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

ESCAPE
Data Lake for Open Science

Xavier Espinal (CERN, ESCAPE WP2 Lead)

Workshop on Open Science practices in Nuclear Physics, 5th-6th December

ESCAPE - The European Science Cluster of Astronomy & Particle Physics ESFRI Research Infrastructures has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement no. 824064.

The ESCAPE H2020 consortium

- 31 partners including 2 SMEs
- 10 ESFRI projects & landmarks: **CTA, EST, FAIR, HL-LHC, KM3NeT, SKA, LSST, VIRGO, ESO, JIVE**
- 2 pan-European International Organizations: **CERN**, **ESO** with their world-class established infrastructures, experiments and observatories
- 2 European Research Infrastructures: **EGO** and **JIV-ERIC**
- 1 involved initiative/infrastructure: **EURO-VO**
- 4 supporting European consortia: **APPEC, ASTRONET, ECFA and NuPECC**

- Budget: **15.98 M€**
- Started: **1/2/2019**
- Duration: **48 months** (end date 31/1/202
- Coordinator: **CNRS-LAPP**

ESCAPE, ESFRIs and world-class Research Infrastructures

Connecting ESFRI RIs through Cluster projects

The H2020 cluster concept was aimed at supporting: "**Open data intensive driven science**" in order to "**rise productivity of researchers and to lead to new insights and innovation**"

The approach:

- Foster the establishment of **cross-border** open innovation environment
- Develop **synergies** and complementarity between involved (ESFRI) research infrastructures
- Adopt global standards and **common solutions** to the data management for economies of scale

The scope:

- Commit in **Open Science** that means implement the FAIRness of scientific data
- Link the ESFRI and other world-class RIs to **EOSC**

Active collaboration across the 5 Science Clusters

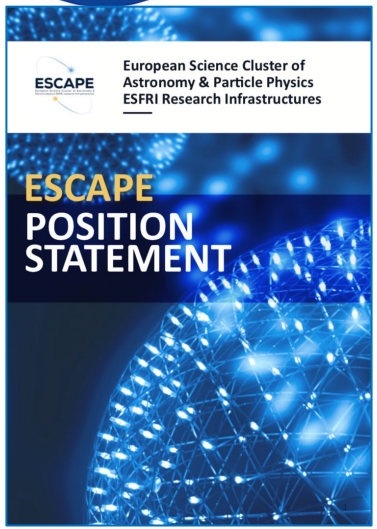

https://www.projectescape.eu/sites/default/files/Escape_position_statement_web.pdf

