ESCAPE
European Science Cluster of Astronomy & Particle physics ESFRI research Infrastructures

Distributed Computing for Open Science: the ESCAPE collaboration

Xavier Espinal (CERN)

ET-PP/ET-EIB workshop @ Geneva: Computing and Data Requirements
Department of Astronomy, University of Geneva, Oct 26 – 27, 2023

Outline

- Need for Commonalities in Scientific Computing
- The ESCAPE project: achievements and model
- The ESCAPE Open Collaboration Agreement and its work program
- The EC-funded project opportunities

From the European Strategy for Particle Physics

D. Large-scale data-intensive software and computing infrastructures are an essential ingredient to particle physics research programmes. The community faces major challenges in this area, notably with a view to the HL-LHC. As a result, the software and computing models used in particle physics research must evolve to meet the future needs of the field. ***The community must vigorously pursue common, coordinated R&D efforts in collaboration with other fields of science and industry, to develop software and computing infrastructures that exploit recent advances in information technology and data science. Further development of internal policies on open data and data preservation should be encouraged, and an adequate level of resources invested in their implementation.***

Deliberation Document
on the 2020 update of the European Strategy for Particle Physics

Astroparticle Physics European Consortium (APPEC)

APPEC

APPEC Contribution to the European Particle Physics Strategy

December 17, 2018

Editorial Board:
S. Katsanevas, A. Masiero, T. Montaruli, J. de Kleuver, A. Haungs

Contact Person:
T. Montaruli (APPEC Chair from Jan. 1, 2019)
Email: teresa.montaruli@unige.ch

Website : http://www.appec.org

The scientific outcomes of particle physics experiments are made possible by the development of an efficient computing and software infrastructure. Computing and software are profound R&D topics in their own right and are essential to sustain and enhance particle physics research capabilities. There is a need for strong community-wide coordination for computing and software R&D activities, and for the development of common coordinating structures that will promote coherence in these activities, long-term planning and effective means of exploiting synergies with other disciplines and industry. Some recently initiated examples are the HEP Software Foundation addressing the common computing and software challenges related to particle physics, and ESCAPE (European Science Cluster of Astronomy & Particle physics ESFRI research infrastructures) exploring the synergies in the areas of astronomy, astroparticle and accelerator-based particle physics.

B. The particle physics community and the European Commission have a strong record of collaboration. ***The relationship between the particle physics community and the European Commission should be further strengthened, exploring funding-mechanism opportunities for the realisation of infrastructure projects and R&D programmes in cooperation with other fields of science and industry.***

C. European science policy is quickly moving towards Open Science, which promotes and accelerates the sharing of scientific knowledge with the community at large. Particle physics has been a pioneer in several aspects of Open Science. ***The particle physics community should work with the relevant authorities to help shape the emerging consensus on Open Science to be adopted for publicly-funded research, and should then implement a policy of Open Science for the field.***

APS Division of Particles and Fields Response to European Strategy Group Call for White Papers: Tools for Particle Physics

DPF Executive Committee and Strategy Whitepaper Editing Group
dpfstrategy@fnal.gov

December 18, 2018

Abstract

The U.S. particle physics strategy process is summarized in a companion white paper that also describes U.S. activities related to the five P5 science drivers. Additional activities within the U.S. particle physics program that are critical to progress in our field are described here.

ESCAPE acknowledged as an example of a real synergistic cooperation scheme in the 2020 European Strategy for Particle Physics

Large commonalities in Scientific Computing needs

Computing scale increasing in all scientific domains, **large commonalities** across scientific communities. Common needs for:

- **Large Scale** Data Management and Data Transfer Services
- Integration and orchestration of **distributed** Storage and Compute services
- Facilitate data **analysis** to large user communities (data, software and code)
- **Reproducibility** and reusability of analysis workflows
- Willingness to promote **Open Science** and **Open Data**
- Outreach and **citizen science** activities

ESCAPE: The European Science Cluster of Astronomy and Particle Physics

Consortium of 31 members, including:

- 10 ESFRI projects & landmarks: **CTA, EST, FAIR, HL-LHC, KM3NeT, SKA, LSST, VIRGO, ESO, JIVE**
- 2 pan-European International Organizations: **CERN and ESO**
- 2 European Research Infrastructures: **EGO and JIV-ERIC**
- 4 supporting European consortia: **APPEC, ASTRONET, ECFA and NuPECC**

Budget: **15.98 M€**

Duration: **48 months** (1/2/2019 -31/1/2023)

ESCAPE, a domain-based EOSC “Cell” (a thematic node?)

