The ET Computing Model and next steps

Paul Laycock (Université de Genève) for the EIB + WP8

ET-PP Final Meeting - Tuesday 3rd November 2026

UNIVERSITÉ DE GENÈVE
FACULTÉ DES SCIENCES
Département d'astronomie

GWSC GRAVITATIONAL WAVE SCIENCE CENTER

The ET computing model has been delivered

EINSTEIN TELESCOPE

Contents

I Introduction 10

1 Scope 11
1.1 Executive Summary . . . 13

2 ET Science Goals 15

3 Gravitational Wave Detection 20
3.1 Basic Principles of GW Detectors . . . 20
3.1.1 Interferometers . . . 20
3.1.2 Data output . . . 22
3.2 Analysis of Gravitational Wave Signals . . . 22
3.2.1 Gravitational wave events . . . 22
3.2.2 Analysis methods . . . 24

4 Overview of ET Computing Use Cases 26
4.1 Computing and Data Resource Estimates . . . 26
4.2 ET Collaboration Structure . . . 28
4.3 Data Format . . . 29
4.3.1 GW detector data format . . . 29
4.4 Detector Operations . . . 30
4.5 Rapid Gravitational Wave Alerts . . . 30
4.5.1 GW event/alert format . . . 31
4.6 Offline Workflows . . . 32

5 Computing Model Design Principles 34
5.1 FAIR Principles . . . 35
5.1.1 A FAIR implementation for the Einstein Telescope . . . 36
5.1.2 Beyond FAIR . . . 37
5.2 Software Management . . . 37
5.2.1 Software best practices and metrics for success . . . 38
5.2.2 Organisational principles . . . 38
5.2.3 Critical engineering practices . . . 39
5.2.4 Metrics to track (DORA) . . . 39
5.2.5 Release management . . . 40
5.3 Software Framework . . . 41

- ***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 **WLCG/CERN/HSF/HEP**

  - 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; ET’s site and geometry have not been decided

What we know about computing in 2040

EINSTEIN TELESCOPE

Thinking…

Combobulating…

Clauding…

Frolicking…

Aborting…

🤷

Two years ago, I knew what coding and analysis would be like in 2030…

(what a horrible time to write a computing TDR)



Computing Challenges for Einstein Telescope

EINSTEIN TELESCOPE

Data

- Strain data will still be strain data
  - Over-estimate **~100 PB/year** of raw detector data (many auxiliary channels)
  - *This is significant, but largely for archive*
- Assume to first order, most analysis only needs h(t), plus data quality and metadata
  - about **10-100 TB/year**
  - *Data management at this scale is a solved problem*

Compute

- Computing needs will increase, the bare minimum we would be required to do is process the data once in low latency
  - Today, LVK uses **20k cores**
  - *This is significant, but many large computing centres can handle that today*
- Scaling up by **1000** makes that a scary number
  - ***We need measurements of end-to-end workflows to measure the scaling factor and give funding agencies an answer***