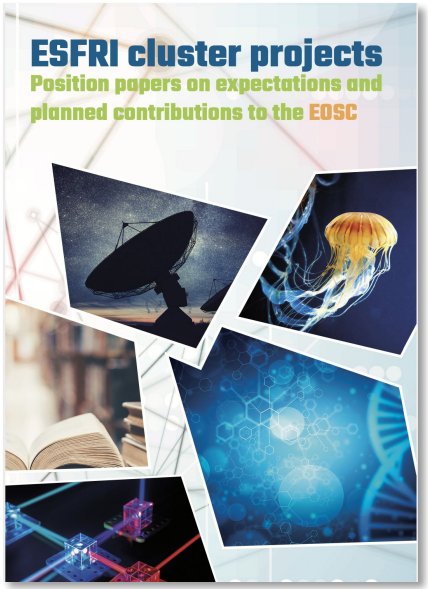

https://zenodo.org/record/3675081 - .X2R2PJNLhTY

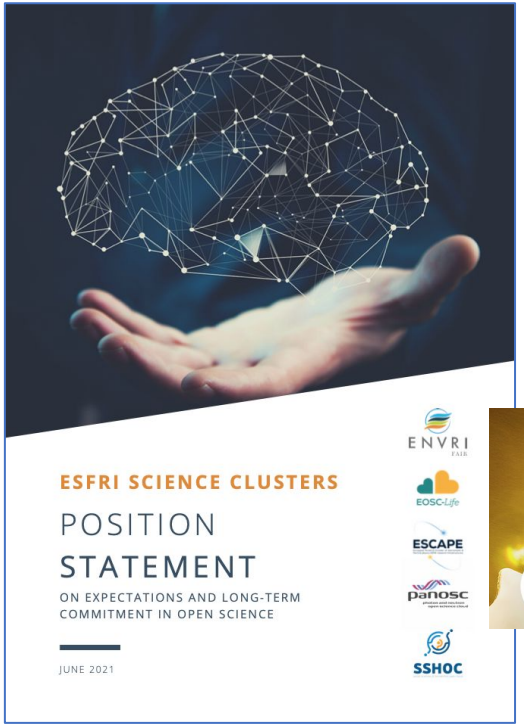

https://zenodo.org/record/4889503

https://indico.in2p3.fr/event/24327/

Active collaboration across the 5 Science Clusters

Fostering synergies in EOSC Future
between Science Clusters & e-infrastructures

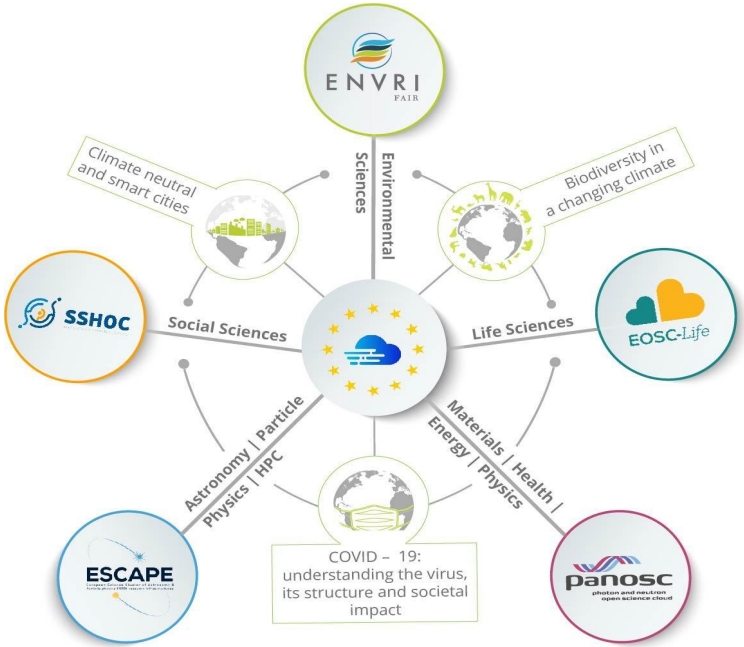

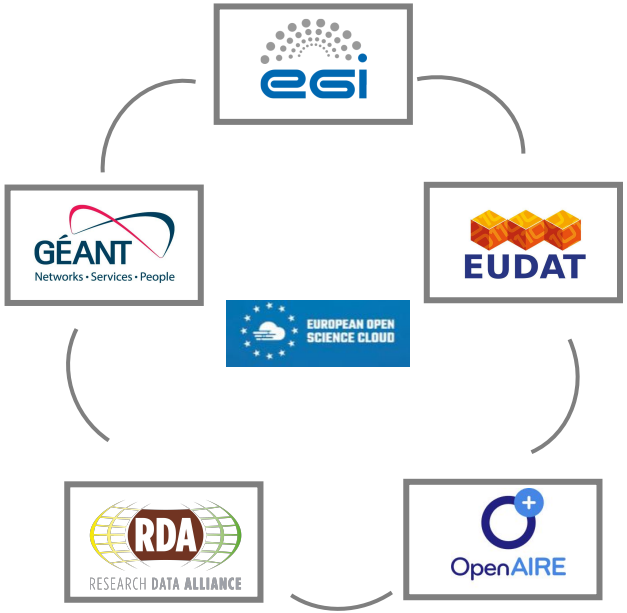

The ESCAPE work programme

Data Lake:
Build a scalable, federated, data infrastructure as the basis of open science for the ESFRI projects within ESCAPE.

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

Science Platforms:
Flexible science platforms to enable the open data analysis tailored by and for each facility as well as a global one for transversal workflows.

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

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

The ESCAPE Data Infrastructure for Open Science

ESFRI RIs

- Address RI's needs in Data Management, Access and Analysis for **Astro-particle, Radio-astronomy, Gravitational Waves, Cosmology** and **Particle Physics.**
- Provide a fully working **common data infrastructure** "The ESCAPE Data Lake" to test novel data management tools and models, giving the opportunity to influence and steer its development.
- Expand **collaborations** and foster involvement with other Scientific Communities. Maintain and strengthen collaborations with related EC initiatives and projects.

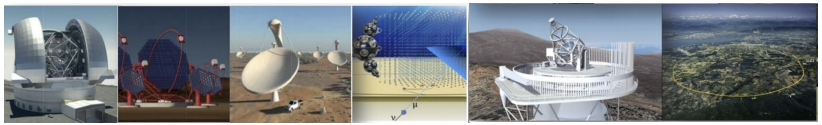

Data centres

A reliable **distributed** data infrastructure capable of managing **Exabyte-scale** data sets, able to deliver data **on-demand** at l**ow latency** to all types of data processing facilities

Services operated by the ESCAPE partner institutes

Petabyte scale storage: DESY, SURF-SARA, IN2P3-CC, CERN, IFAE-PIC, LAPP, GSI and INFN (CNAF, ROMA and Napoli)

Data management and storage orchestration (Rucio)

File transfer and data movement services (FTS)

Global Data Lake Information System (CRIC)

ESCAPE IAM: common Auth/Authz/IM (AAI)

Content delivery and latency hiding: XCache

Widening access with several data access protocols: http, xrootd and gridftp.

