Status of e-Infrastructure Board

Einstein Telescope Symposium
Maastricht – May 9th, 2024

Stefano Bagnasco, Patrice Verdier (EIB)
Nadia Tonello, Achim Stahl (WP8)

Introduction

«…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), 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: Anastasios Fragkos (U. 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 (EGO)

Joint WP8+EIB weekly call for coordination

Division 1: Software, frameworks, and data challenge support

MDC input data distribution

- EIB Div 1 is providing support for ET Mock-Data Challenge by hosting the input data, distributing them to computing centres and offering means to download the data
  - Installed OSDF origin service at UCLouvain to host the MDC input data
  - MDC input data files metadata are published into CVMFS
    `/cvmfs/et-gw.osgstorage.org/et-gw/PUBLIC/MDC1/`
    - Requires CVMFS client (available in WLCG and IGWN computing centres)
    - When a file is being accessed, it is downloaded from OSDF
  - Data are publicly accessible
  - Download possible from UCLouvain's OSDF origin
    - Requires XRootD client (available in WLCG and IGWN computing centres)
  - Installed web server on UCLouvain's OSDF origin to allow download without need to install any software
  - Documentation on how to access MDC data is in EIB Div 1 wiki
  - Created MDC Computing Helpdesk (gitlab issues) for user support (0 issues so far!)

7

Andres Tanasijczuk & Steffen Hahn

Open Science Data Federation* (OSDF) infrastructure from Open Science Grid (OSG)

Resources in EU provided by institutions affiliated to Virgo

- Virgo "origin" at UCLouvain
- Caches at PIC, IN2P3, CNAF, Amsterdam

Added a new "origin" for ET at UCLouvain

Division 1: Software, frameworks, and data challenge support

Using ESCAPE (c.f. next slides) Virtual Research Environment :

- **VRE consists of:**
  - (Rucio) Data Lake
  - Jupyterhub with extension to integrate with Data Lake
  - REANA data analysis platform (job submission)
  - Auth: integration with ESCAPE Indigo IAM

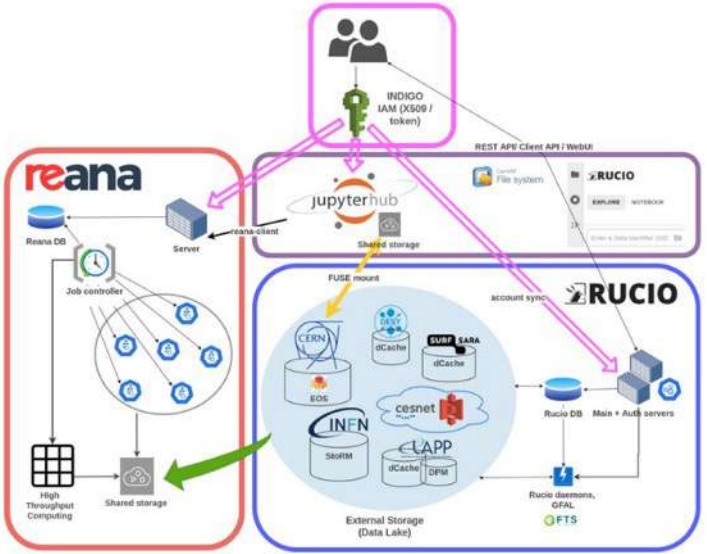

Uploaded a few MDC files to the Data Lake

- Successfuly ran a (simple) stochastic GW background (pygwb) pipeline on REANA
- Activity done together with Div 3, since relevant for compute resources estimation
  Georgy Skorobogatov (U. Barcelona), Andres Tanasijczuk (UCLouvain)
- Became the main users and testers of the VRE

Further developments in OSCARS proposal

MDC-code rewrite - what has been done?

Why should we do a rewrite?

- monolithic codebase
- code duplication
- reused variables (side effects)
- polluted repository

rewrite

Basic ideas

- C?? → C++(21/23)
- enforce modularization
- (automatic) validation/memory tests
- build wrappers around C-code from LAL

What has been done (so far)?

