ET Symposium EIB session

Div3 - Computing model and computing resources

G. Merino

Maastricht, 7 May 2024

D8.1 ET-PP Computing and Data Requirements

Document generated by ET-PP with input from the Geneva workshop Oct'23 and a lot of work from Paul (thanks!) - Our baseline reference for computing & data requirements.

- Greatest challenge will be training and retaining sw & computing expertise.
- Modest computing challenge compared to HL-LHC.
  - In terms of resources, computing more challenging than data.

The role of MDCs as a tool to make measurements and extrapolations that we will use to build the computing model is clearly referenced.

> *"The computing requirements summarised here are based on the information available at the time of writing. These will evolve over time and **will be measured using a rolling program of MDCs** that are then curated by the ET collaboration."*

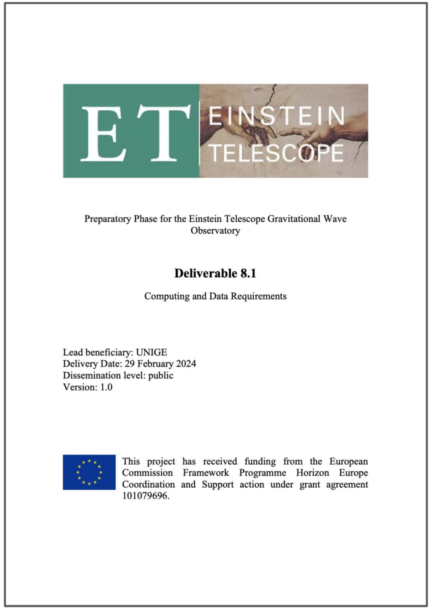
ET EINSTEIN TELESCOPE

Preparatory Phase for the Einstein Telescope Gravitational Wave Observatory

Deliverable 8.1

Computing and Data Requirements

Lead beneficiary: UNIGE
Delivery Date: 29 February 2024
Dissemination level: public
Version: 1.0

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

D8.1 - ET Computing Domains

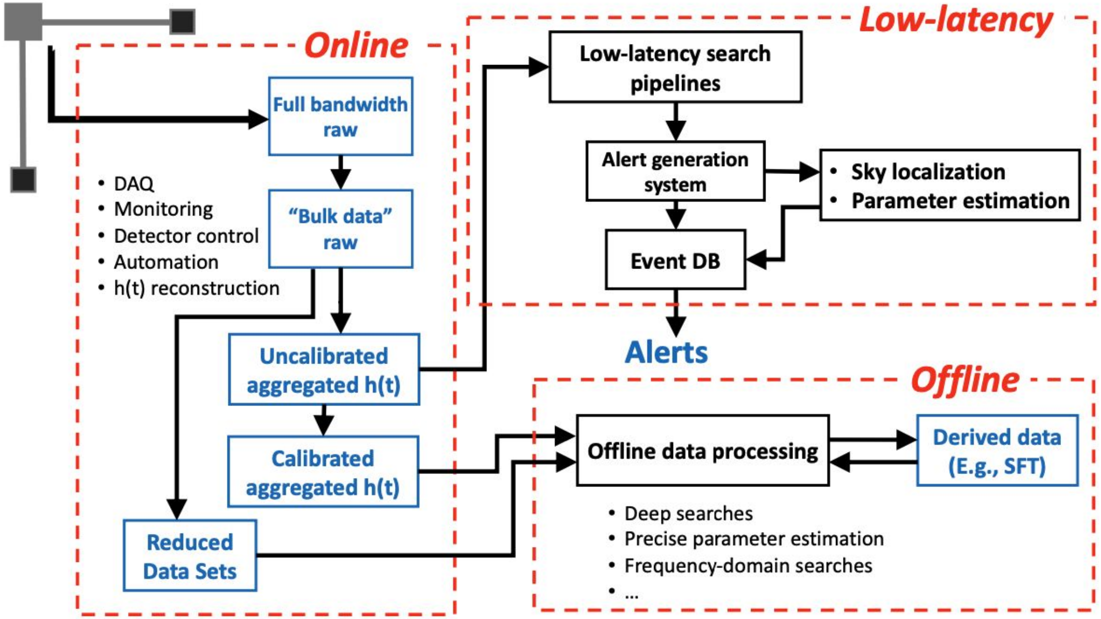

Fig. 5. Schematic of the ET Computing domains (red) and data types (blue), slightly simplified for clarity.