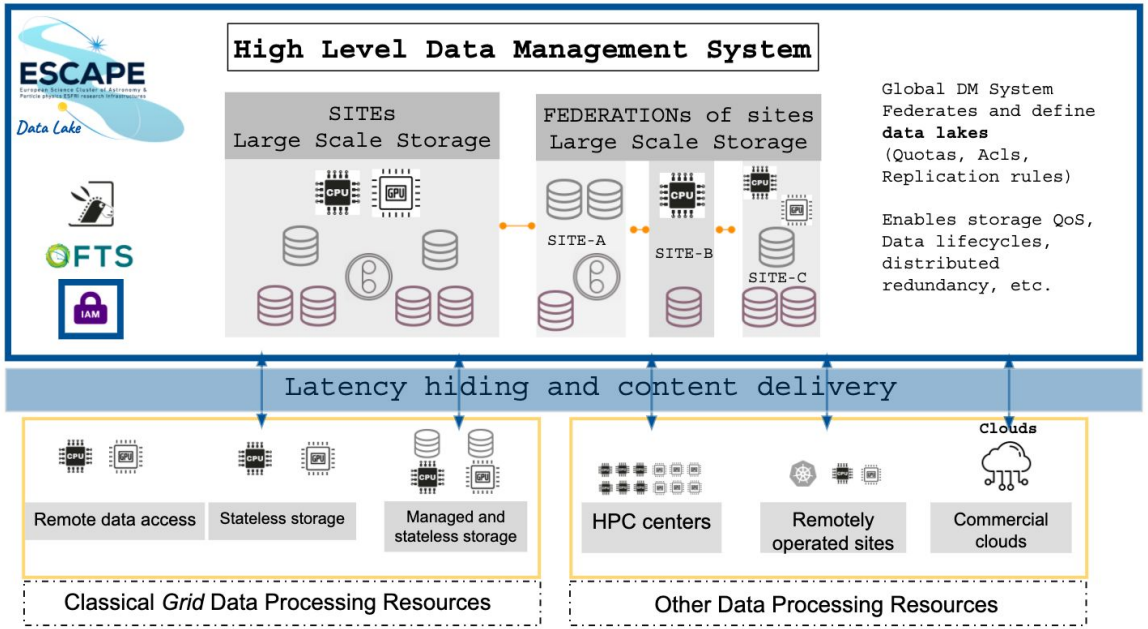

The ESCAPE Data Lake at work (Data Challenges)

Run realistic experiment use cases: f**ull data life cycles**

Several RIs implemented their **dedicated** *Data Lake*

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

Experiments successfully performed:

- **Ingestion** of raw data from data taking sources: detectors/telescopes/antennas
- Extreme distance (long haul) data management
- Rule based data replication at scale
- Data processing on several facilities, including large batch systems and user notebooks

FAIR 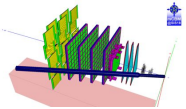	**Ingestion and replication** of simulated R3B data of mCBM detector on FAIR-ROOT. Ingestion and replication of simulated and digitised raw PANDA fallback data.
KM3NeT 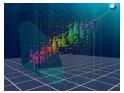	**Offload data** from the storage buffer in the coast, **replicate** across sites, run **data calibration, store** back. Data product ready for user consumption
cta cherenkov telescope array 	CTA-RUCIO @PIC: non-deterministic (La Palma) and deterministic (PIC) RSEs
HiLumi HL-LHC PROJECT ATLAS EXPERIMENT	2000 files uploaded in the Data Lake. Two copies of such files (rules) into at least two RSEs. 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 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 a multiwavelength image.**
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.

Data Lake at work for Nuclear Physics: R3B and CBM

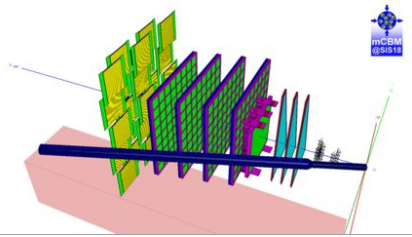

- **Registration of RAW data** acquired by the mCBM detector on FAIR-ROOT
- Ingestion and replication of **simulated R3B** data
- Ingestion and replication of **simulated and digitised raw PANDA** fallback data
- Particle-transport and digitisation of Monte-Carlo events
- **Live ingestion** of simulated data
- **Retrieval of stored RAW data** from the Data Lake, **processing** of the data and **storing** the processed data back to the Data Lake
- **Retrieve raw mCBM data** from the data lake, **run reconstruction** on it and **store the results**
- **Analyse simulated R3B data** stored in the data lake, **upload resulting histograms** and bitmaps to the DL

November 2021!

