

ESCAPE Data Lake

Next-generation management of cross-discipline Exabyte-scale scientific data

Riccardo Di Maria^{1,*} and Rizart Dona^{1,**} on behalf of the ESCAPE project

¹European Organization for Nuclear Research (CERN), Geneva, Switzerland

Abstract. The European-funded ESCAPE project (Horizon 2020) aims to address computing challenges in the context of the European Open Science Cloud. The project targets Particle Physics and Astronomy facilities and research infrastructures, focusing on the development of solutions to handle Exabyte-scale datasets. The science projects in ESCAPE are in different phases of evolution and count a variety of specific use cases and challenges to be addressed.

This contribution describes the shared-ecosystem architecture of services, the Data Lake, fulfilling the needs in terms of data organisation, management, and access of the ESCAPE community. The Pilot Data Lake consists of several storage services operated by the partner institutes and connected through reliable networks, and it adopts Rucio to orchestrate data management and organisation. The results of a 24-hour Full Dress Rehearsal are also presented, highlighting the achievements of the Data Lake model and of the ESCAPE sciences.

1 Introduction

The main goal of the European Science Cluster of Astronomy & Particle physics ESFRI research infrastructures (ESCAPE) project [1, 2] is to address Open Science challenges in Astronomy and Particle Physics. These challenges are shared by European Strategy Forum on Research Infrastructures (ESFRI) facilities [3] - i.e. CTA, ELT, EST, FAIR, HL-LHC, KM3NeT, SKA - and pan-European research infrastructures [4] - i.e. CERN, ESO, JIVE-ERIC, EGO-Virgo - for which ESCAPE aims at developing and delivering solutions to ensure optimised access, management, and organisation of the enormous amount of data handled by the partners.

The objectives of such solutions are threefold and in support of Findable, Accessible, Interoperable, and Reusable (FAIR) data principles [5]: (1) connect ESFRI projects to the European Open Science Cloud (EOSC) [6] ensuring integration of tools and data; (2) support shared approaches for open data management; (3) establish interoperability within EOSC as an integrated multi-messenger facility for fundamental science. The diverse expertise in data management and computing of Astrophysics and Particle Physics ESCAPE communities had a crucial role in setting up an infrastructure that is beyond the current state-of-the-art: the Data Lake.

*e-mail: riccardo.di.maria@cern.ch

**e-mail: rizart.dona@cern.ch

This paper describes the infrastructure of the ESCAPE Pilot Data Lake supporting Open Science for the ESFRIs being part of the project. The overall architecture and the various functional elements are detailed, and the achievements and challenges for this Pilot phase, as well as plans towards the forthcoming Prototype phase, are illustrated.

2 The Pilot Data Lake

The Data Lake is a modular ecosystem of services and tools shaped around the ESCAPE scientific communities. Available existing implementations of each functional element were analysed and adapted according to ESFRI partners' needs, identifying the most suitable options. The adopted open-source community-supported products have required specific R&D effort to cover all ESCAPE use cases, even though EU-funded initiatives, e.g. INDIGO-DataCloud [7], AARC/AARC2 [8] and FIM4R [9] projects, have been leveraged extensively when possible.

The mandate for the Pilot phase was to demonstrate at a functional level the choices made in terms of architecture and reference implementations, assessing core functionalities of data transfer, access, and storage, and of the authentication and authorisation model.

The upcoming Prototype phase will enlarge the Pilot to all partner sites and expose the various functionalities to the science projects. Real use cases and applications from ESFRIs will drive scalability tests, and fully prove that the ESCAPE Data Lake implementation can serve as an infrastructure for FAIR data.

2.1 Architecture

Different functional elements constitute the architecture of the Data Lake, designed to provide a flexible and scalable infrastructure in terms of data organisation, management, and access. The components and services, discussed in the following sections, are also available to ESCAPE sciences as stand-alone services, modulo R&D effort required to integrate them into external specific environments.

Storages

Storage services are the core of the infrastructure and they are hosted at different facilities which are diverse in size and expertise. As a consequence, different classes of service are provided, from large national centres operating archive storage to smaller disk-based systems. The Data Lake counts dCache [10], DPM [11], EOS [12], StoRM [13], and XRootD [14] as software technologies for its storage services, using diverse backends such as classic RAID systems, Erasure Coding, CEPH, and multi-replica. The reliability and capability of such solutions at hundred-Petabyte scale have been demonstrated by their decade deployment in the Worldwide LHC Computing Grid (WLCG) infrastructure [15].