- original features reimplemented
- modularized code (one file ≈ one task)
- improved configurability
- output validated (simple tests)

The code seems to do what the old code did. Therefore, first phase of rewrite is more-or-less complete. However, there are still a couple of "issues" before the first version.

Division 2: Services and Collaboration Support

Provide services to the ET users: strong support from EGO IT Dpt

- ET Member Database & Active Directory accounts
- Indico agenda server
- ET Wiki & web page
- …

ET Authentication and authorization infrastructure (AAI)

- Implementing the AAI solution developed within ESCAPE and supported by major European centers/labs in HEP and astroparticle phsyics (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
  - Interface the ETMD with Indigo IAM

Division 3: Computing and data model, Resource Estimation

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."*

Gonzalo Merino

2

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

Gonzalo Merino

12

Division 3: Computing and data model, Resource Estimation

Need for HPC resources

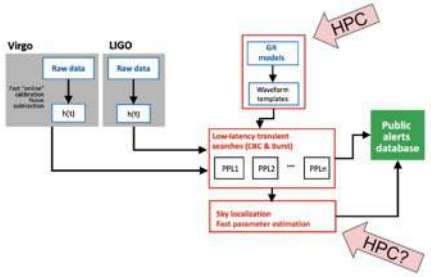

Fig. 3: Simplified conceptual representation of 2G low latency data analysis workflows.

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 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).

13

What about the Cloud?

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.

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

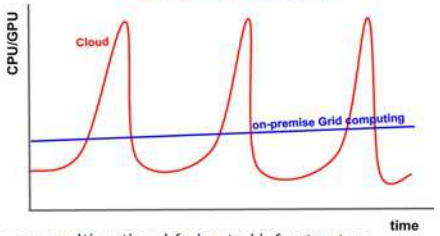

On-premise Grid has clear advantages as a multi-national federated infrastructure. Mostly, it provides a simple mechanism for countries to contribute.

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? ...*

Gonzalo Merino

15

Division 4: Multimessenger alerts infrastructure

- **Division 4, multimessenger infrastructure, relies on close collaborations with other groups**
- **Within ET: close connections to OSB Div4 and Div10**
  - Jointly organised a mini-workshop with OSB Div4 as part of the ET annual meeting in November
  - Investigating the MDC in the context of low-latency analysis for computing requirements
- **External to ET: need to participate in the larger multimessenger community**
  - Participating in various cross-RI workshops and discussions as opportunities arise
  - Concrete involvement in the recent INFRA-TECH proposal, M2Tech, submitted 03/2024

    => Next slide
- **As input-gathering continues, independent Div4 activities will grow**
  - Please get in touch if you want to contribute, or have ideas on how to best proceed

Steven Schramm

Infratech : M2TECH proposal

M2TECH : Technologies for Multimessengers Astrophysics

HORIZON-INFRA2022-TECH-01-01

R&D for the next generation of scientific instrumentation, tools and methods

At least 3 world class research infrastructures RIs:

CTAO and **MAGIC, ET** and **Virgo, KM3NeT, ELI**

- First step to structure R&D for ET
- Computing WP6 lead by Steven Schramm
  - Efficient data processing
  - Sustainable large-scale computing
  - Multimessenger alert tools
- Resubmitted in March 2024