Data Lake at work for Nuclear Physics: R3B and CBM

Experiment and point of contact	Test name	Description and the running method (cron, script,...)	Method (cron script, interactive) Cron Offset	Anticipated timeline Success metric	Estimated #Files, data volume	QoS requirements	replicas / RSEs	Data access/ analysis plans?	Completed? Link to summary of each test
FAIR Marek Szuba	mCBM ingestion	Registration of data acquired by the mCBM detector on FAIR-ROOT	Python script, run as ~~an inotify-driven persistent service~~ a cron job + interactively	Raw mCBM data is successfully registered in the data lake.	TBC. Note that this will be zero-copy ingestion, i.e. storage will effectively expand as new files are added	N	FAIR-ROOT	N	Completed successfully
FAIR Marek Szuba, Maisam Dadkan	R3B ingestion	Ingestion and replication of simulated R3B data	Combination of shell and Python scripts, run interactively or by cron	• Data is successfully registered • Data is successfully replicated • The associated metadata is assigned • Data can be discovered using the configured ESCAPE-Rucio client	Main: ~300 files x ~5 GB Aux: ~300 files of negligible size # runs: 1 # output replicas: 2 Total: **4 TB**	Y	FAIR-ROOT + 1x QOS=SAFE	Y	Completed successfully
FAIR Marek Szuba, Ralf Kliemt	PANDA ingestion	Ingestion and replication of simulated and digitised raw PANDA fallback data	Shell script	Data is successfully registered	Fallback input data: 500 GB x 1 replica	Y	FAIR-ROOT, QOS=SAFE, QOS=FAST	Y	Completed successfully
FAIR Marek Szuba, Ralf Kliemt	PANDA simulation	Particle-transport and digitisation of Monte-Carlo events. Live ingestion of	Batch jobs	Data successfully processed and sent to the data lake	Output size: 1150 TB per run # runs: 1 # output replicas: 3 Total: **3.95 TB**	Y	FAIR-ROOT, QOS=SAFE, QOS=FAST	Y	Completed successfully

November 2021!

