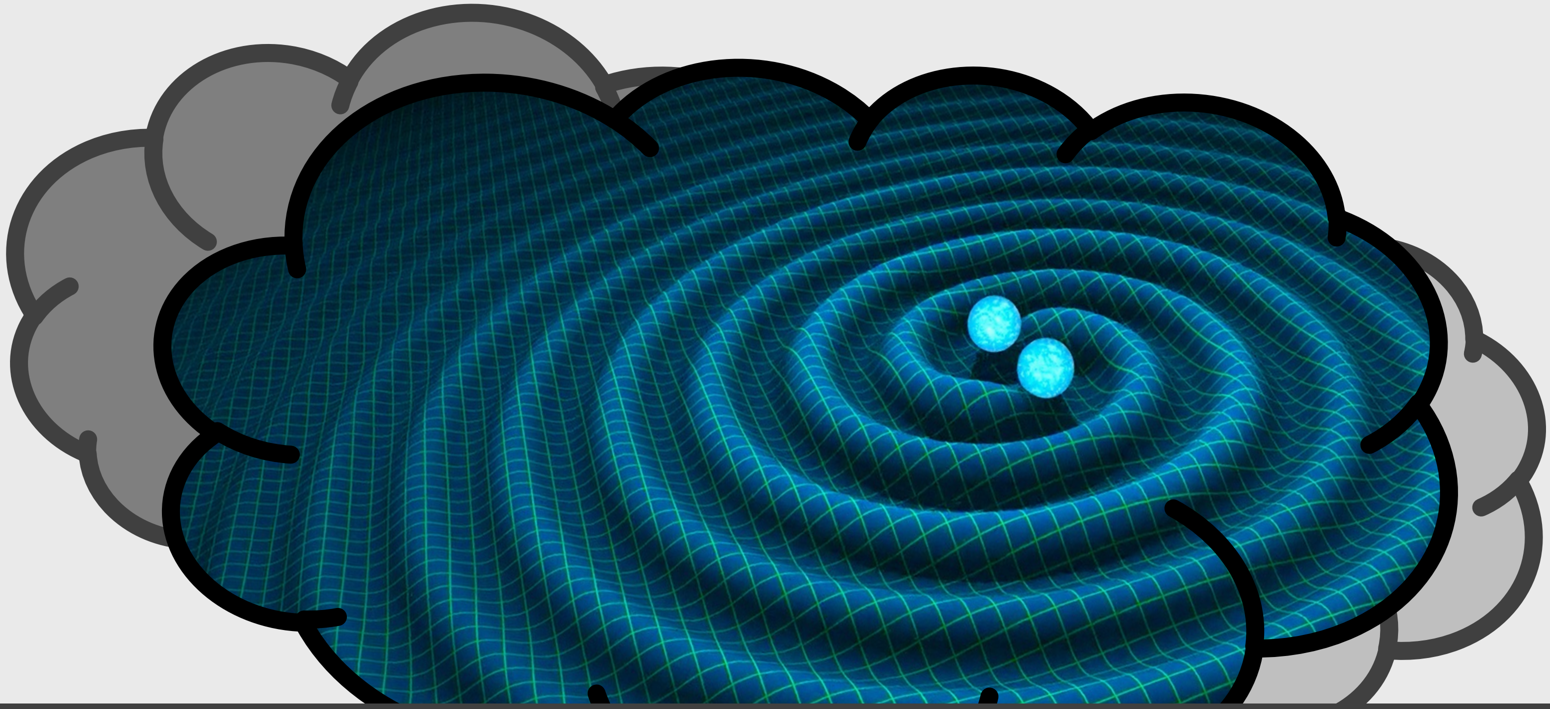

# E-INFRASTRUCTURE BOARD


**Stefano Bagnasco, INFN**

For the e-Infrastructure Board


**XV ET Symposium, Bologna**

EINSTEIN TELESCOPE

May 29, 2025

# DRAMATIS PERSONAE

**EIB Chairs**: Stefano Bagnasco (INFN), 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 (UC 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 (Geneva)

**Task 8.3: Resources**
Leader: Silvio Pardi (INFN Napoli)

**Task 8.4: Data Access Implementation**
Leader: Nadia Tonello (BSC)

**Liaison with OSB Div. 10:** John Veitch (University of Glasgow), Elena Cuoco (Bologna)

Joint ET-PP WP8 & ETC-EIB management (e.g., weekly call for coordination)

EINSTEIN TELESCOPE

 

# 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!

 

# THE ESCAPE OPEN COLLABORATION


# AT THIS SYMPOSIUM

**EIB**

📍 Room 213, CNR

| | |
|---|---|
| 09:00 – 09:10 | *Introductions* s Patrice Verdier, Stefano Bagnasco |
| 09:10 – 09:30 | *Div 3 report* |
| 09:30 – 09:50 | *ET-PP WP8* |
| 09:50 – 10:10 | *Challenges in Multimessenger Astronomy in the ET Era: From Interoperability to Multimodal Generative AI Systems* Giuseppe Greco |
| 10:10 – 10:30 | *Update on the CTLab4ET deployment* Lia Lavezzi |
| 10:30 – 10:50 | *GW data analysis on the ESCAPE Virtual Research Environment* Alberto Iess |
| 10:50 – 11:10 | *Discussion* |

**EIB**

📍 Room 213, CNR

| | |
|---|---|
| 11:30 – 11:50 | *The ET Members Database* Gary Hemming |
| 11:50 – 12:10 | *ETAP + MADDEN status* |
| 12:10 – 12:30 | *The UK effort for 3G computing* |
| 12:30 – 13:00 | *Discussion* |

**OSB & EIB joint session**

📍 Rooms 215 + 216, CNR

| | |
|---|---|
| 14:30 – 14:45 | *Report from EIB Div1* |
| 14:45 – 15:00 | *Report from OSB Div10* |
| 15:00 – 15:15 | *Results from MDC1* |
| 15:15 – 15:30 | *MDC data generation code management* |
| 15:30 – 15:45 | *EIB Div4 and LL computing model* |
| 15:45 – 16:00 | *Discussion on MM science and the LL computing model* |

# WP 8: Deliverables and milestones

EINSTEIN TELESCOPE

| Content | Type | Status | Date |
|---|---|---|---|
| M8.1 Workflows requirements collection and constraints. Workshop Uni Geneva | Milestone | Done - indico | Oct 2023 |
| D8.1 Computing and Data requirements submitted (UniGe). Reviewed and updated | Deliverable | Delivered - TDS | Feb 2024 |
| M8.2 Computing infrastructure availability for ET workflows characteristics. Workshop Napoli | Milestone | Done - indico | July 2024 |
| 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 |
| **D8.2 Computing and data model for ET** | **Deliverable** | | **Feb 2026** |
| M8.5 data management, data access policy and implementation | Milestone | | July 2026 |
| **D8.3 Data access policy implementation** | **Deliverable** | | **July 2026** |

XV ET Symposium - Bologna 26/05/2025

Nadia.Tonello@bsc.es

# FIRST DELIVERABLE DELIVERED: ET-PP D8.1

- "Computing and Data Requirements"
  - https://apps.et-gw.eu/tds/?r=19444
  - Essentially a description of how we do things today

- Resubmitted after one iteration
  - Lack of predictions about LL requirements
  - Reasonable answers on scaling made thanks to the Blue Book having been released in the meanwhile
  - A good example of the difficulties, we have to plan for a moving target