Virtual Observatory:
Extend the VO *FAIR* standards, methods to a broader scientific context; prepare the VO to interface the large data volumes of next facilities.

Science Platforms:
Flexible science platforms to enable open data analysis tailored by and each facility as well as a globally ansversal workflows.

Citizen Science:
Open gateway for citizen science on ESCAPE data archives and ESFRI community

ESCAPE OSSR
Catalogue & Repository of resources

Datasets
Software & services
Tutorials
Training
Publications

Software Repository:
Repository of "scientific software" as a major component of the “data” to be curated in EOSC.

ESCAPE VO — Virtual Observatory
Astronomy Data centres
VO Registry
VO Registry
Analysis Tools
VO Services

TSP’s
RI-Specific Science Platforms
ESCAPE ESAP — Science Platforms
Workflows, notebooks, deployment platforms, packaging

ESCAPE CS — Citizen Science

ESCAPE DIOS — Data Lake
FAIR data management
Content discovery and delivery

HPC
PRACE
EuroHPC Joint Undertaking

HTC
Grid clusters, etc

Private/public clouds

Commercial clouds

GÉANT

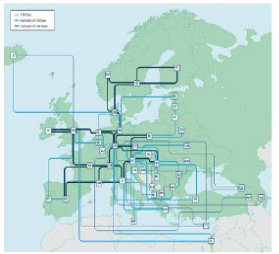

Key high-level ESCAPE project results

- Provided a full **Exabyte-scale Data Management** prototype (Storage, Transfer and Access) service with **common AAI** framework for distributed scientific computing
  - validated by ESFRIs through Data Challenges
- Developed a **catalogue** to publish digital scientific products of ESCAPE communities;
  - onboarded into EOSC portal
- Developed i**nteroperability standards** for astronomical data services;
  - VO services prepared to be integrated with EOSC
- Produced a **reusable analysis toolkit** for integrating diverse service offerings, and a **Virtual Research Environment** prototype that integrates ESCAPE services
  - A number of RI's using tools in their analysis frameworks
- Several **Citizen Science** projects built upon ESCAPE developments,
  - and attracted significant interest/involvement
- Dark Matter and Extreme Universe Science Projects as demonstrations of Open Science in ESCAPE
  - and integration with EOSC platform through EOSC-Future

The ESCAPE Data Lake

- The Data Lake **harness heterogeneous facilities** different types of storage systems (eg. http-enabled) provided by partner sites

- Able to **serve data** to heterogeneous facilities, eg. Grid sites, HPC, Cloud providers

- **Data accessible from single-user nodes** (eg. notebooks) via standard protocols (http)

- **Data Management** and orchestration: **Rucio**

- **File Transfer** and data movement: **FTS**

- **AAI**: Indigo **IAM** (OAuth2/OIDC and legacy x509 support)

