interferometer.

Computing and Data requirements have been estimated in [D8.1](#)

The scaling model *per interferometer* is

$$ET = \text{Virgo} * 5$$

- This accounts for more auxiliary channels and higher sampling freq.
- The final number will depend on geometry

We think this is very conservative

Onsite Computing needs for ET

From ET-PP Deliverable 8.1 Computing and Data Requirements

	Minimal scenario	Operational safety margin
Operations storage buffer (TB)	800	1600
Long-term storage (PB)	20	40
CPU cores	2150	6450
RAM (GB)	6300	18900
Network	100 Gb/s	2 * 100 Gb/s

Table 7: Baseline online computing requirements per ET interferometer for data storage capacity, processing power, RAM memory and network speed.

The scaling model per IF is

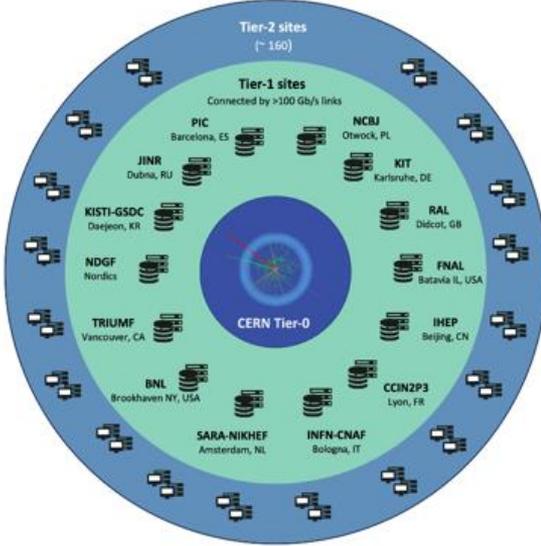
$$ET = \text{Virgo} * 5$$

The number depends on geometry, and allows for more auxiliary channels and higher sampling freq.

We think this is very conservative

While the computing power needed is modest, there is still a relatively large amount of data needed onsite and, particularly during commissioning, an efficient data reduction strategy will be a priority