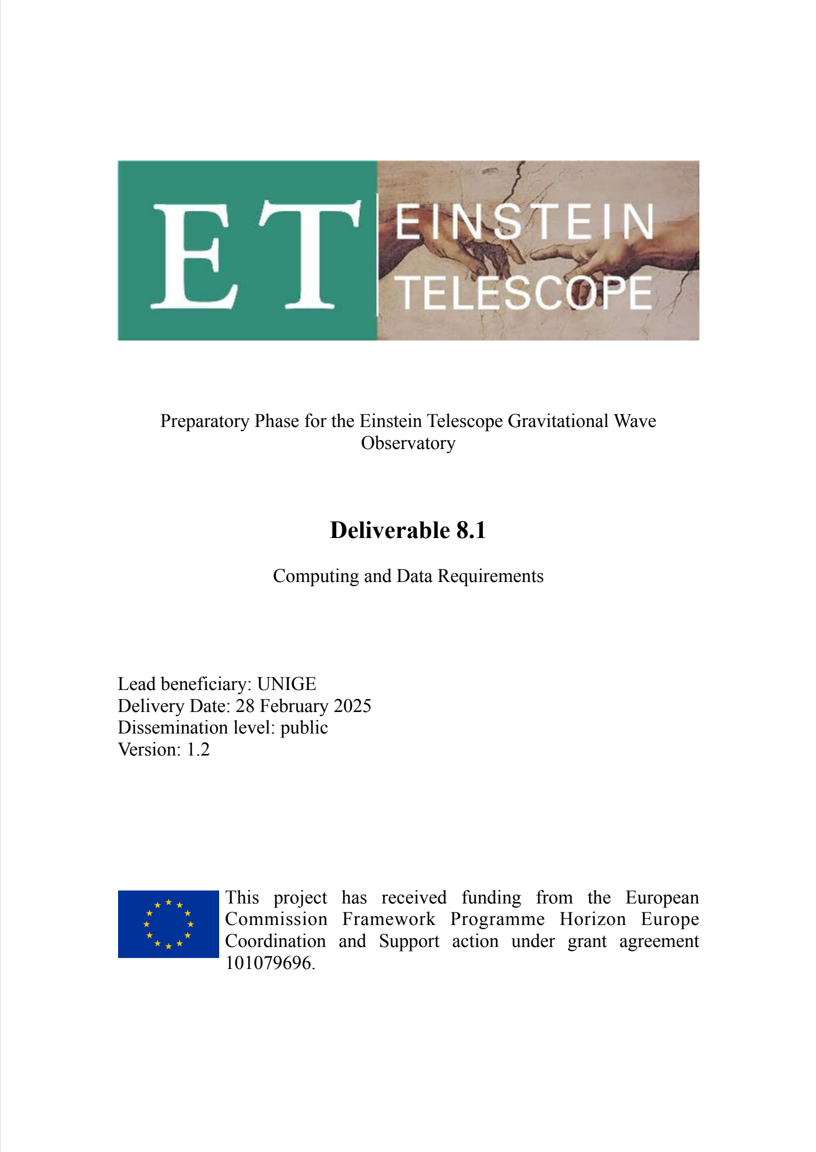
ET EINSTEIN TELESCOPE

Preparatory Phase for the Einstein Telescope Gravitational Wave Observatory

**Deliverable 8.1**

Computing and Data Requirements

Lead beneficiary: UNIGE
Delivery Date: 28 February 2025
Dissemination level: public
Version: 1.2

This project has received funding from the European Commission Framework Programme Horizon Europe Coordination and Support action under grant agreement 101079696.

# WP 8: Deliverables and milestones

EINSTEIN TELESCOPE

| Content | Type | Status | Date |
|---|---|---|---|
| M8.1 Workflows requirements collection and constraints. Workshop Uni Geneva | Milestone | Done - indico | Oct 2023 |
| D8.1 Computing and Data requirements submitted (UniGe). Reviewed and updated | Deliverable | Delivered - TDS | Feb 2024 |
| M8.2 Computing infrastructure availability for ET workflows characteristics. Workshop Napoli | Milestone | Done - indico | July 2024 |
| 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 |
| **D8.2 Computing and data model for ET** | **Deliverable** | | **Feb 2026** |
| M8.5 data management, data access policy and implementation | Milestone | | July 2026 |
| **D8.3 Data access policy implementation** | **Deliverable** | | **July 2026** |

XV ET Symposium - Bologna 26/05/2025

Nadia.Tonello@bsc.es

# THE COMPUTING MODEL

- The **overall architecture of the e-Infrastructure**, either as a single integrated system or as a few separate systems (e.g. instrument control and DAQ, low-latency, and offline)
- A **documented way of evaluating the required computing power** and storage space from the evolving scientific program of the collaboration
- Estimates of the **involved costs** and **growth timelines**
- A **description of the data flows**, with estimates for the needed network performances
- A description of the **User Experience and workflows** for relevant activities
- A description of the **tools** to be chosen or developed to provide all the required functionalities (foundation libraries, frameworks, middleware,...)
- Subsequent "Work Breakdown Structure" and "Implementation Plan" documents

# THE COMPUTING MODEL AS D8.1

- The **overall architecture of the e-Infrastructure**, either as a single integrated system or as a few separate systems (e.g. instrument control and DAQ, low-latency, and offline)
- A **documented way of evaluating the required computing power** and storage space from the evolving scientific program of the collaboration
- ~~Estimates of the **involved costs** and **growth timelines**~~
- A **description of the data flows**, with estimates for the needed network performances
- A description of the **User Experience and workflows** for relevant activities
- A description of ~~the **tools** to be chosen or developed to provide all~~ the required functionalities (foundation libraries, frameworks, middleware,…)
- ~~Subsequent "Work Breakdown Structure" and "Implementation Plan" documents~~