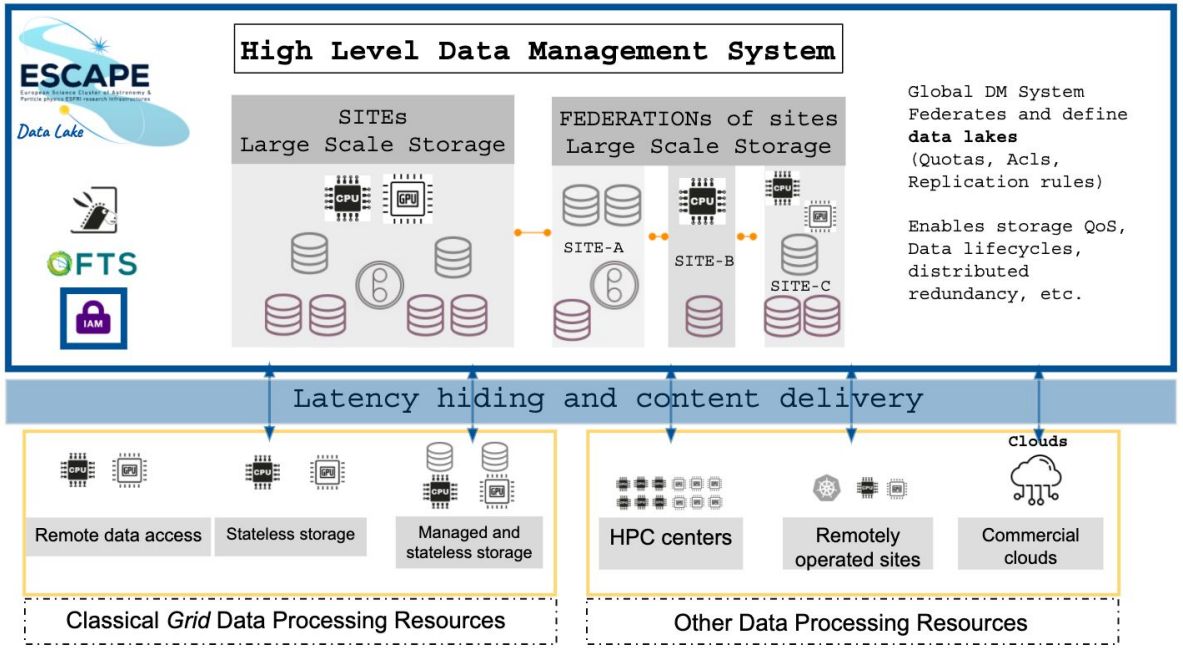

The ESCAPE Data Lake

- The Data Lake **harness heterogeneous facilities** different types of storage systems (eg. http-enabled) provided by partner sites

- Able to **serve data** to heterogeneous facilities, eg. Grid sites, HPC, Cloud providers

- **Data accessible from single-user nodes** (eg. notebooks) via standard protocols (http)

- **Data Management** and orchestration: **Rucio**

- **File Transfer** and data movement: **FTS**

- **AAI**: Indigo **IAM** (OAuth2/OIDC and legacy x509 support)