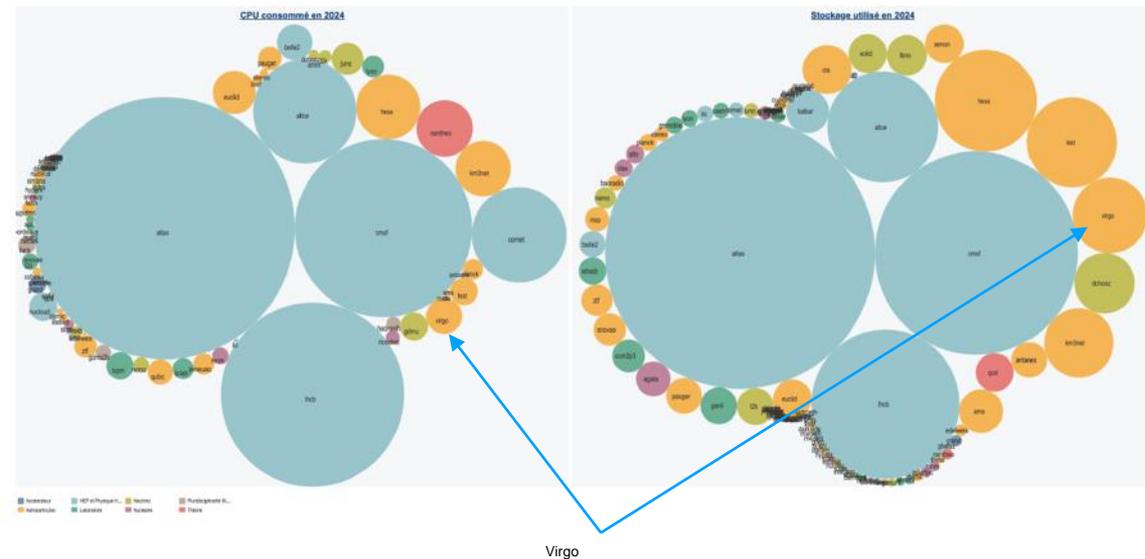
Available Computing Centers



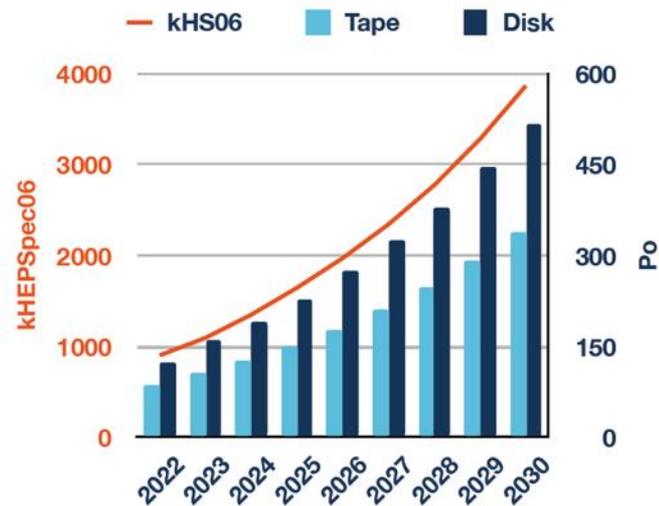
Switzerland's participation is contingent upon the ratification of its accession to Horizon Europe.
The boundaries and names shown and the designations used on this map do not imply official endorsement or acceptance by the European Union.
This designation shall not be construed as recognition of a State of Palestine and is without prejudice to the individual positions of the Member States on this issue.
Administrative boundaries: © EuroGeographics © OpenStreetMap
Cartography: Eurostat - IMAGE, 05/2025

Experiment resource usage at CC-IN2P3

- LHC is still the biggest user and will stay but others are growing => important to prepare and organize



Projection of resource evolution



ET Mock Data Challenge

ET Data analysis workshop in Bologna (2025)



@ Elena Cuoco



Discussion and actions



- Define a set of standard software distribution(s)
- Provide a CVMFS repository for distribution and policies
- Call for computing resources requests
- ET Computing Cloud MoU
- Data/metadata catalogue (annotations) for datasets
- Define the requirements for metadata
- New MDC data production code review for modularity

This event is designed as a collaborative meeting for individuals involved at various levels in the data analysis of the Einstein Telescope, as well as in the development of the infrastructure for data analysis and generation.

- Prioritized list of requested datasets
- Resources (CPU, storage, personpower required both for development & running)
- Strategies for reuse and "modular" generation
- Set up MDC analysis Code working group
- Start definition of best practices for code
- Organize MDC timeline definition meetings (or whatever)
- Define a plan for user-facing portal



Sustainability

EU analysis on RI sustainability

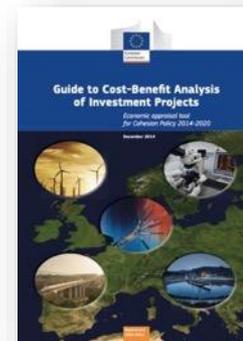
https://www.esfri.eu/sites/default/files/ESFRI_SCRIPTA_SINGLE_PAGE_19102017_0.pdf

Sustainability studies are supported by EU calls

<https://data.europa.eu/doi/10.2777/76269>

Cost-Benefit analysis:

Research Infrastructure guides foresee this approach
Required by ESFRI



<https://greendigit-project.eu>

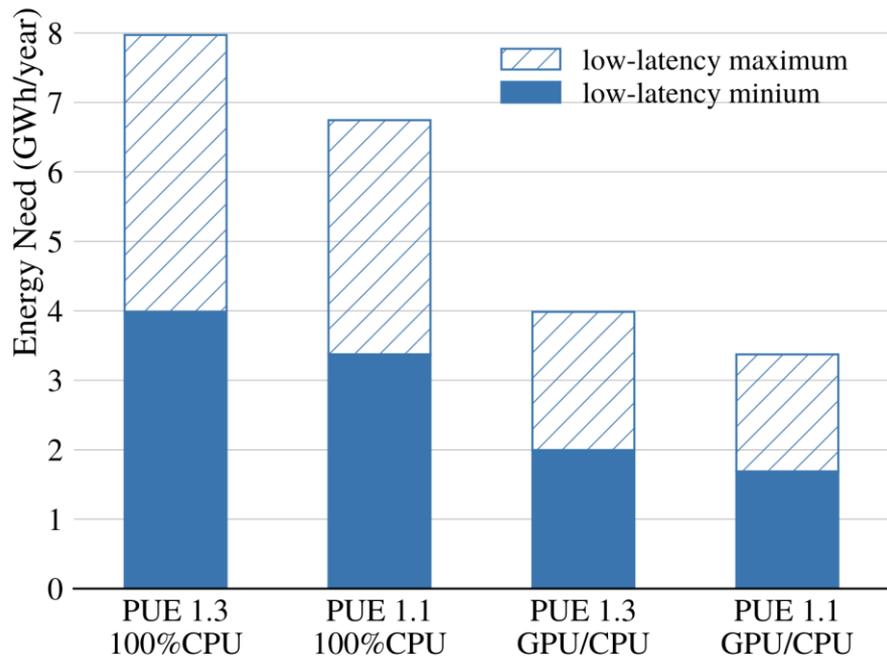
Ex. Sustainability in software development

- optimizing code for energy efficiency
- utilizing energy-efficient algorithms
- integrating low-power hardware



EGI, SoBigData
SLICES, EBRAINS

ET Computing - energy footprint



In the ET computing model, we present an estimate of the energy footprint of ET low latency computing. This is expected to dominate assuming the **full data stream** is only processed by low latency computing

This follows the [methodology used by the WLCG](#) for worldwide computing needs for CERN experiments (beyond the CERN sites)

Power Usage Efficiency (PUE) makes a big difference, as do strategies that leverage GPUs

Training & Open Science



The JENA communities should aim to create coherent, stable, and predictable training programs and events that harness the strengths of the European community on Federated Computing.