D8.1 - Offline computing requirements

ET will generate 100.000 evts/yr, BUT only considering "golden events" for the dominant part of computing needs: 1000 evts/year

⇒ x10 more CPU than O3 for CBC PE

CPU and storage requirements are scaled as described in Sec. 3.1; other parameters are baseline O3 values.

Custodial storage	2 * 10 PB/ ET interferometer / year of observation
CPU	5×10^9 HS06 hours/year
RAM (GB)	2GB/core
Throughput to WNs	25 kbps per core

Table 9: Baseline offline computing requirements. Similarly to what happens for low latency, GPU acceleration is currently used by a growing number of pipelines and we expect this to increase in ET.

D8.1 - Offline computing requirements

ET will generate 100.000 evts/yr, BUT only considering “golden events” for the dominant part of computing needs: 1000 evts/year

⇒ x10 more CPU than O3 for CBC PE

CPU and storage requirements are scaled as described in Sec. 3.1; other parameters are baseline O3 values.

Custodial storage	2 * 10 PB/ ET interferometer / year of observation
CPU	5×10^9 HS06 hours/year
RAM (GB)	2GB/core
Throughput to WNs	25 kbps per core

~50k cores, flat

Table 9: Baseline offline computing requirements. Similarly to what happens for low latency, GPU acceleration is currently used by a growing number of pipelines and we expect this to increase in ET.

D8.1 - Low latency computing requirements

The following baseline requirements reflect the situation during the LIGO-Virgo O4 observation period, and are discussed in Sec 3.1:

Caching storage	Negligible
CPU	15000 cores
RAM (GB)	2GB/core
Latency	<10s

Table 8: Baseline low latency computing requirements. In addition to the CPU, a diverse set of GPUs are also used by LIGO and Virgo to accelerate computations. The usage of GPUs is expected to increase and will be particularly important for low latency computing.

	Pipelines	**CPU cores**	**GPUs**
Caltech	GstLAL, PyCBC Live, SPIIR oLIB RAVEN, LLAMA Omega	14900	438
EGO	MBTA, cWB Omicron	1000	-

Table 5: Resources used by the 2G network for low latency data analysis during O4.

D8.1 - online computing requirements

	Minimal scenario	Operational safety margin
Operations storage buffer (TB)	800	1600
Long-term storage (PB)	20	40
CPU cores	2150	6450
RAM (GB)	3780	11340
Network	100 Gb/s	2 * 100 Gb/s

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

PE computing estimation needs

"We need to find out how these numbers multiply together to find/define the ceiling for PE" (J. Veitch, Geneva workshop,Oct 2023)