ESCAPE storages, in this context also known as Rucio Storage Elements (RSEs), are 15 and are spread across Europe as shown in Figure 1. The total quota and watermark (WM) these RSEs contribute to the Data Lake are 821.2 TB and 706.7 TB, respectively, where WM indicates the maximum amount of space that should be occupied. The WM is set per-RSE at discretion of the site itself. This allows scenarios such as the utilisation of an RSE for data injection buffering where the WM is enforced to be lower than 1% of the total storage capacity. Specifically, MAGIC experiment [16] pioneered this controlled data injection model in ESCAPE during the 24-hour assessment of the Data Lake as described in the following, whose outcome proved to be highly valuable for the other scientific communities. The DPM

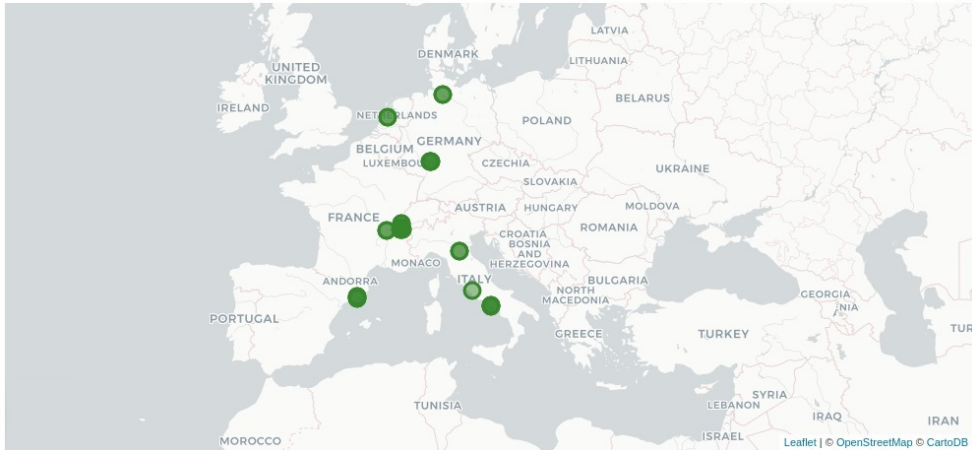


Figure 1. The RSEs distribution of the Pilot Data Lake from ESCAPE monitoring.

storage solution is also used in federated and geographically-distributed data centres under a unique storage namespace for specific ESCAPE institutes, offering a single user entry point.

Data are accessed through three protocols, depending on the implementation at site level: gridFTP, HTTP(S), and xroot. The latter supports predominantly High Energy Physics (HEP) use cases since files are stored in ROOT format, whereas gridFTP is present for legacy reasons. The HTTP protocol is widely adopted within the open-source community and covers the vast majority of ESCAPE science projects. These protocols offer download, upload, and streaming file functionalities, as well as third-party transfers and deletion. X.509 is the standard authentication method for the Pilot phase, even though a proof-of-concept (PoC) for token-based authentication support has been carried out as a precursor of the Prototype stage objective.

The varied availability of storage systems and disk backend configurations in the ESCAPE Data Lake represents the perfect playground for the implementation of storage Quality of Service (QoS) intelligence. Consequently, experiments can optimise the use of storage and can define specific data life-cycles. A first PoC implementation of cost and performance optimised storage usage results in different QoS denominations for each storage element in the Data Lake. The metrics under consideration and the related set of capabilities at storage level are evolving throughout the ESCAPE mandate, but are summarised in the reliability, performance, and cost categories, including their combinations.

Data Orchestration

The open-source data management system Rucio [17] represents the reference implementation of the ESCAPE storage orchestration service, whose core consists of main and authorisation servers in redundancy deployment. Together with other Rucio components, they interact with an Oracle database hosted at CERN in a shared production environment, consequently providing backup in case of major disruption of services. APIs for data management and access hides the complexity of the machinery to the end-user, for which the web user interface service `escape-rucio.cern.ch` is also deployed.

Data management is driven by user-defined policies named rules which are enforced through asynchronous agents, or daemons in Rucio jargon. Data injection, organisation,

replication, and deletion follow these rules by performing operations on Rucio data objects, addressed with Data IDentifiers (DIDs) and composed of a scope and a name. ESCAPE Rucio exploits a specific subset of daemons, whose redundancy and working-threads are designed to fulfil the Data Lake needs. These agents are responsible for:

- updating the usage of accounts, collection replicas, and RSEs;
- managing file transfers by preparing and submitting jobs, checking their status, and updating Rucio internal state for finished transfers;
- delivering messages via STOMP to a message broker;
- managing replication rules by asynchronously injecting, executing, evaluating, repairing, and cleaning, as well as creating rules for DIDs matching a subscription;
- replica deletion and management of expired DIDs.

FTS [18] and GFAL library [19] are used for file transfer and upload/deletion, respectively, and the system relies on gridFTP, xroot, and HTTP protocols.