EINSTEIN TELESCOPE

# AS IT IS DONE TODAY

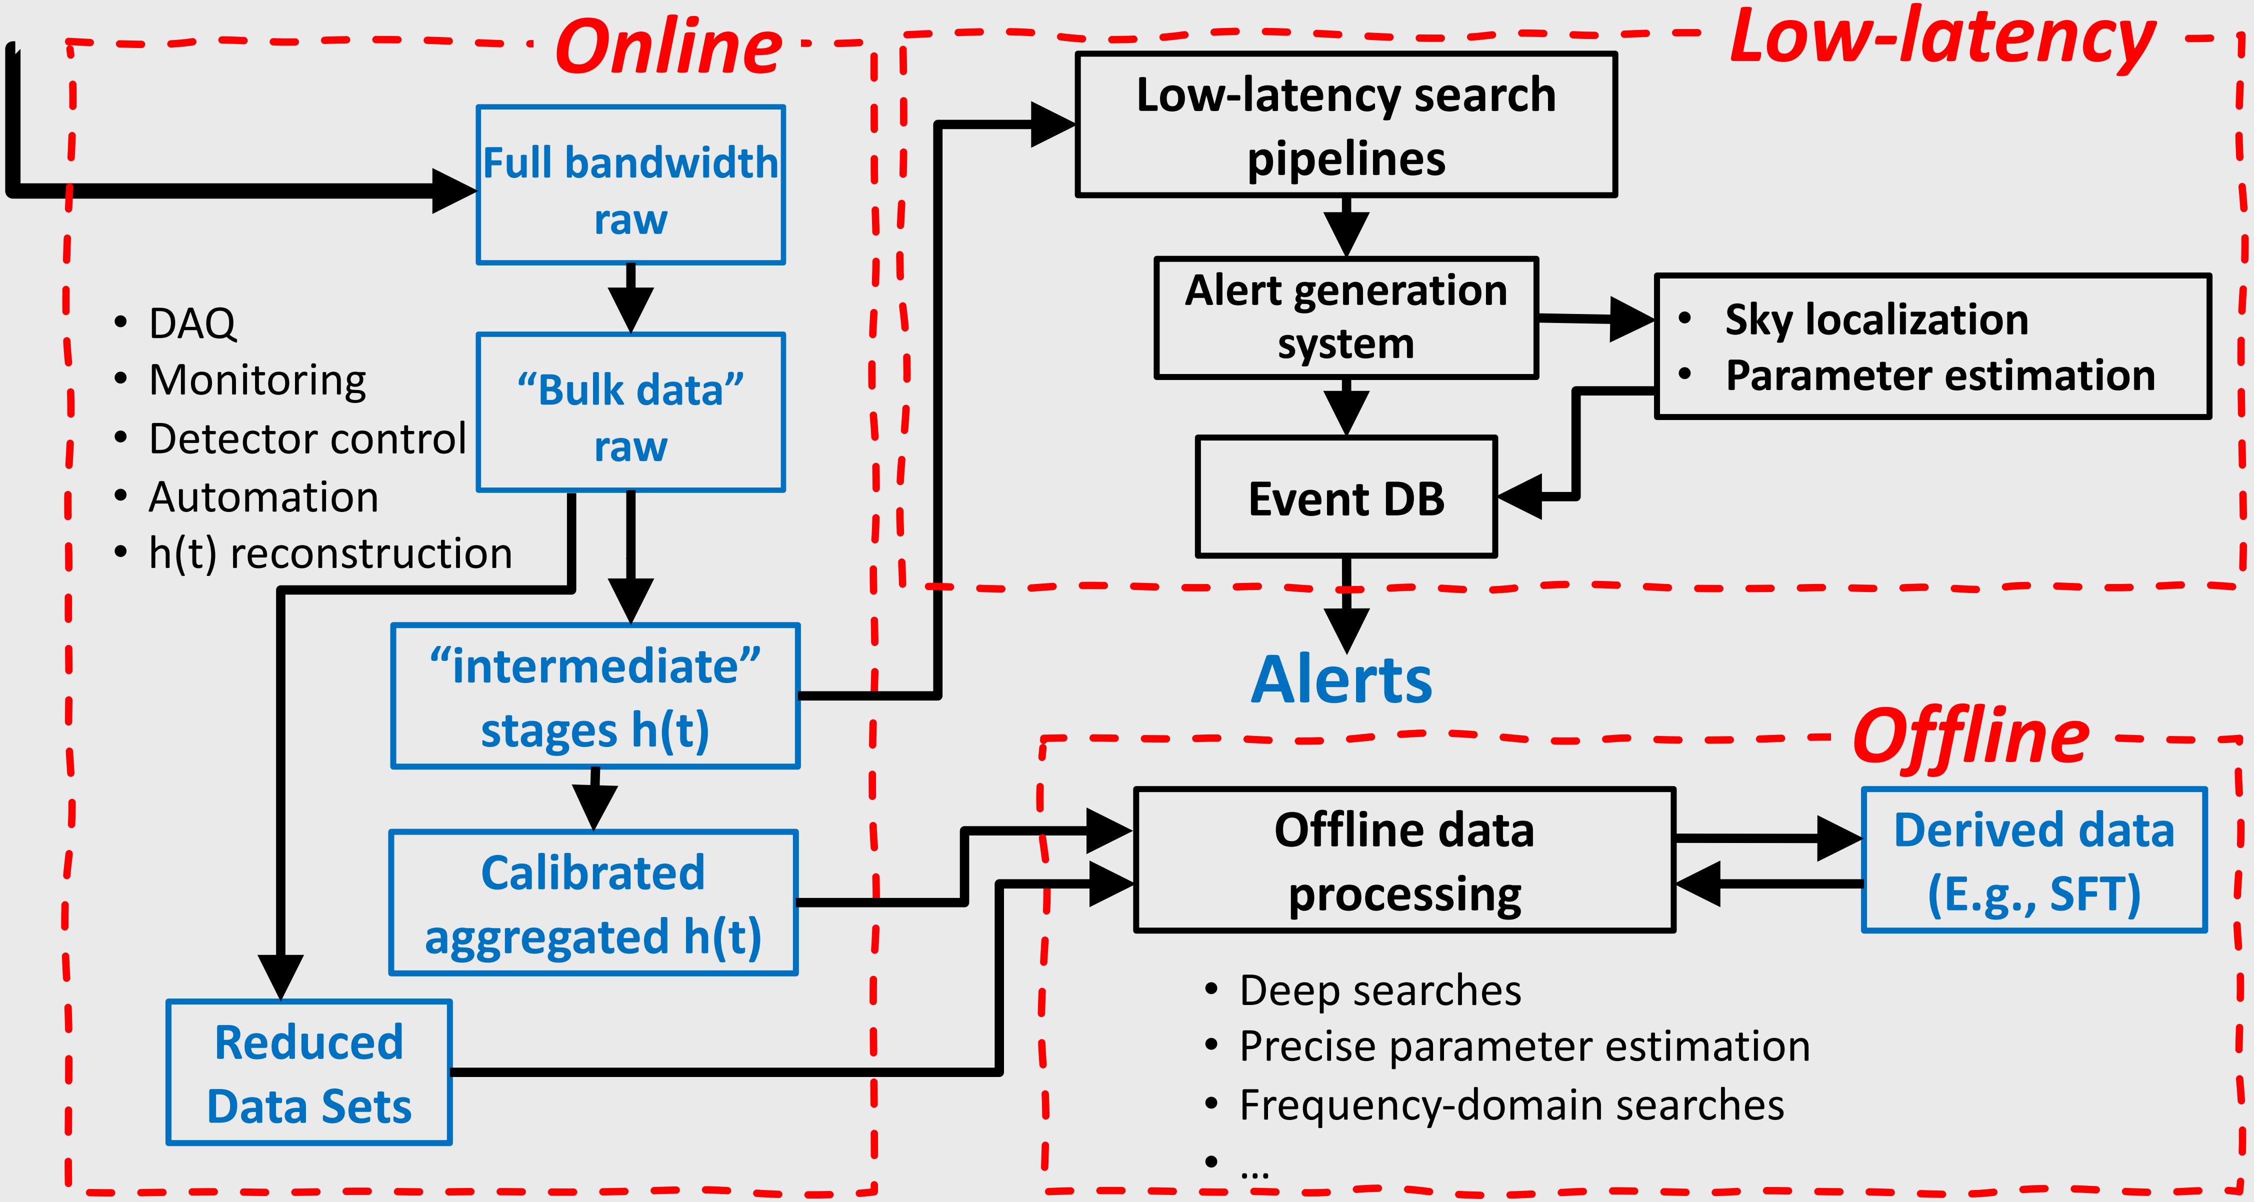


EINSTEIN TELESCOPE

- 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

# Toward an ET Computing Model

EINSTEIN TELESCOPE

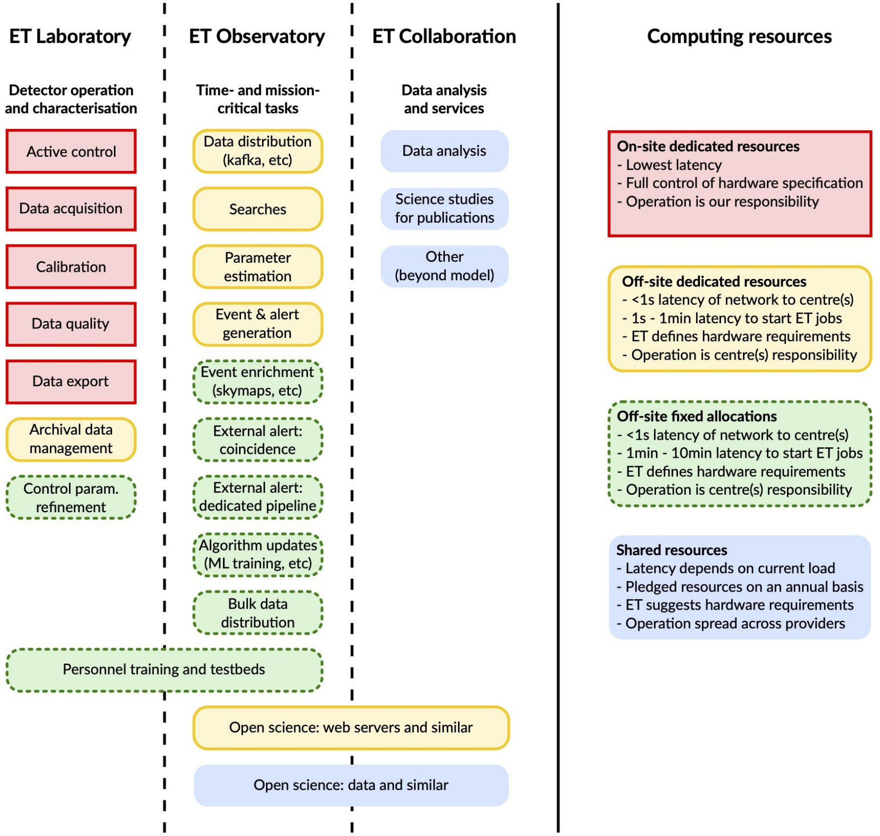


We can move forward and propose an ET Computing Model that performs this work

- Based on evolving the IGWN computing model used today
- Evolution includes making more use of experience and tools from the LHC

Computing estimates will evolve, we cannot provide a better estimate in the next 6 months

***However, as part of the Computing Model document, we will need to provide a plan of how ET aims to systematically demonstrate we can deliver the science***

UNIVERSITÉ DE GENÈVE