Distributing ET Computing tasks

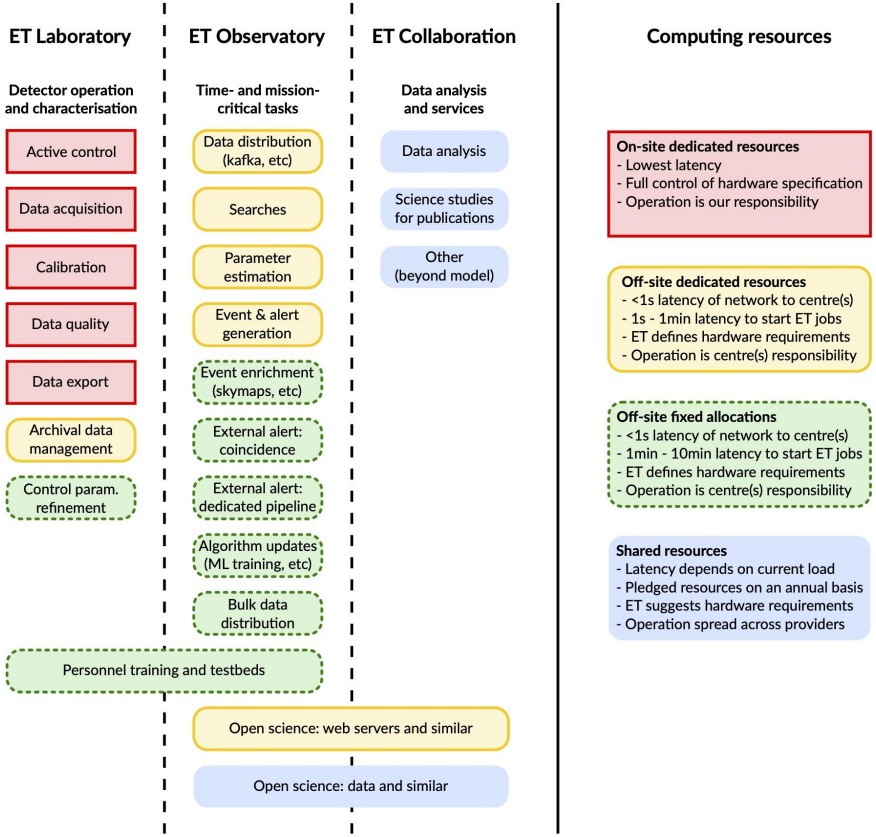

Like LVK, we assume ET will share data with other GW detectors and use distributed computing resources

The ET Computing Model implementation may depend on the site decision

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

“On-site” refers to things that are strictly ***required to be at detector site(s)***

On-site computing for ET

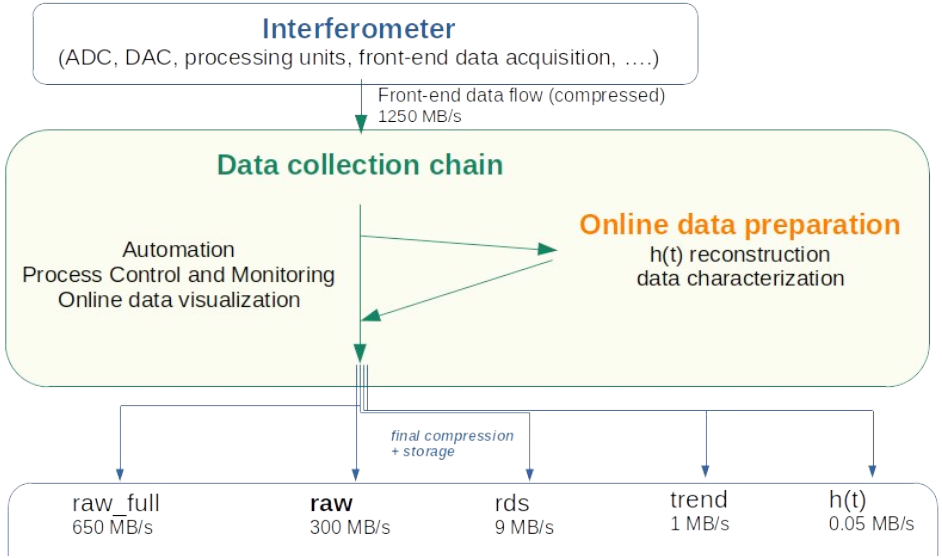

	Minimal Scenario	Operational Safety Margin
Operations Storage buffer (TB)	800	1600
Long-term storage (PB)	10	20
CPU Cores	2150	6450
RAM (GB)	6300	18900
Network (per site)	100 Gb/s	2 * 100 Gb/s

The conservative scaling model ***per IF***:

ET = Virgo * 5 (storage)

ET = Virgo * 3 (CPU and RAM)

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

Computing power needs are significant; the amount of data needed on-site, particularly during commissioning, means efficient data reduction will be a priority