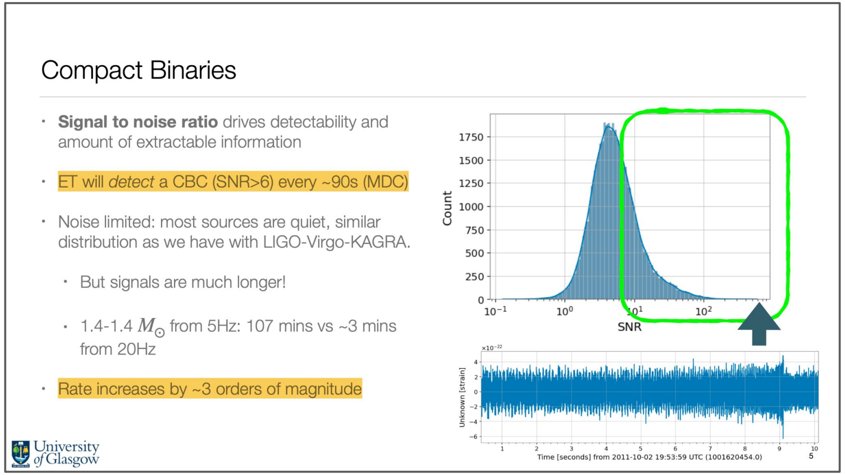

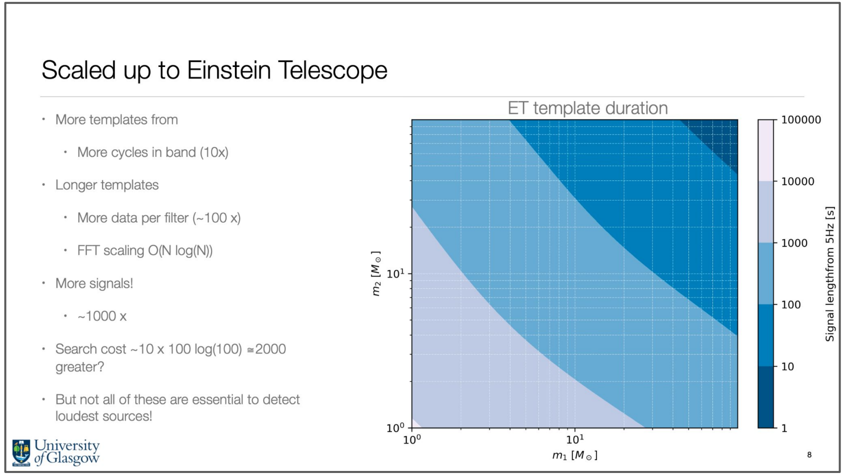

We will need to keep an eye on this $5x10^9$ HS06hours ceiling, to detect if major deviations appear as new analyses activities unfold.

GPUhours needs continues to be the "elephant in the room" …

HEPIX Benchmarking Working Group

Focused work the last years to develop HEPScore benchmark and HEP Benchmark Suite.

- Modular mechanism to add realistic workload based benchmarks
- Framework to automate setup, execution and results data collection

If **GPU workloads** are going to play a key role in GW pipelines, having a better understanding on how to measure & compare performance seems a needed building block.

- Is this a useful development? (trying to check with the HEPIX Benchmark WG chairs)

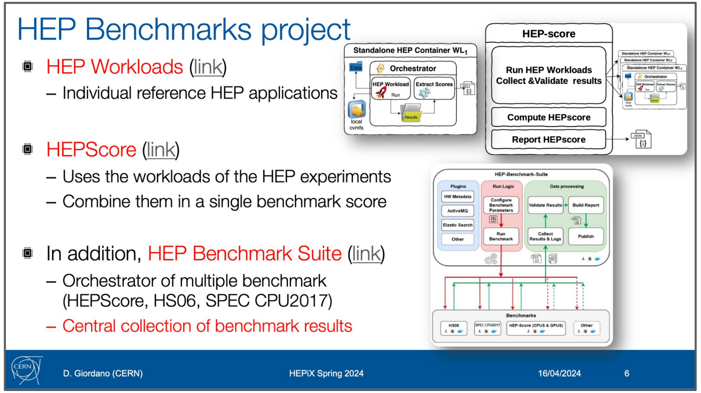

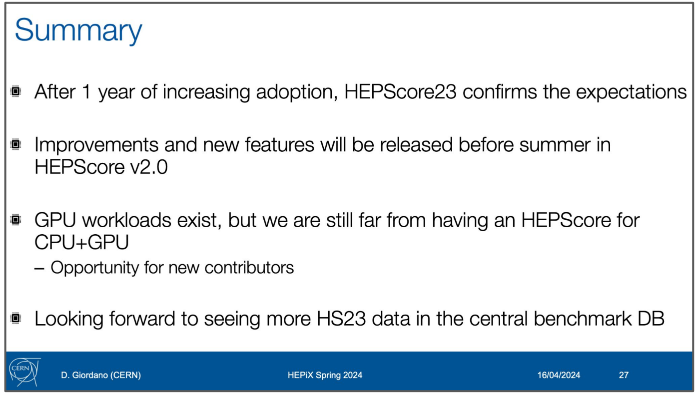

slides from D. Giordano, HEPIX Paris Apr 2024

Contribution: Georgy Skorobogatov from ICCUB

Georgy expressed interest in contributing to Div3 last year (thanks!)

Worked on setting up and executing a pygwb workflow on ET MDC data using the ESCAPE VRE/Reana infrastructure

- Aim: Set up ET MDC runtime environment that can be used to extract computing resources information.