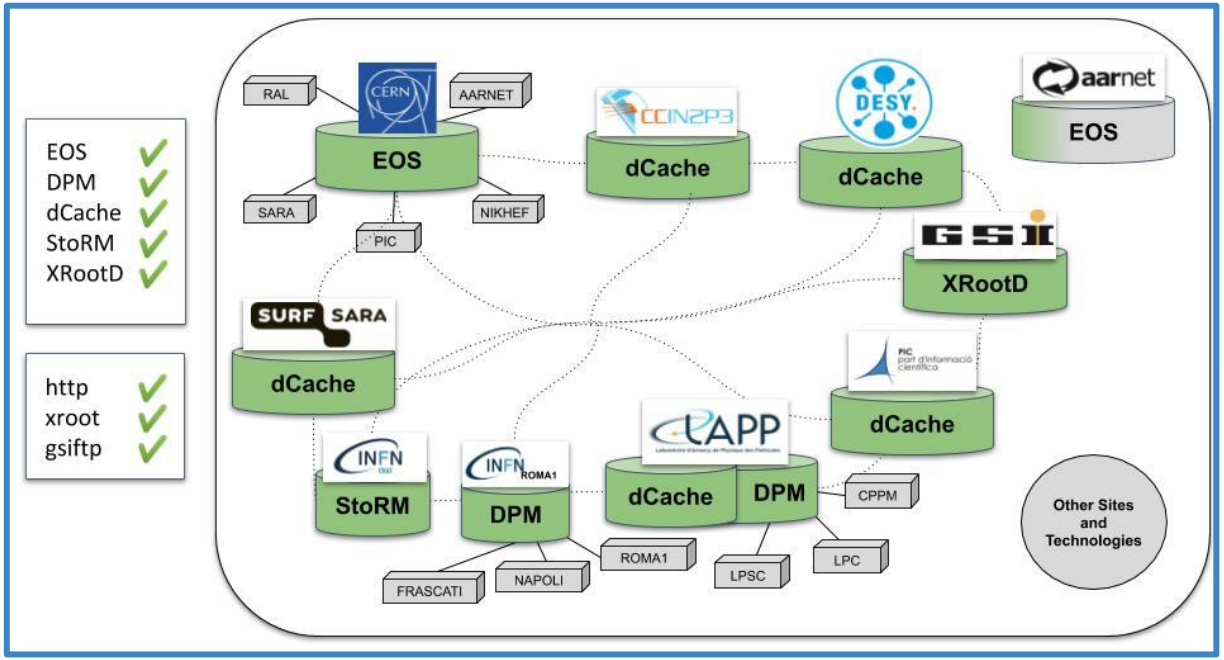

The ESCAPE Data Lake at work

Experiment's realistic data workflows:

- **Injection** of raw data, from the source: detectors, telescopes and antennas
- **Data management**, including Extreme distance data workflows
- Full **Data Life-Cycles**
- **Data processing** on several facilities, including large batch systems and user notebooks (browsing Data Lake files)
- Enabled **multi-wavelength** analysis

Some RIs deployed their **own Data Lake**

- Certifying the successful cross-operation of several *Data Lakes* in a **shared and common** storage infrastructure

FAIR 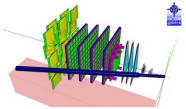	**Ingestion and replication** of simulated R3B data of mCBM detector on FAIR-ROOT. Ingestion and replication of simulated and digitised raw PANDA fallback data
KM3NeT 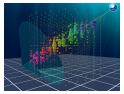	**Offload data** from the storage buffer in the coast, **replicate** across sites, run **data calibration, store** back. Data product ready for user consumption
HiLumi HL-LHC PROJECT ATLAS EXPERIMENT	User **analysis pipeline** tests on experimental particle physics by using augmented open data
MAGIC Major Atmospheric Gamma Imaging Cerenkov Telescopes	Long haul raw data ingestion and replication. Data is successfully transferred from the telescope station (La Palma, Canary Islands) and replicated to the Data Lake, **file deleted on the telescope storage buffer**
LOFAR	Process data in the Data Lake at an external location, **combine results with other astronomical data to produce multiwavelength images**
HiLumi HL-LHC PROJECT CMS	DL interface with local and heterogeneous resources: multi purpose **Analysis Facility PoC** with data access via DASK leveraging HPC and batch clusters. Content delivery and caching. Performance comparison
SKAO	**Long haul data replication**: data in correct place in timely manner using SKAO Rucio for Australia and South-Africa to UK RSEs. End-to-end proof of concept **data lifecycle** test, AUS/SA to northern hemisphere sites

The Virtual Research Environment (link to demo)

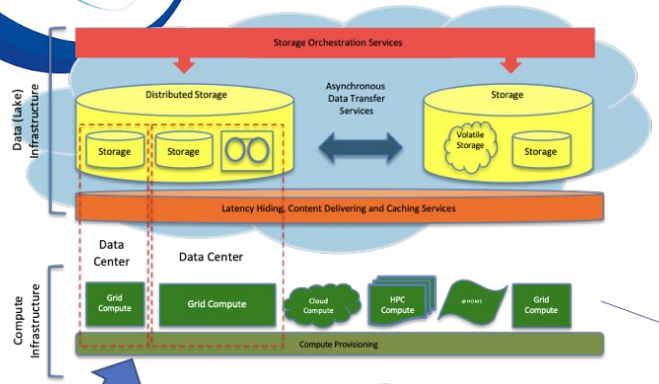

Data Lake: RUCIO

VRE: jupyter reana kubernetes

ESCAPE
European Science Cluster of Astronomy & Particle physics ESFRI research infrastructures
Virtual Research Environment

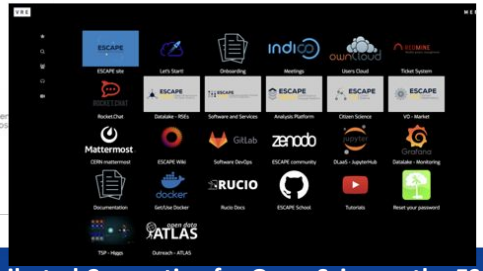

Software Repo: ESCAPE OSSR | Open-source Science and Service Repository; zenodo

ESCAPE AAI

EOSC Future

AAI Federated with EOSC

ESCAPE Data Lake (Federated storage)