ET Computing - energy footprint

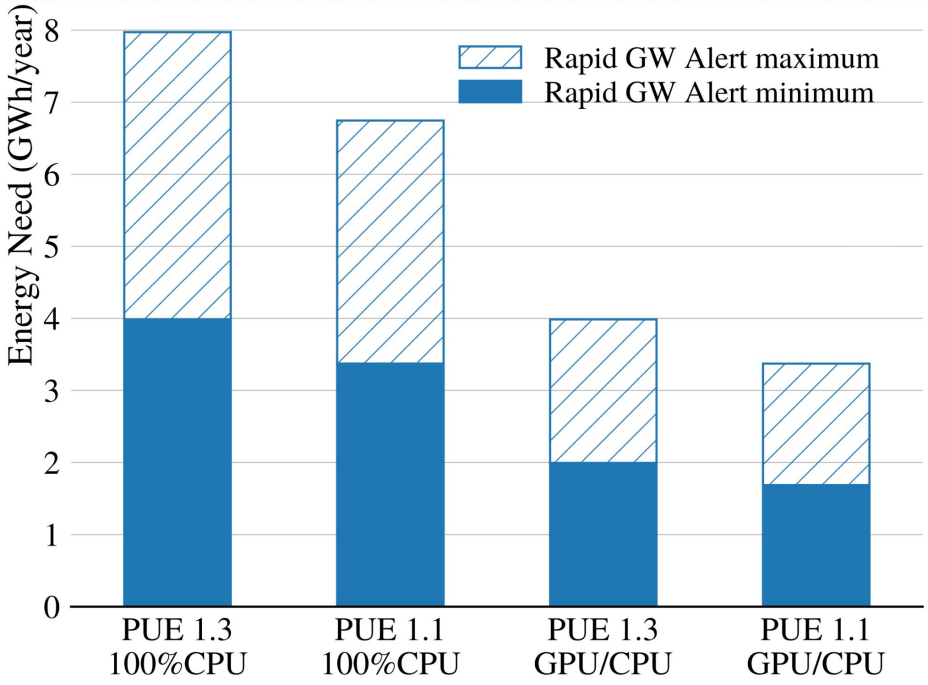

In the ET computing model, we present an estimate of the energy footprint of the computing power needed for **Rapid GW Alerts** *(fed to WP9)*

This assumes the Rapid GW Alert workflow is the only time the entire data stream is processed

The estimate 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 differences, as does leveraging GPUs