R³B
Reactions with Relativistic Radioactive Beams

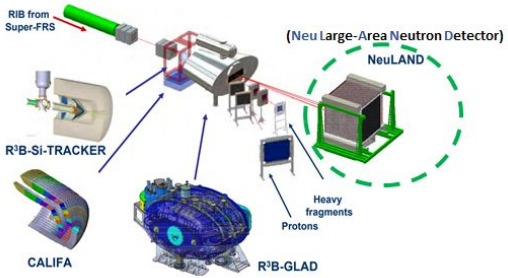

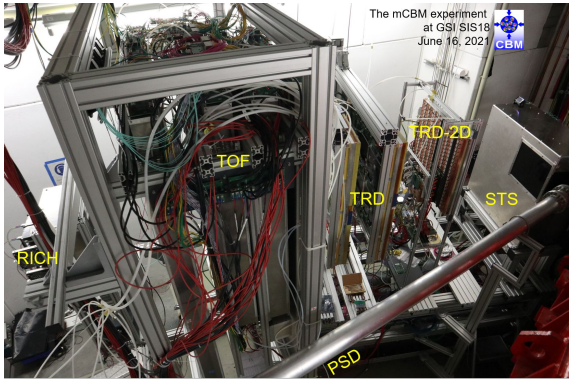

Data Lake at work for Nuclear Physics: R3B and CBM

Experiment and point of contact	Test name	Description and the running method (cron, script,...)	Method (cron script, interactive) Cron Offset	Anticipated timeline Success metric	Estimated #Files, data volume	QoS requirements	replicas / RSEs	Data access/ analysis plans?	Completed? Link to summary of each test
FAIR Marek Szuba, Eoin Clerkin	CBM simulation	Particle-transport and digitisation of Monte-Carlo events	Cron or batch jobs	Success may be determined via the successful retrieval and upload of material. Suggest hashes of datasets are stored locally at each step and compared to determine success.	Fallback input data: 1 run - ~ 3 MB x 1 replica Output size: ~5 GB/run # runs: 84 Lifetime of output data: 48h # output replicas: 1 Total: **122 GB** peak	N	FAIR-ROOT	N	Completed Successfully
FAIR Marek Szuba, Ralf Kliemt	PANDA reconstruction	Retrieval of stored RAW data from the data-lake, processing of the data and storing the processed data back to the data lake	Batch jobs	Data successfully retrieved, processed and returned to the data lake	Fallback input data: 500 GB x 1 replica Output size: 30 MB per run # runs: 84 # output replicas: 2 Total: **505 GB**	Y	FAIR-ROOT, QOS=OPPORTUNISTIC	N	Completed Successfully
FAIR Marek Szuba, Eoin Clerkin	mCBM reconstruction	Retrieve raw mCBM data from the data lake, run reconstruction on it and store the results	Cron or batch jobs	Success determined by the successful completion of event reconstruction as well as the successful upload/retrieval from data lake	Fallback input data: 10 files x 10 GB x 1 replica Output size: ~1 GB/run # runs: 84 Lifetime of output data: 48h # output replicas: 1 Total: **125 GB** peak	N	FAIR-ROOT	N	Completed Successfully
FAIR Marek Szuba, Maisam Dadkan	R3B data analysis	Analyse simulated R3B data stored in the data lake, upload resulting histograms and bitmaps to the DL	JupyterHub (kernels: Python, possibly C++)	• Successful authn/authz for user • Gain access to the data lake and perform the analysis • Write results back to the data lake	Failsafe input data: 1 run - 100 GB x 2 replicas (max.) Analysis output: ~10 GB per run # runs: 30 # output replicas: 2 Total: **201 GB**	Y	*FAIR-ROOT*, QOS=CHEAP-ANALYSIS, QOS=FAST, QOS=SAFE	N	Was not able to download files from DL using OIDC. completed successfully using x509.

November 2021!