Hosts/ingests open data sources into a common store; Includes EOSC Exchange provisioned storage

EOSC Future

Virtual Research Environment (VRE)

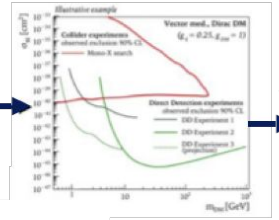

Publish results to Zenodo

zenodo

Cloud computing using EOSC Exchange provisioned compute, & HPC

EOSC Future

Software Repository

Archive of reusable pipelines: Onboarded to EOSC marketplace

EOSC Future

EOSC Core:
+ Helpdesk as a service
(+monitoring)
(+accounting)

EOSC Future

Multi-wavelength analysis in the notebook (3C196 quasar)

1. Data **injected** to the Data Lake from three **radio** source observations in **external** locations

2. User **download** the data, **process** and **store** results back to the Data Lake

3. User interested in **combining** results stored with other public data to cover also **visible** spectrum

4. Combined **optical** data from the Hubble located via the VO plugin available in the **notebook**

5. **Optical** and **radio** data aggregate in the **Analysis Platform combined analysis done.** Results uploaded back to the DL.

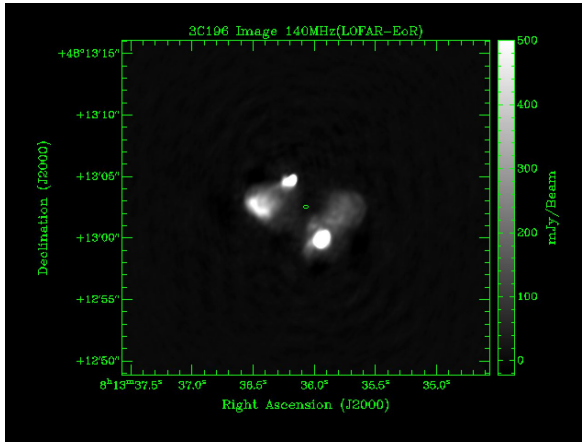

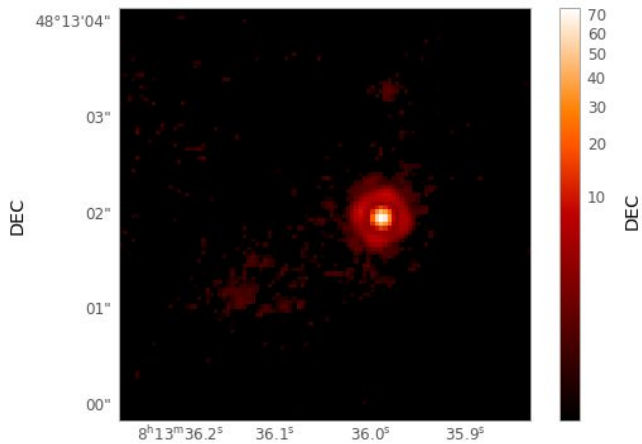

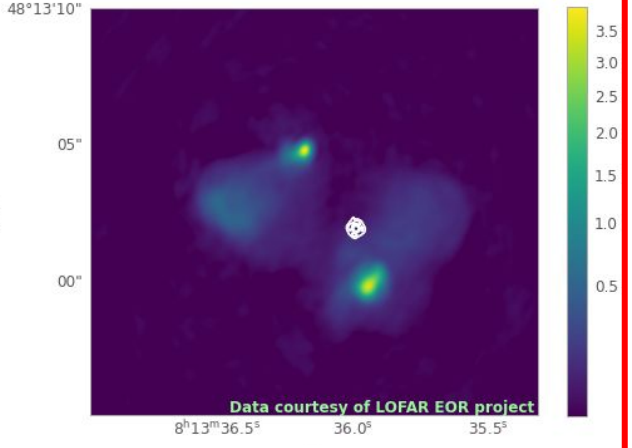

The central core of 3C196 is sitting right in the middle of the radio component, only visible on optical images.