FACULTÉ DES SCIENCES
Département d'astronomie

Paul Laycock
13

GWSC GRAVITATIONAL WAVE SCIENCE CENTER

Paul.Laycock@unige.ch

EINSTEIN TELESCOPE

# The scale of ET computing

EINSTEIN TELESCOPE

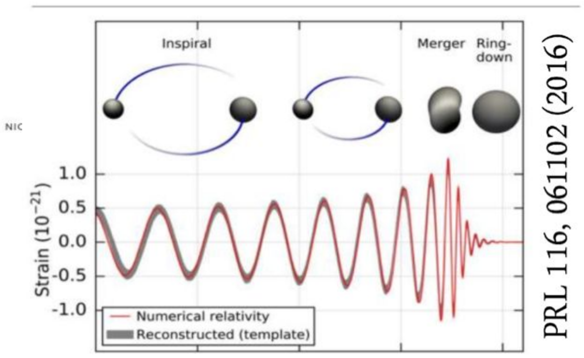


- Increased signal sensitivity for ET means signals are also in-band for much longer (minutes, hours, days)
  - Signal “pile-up” complicates things
  - ***Naively, ET would need 40M cores just for low latency***
  - Naive here means assuming we only need to double the compute to handle pre-merger analysis

**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

EINSTEIN TELESCOPE

 