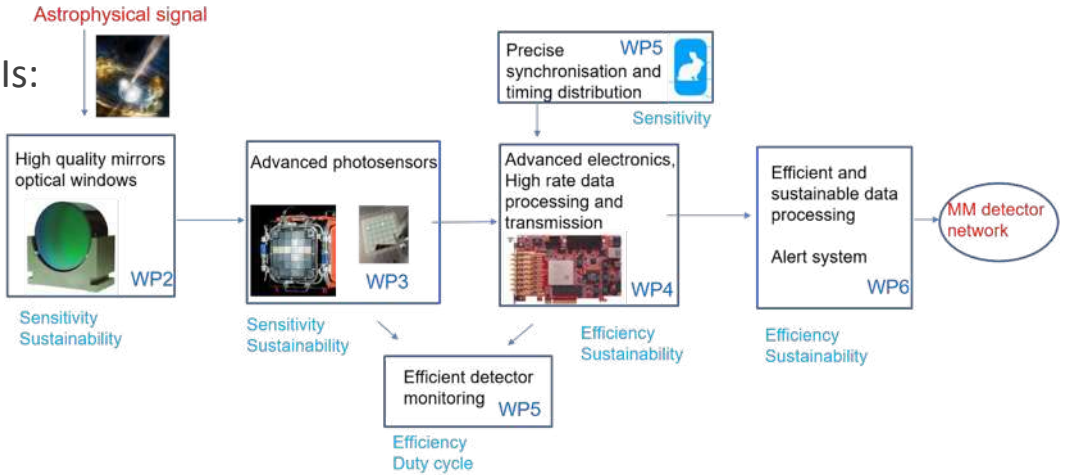

ESCAPE and Einstein Telescope

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

OSCARS

ESCAPE

HORIZON EUROPE

https://oscars-project.eu/

The OSCARS – ***O**pen **S**cience **C**lusters' **A**ction for **R**esearch & **S**ociety* proposal was submitted in March in response to the call HORIZON-INFRA-2023-EOSC-01-01 for a budget of EUR 25.00 million.

Two activities:

1) Consolidation of the Science Cluster approaches between the RI communities, where ESCAPE focuses on Data Lake, VRE and training. As part of this activity, we were also asked to study and implement an ESCAPE "community-based competence centre".
2) Demonstrate and pilot the use of EOSC resources by multiple research communities through cross-RI and/or cross-domain open science projects and services, by running 'open calls with cascading grants' with a budget of approximately EUR 16.00 million.

OSCARS is based on the co-participation of the five science cluster collaborations in Europe ENVRI-FAIR (Environment and Earth Sciences), EOSC-LIFE (Biomedical Sciences), ESCAPE, PANOSC (Light and Neutron Facilities), SSHOC (Social Sciences and Humanities).

5 Giovanni Lamanna 07/07/2023

OSCARS call

Several elements of the computing and software ecosystem are very interesting (notably Rucio for data management)

What we propose

Two complimentary proposals

1a) Use Rucio for multiple GW detector data management

1b) Enhance GW data handling tools, building on RucioFS

2a) Deploy a virtual research environment (VRE) for MDCs

2b) Monitor computing resource needs

2c) Provide access to GW experiment metadata

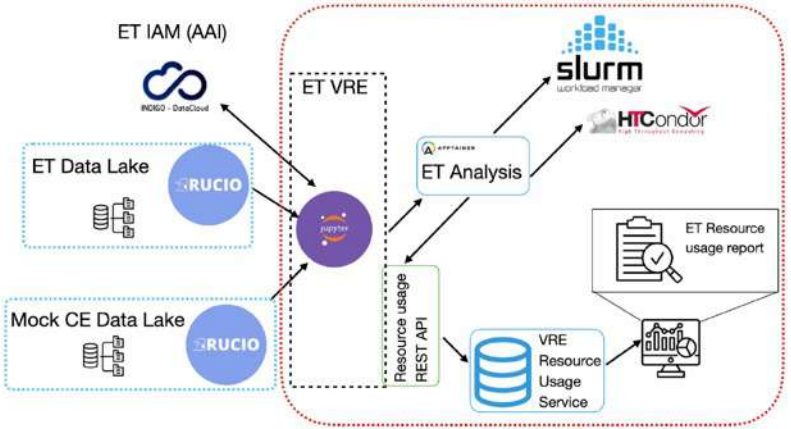

ET-PP

ET-PP WP8 Computing and Data models

WP8 milestones

	Workshops	WPs	Due date	
M8.1	Workflows Requirements collection and constraints: computing and data		Sep 2023	Oct 2023
M8.2	Computing Infrastructures availability for ET workflows, characteristics	WP9	Aug 2024	July 2024
M8.3	On site infrastructure, computing and data model	WP6	Aug 2025	
M8.4	Low latency and offline workflows and computing model	WP6	Dec 2025	
M8.5	Data management, access, policy and implementation	WP2, WP6	July 2026	