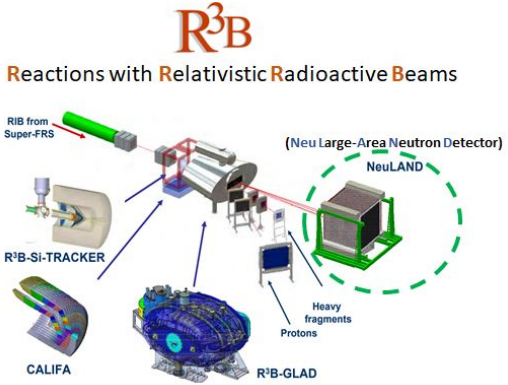

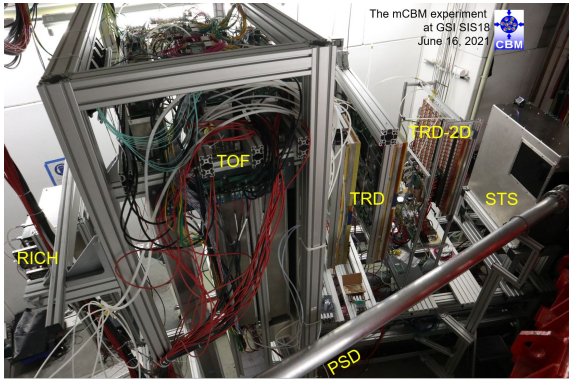

Data Lake enabling user analysis: **The Virtual Research Environment**

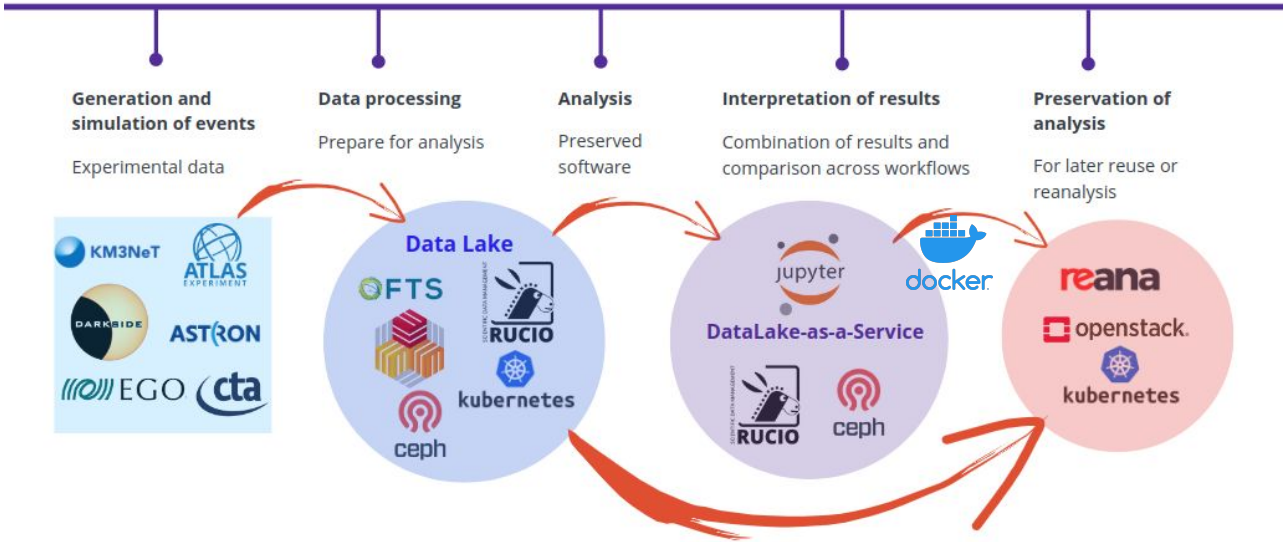

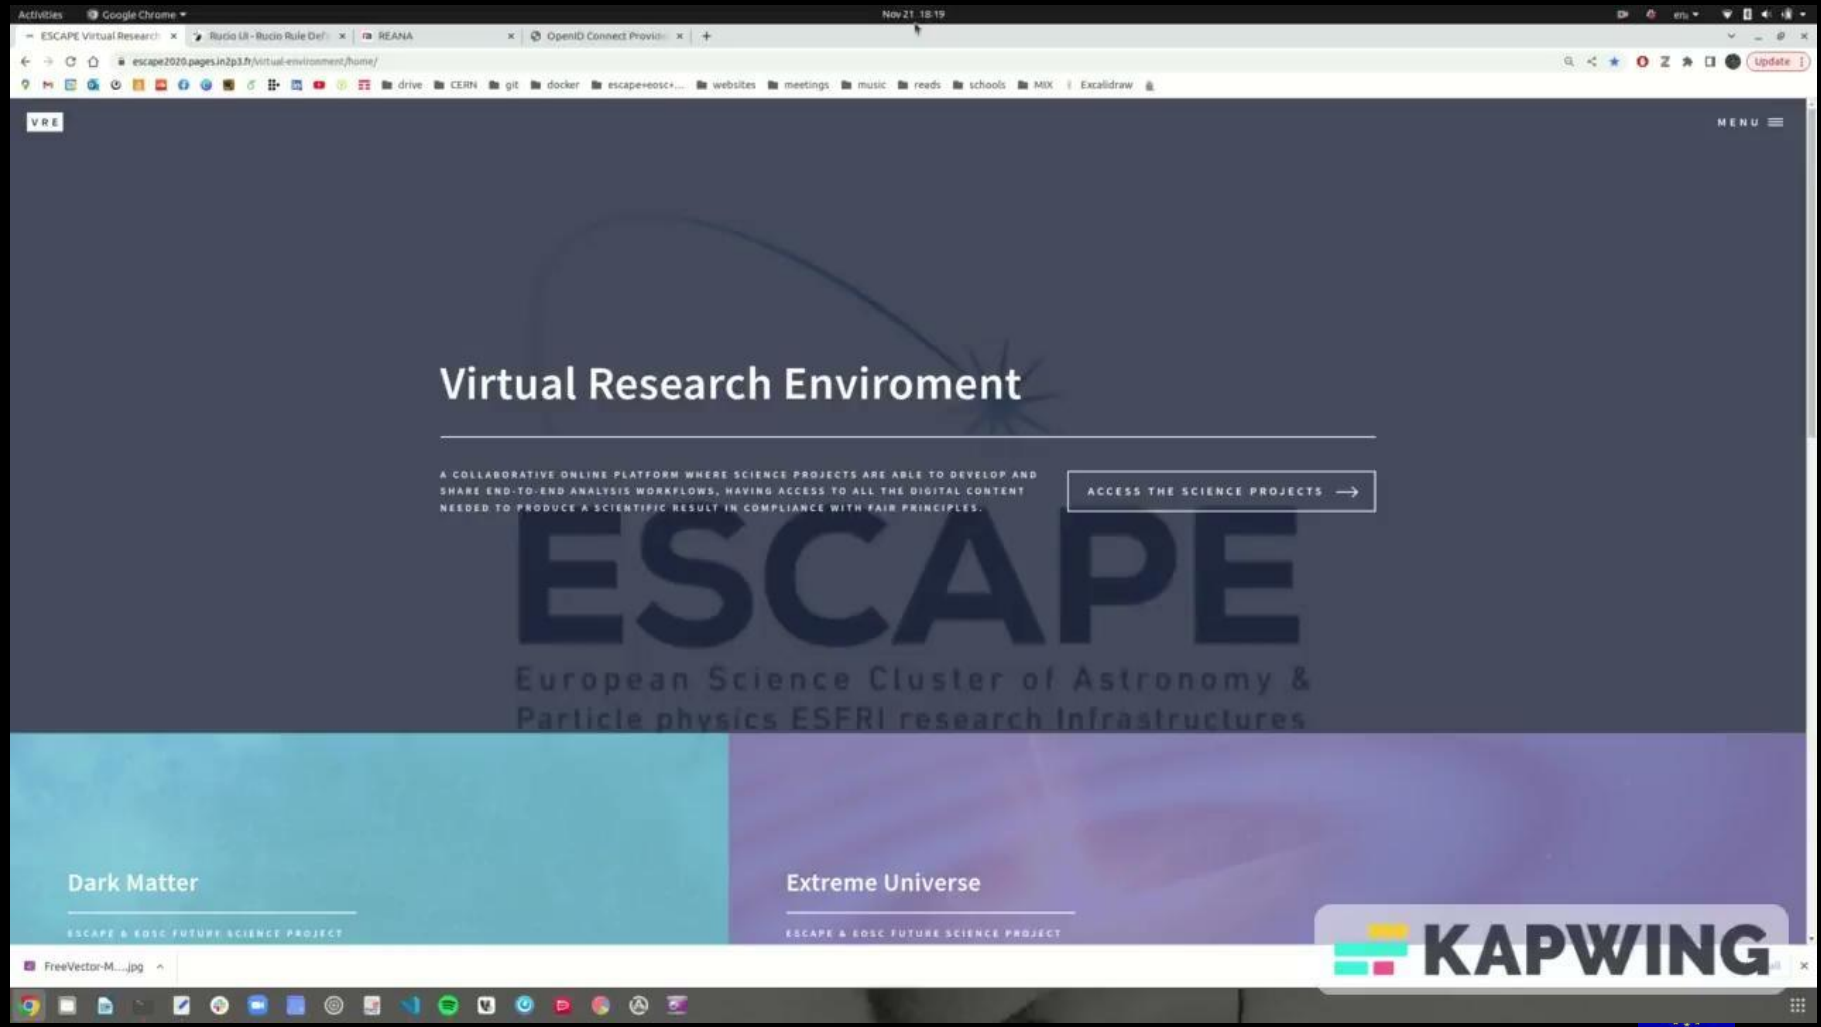
Activities
Google Chrome
Nov 21 18:19
ESCAPE Virtual Research
REANA
escape2020.pages.in2p3.fr/virtual-environment/home/
VRE
MENU
Virtual Research Enviroment
A COLLABORATIVE ONLINE PLATFORM WHERE SCIENCE PROJECTS ARE ABLE TO DEVELOP AND SHARE END-TO-END ANALYSIS WORKFLOWS, HAVING ACCESS TO ALL THE DIGITAL CONTENT NEEDED TO PRODUCE A SCIENTIFIC RESULT IN COMPLIANCE WITH FAIR PRINCIPLES.
ACCESS THE SCIENCE PROJECTS
ESCAPE
European Science Cluster of Astronomy & Particle physics ESFRI research Infrastructures
Dark Matter
ESCAPE & EOSC FUTURE SCIENCE PROJECT
Extreme Universe
ESCAPE & EOSC FUTURE SCIENCE PROJECT
KAPWING