These efforts are supported at European level through the EOSC association (<https://eosc.eu/>) which aims at putting in place a system to help researchers store, share, process, analyse and reuse FAIR research outputs within and across disciplines and borders.

The European Science Cluster of Astronomy & Particle physics ESFRI research infrastructures



Fostering the uptake of Open Science in Europe

In response to the EU call on EOSC HORIZON-INFRA-2023-EOSC-01-01



2 ET projects funded in the OSCARS cascading grants calls

Training Key Initiatives

HEP Software Foundation (HSF) - 2800+ trained participants to facilitate coordination and common activities of software and computing in High Energy Physics and beyond

<https://hepsoftwarefoundation.org/activities/training.html>



EGI federation of computing and storage resource providers

Large training catalog: <https://www.egi.eu/trainings-and-webinars>



ESCAPE - Open Science training for data analysis

Gray Scott School : <https://projectescape.eu>



EVERSE - Research software excellence framework

European Virtual Institute for Research Software Excellence

The EVERSE project aims to create a framework for research software and code excellence.



Conclusions on Computing Model

- **The existing e-Infrastructure developed for LIGO-Virgo-KAGRA/IGWN is used as a starting point for the Einstein Telescope Computing Model**
- **However**, there are many interesting and new challenges for 3G GW interferometers:
 - A huge number of alerts in the era of multi-messenger physics, i.e. multi-Ris
- The complexity of data processing will require new techniques and technologies, including:
 - Real time monitoring of the infrastructure, the detector and the environment enabling background/noise subtraction and Digital Twin solutions
 - AI/ML techniques (everywhere !)
- **Sustainability, and especially energy consumption, is one of the critical aspect of the ET computing model** : it require new tools and methods to be centrally integrated in the project organization and a small but dedicated team of software engineers and computing experts
- **ET Computing Model and further developments must build on cross-domain expertise developed especially within the European landscape (EOSC, ESCAPE, EVERSE, SPECTRUM, GreenDigit, ...). Tight collaborations with HEP/CERN, KM3NeT, CTA, SKA, LSST is mandatory**

Many possibilities of contribution and collaboration with ET

Conclusion and next steps

- **ET Data Access Policy:**

Delivered end of 2025 (ET-PP Deliverable): based on LVK model with embargo period

- **ET Open Source Policy:**

Proposal Introduction and Discussion: Inspired by LIGO and Virgo, a draft proposing an open source policy was discussed to address licensing and authorship issues.

- **Code development Policy:**

Preparation of a code development, management and quality “best practices” document to be proposed for all software development within the collaboration

To be included in ET
Data Management Plan

- **Use Mock-Data-Challenges** as bench tests for the development of the ET e-Infrastructures

TTG: Technology Tracking Working Group

ET INDUSTRY COMPUTING WORKSHOP

14 - 15 April, 2026
Nikhef, Amsterdam

- AI-enabled algorithms, tools and approaches for data analysis
- AI tools for software development and more
- Evolution of processor architectures (accelerators, stream processors, neuromorphic...) incl. Quantum
- Storage and data management technologies
- Network technologies and cybersecurity
- Infrastructure and sustainability



Overview
Timetable
Participant List
Workshop venue
How to reach the venue
Dinner Tuesday 14th April
Accommodation

INDUSTRY COMPUTING WORKSHOP 2026

The Einstein Telescope Preparatory Phase European Project (ET-PP) and the ET Collaboration e-Infrastructure Board are organising a workshop on ET computing, to be held at Nikhef on April 14-15, 2026.

The objective of the workshop is twofold: firstly, to examine the opportunities that the development of information technologies will offer for realizing the scientific potential of the Einstein Telescope; and secondly, to explore possible approaches to collaboration with industry that can be effective in enabling beneficial synergies.

Topics for the workshop will be hardware and software technologies, future trends and sustainability issues related to:

- AI-enabled algorithms, tools and approaches for data analysis
- AI tools for software development and more
- Evolution of processor architectures (accelerators, stream processors, neuromorphic...)
- Storage and data management technologies
- Network technologies

<https://indico.cern.ch/event/1605739/>

*This Workshop is being organized by **ET-EiB** and **ET-PP WP8**, in collaboration with **ET-PP WP7***

Reserve

Services and Collaboration Support



EINSTEIN
TELESCOPE

Welcome to **et**

Sign in with your et credentials

Username

Password

Sign in

[Forgot your password?](#)

Or sign in with

eduGAIN

Not a member?