WP8 Documents

	Deliverables	Lead	Due date	
D8.1	Computing and Data Requirements	UniGe	Feb 2024	Feb 2024
D8.2	Computing and Data Model	UniGe	Feb 2026	
D8.3	Data Access Implementation Guidelines	IFAE	July 2026	

Barcelona Supercomputing Center

ET Symposium - May 2024

Nadia Tonello & Achim Stahl

EiB-WP8 workshops

M8.1 Computing and data requirements
Host: Geneva Date: 26th-27th October
Agenda: https://indico.ego-gw.it/event/590/

Sustainability workshop participation (WP9)
Co-located with the ET annual meeting 16-17 Nov 2023
Paul's report
Focus on computing aspects: optimization of software, workflows and computing resources, reproducibility, automation. Carbon emissions related to computing is identified as challenging but crucial. Expertise is needed.

The big Picture

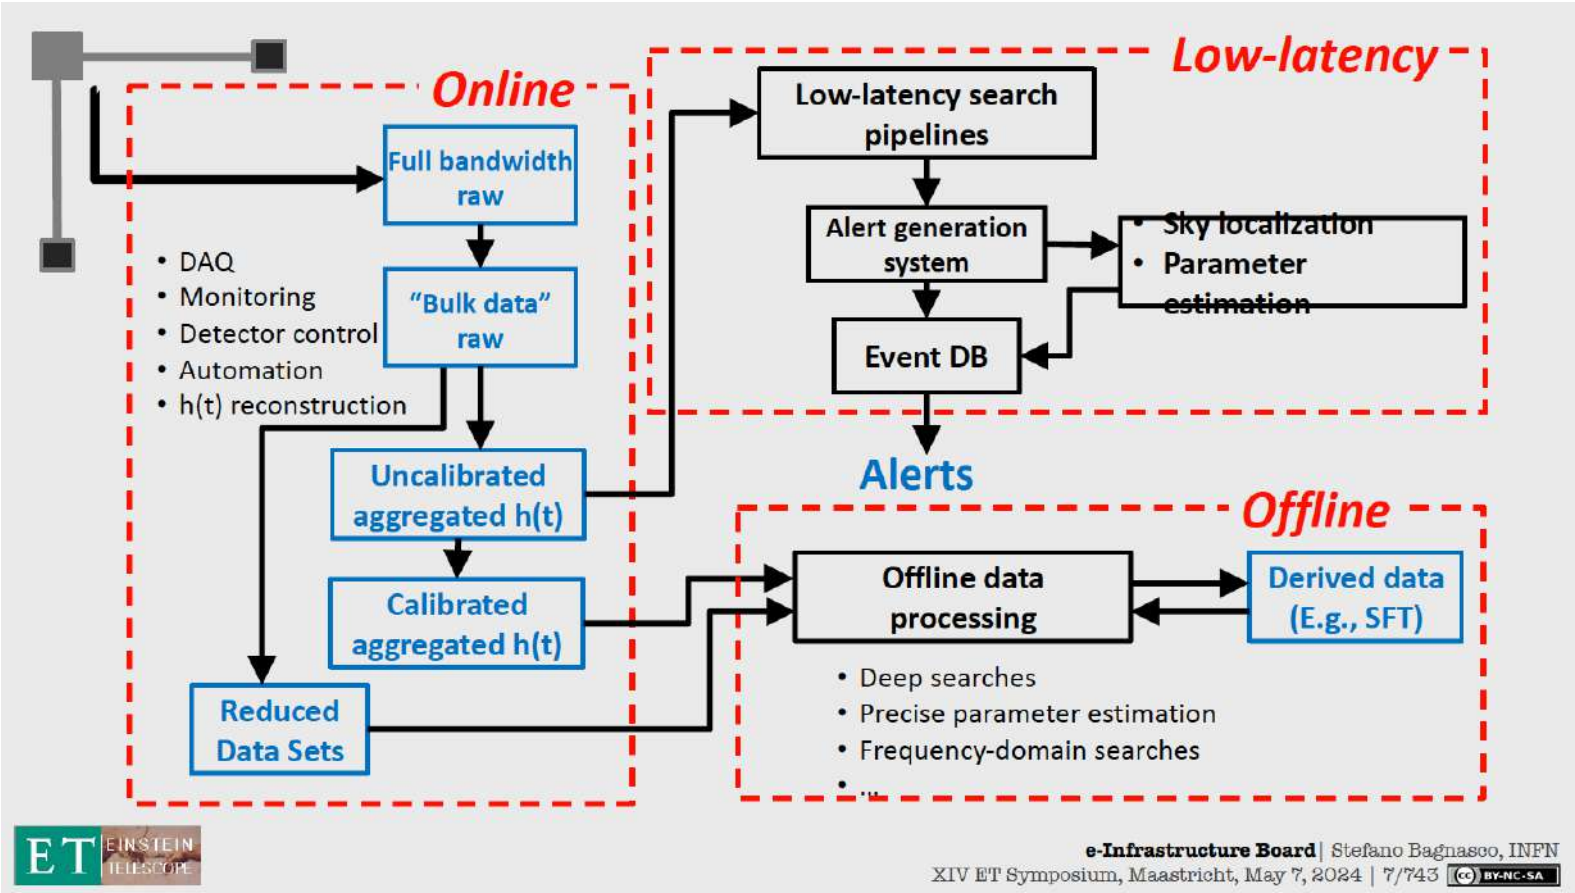

The big Picture

THE "USER"

- Write some simple code to do an analysis
  - Use common and well-documented libraries…
  - …packaged in easy-to-use environments
- Find the data you need
  - By querying a database, possibly with complex metadata requirements including "external" annotations
  - E.g., "2 hours around each external EM trigger of the type xyz, with the following data quality requirements"
  - Even better: "get also the EM public data from this detector"
- Fire up a Jupyter Notebook
  - Or something like that
  - Do not worry at all about data locality
  - Do not worry about workload management
- Run your analysis
  - A wonderful paper is out!

THE "POWER USER"

- Write some complex code to do a computing-intensive analysis
  - Use common and well-documented libraries