A new ESCAPE Collaboration Agreement

Brussels, 26 October 2022

After the successful experience of the H2020 ESCAPE Project, the nine partner Research Infrastructure Directorates have signed a new "**Open Collaboration Agreement**", which consolidates their action for the benefit of Open Science, the implementation of the EOSC and the establishment of new sustainable cooperative schemes within Horizon Europe, the European Strategy for Data and Excellence Science.

Paolo Giubellino (FAIR DG) as a member of ESCAPE is one of the parties of the new collaboration

New ESCAPE Open Collaboration Agreement

- It will continue to work as a **domain-based Science Cluster** and open to new RIs
- **Inclusive** potential partnership, embracing the concerned scientific communities at large at pan-European level
  - Giovanni Lamanna statement during one recent NuPPEC general assembly: "*the ESCAPE collaboration is potentially open to nuclear physics world-class RIs*"
- Its parties will convene their **national partner institutes** to join and support the ESCAPE work programme
- It will leverage the existing **international bridges** and consortia to operate more globally - many RIs are international.
- The **external advisory board** played and will play a key role in supporting ESCAPE, guaranteeing the link with JENA (Joint ECFA-NuPECC-ApPEC), national institutes at large and the concerned community

Backup slides

Virtual Research Environment

Video/Demo workflow:

- VRE website to showcase science projects
- Look up the code on gitlab
- Log into the jupyterhub via Docker images to rerun the analysis
- Rucio client installed, and Rucio jupyterlab extension enabled to browse through scopes and files
- The cloud S3 (EOSC resources) RSE has been added to the Rucio Data Lake and the files are uploaded there (using the ESCAPE IAM authentication
- The correct state of uploaded files can be checked via the Rucio webUI interface
- The preliminary analysis can be run in notebook cells – see the plots
- Otherwise, the analysis - whenever computationally heavier - can be dispatched to multiple nodes of a cluster that has been spawned on EOSC resources at CNAF bologna (by interacting with the CNAF Openstack thanks to the integration between egi oidc and ESCAPE IAM)
- Reana web interface to follow up the workflow running
- Documentation on onboarding can be found on the VRE webpage
- Files are fetched from Rucio to the Reana workspace (object storage close to all nodes), analysis is run across nodes and the plots are produced.