Apply for an account

Register an account with eduGAIN

[Info and Privacy Policy](#)

ET Authentication and authorization infrastructure (AAI)

- Implementing the AAI solution developed within ESCAPE and supported by major European centers/labs in HEP and astroparticle physics (AAI solution adopted by CERN, CC-IN2P3, CNAF...) : **Indigo IAM: <https://indigo-iam.github.io/v/v1.8.3/>**
- **Working Group in place: CNAF, EGO, IJCLab**
 - Deploying an Indigo IAM instance for ET at CNAF: done - <https://iam-et.cloud.cnaf.infn.it/>
 - Interface the ETMD with Indigo IAM : in progress
 - Implement first service/tool with this AAI system : to come

Fuller details available here:

<https://apps.et-gw.eu/tds/ql/?c=18244>

ESCAPE/OSCARs and Einstein Telescope

CASCADING-GRANT CALLS FOR OPEN SCIENCE PROJECTS 



Open Call for Open Science Projects

Launch event

15 March 2024
Online



GOAL:

Build on the science cluster approach to ensure the uptake of EOSC, i.e., consolidate FAIR services of the five Science Clusters and, more broadly, perform excellent science and pursue societal benefits by leveraging an Open Research approach.

TARGET USER COMMUNITIES:

Science Clusters and wider community (RIs, Universities, Institutes, either consortia, or individual researchers)

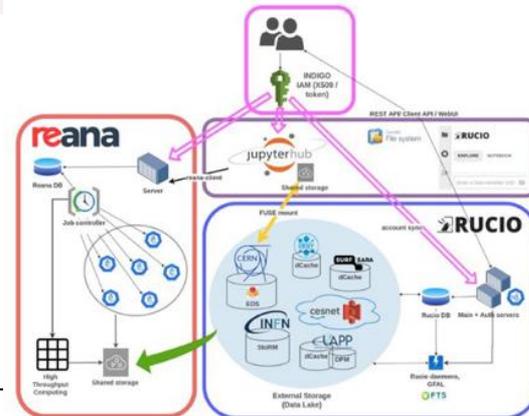
Evaluation criteria for the independent expert panel

- Opens: ~ March 2024 / Nov. 2024
 - Submission within 60 days
 - Project start: Sept-Dec. 2024 / Aug-Oct. 2025
 - Budget: 100 - 250 k€ / project
 - Duration: 1 - 2 years
- Project description: clear objectives, towards FAIR and open
 - Scientific impacts: multiple RIs / cross-cluster
 - Digital resources: use of EOSC services / new EOSC service
 - Implementation: realistic within budget



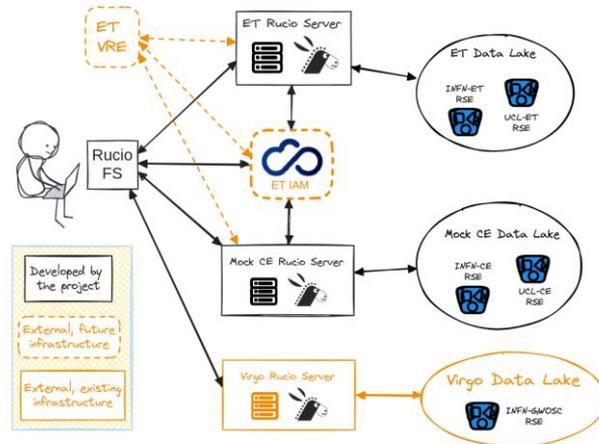
July 2023:
ET application
to join the
ESCAPE
consortium was
approved

G.LAMANNA



Madden project

MADDEN Multi-RI Access and Discovery of Data for Experiment Networking



Participating organizations:

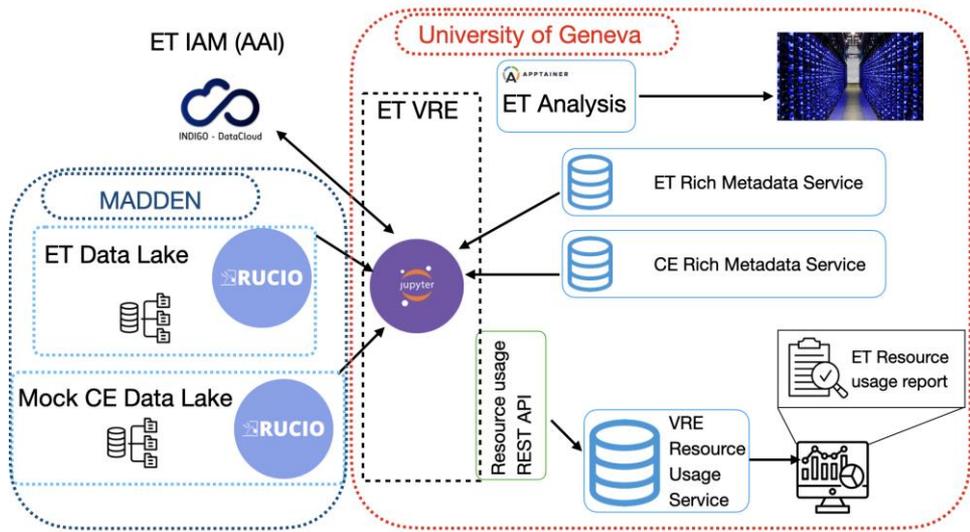
- INFN Torino (PI: Federica Legger)
- UC Louvain (coord.: Andres Tanasijczuk)
- Targeted start date: January 2025
- Duration of the project: 24 months
- Overall funding: 210 K€