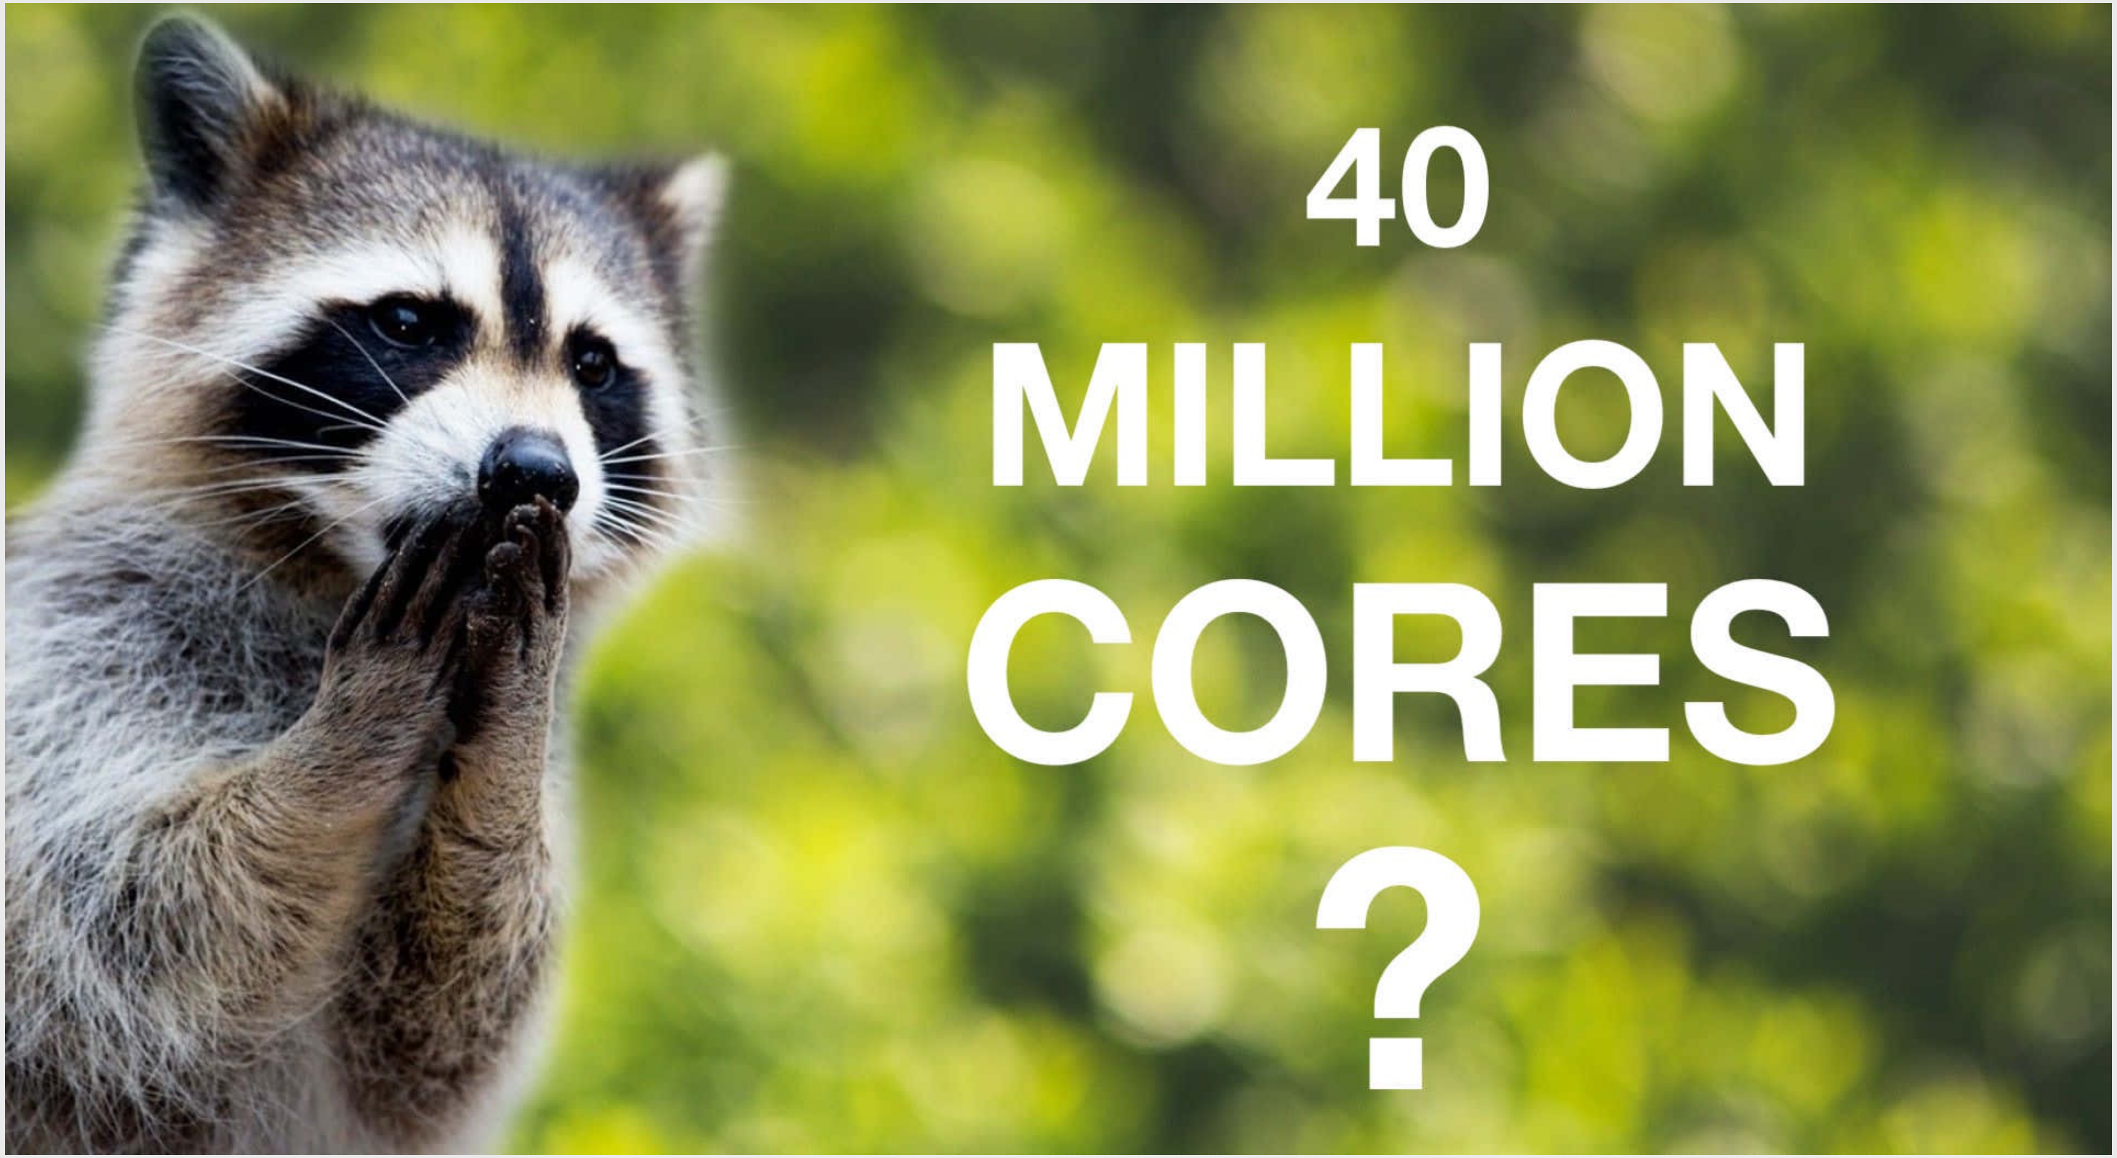


Sara.Vallero@to.infn.it

# 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.

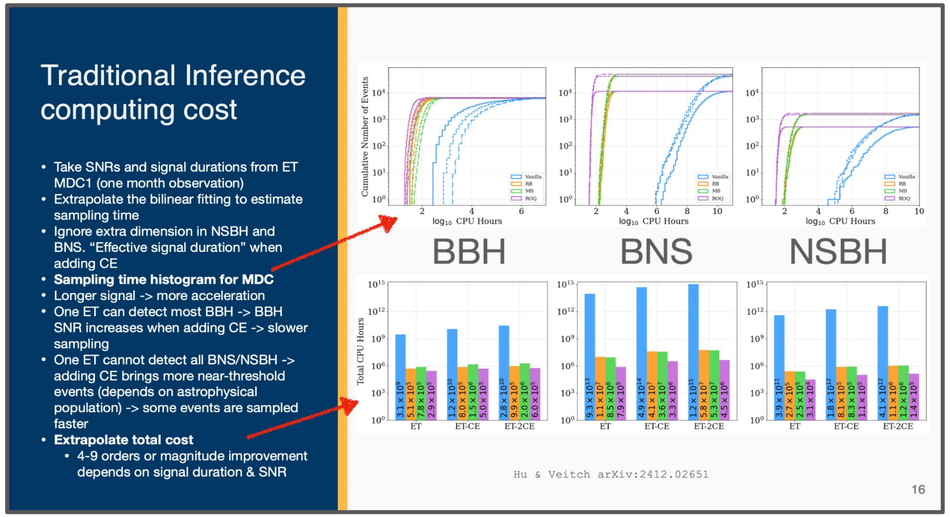


arXiv > gr-qc > arXiv:2412.02651

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

# 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
  - See also Steven’s and Paul’s presentation tomorrow

EINSTEIN TELESCOPE

# NEXT UP

ET-PP WP8 Workshop: On sit

indico.ego-gw.it/event/820/overview

MNTN | Amm.ne INFN | MonALISA | OpenNebula Sunst... | OCCAM | TOCLOUD Jira | GW Computing | Biblio | CloudTO | PNRR | All Bookmarks

Public | Europe/Paris | S. Bagnasco

## ET-PP WP8 Workshop: On site infrastructure, computing and data model

EINSTEIN TELESCOPE

30 June 2025 to 1 July 2025
Lyon
Europe/Paris timezone

- Overview
- Timetable
- Registration
- Participant List
- Link to CC-IN2P3

**Contact**

patrice.verdier@in2p3.fr
+33 6 13 49 57 51

CC-IN2P3 will host the workshop dedicated to "On site infrastructure, computing and data model" for ET, as part of WP8 of the Einstein Telescope Preparatory Phase Project.