From left to right: Radio image, Optical image and the Combined image (LOFAR with optical contours)

Thoughts after the ESCAPE project

- Understand what is **useful** for us, where can we bring benefits
  - Establish several competence centres, to share expertise etc.
- A working model for a **long term ESCAPE activities:**
  - Work programme **driven by scientific needs**. Ensuring our RIs are fully involved in the planning process
  - Use of relevant **funding opportunities** to address specific and well understood needs
- **Strengthen the work with the Science Clusters** to understand how we all fit in EOSC evolution
  - We really need to re-state what we want (science usually a bit ignored in some of the EC discussions)

The ESCAPE Open Collaboration Agreement

The ESCAPE Open Collaboration Agreement

- Agreement signed by the **directors** of all partner Research Infrastructures
- Recognising the many **synergies**: communities, technical, coordination, political and funding
- **Common** facilities (data centres, networks)
- **Common** scientific goals
- **Common** challenges both in scale and integration of new technologies/facilities (AI/ML, HPC, Q, etc.)
- Representing a single community as a **consolidated political force** towards the EC, EOSC, and other bodies

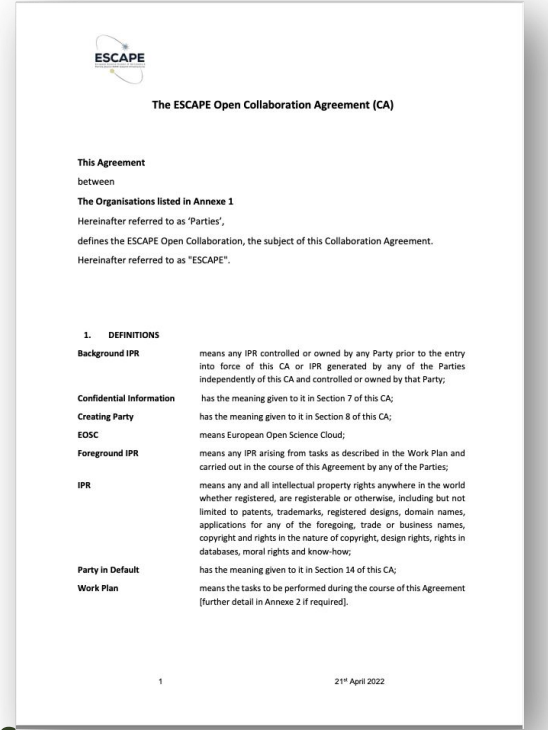

ESCAPE

The ESCAPE Open Collaboration Agreement (CA)

This Agreement

between

The Organisations listed in Annexe 1

Hereinafter referred to as 'Parties',

defines the ESCAPE Open Collaboration, the subject of this Collaboration Agreement.

Hereinafter referred to as "ESCAPE".

1. DEFINITIONS

Background IPR	means any IPR controlled or owned by any Party prior to the entry into force of this CA or IPR generated by any of the Parties independently of this CA and controlled or owned by that Party;
Confidential Information	has the meaning given to it in Section 7 of this CA;
Creating Party	has the meaning given to it in Section 8 of this CA;
EOSC	means European Open Science Cloud;
Foreground IPR	means any IPR arising from tasks as described in the Work Plan and carried out in the course of this Agreement by any of the Parties;
IPR	means any and all intellectual property rights anywhere in the world whether registered, are registerable or otherwise, including but not limited to patents, trademarks, registered designs, domain names, applications for any of the foregoing, trade or business names, copyright and rights in the nature of copyright, design rights, rights in databases, moral rights and know-how;
Party in Default	has the meaning given to it in Section 14 of this CA;
Work Plan	means the tasks to be performed during the course of this Agreement [further detail in Annexe 2 if required].

1 21st April 2022

Example 1: Data Management

The pursuit of commonalities in Data Management and Data Access is key. To ensure **sustainable and efficient delivery of resources and services** to the Research Infrastructures (RIs)

- Foster an holistic approach favouring **economy of scale** to globally improve hardware usage, software tools. Maximise manpower impact on computing.
- **Privilege common approaches**, identify and adopt best practices. Facilitate the organisation and management of scientific data, compliant with the FAIR principles.
- Pursue a global **common understanding** of Data Management and Access. Easing the implementation of data life-cycles, replication policies and access rules. Facilitate the transition to Open Data demands.
- **Expand** the data access and data processing possibilities by ensuring integration with growing demands for **heterogeneous resources and facilities**, ie. Clouds, HPC, GPUs, FPGA, ML/AI clusters, etc.