Sustainable Computing for ET

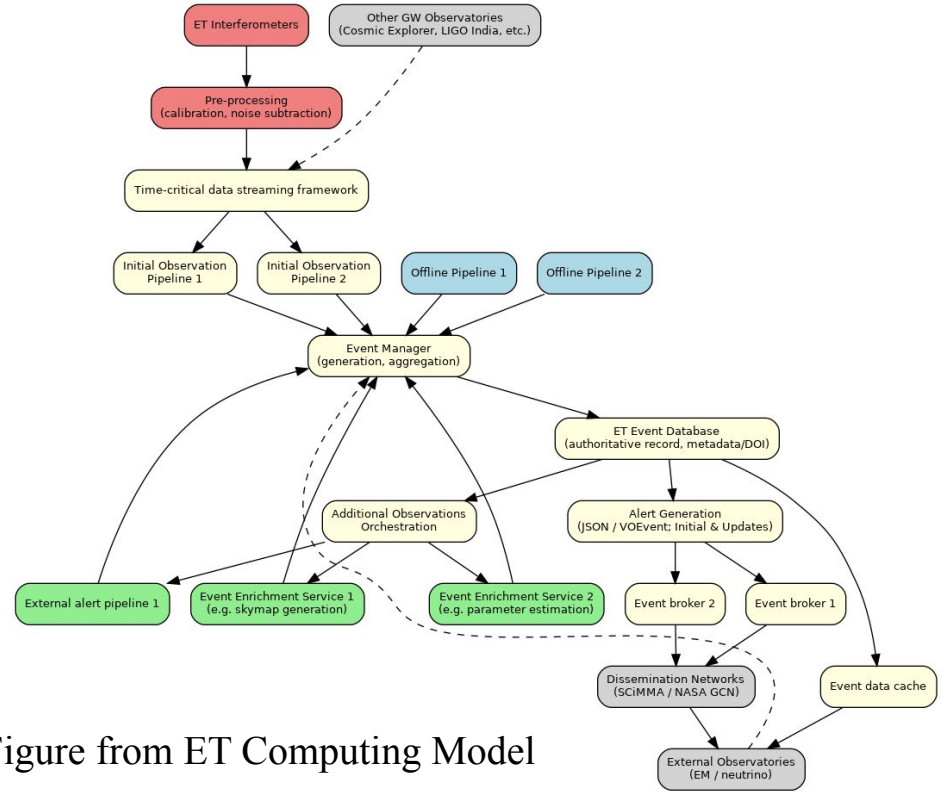

Figure from ET Computing Model

*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 minimising duplication

Scalable, portable computing infrastructure

Why software engineering matters

EINSTEIN TELESCOPE

LVK circa 2015

- Scientists write and run the code
- It produces excellent science
- The results get deserved credit, LVK software helped deliver a Nobel-prize

LVK today

- Ad hoc software becomes difficult to maintain (bus factor, people move on and even out of the field)
- Ad hoc deployment strategies become technical debt and increasing operations burdens
- No end-to-end testing means updates become operational risks

What ET needs

- Software that has a long-term support model, can be rerun by others and adheres to best practices
- A common deployment model that can run anywhere ET has resources
- A common development environment enabling end-to-end testing as standard

The issue is not whether the code (once) worked; it is whether ET can safely depend on it for decades

Software engineering best practice must be the foundation of ET software development

A new software engineering division for the EIB

EINSTEIN TELESCOPE

Div 1
Data Challenge Support

End-to-end MDC workflow evaluation
Resource-footprint measurements

X-coord: OSB-Div10

Div 2
Collaborative Services

IAM, helpdesk, onboarding
collaboration tooling
(Security incubator)

Div 3
Distributed Computing

Federated computing sites and services (data, WMS, accounting, monitoring…)
Resource estimates

Div 4
Multi-messenger Alerts

Low-latency alert infrastructure

X-coord:
MMMB Coordinator
OSB-Div4

Technology Tracking WG

Ensuring a place for R&D

Future divisions

Div 5 Security | Div 6 Open Science | Div 7 Site Operations

Division 0 — Core Software and Engineering Foundations

Best practices | FAIR-by-design | CI/CD | **Software framework** | **versioned software releases** | support | training