The main objectives of this project are:

- Build a multi-RI Data Lake managed with Rucio.
- Develop and test RucioFS, a tool to provide a POSIX-like view of the Rucio catalogue in a multi-RI environment.
- Extend RucioFS to support advanced querying capabilities using metadata.

OSCARs call: ETAP

Einstein Telescope Analysis Portal (ETAP)



- Deploy the CERN ESCAPE VRE at University of Geneva
- Connect to multi-RI Data Lakes managed by Rucio (MADDEN)
- Deploy multi-RI Metadata services from the HEP Software Foundation (HSF)
- Design a flexible computing resource monitoring service

- University of Geneva
- Start date January 2025
- Duration 18 months
- Funding ~250k

Paul.Laycock@unige.ch

The Virtual Research Environment



Running an analysis on the ESCAPE VRE

Example: notebook running multiprocess Wavelet Detection Filter pipeline for burst signal detection (Cuoco et al. 2018, Cuoco et al. 2001)

The screenshot shows a JupyterLab environment with a file browser on the left and a code editor on the right. The file browser shows a directory structure with files like 'ARcoeff-AR300...', 'LVcoeff-AR300...', 'parametersUsed...', and 'WDFTriggers-E1...'. The code editor shows a notebook cell with the following code:

```
[15]: wdf=wdfUnitDSWorker(par,fullPrint=1)

[*]: # ---- Run multiprocess wdf (without state vectors) ----
with mp.Pool(par.nproc) as p:
    wdf=wdfUnitDSWorker(par,fullPrint=1)
    p.map(wdf.segmentProcess, par.segments)
    p.close()

#wdf.segmentProcess(par.segment)
```

The output of the code is a log of analysis results for multiple segments, including channel information, sampling rates, and estimated sigma values.

```
2025-05-27 00:20:51 jupyter-alberto@root[2088] INFO Analyzing segment: 1000002048.0-1000004096.0 for channel E1:STRAIN downsampled at 4096Hz
2025-05-27 00:20:51 jupyter-alberto@root[2089] INFO Analyzing segment: 1000004096.0-1000006144.0 for channel E1:STRAIN downsampled at 4096Hz
2025-05-27 00:20:51 jupyter-alberto@root[2090] INFO Analyzing segment: 1000000000.0-1000002048.0 for channel E1:STRAIN downsampled at 4096Hz
2025-05-27 00:20:51 jupyter-alberto@root[2088] INFO Load AR parameters
2025-05-27 00:20:51 jupyter-alberto@root[2087] INFO Load AR parameters
2025-05-27 00:20:51 jupyter-alberto@root[2090] INFO Load AR parameters
2025-05-27 00:20:51 jupyter-alberto@root[2088] INFO Load AR parameters
2025-05-27 00:20:51 jupyter-alberto@root[2089] INFO Load AR parameters
2025-05-27 00:20:51 jupyter-alberto@root[2088] INFO Estimated sigma= 2.48899e-23
2025-05-27 00:20:51 jupyter-alberto@root[2090] INFO Estimated sigma= 2.48976e-23
2025-05-27 00:20:51 jupyter-alberto@root[2089] INFO Estimated sigma= 2.49079e-23
2025-05-27 00:20:51 jupyter-alberto@root[2087] INFO Estimated sigma= 2.49165e-23
2025-05-27 00:21:01 jupyter-alberto@root[2087] INFO Starting detection loop
2025-05-27 00:21:02 jupyter-alberto@root[2088] INFO Starting detection loop
2025-05-27 00:21:02 jupyter-alberto@root[2089] INFO Starting detection loop
2025-05-27 00:21:03 jupyter-alberto@root[2090] INFO Starting detection loop
```

The bottom of the screenshot shows a blue bar with the number '10' on the left and the date '27/05/2025' on the right.