Example 2: Data Access. Analysis.

Minimise time-to-science. Provide ergonomic access to data, software and code.

- **Analysis Facilities**, Analysis **Platforms** or **Virtual Research Environments** to provide access to the infrastructure offering the digital content necessary to perform a scientific analysis
  - Users spend less time in (*non-interesting*) pre-analysis environment setup, eg. look for data, software access, code, access and permissions, etc.
- Flexible and **scale-out** access to computing resources, favouring integration of **heterogeneous** computing backends: local clusters, clouds and HPCs. Offload analysis products from the platform to large infrastructures (from the laptop to Exascale)
  - Allow analysis workflow **encapsulation** to ease analysis preservation, reproducibility, reusability and shareability
- Evaluate potential benefits in the community to adopt **new trends** in user-analysis
  - Visual and interactive wrt. terminal-text based job-submit.
  - Breaking the barrier between outreach and large scale analysis with notebooks?

Opportunities in the ESCAPE Open Collaboration

- Funding calls in Horizon Europe; highly relevant and potentially important to us
  - Building on the science cluster approach to ensure uptake by RIs and research communities
  - Consolidation of activities and new Science Projects
  - Bring in new world-class RI and others via open calls and cascading grants
  - Clusters expected to work together

HORIZON-INFRA-2024-DEV-01-01: Research infrastructure concept development

Call - Next generation of scientific instrumentation, tools, methods, and advanced digital solutions for RIs (2024) 105
Conditions for the Call 105
HORIZON-INFRA-2024-TECH-01-01: R&D for the next generation of scientific instrumentation, tools, methods, solutions for RI upgrade 106
HORIZON-INFRA-2024-TECH-01-02: New technologies and solutions for reducing the environmental and climate footprint of RIs 108
HORIZON-INFRA-2024-TECH-01-03: Development of tools, solutions, modules to enable R&I on the social aspects of the green transition 110
HORIZON-INFRA-2024-TECH-01-04: New digital twins for Destination Earth 112
HORIZON-INFRA-2024-TECH-01-05: AR/VR-empowered digital twins for modelling complex phenomena in new RI application areas 115

Call - Enabling an operational, open and FAIR EOSC ecosystem (2024) 59
Conditions for the Call 59
HORIZON-INFRA-2024-EOSC-01-01: FAIR and open data sharing in support of the mission adaptation to climate change 60
HORIZON-INFRA-2024-EOSC-01-02: Supporting the EOSC Partnership in further consolidating the coordination and sustainability of the EOSC ecosystem 63
HORIZON-INFRA-2024-EOSC-01-03: Enabling a network of EOSC federated and trustworthy repositories and enhancing the framework of generic and discipline specific services for data and other research digital objects 67
HORIZON-INFRA-2024-EOSC-01-04: Long-term access and preservation infrastructure development for EOSC, including data quality aspects 70
HORIZON-INFRA-2024-EOSC-01-05: Innovative and customizable services for EOSC Exchange 73

https://www.science-clusters.eu/

Horizon Europe: The OSCARS Project

- The OSCARS – Open Science Clusters' Action for Research & Society proposal was submitted in March in response to the call *HORIZON-INFRA-2023-EOSC-01-01*
- Two main activities:
  - **Consolidation** of the Science Cluster approaches between the RI communities. ESCAPE will focus on **Data Lake, VRE and training**
    - Study and implement an ESCAPE "**community-based competence centre**".
  - 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'
- *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) and **SSHOC** (Social Sciences and Humanities)*

OSCARS: Computing joint roadmaps

Community based **Competence Centers** to strengthen the role of the **Science Clusters** in the **European Research Area**

- ESCAPE Open CA and the CCCs are a recognised umbrella where to anchor **joint activities, projects and working groups** (eg. ongoing analysis working group)
- The CCCs will provide the opportunity for **new generation of scientific computing experts** to build-up expertise in computing models, infrastructure operations and distributed computing tool.
  - Also as a discussion forums and operations “hub”
- Next two slides are **preliminary** ideas, gathered from the ESCAPE project, as a *strawman* guideline to start building up the roadmap/work program of the CCCs
  - The roadmap needs to be kept open and agile enough to address and fit new requirements coming from the communities (some of them new… like ET!)