Code in EIB/Div1 gitlab

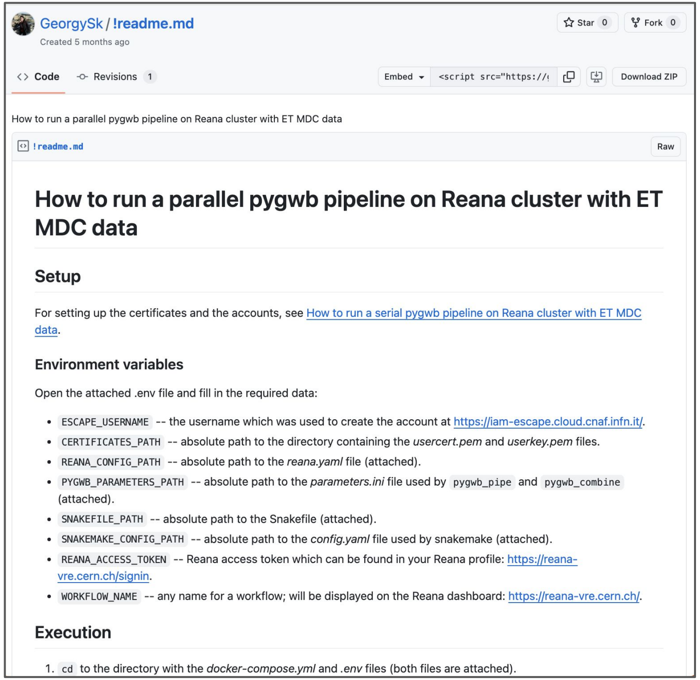

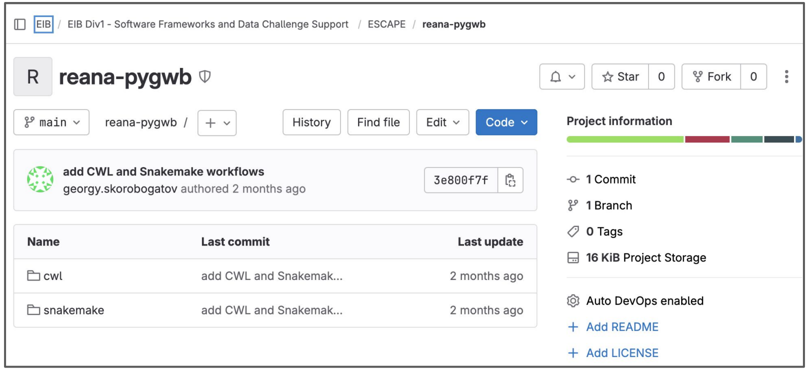

Need for HPC resources

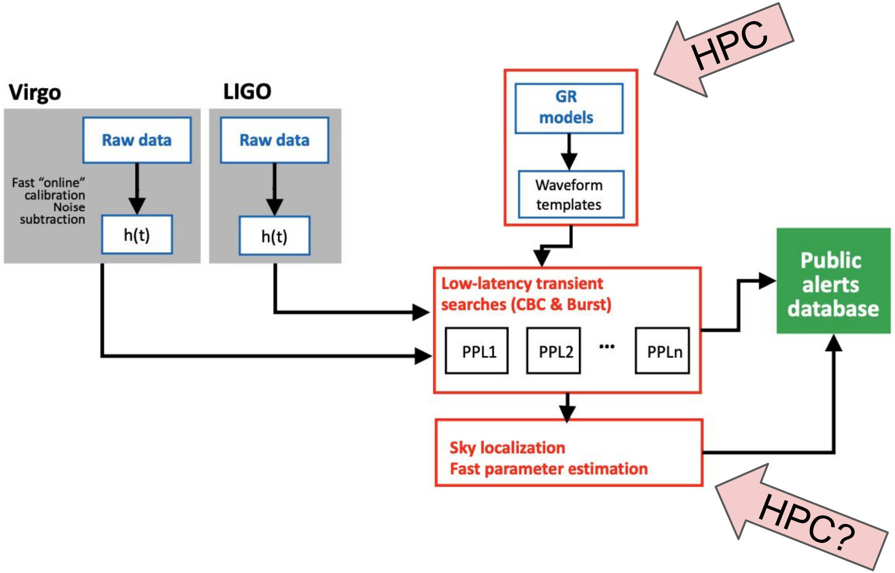

Fig. 3: Simplified conceptual representation of 2G low latency data analysis workflows.

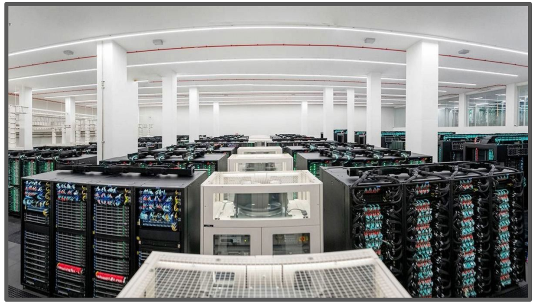

From Spain's HPC network public accounting → one single researcher in GR models gets consistently around 30Mhrs/yr.