Local organizing team: Eric Fede, Patrice Verdier

# </rant>

# DEFINING A STANDARD BENCHMARK

## Current Status and Future Work

- Working with the HEPiX development team on integration of RIFT to the HEPScore
- Developed a draft of the workload:
  - Created a Dockerfile.append that sets up dependencies, clones the RIFT, and generates synthetic data – to be validated
  - Updated the script for running a single job – works only with one copy
  - Score calculation is implemented but needs to be reviewed
- Next steps:
  - To verify what other metrics, logs, or results need to be collected
  - Whole workload needs to be rewritten using a different simpler template
  - Input data, parameters, and constants have to be updated so that they make sense for ET

EINSTEIN TELESCOPE

skorobogatovgeorgy@icc.ub.edu

# MDC1

- Data produced by OSB-div 10 with a revised version of a code developped in 2012
    - ET-0148A-25: *A mock data challenge for next generation detecto*
- Data contains colored Gaussian noise+ GW signal fr … gle.
- Challenges include detection and parame … ce of a large number of long and overlapping s…

See Elena's plenary presentation tomorrow

- **In fieri**:
    - *EIB is rewri … re waiting for access to the code after which we w…*
        - *Div10 cha … a checklist to test the new code.*
        - *A review team will be in charge of comparing the outputs of the two codes.*

Gianluca.Guidi@uniurb.it

# ET-IAM DEPLOYMENT AT CNAF

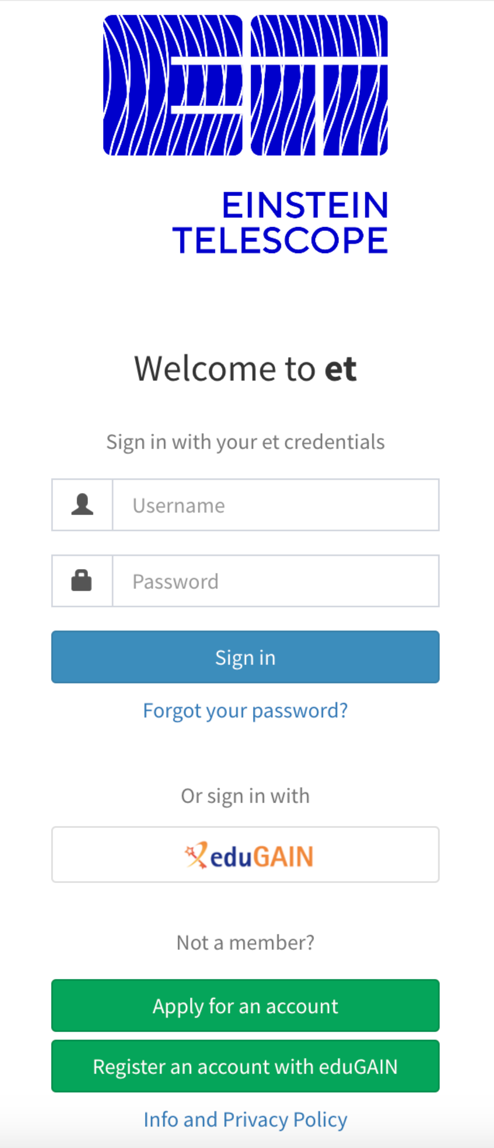


- Installed, configured and working!
  - https://et-iam.cnaf.infn.it
  - Thanks to Michel Jouvin, after specific MoU with CNAF
  - Plan available to develop interface to ETMD
- For now, local temporary accounts for testing
  - Most functionalities for testing downstream applications available
  - Already being used (WIP) by ETAP & MADDEN projects

# ETMD DEVELOPMENTS - Docs


ET Docs

Document list

ET Wiki available here | + Add a new document

Displaying 1-51 of **2565** documents

| Code | Access | Title | Date | Author(s) | Review status |
|---|---|---|---|---|---|
| ET-0197A-25 | ET | Presentation in ISB-SUSP: ET-LF TM Suspension Conceptual Design Update of 12m suspension - Concept B | 2025-05-23 | R. J. Cornelissen, N.A. Holland, J. van Dongen, P. Saffarieh, M. Valentini, H. Verkooijen, A. L. Mitchell, S. Sijtsma, C. M. Mow-Lowry, J. V. van Heijningen, A. Bertolini | |
| ET-0196A-25 | ET | Photoinduced Effects in Mechanical Loss and Elasticity of GaAs for Ultra-Stable Laser Interferometers | 2025-05-23 | Nico Wagner, Michael Zimmer, Michael Jetter, Stefanie Kroker | |
| ET-0195A-25 | ET | Innovative On-Site Continuous Production and Welding of the Einstein Telescope Vacuum Pipes File(s) not yet attached | 2025-05-23 | Charlotte Benning, Robert Joppe, Oliver Pooth, Achim Stahl | |
| ET-0194A-25 | ET LKV | Characterisation of Crystalline Silicon Fibres: A material study for the ET-LF payload File(s) not yet attached | 2025-05-23 | Ardiana Nela [a.nela.1@research.gla.ac.uk], Andrew Spencer, Alan Cumming, Flavio Travasso, Karl Toland, Giles Hammond, Gregoire Lacaille, Karen Haughian, Victoria Graham, James Hough, Russel Jones, Sheila Rowan, Iryna Buchovska, Robert Menzel, Charlie Patterson, Claire Wilson, Sarah Dugmore | |
| ET-0193A-25 | ET | First eccentric inspiral-merger-ringdown analysis of neutron star-black hole mergers | 2025-05-23 | Lluc Planas, Sascha Husa, Antoni Ramos-Buades, Jorge Valencia | |
| ET-0192A-25 | ET ET-CB ET-EB ET-INFRA-DEV ET-SAR ET-SC PUBLIC LKV | AEI 10m Prototype - Recent developments towards the Sub-SQL Interferometer | 2025-05-23 | Sara Al-Kershi | |
| ET-0191A-25 | ET | Impact of the (l=2,m=0) spherical harmonic mode with memory on parameter estimation for ground-based detectors | 2025-05-23 | Maria Rosselló-Sastre, Sascha Husa, Sayantani Bera, Yumeng Xu | |
| ET-0190A-25 | ET | Trimble Connect Quick Guideline for ET Detector Layout(s) File(s) not yet attached | 2025-05-23 | J. Bratanata, W. Wahbeh, M. Marsella | |
| ET-0189A-25 | ET VIRGO LKV | Einstein-Podolsky-Rosen conditional squeezing for future GW detectors File(s) not yet attached | 2025-05-22 | Francesco De Marco on behalf of the EPR working group | |

- Work in progress
- Provides list of members for Document Review in Docs
- Working with the Editorial Committee, in particular with Paola Leaci (ET-Roma1)



Gary.Hemming@ego-gw.it

## Digital infrastructure

Develop new computational algorithms and infrastructure designed to process the hundreds of thousands to millions of black holes and neutron star mergers each year observed by the next-generation GW observatories.