Horizon Europe: The EVERSE Project

- EVERSE aims to create a **framework** for research **software and code excellence**, collaboratively designed and championed by the research communities
  - In pursuit of building a **European network of Research Software Quality** and setting the foundations of a future Virtual Institute for Research Software Excellence.
- EVERSE will establish a **network** amongst the five EOSC Science Clusters (ENVRI-FAIR, EOSC-Life, ESCAPE, PANOSC, SSHOC) who develop and make research software for their respective European communities
- Community connections:
  - High **Energy Physics Software Foundation (HSF)** collaboration to foster best practices, training, summer schools, common tools and reusable software solutions in particle physics.
  - ESCAPE OSSR (software repository, training collection, best practices in software). Existing efforts in career rewarding and development.

Science Cluster Competence Center :: Data Management

what	*why*
Storage Federation, Data Lake and Data Management system (Rucio)	Provide to the experiment(s) and clusters an integral Data Management solution. From raw data recording/injection to global replication rules, access policies and data life-cycles.
File Transfer Service (FTS)	Global file transferring machinery connecting the different storage systems. File movement “executor backend” of the data management system. Able to speak transfer protocols required by the infrastructure providers.
Common Trust Layer (IAM)	Common layer of trust to ensure single (and shared) identity across data management services, analysis facilities and resources.

Science Cluster Competence Center :: Data Management

what	*why*
Content Delivery and Latency Hiding	1. **DAQ buffers**. RIs and labs needing raw data injection tools and/or data download. Asses deployment of **edge services** addressing CDN and latency hiding, allowing to interface storage buffers with the Data Lake approach.
	2. **HPCs**. Edge services to blend HPC internal storage and RI data/ Accelerate collaboration with FENIX and EuroHPC Joint Undertaking to ensure new federated models of HPC are in sync with EU Science Clusters strategies.
	3. **Public**, **Hybrid and private clouds**. Ability to deploy Content Delivery and Latency Hiding as edge services to exploit opportunistic resources.
Multi-messenger and **multi-wavelength** ready infrastructure	To be able to combine data from different experiments (detectors, telescopes, radio-antennas, etc.), the data Management systems should able to provide extended metadata fields and/or interface to experiment’s metadata DDBBs. Examples are the current Rucio-IVOA integration.

Science Cluster Competence Center :: Data Access and Processing	
what	*why*
Tools, services and plugins	Ensure **sustainability of the key components** (tools and services) to enable data platforms and research environments access to the RIs data in the Data Lake: eg. Rucio plug-in maintenance and development, metadata extensions, integration of analysis preservation, code and software access, etc.
Content Delivery Network and **Latency Hiding** services	Analysis Facilities, Analysis Platforms and Virtual Research Environments to be able to connect to content delivery and latency hiding services. Continue investigations with XCache (xroot) service via the VRE environment.
Shared Filesystem integration	AF/VRE ability to leverage the facility's internal file systems and possible links to “local” Data Lake storage via shared file systems (posix)

Science Cluster Competence Center :: Data Access and Processing	
what	*why*
Software distribution and FS mount	Common need for a system that can handle software distribution: accessed by the RIs librarians to populate and automatically distributed to the computing nodes. Some level of posix-ness needed.
Analysis Preservation	Facilitate the immediate analysis “encapsulation“ (=> preservation, reproducibility, reusability and shareability)
Job *“capsules”* **distribution**	Ability to send offload the the notebook by wrapping jobs in a *capsule* (ie. REANA). These analysis capsules are self-standing jobs able to run elsewhere having pointers to data, software and code.
Analysis job **parallelization**	Evaluate tools for local cluster distribution, job scheduling and chunking with novel data patterns and access, eg RNtuples

- Early Mock Data Challenges to develop and validate everything