MDCs for estimating ET computing needs

EINSTEIN TELESCOPE

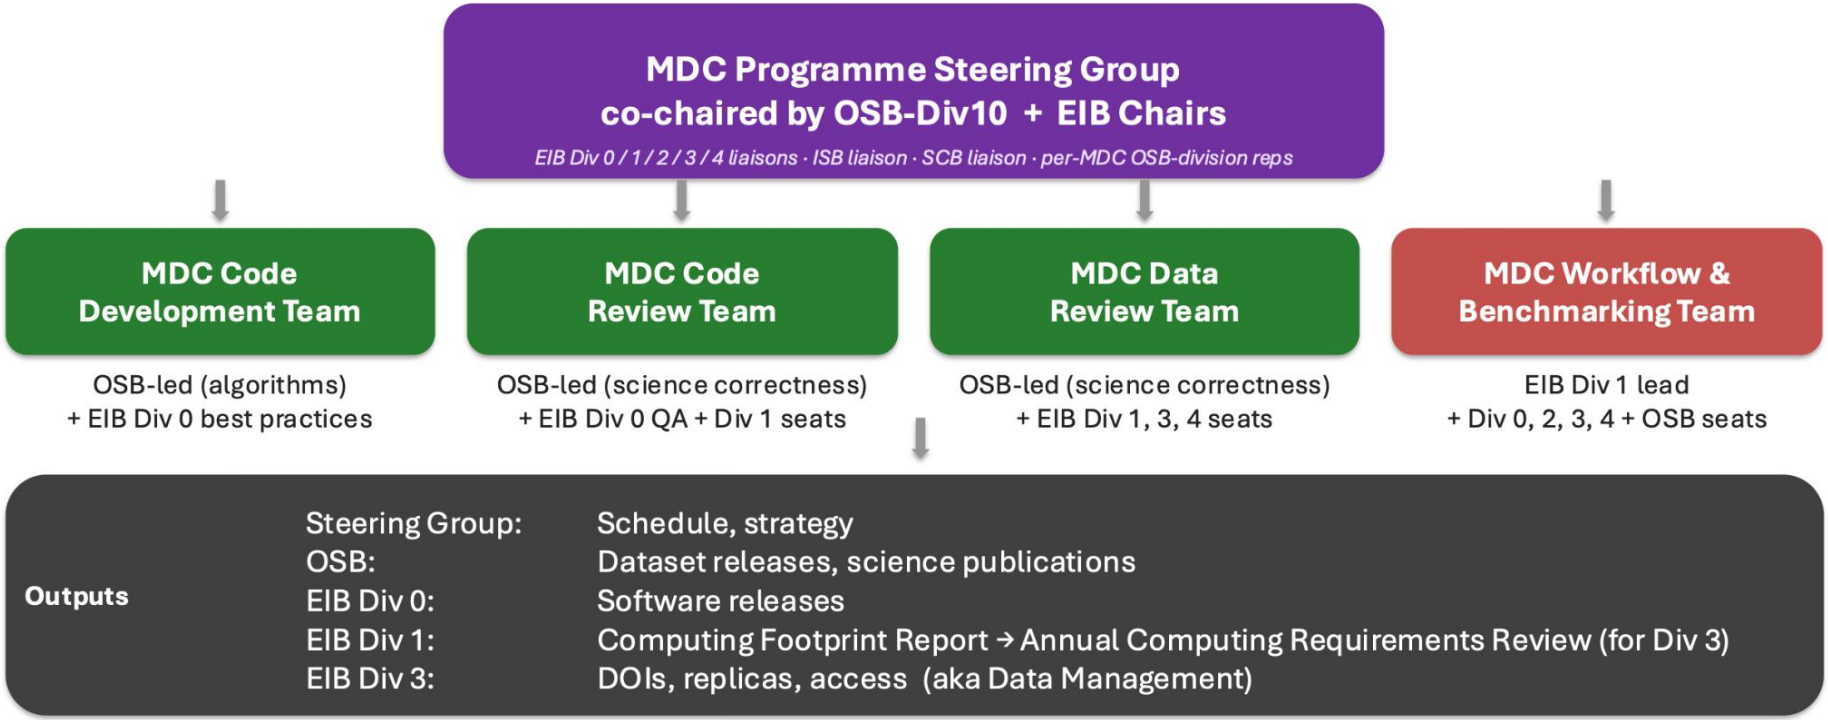

The EIB board needs the ET collaboration

EINSTEIN TELESCOPE

- EIB has low participation, somewhat inheriting the lack of computing expertise in Virgo
  - The handful of existing experts* are ***heavily*** over-subscribed

- And yet we need to on-board new computing sites to support Mock Data Challenges
  - **AI-4-ET** will be a fantastic accelerator, but we still need people!

- and real effort in the EIB for the services that AI-for-ET catalyses

**Nota Bene: Gary Hemming is a real human being, we don't have more*

MDCs and computing resources

ETAP successfully concluded in July, delivering the ESCAPE VRE helm chart:

https://github.com/vre-hub/helm-charts

MADDEN continues and we have several Rucio “Storage Elements”

We on-boarded PIC using Indigo IAM for high priority MDC review work

And we revived the ET ***zenodo*** community for publishing data (a magnet for SPAM unfortunately, but now up!)