- Computational infrastructure and algorithms
  - **Waveform generation**, led by the University of Birmingham
  - **Real time searches**, led by the University of Portsmouth
  - **Signal and population inference**, led by the University of Glasgow
- **Prototype event database**, led by Cardiff University

RaymondV@cardiff.ac.uk

# Next-Gen event database

- 3G event rate 1000x higher than 2G
- High-SNR negative-latency detection allows early warning alerts
- Challenge to provide a robust, scalable, event database and alerting infrastructure to serve the GW observatory network and EM partners
- Cardiff leading UK effort to scope requirements and conceptual design
  - Requirements link
- Needs detailed understanding of access patterns
- Introductory study of prospective database technologies and data placement algorithms
  - Link

NGDB Software Requirements Specification

Development version 6dbb5bac

2024-12-13

## Contents



RaymondV@cardiff.ac.uk

# General overview

Karlsruhe Institute of Technology

## mdcpp

- C++ library + bindings to Python
- standard application to simulate all implemented components

**Example:** simulate noise and blips [for a build using libtorch, job 0] with standard application `detsim` for a ET-triangle-like dummy network:

```
./bin/detsim -j 0                                         \
    --add_ET_dummy --gen_gcnoise --gen_blipgl             \
    --output_dir ./validation/output                      \
    --input_dir ./cfg                                     \
    --start_time 1000000000 --segment_duration 2048 \
    --output_format gwf
```



Steffen.Hahn@kit.edu

# 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)

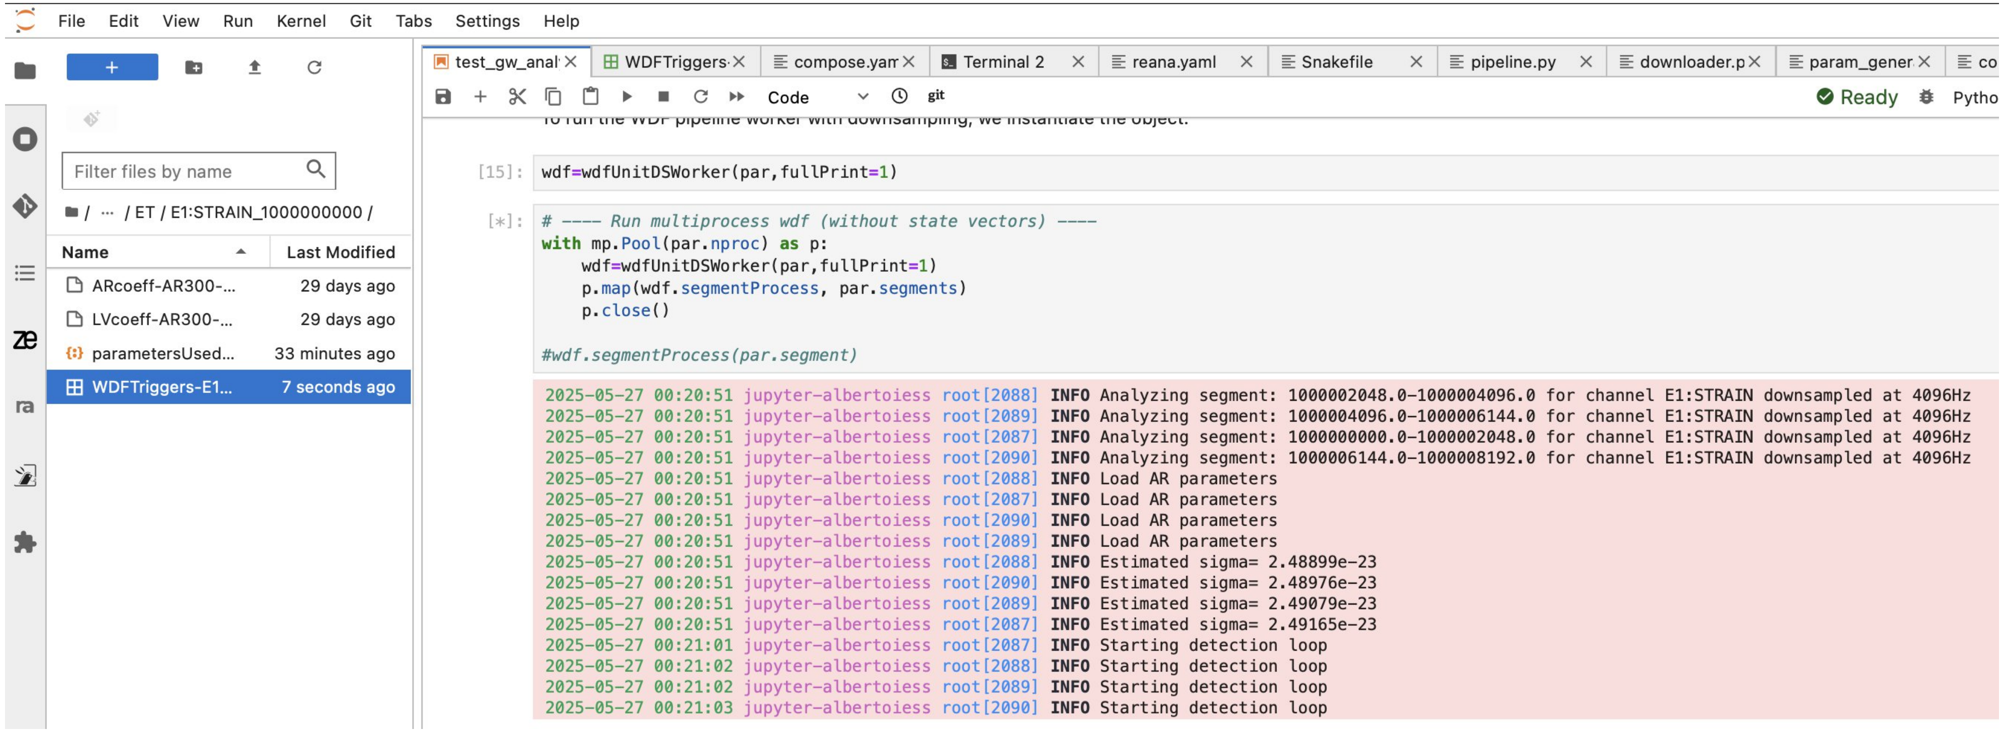


Alberto.Iess@lapp.in2p3.fr

# MADDEN: DATA MANAGEMENT AND DISTRIBUTION

## Proposed Setup

OSCARS

- The **Multi-RI data lake** for ET and CE is the innovative development
  - **OpenID Connect (OIDC)** for authz
  - CE Rucio to trust ET IAM as an IdP
  - allow Rucio client to connect to more than one Rucio server
- POSIX like view of the data with **RucioFS**

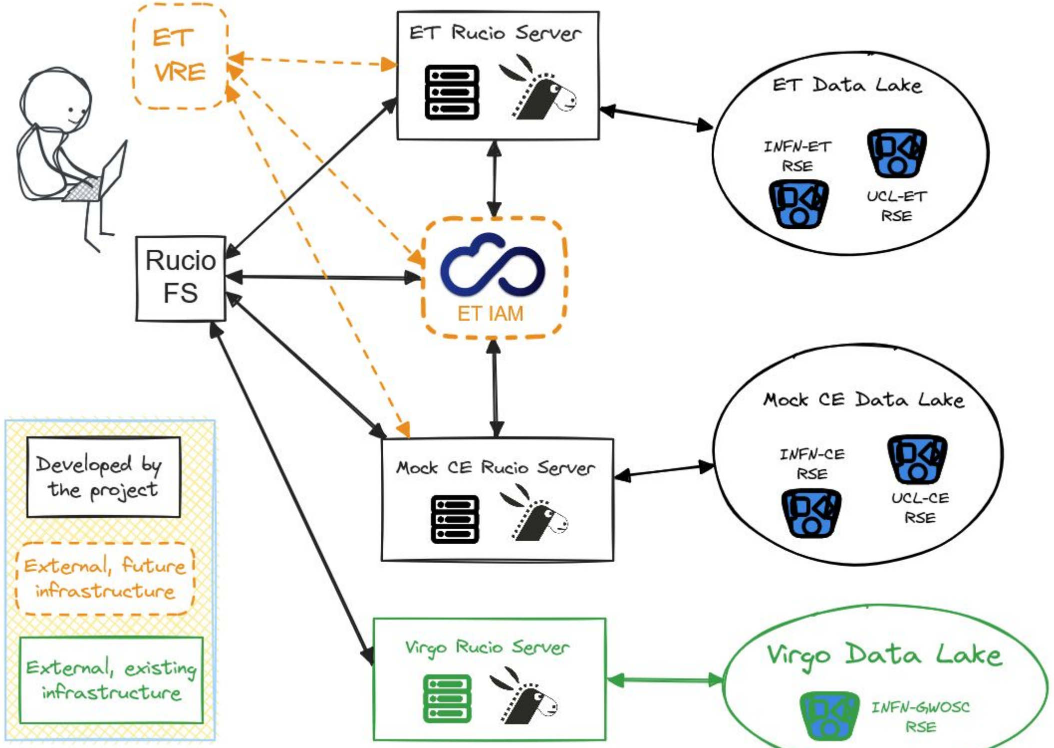


ENVRI community | LIFE SCIENCE RI | ESCAPE | panosc photon and neutron open science cloud | SSHOpenCluster Social Sciences & Humanities Resources