Funded by the European Union's
Horizon 2020 - Grant N° 824064

Content delivery and latency hiding

- Streaming caches demonstrate potential on **latency hiding and file re-usability** in Particle Physics workflows
- Understanding whether caching can also help on non-event based formats, e.g. images, data-cubes,...
- Caches can facilitate ingress/egress of data with heterogeneous computing resources: **Commercial Clouds and HPCs**
- Effort made towards a vanilla installation (experiment-unbiased) caching service
  - Based on XCache technology
  - Easy deployable by the partners

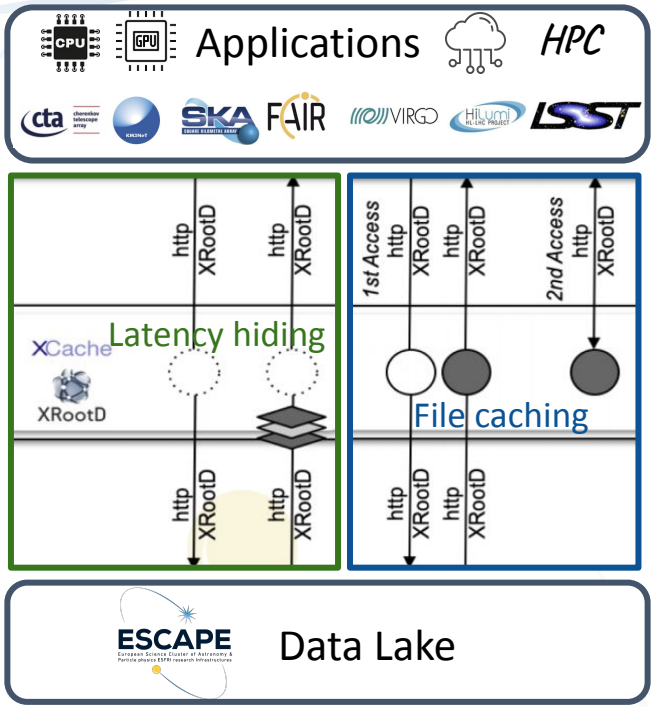

Funded by the European Union's Horizon 2020 - Grant N° 824064

Data Lake integration with notebooks (1/2)

- Developed a JupyterLab extension to **enable data access** from a notebook platform
- Rucio JupyterLab Extension: **notebook datalake integration**
- **Browse/download/replicate** datalake data from the notebook sidebar
- Remote storage fuse mount, XCache integration and multi VO support
- Technology being implemented: fruitful joint work with WP5

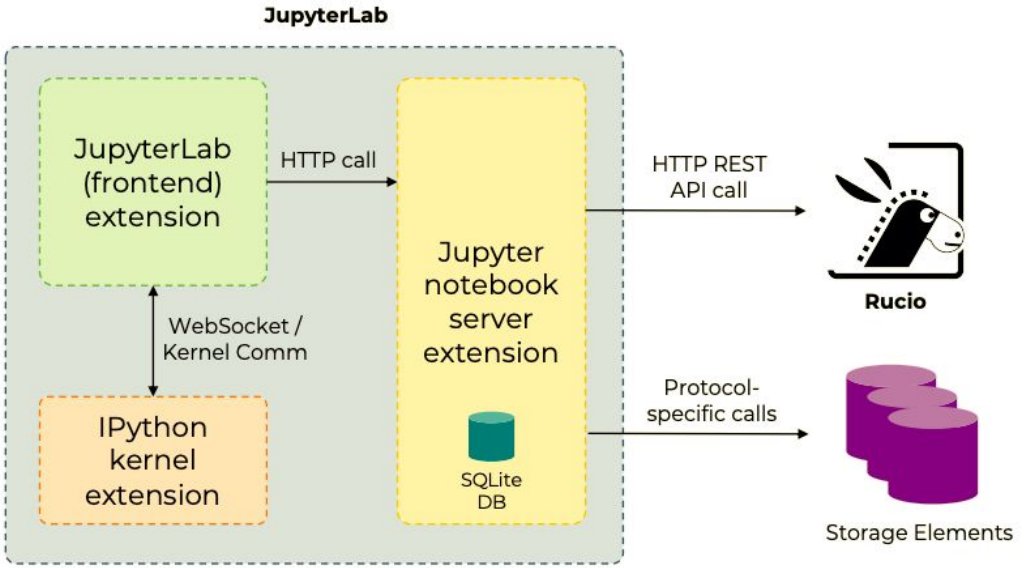

Funded by the European Union's Horizon 2020 - Grant N° 824064