Alberto.Iess@lapp.in2p3.fr



@ Art Institute Of Chicago

=> ET-PP WP9

EU analysis on RI sustainability

https://www.esfri.eu/sites/default/files/ESFRI_SCRIPTA_SINGLE_PAGE_19102017_0.pdf

Sustainability studies are supported by EU calls

<https://data.europa.eu/doi/10.2777/76269>

Cost-Benefit analysis:

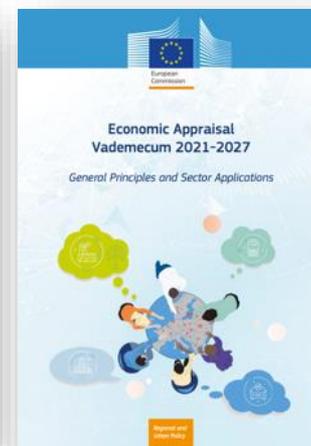
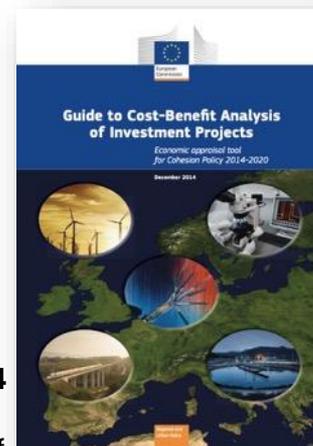
Research Infrastructure guides foresee this approach

Required by ESFRI

J. Gutleber (CERN), « Future Circular Collider Sustainability Aspects », May 2024

ESFRI, « Long-Term Sustainability for Research Infrastructures », 2017

https://www.esfri.eu/sites/default/files/ESFRI_SCRIPTA_SINGLE_PAGE_19102017_0.pdf



JENA - Joint ECFA, NUPECC, APPEC Activities

The JENA Symposium in May 2022 in Madrid (<https://indico.cern.ch/event/1040535/>) revealed an increased need for discussions on the strategy of EU federated computing at future large-scale research facilities.

Focused workshop on the strategy of computing in [Bologna June 2023](#) aimed to define computing requirements in the next decade and to try and find synergies.

Outcome: creation of 5 WGs to generate input (whitepapers) for JENA Symposium in 2025:

- WG1: HTC, WLCG and HPC
- WG2: Software and Heterogeneous Architectures
- WG3: Federate Data Management, Virtual Research Environments and FAIR/Open Data
- WG4: Machine Learning and Artificial Intelligence
- WG5: Training, Dissemination, Education



12

@Gonzalo Merino

CC-IN2P3

Description and missions

- National research infrastructure for IT resources dedicated to our research fields
 - main experiments LHC/HL-LHC T1, LSST, Belle II, CTA, KM3NET, DUNE,...
- Provides storage (disk+tape) and computing resources with the more appropriate architecture
 - Mainly HTC but increasing part of GPU and small HPC resources included
- Provides IT related services
 - IN2P3 sites connexion in relation with Renater
 - Tools for software developments, set of collaborative tools



CC-IN2P3 Datacenter

- 2 computer rooms: 1700 m² up to 390 racks
- 929 kHS23 ~ 57 664 cores
- 80 PB disks + 180 PB tapes

- 80 staffs

- Users :
 - CPU 104 teams/930 users
 - STO 221 teams/4596 users



CC-IN2P3 and experiments



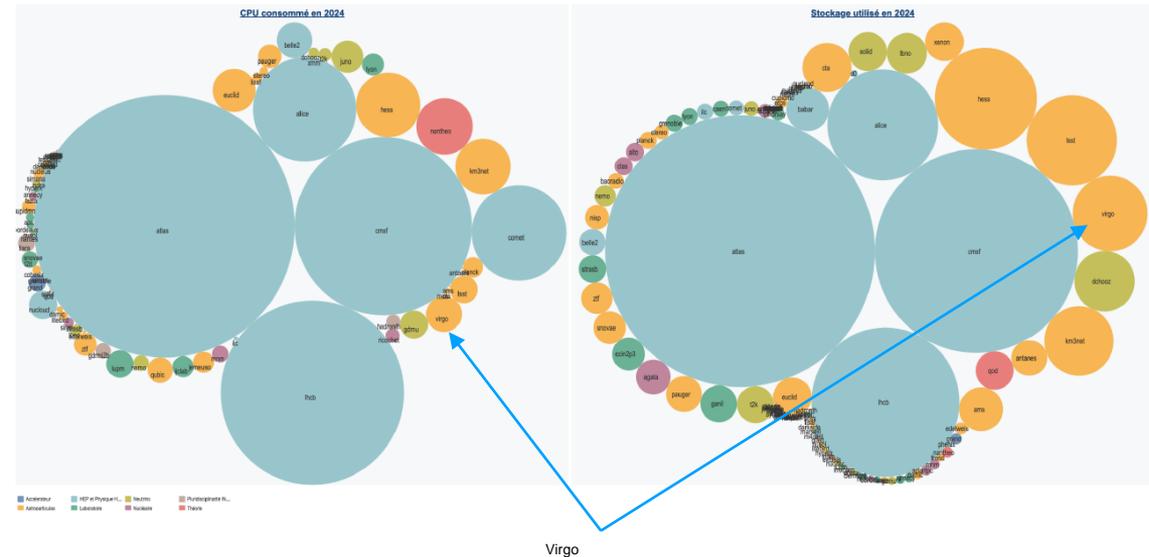
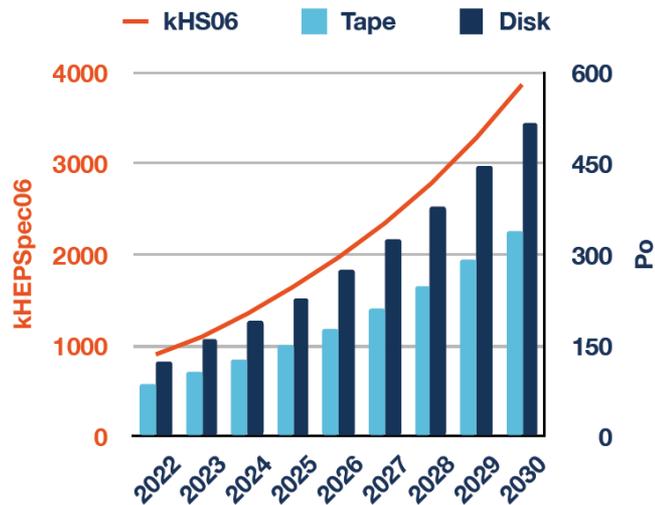
CC-IN2P3 supports 80 experiments