- Find the data you need
  - Same as before
- Plug your algorithm in a framework
  - With well-defined, standardised APIs
  - Express you computing needs (HPC, GPUs, whatever)
  - Do not worry about data locality
  - Do not worry about workload management and resource availability
- Run your analysis
  - One more wonderful paper out!

THE "MULTIMESSENGER USER"

- Write some code to do a natively multimodal analysis
  - Use common and well-documented libraries
- Find the data you need
  - Querying across multiple ineroperable metadata-rich catalogues
  - …using federated AuthN/AuthZ: not only public data!
- Plug your algorithm in a framework
  - Same as before, but more so
- Run your analysis
  - This will be a really stunning paper!

THE LOW-LATENCY EXPERT

- Write some complex code to do a search, sky localization, PE, or whatever
  - Use common and well-documented libraries (the same used for offline)
- Have the data available in the facility you need
  - Possibly even large HPC facilities?
- Plug your algorithm in a framework
  - That API-wise looks a lot like the one for offline
- Run the alert management services on a standard orchestration platform
  - On owned or shared cloud resources
- Broadcast your alerts

THE DATA LIBRARIAN

- Manage large-scale reprocessing campaigns
  - E.g., re-calibration of data
- Be able to annotate data with rich metadata
  - Detector and environment status ("conditions"), identified noise transients, event candidates, external triggers,…
- Manage data "cycles" between custodial storage and data distribution network
- (More…)