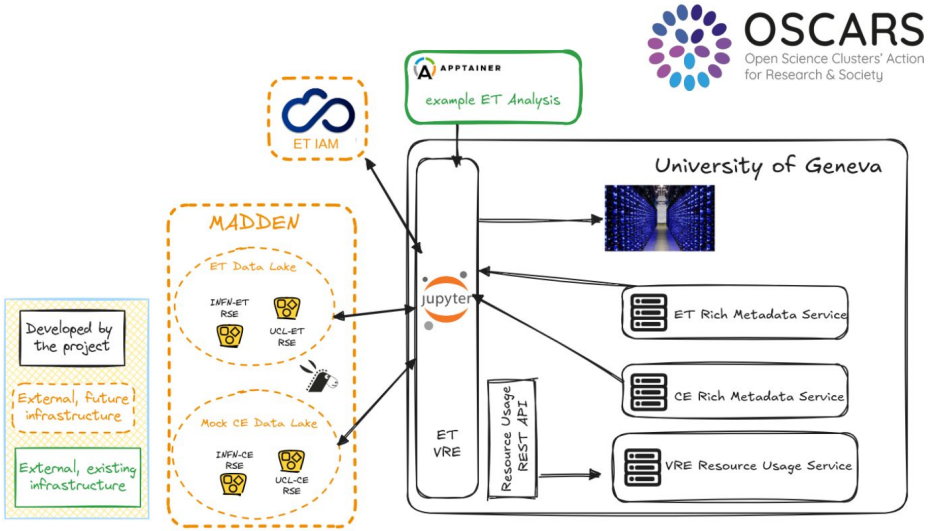

One collaboration, one computing ecosystem

EINSTEIN TELESCOPE

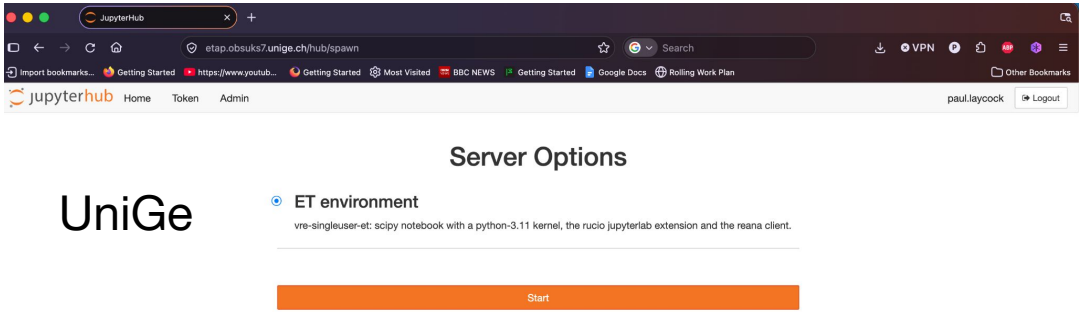

UniGe

PIC port d'informació científica

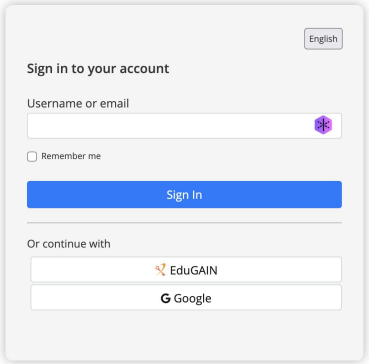

PIC

Goal: **one** ET account gives you access to **all** ET data and ET computing resources, wherever they are

Requires a scheme for recognising computing contributions (pledges)

Successes: Data Management with Rucio

A success story: ET engagement with the Rucio data management software is excellent, well represented at the Rucio workshop last week

Largely driven by the Torino group and thanks to funding through OSCARS

Uncovered: the rest !



EIB Divisions

EINSTEIN TELESCOPE

Div 1
Data Challenge Support

End-to-end MDC workflow evaluation
Resource-footprint measurements

X-coord: OSB-Div10

Div 2
Collaborative Services

IAM, helpdesk, onboarding
collaboration tooling
(Security incubator)

Div 3
Distributed Computing

Federated computing sites and services (data, WMS, accounting, monitoring…)
Resource estimates

Div 4
Multi-messenger Alerts

Low-latency alert infrastructure

X-coord:
MMMB Coordinator
OSB-Div4

Technology Tracking WG

Ensuring a place for R&D

Future divisions

Div 5 Security | Div 6 Open Science | Div 7 Site Operations

Division 0 — Core Software and Engineering Foundations

Best practices | FAIR-by-design | CI/CD | **Software framework** | **versioned software releases** | support | training

EIB Divisions

EINSTEIN TELESCOPE

Div 1
Data Challenge Support

End-to-end MDC
workflow evaluation
nt

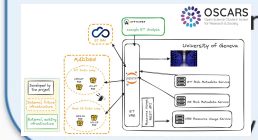

/10

Div 2
Collaborative Services

lesk,
g
ion tooling
ncubator)