- each collaboration is asked by IN2P3 to provide plan for ressource needs for their lifespan and DMP updated regularly
- resource requests reviewed each year
- computing needs discussed and evaluated for all new experiments before approval

Experiment resource usage at CC-IN2P3

- LHC is still the biggest user and will stay but others are growing => important to prepare and organize

Projection of resource evolution



FITS project

A new computing room at CC-IN2P3

2025



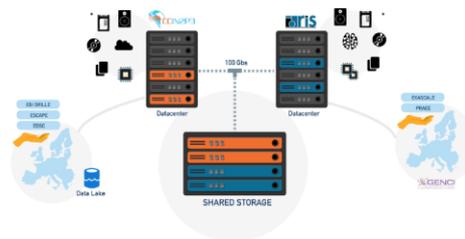
05/01: start of work

15/05: end of structural work and beginning of second work

December: qualification and delivery

This project aims to federate these two computing centers, respecting their specific missions and know-how through the implementation of a distributed infrastructure for data storage, processing and provision as well as their dissemination and enhancement. Hosted in environmental conditions with a low carbon footprint, this infrastructure will provide research infrastructures with computing and data processing resources combining high-performance computing and high-throughput computing, supercomputing with graphics accelerators and data processing farms, containerization and cloud computing.

Upgrade hosting/housing capacities of both sites
 Shared storage infrastructure
 Access portal
 Sustainability



FITS

CNRS Federated IT services
 for Research Infrastructures



Conclusion

EIB has a general strategy of iterative “early” deployment of the e-infrastructure

- Using MDCs as multipurpose tools:
 - Informal milestones
 - Assess the suitability of tools and infrastructures with feedback from the community
 - Evaluate the parameters of the Computing Model
- Prototypes for some of the functionalities are being developed
 - Also thanks to the OSCARS funding, we try not to develop what can be shared
 - Not yet all of them (frameworks!)
 - And some interesting work on possible future advanced developments
- Work on the computing model is ongoing
 - Deliverable D8.2 due Feb 2026
 - Draft released to the collaboration Dec 2025
 - Will include plans for resource requirements evaluation and a first estimate of personpower
- Still many (but not too many!) years to go
 - We still don't know a lot of things
 - We have time to learn and develop tools
 - Technology will evolve, possibly in unexpected directions

Conclusion and next steps

- **ET Data Access Policy:**

Delivered end of 2025 (ET-PP Deliverable): based on LVK model with embargo period

- **ET Open Source Policy:**

Proposal Introduction and Discussion: Inspired by LIGO and Virgo, a draft proposing an open source policy was discussed to address licensing and authorship issues.

- **Code development Policy:**

Preparation of a code development, management and quality “best practices” document to be proposed for all software development within the collaboration

To be included in ET
Data Management Plan

TTG: Technology Tracking Working Group

ET INDUSTRY COMPUTING WORKSHOP

14 - 15 April, 2026
Nikhef, Amsterdam

- AI-enabled algorithms, tools and approaches for data analysis
- AI tools for software development and more
- Evolution of processor architectures (accelerators, stream processors, neuromorphic...) incl. Quantum
- Storage and data management technologies
- Network technologies and cybersecurity
- Infrastructure and sustainability



Overview

Timetable

Participant List

Workshop venue

How to reach the venue

Dinner Tuesday 14th April

Accommodation

INDUSTRY COMPUTING WORKSHOP 2026

The Einstein Telescope Preparatory Phase European Project (ET-PP) and the ET Collaboration e-Infrastructure Board are organising a workshop on ET computing, to be held at Nikhef on April 14-15, 2026.