THE COMPUTING TEAM

- Monitor the infrastructure operations across heterogeneous resources from different providers
  - Alterts forwarded by underlying infratructures
  - Custom application-specific metrics and alerts
- Distribute the code to heterogeneous etc.
  - CI/CD provides automated ways of building/testing/packaging for different resources
- Manage services via simple interfaces on industry-standard orchestration
  - Which todays means Kubernets, tomorrow?
  - On owned and shared Cloud resources, with uniform tools
- Gather consistent and high-granularity accouting information
  - And be able to enforce priorities across computing tasks

Stefano Bagnasco

The big Picture

SHOPPING LIST

Data transfer and storage: safely and efficiently transfer all data to custodial storage and processing centres, including low-latency transfers;

Software packaging and distribution: manage software lifecycle, and make packages available ubiquitously;

Computing power: provide and manage computing resources (HTC and HPC) for the processing of data, in all computing domains;

Data distribution: make data available to worker nodes in computing centres anywhere, and possibly also to single workstations, including support to public releases of data;

High-availability service management: provide a platform for running the collaboration's services (e.g. alert generation services, event databases,...)

Data cataloguing and bookkeeping: organise all data and metadata and provide querying and discovering capabilities;

Job lifecycle management: provide a uniform job submission and runtime environment to research groups;

High-level workload management: keep a database of all jobs and allow the enforcement of priorities and scheduling strategies; provide support for organized large-scale data processing campaigns;

Monitoring and accounting: monitor local and distributed computing, checking performance and looking for issues, and provide reliable accounting both at the user/job and site level;

Authentication, Authorisation and Identity management: provide consistent AAI across all domains and activities.

Collaboration services: provide tools for efficient collaboration management, coordination, and outreach (e.g. document repositories, collaborative tools, administrative databases, communications,...)

ET EINSTEIN TELESCOPE

The big Picture

WORKFLOW EVALUATION KITS

- Independent packaged parts of the final architecture
  - Providing limited functionalities, possibly some as mere demonstrators
  - But actually to be released to users (i.e., they MUST be functional)
  - Different implementations may exist, with different tools/technologies used to provide same functionality
  - Integration of existing tools, not excessive bespoke developments to map "kits" onto smallish projects
- Examples:
  - ESCAPE Datalake + RucioFS for data distribution
  - IAM-based AAI
  - ESCAPE Datalake + VRE interactive data analysis
  - OSDF + INFNCloud interactive data analysis
  - "Packaged" and quality-tested MDC data generation tool
  - Paul's rich metadata tool

Conclusion and next steps

- **With current and future Mock Data Challenges:**
  Iteratively improved infrastructure prototypes and usage data collection tools

- **ET Representation and Collaboration**
  ET currently informally represented in WLCG through Virgo. Likely to evolve in a formal “general” GW representation (IGWN + ET)
  Effort to showcase ET computing case in major physics computing venues (CHEP, ACAT)
  https://www.epj-conferences.org/articles/epjconf/abs/2024/05/epjconf_chep2024_04015/epjconf_chep2024_04015.html

Computing Challenges for the Einstein Telescope project
CHEP2023 Proceeding, arXiv:2312.11103

- **ESCAPE/OSCARS:**
  EIB-WP8 collaboration, engage with various communities including Rucio developers, WLCG data challenge participants, and those involved in the M2Tech proposal.

- **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

Conclusion and next steps

- **ET-PP WP8 – 2024 milestone:**
  Computing infrastructures availability for ET workflows workshop, in collaboration with WP9 (sustainability):
  July 2024: Workshop on “Computing Infrastructures availability for ET workflows, characteristics” to held in Napoli

- **Future organization of the IGWN collaboration:**
  - Current (2G) data/software model is IGWN (OSDF+HTCondor,…) extended to European CCs
  - Follow carefully the discussion on the evolution of the IGWN collaboration to prepare the 3G era with CE
  - A common GW computing infrastructure and/or coordination body (as WLCG for LHC) seems like a good target

Right time to get involved in Computing & Software activities for ET:
please contact us if you are interested to contribute