4

Federica.Legger@to.infn.it

## Proposed Setup

OSCARS

At UniGe, deploy on k8s:

- Jupyterhub
- REANA
- HSF Metadata DB
- Monitoring and accounting
- Deliver [integrated helm chart]

Connect to:

- Indigo IAM@CNAF
- MADDEN@Torino

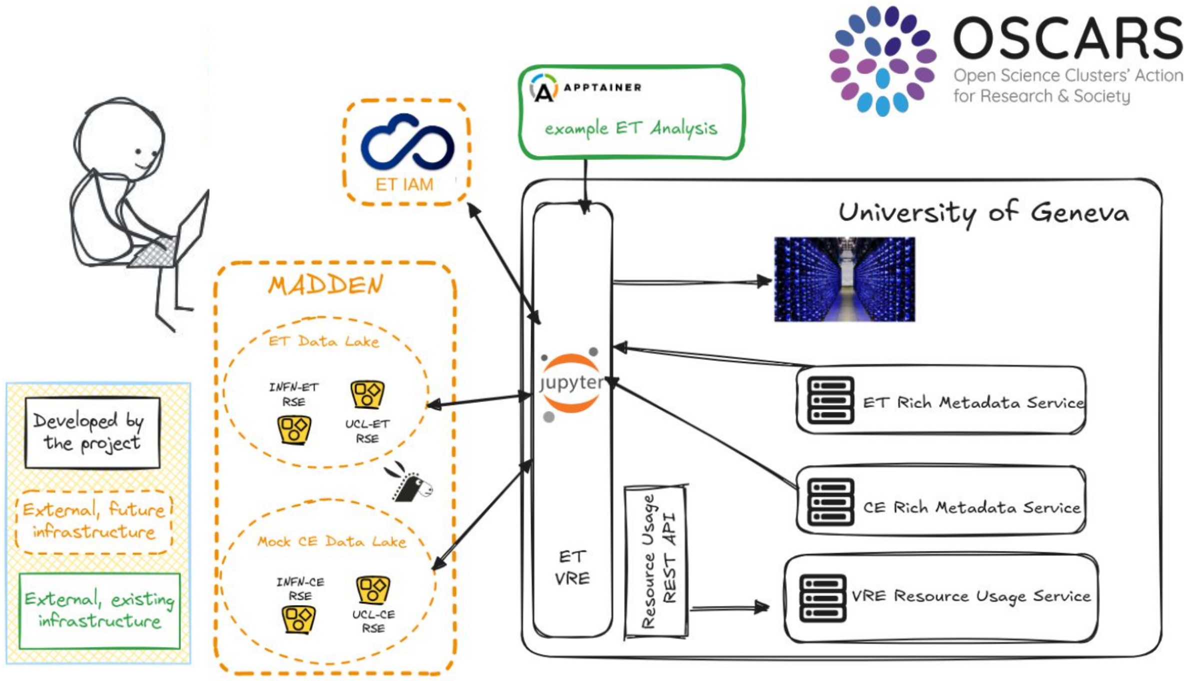


ENVRI community · LIFE SCIENCE RI · ESCAPE · panosc photon and neutron open science cloud · SSHOpenCluster Social Sciences & Humanities Resources

4

Paul.Laycock@unige.ch

# Inauguration on April 14th, 2025


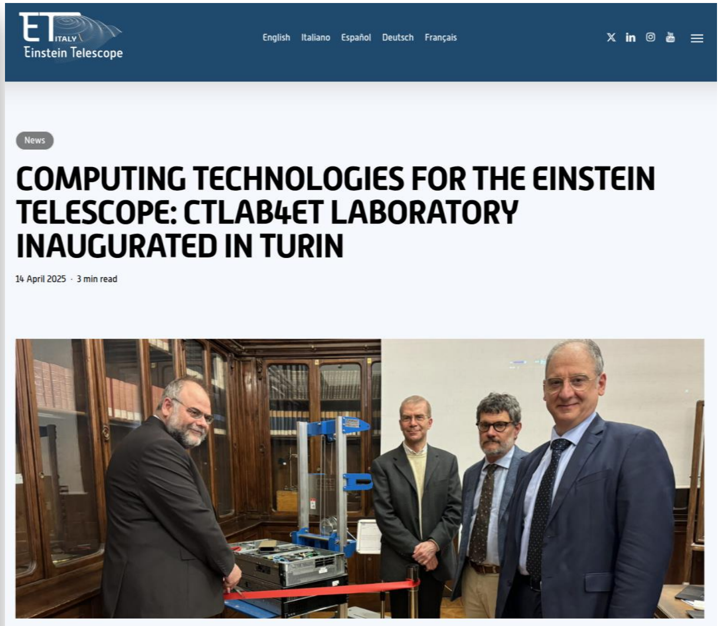


from www.einstein-telescope.it


Home / NEWS / Computing Technologies for the Einstein Telescope: CTLAB4ET Laboratory inaugurated in Turin

NEWS

COMPUTING TECHNOLOGIES FOR THE EINSTEIN TELESCOPE: CTLAB4ET LABORATORY INAUGURATED IN TURIN

14 April 2025

from www.infn.it

Lia.Lavezzi@to.infn.it

## Next...

**K8s cluster TO DO list:**

- ✅ Three control planes in High Availability
- ✅ /home directory on GlusterFS distributed filesystem
- ✅ Dedicated storage for the data
- ✅ Deployment of the nodes
- ✅ Authentication via ET-IAM
- ✅ Monitoring

**Desiderata**

Installation of Jupyter Notebook

Installation of Rucio client

Installation of VRE

...

**Possibility**

Possible integration with INFN DataCloud

**From Bologna WS last February**

What additional tool do we want/need?

- Only interactive analysis or also pipelines?
- Dask https://www.dask.org/
- VRE, from OSCARS project ETAP
- REANA https://reanahub.io/
- Snakemake https://snakemake.github.io/
- Spark https://spark.apache.org/
- Rucio https://rucio.cern.ch/
- ML tools

Lia.Lavezzi@to.infn.it

# Experimental Standards

## Text and Semantic MOCs

Encoding sky regions and textual piece of information for simultaneously semantic and sky space operation – including multimodal Generative AI.

Basic JSON structure of a Space MOC

```
{
  "order_n": [npix_i | where npix_i uniquely
     defined integers],
  // ... continue for other orders ...
}
```

Adding new entries in the JSON MOC serializations

```
"text": "Your textual
description here",
“multimedia” : “URL”
```

7

Giuseppe.Greco@pg.infn.it

# SUMMARY

**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!)
  - Ando 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

# THANKS!

- Questions?