The objective of the workshop is twofold: firstly, to examine the opportunities that the development of information technologies will offer for realizing the scientific potential of the Einstein Telescope; and secondly, to explore possible approaches to collaboration with industry that can be effective in enabling beneficial synergies.

Topics for the workshop will be hardware and software technologies, future trends and sustainability issues related to:

- AI-enabled algorithms, tools and approaches for data analysis
- AI tools for software development and more
- Evolution of processor architectures (accelerators, stream processors, neuromorphic...)
- Storage and data management technologies
- Network technologies

<https://indico.cern.ch/event/1605739/>

*This Workshop is being organized by **ET-EiB** and **ET-PP WP8**, in collaboration with **ET-PP WP7***

ET-PP: Computing Model

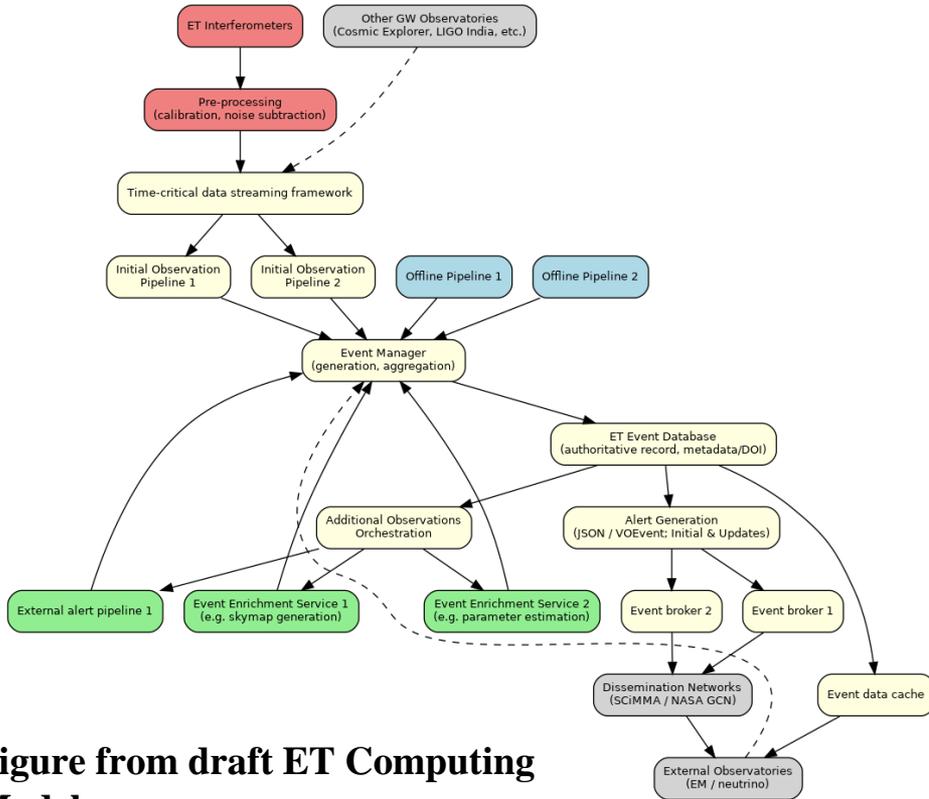
@ Nadia, Oscar, Paul

Successful coordination and organization of the activities of WP8 with the ET e-Infrastructure Board (EIB).
Chairs and all-hands online regular meetings.

		WP8 Tasks			
		T8.1 T0 data center	T8.2 Computing and data model	T8.3 Resources	T8.4 Data access implementation
ET EIB Divisions	D1 SW frameworks and Data Challenges		Computing frameworks domains and data formats	Resources for frameworks execution and data storage availability	Data availability Data releases format
	D2 Services and collaboration support				Tools for monitoring, AAI (IAM) data access
	D3 Computing and data models, resources estimation	T0 storage and computing resources estimation	Computing model Data model	Resources estimation	
	D4 Multimessenger alerts infrastructure				Tools for multi-messenger alerts
		Technology tracking working group			

Content	Type	Status	Date
M8.3 on-site infrastructure, computing and data	Milestone	In preparation	July 2025
M8.4 low-latency and offline workflows, computing and data model	Milestone		Dec 2025
M8.5 data management, data access policy and implementation	Milestone		July 2026
D8.2 Computing and data model for ET	Deliverable		Feb 2026
D8.3 Data access policy implementation	Deliverable		July 2026

Sustainable Computing for ET



*For ET, the biggest challenge for a sustainable computing model is keeping computing resource needs under control for ET's **50 year lifespan***

There are clear needs for a small, dedicated team of software and computing professionals to provide:

A common software framework

Comprehensive end-to-end testing of all ET software

Intelligent analysis platforms that eliminate duplication

Scalable, portable computing infrastructure

Figure from draft ET Computing Model