Can we gather EU HPC accounting data?

Should we invest effort in "lobbying" for getting some guaranteed allocation at EuroHPC for ET(LISA?)?

How relevant is this within the overall ET computing model?

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

JENAA
Joint ECFA-NuPECC-APPEC Activities

JENA HPC Working Group - scope

The overall goal is to try and have a coordinated voice from the three JENA communities towards EuroHPC, the organisation that plans, runs and manages the funding for the large HPC machines in Europe.

Concrete goals:

1. Try to get some **"priority/strategic" long-term allocation** in EuroHPC so that ENA experiments could access a number of CPU/GPU hours/year without the need to submit proposals quarterly.
2. Have a voice in the **planning process for the large HPC in Europe**, both at the design level (e.g. ask for more or less CPU vs GPU or certain network requirements) as well as the operations level (e.g. ask for consistent backfill mechanisms in all the EuroHPC machines so that idle cpu-hours could be used by opportunistic workloads).

SPECTRUM EU HE Project (Jan'24 → Jun'26)

SPECTRUM Computing Strategy for Data-intensive Science Infrastructures in Europe

- Joint effort: RIs (HEP+Radio astronomy) + e-infrastructure providers
- Goal: deliver Strategic Research, Innovation, and Deployment Agenda (SRIDA) along with a Technical Blueprint to facilitate the creation of an exabyte-scale **research data federation and compute continuum**, fostering **data-intensive scientific** collaborations across Europe.

The SPECTRUM Community of Practice (CoP) aims to create a community of scientists and infrastructure managers with mutual understanding of future needs, challenges and possible solutions.

- WG1 Data Management and Access
- WG2 Workflow management and organization
- WG3 Compute Environment
- WG4 SW tools
- WG5 Scientific Use cases
- WG6 Facilities

SPECTRUM CoP

What about the Cloud?

the peaks could be conferences, or CW searches after the run, or ...

Grid computing was born before the Cloud.

Most experiments computing models have evolved on a “flat provisioned resources” environment.

- Very good at keeping any available resource busy.

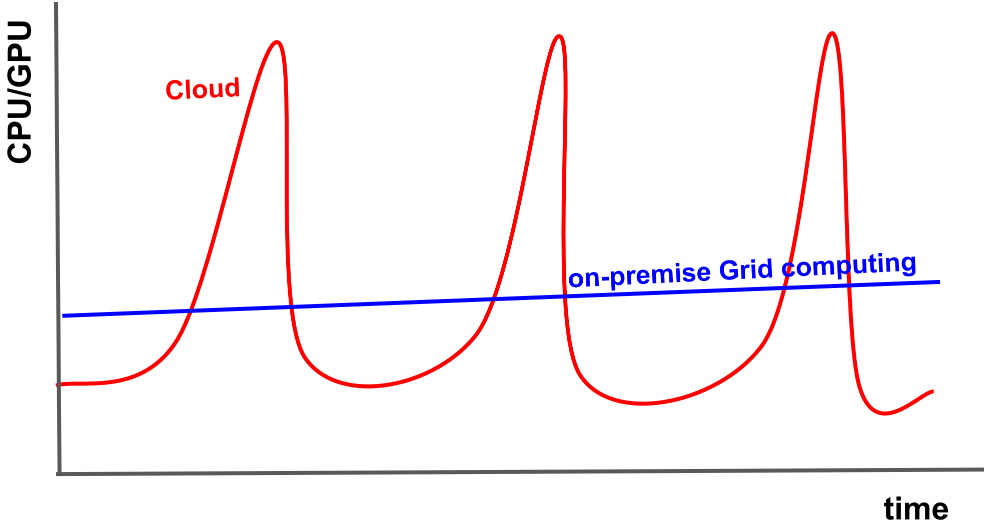

On-premise Grid has clear advantages as a multi-national federated infrastructure. Mostly, it provides a simple mechanism for countries to contribute.

Should we try to explore the Cloud model advantages for ET, if any?

- *which one of the two above is cheaper?*
- *which one minimizes the time to results?*
- *which one consumes less energy? ...*

thank you