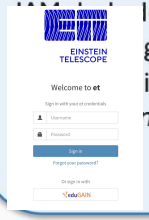

Div 3
Distributed Computing

Federated computing sites
and services (data, WMS,
onitoring...)
hates



Div 4
Multi-messenger Alerts

Low-latency alert
infrastructure

X-coord:
MMMB Coordinator
OSB-Div4

Technology Tracking WG

Ensuring a place for R&D

Future divisions

Div 5 Security | Div 6 Open Science | Div 7 Site Operations

Division 0 — Core Software and Engineering Foundations

Best practices | FAIR-by-design | CI/CD | **Software framework** | **versioned software releases** | support | training

EIB Divisions

EINSTEIN TELESCOPE

Div 1
Data Challenge Support

End-to-end MDC
workflow evaluat

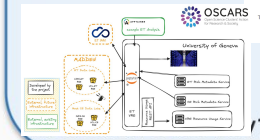

/10

404
Not Found
The resource requested could not be found on this server!

Div 2
Collaborative Services

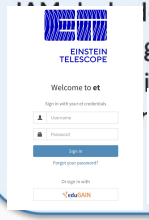

404 Not Found
404 Not Found
404 Not Found

Div 3
Distributed Computing

404 Not Found
404 Not Found
404 Not Found
404 Not Found



Div 4

404
Not Found
The resource requested could not be found on this server!

Technology Tracking WG

Ensuring a place for R&D

Future divisions

Div 5 Security | Div 6 Open Science | Div 7 Site Operations

Division 0 — Core Software and Engineering Foundations

Best practices | FAIR-by-design | CI/CD | **Software framework** | **versioned software releases** | su

404
Not Found
The resource requested could not be found on this server!

Conclusions

We have a computing model concept for ET, based on LVK, leveraging all of the acronyms

(WLCG/CERN/HSF/HEP/EOSC/ESCAPE…)

Distributed computing is largely a solved problem in terms of technology, and Europe expects us to use those existing solutions

We need more ET collaboration members to sign up to integrate, operate and maintain a federated distributed computing system that serves the whole ET collaboration

Writing software is cheap. ET needs software that we can rely on for **decades**

We have a new software engineering division to embody this, and now we need to find the engineers!



EINSTEIN TELESCOPE

Backup

Sustainable Computing for ET

The SCOPE project (DE) aims at developing a realistic, scalable and sustainable prototype for an ET Computing Center inside a shipping container:

- Adaption of computing power to energy availability
- (Close to) 100% self sustaining with own renewable production
- Efficient cooling & waste heat reusability

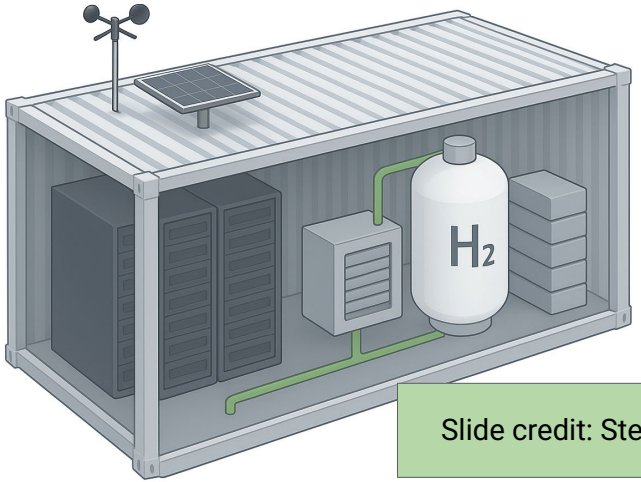

For ET, the biggest challenge for a sustainable computing model is keeping computing resource needs under control for ET's **50 year lifespan**

ET will not require new Data Centres, but there could be reasons why these may be built

There are already R&D activities for Green solutions

See Stefan's talk at CHEP

Slide credit: Stefan Krischer | RWTH Aachen University