The CERN Agile Infrastructure store, organise, and expose events through customised dashboards to ESCAPE partners, exploiting the Rucio built-in messaging capabilities. A series of automated Rucio-level tests have been designed and integrated as health sanity checks of the Data Lake. They are responsible for uploading, downloading, replicating, and deleting a range of different sized files, i.e 1 MB, 10 MB, 100 MB, 1 GB, and 5 GB, at a different rate for all ESCAPE RSEs. Detailed performance studies and operational debugging are carried out using these monitoring data, which have played a major role in the improvement of the infrastructure.

This specific Rucio implementation resulted in a fundamental contribution to the improvement of several components in Rucio itself due to the variety of needs of the ESCAPE science projects. Several ESRIs are widely adopting it as a consequence of its ability to organise data in collections and to manage them at a file granularity-level. Moreover, additional R&D work is being performed for the evolution of storage QoS and for the transition to token-based authentication, representing two of the main objectives of the upcoming Prototype phase.

File Transfer Service, Networking, and Information System

The decade history of the File Transfer Service (FTS) in WLCG demonstrated scalability and performance features required in the context of ESCAPE. Moreover, it exploits mechanisms to avoid congestion and optimise the usage of transfer bandwidth and network. This resulted in FTS as the reference implementation for an asynchronous transfer scheduling service, ensuring data movement across storages and triggering third-party-copy for various protocols and multiple files in parallel. However, the transition from X.509 to tokens is crucial for improved granularity and flexible management [20], and, as transferred data greatly differ in size, it will reduce authentication overheads affecting performance at Kilobyte-scale.

A continuous testing and monitor machinery has been deployed to inspect the health of the technology stack in the Data Lake. As Rucio delegates file transfer management to FTS, a health-check on this service has been designed and deployed. Automated FTS-level tests run continuously, providing comprehensive information to the common ESCAPE monitoring infrastructure. The Pilot phase has profited from the existing FTS monitoring for third-party transfers and the capabilities of its data analytic backend, and dashboards have been developed for information to be exposed to the different ESRIs. However, it has to be taken into account that Rucio-triggered transfers also appear in the FTS dashboard as mediated by FTS itself. Since Rucio and FTS rely on GFAL, an automated GFAL-based tool additionally runs every minute to test upload, download, and deletion on the ESCAPE endpoints, using all

protocols supported per site and not relying on third-party management. This testing suite of FTS and GFAL, together with Rucio-level tests, ensure a quasi-real-time view of the load and performance of the Data Lake, necessary to identify and investigate any potential issue.

The health of network links between sites is monitored by point-to-point transfers and latency tests using the existing perfSONAR infrastructure [21]. Automated transfers probe packet-loss and throughput between ESCAPE perfSONAR instances, and results are routed into ESCAPE dashboards.

Finally, the interoperability of the Data Lake services is ensured by the Compute Resources Information Catalog (CRIC) [22], a central catalogue containing services information and configuration for Rucio, e.g. supported protocols at sites and their prioritisation.

Authentication and Authorisation

Due to the distributed and heterogeneous nature of the infrastructure, a major role is played by the ESCAPE Authentication, Authorisation, and Identity (AAI) layer, based on the INDIGO Identity and Access Management (IAM) service, to manage data access in the Data Lake. As an evolution of VOMS [23], this technology leverages 15-year experience in WLCG and it has demonstrated the required flexibility and security to serve the purpose of multi-Petabyte scale experiments. The ESCAPE IAM instance is deployed on a Kubernetes [24] cluster at INFN-CNAF, supports authentication also through EduGAIN, and implements user registration and management, as well as token and X.509/VOMS attribute provisioning services, acting as a community-scoped authorisation service. This choice accommodates the needs of ESFRI partners and policies for embargoed data, ensuring the proper level of security for various potential scenarios.

Data Access, Content Delivery and Caching

Data access, processing, and content delivery services are not extensively covered by the Pilot phase. However, several PoCs and implementations have been deployed at CC-IN2P3, CERN, GSI, and INFN to prove the validity of latency-hiding and caching solutions. These solutions support X.509 and token-based authentication and authorisation, and address access for both embargoed and open data.

The Data Lake model implies processing resources not necessarily co-located with data. Remote data access represents a challenge at Exabyte-scale that is being studied during the Prototype stage, together with data serving to heterogeneous resources. Caching solutions, in this context based on XRootD-XCache technology, unfold several possibilities to enable data access from non-standard resources, such as opportunistic resource, commercial clouds, or High-Performance Computer centres, for which storage capacity may be limited. A content delivery layer ensures buffering to allow local-like data processing, and caching most-accessed data close to the computing resources for future re-use could optimise access performance and application efficiency.

A vanilla implementation, and following more complex versions based on Docker and Puppet, were tailored to ESFRIs' needs and deployed as multi-experiment service across the Data Lake. A qualitative assessment of XCache has been performed during this stage of the project by running experiment-specific workflows, prominently from ATLAS [25] and CMS [26], resulting in promising quantitative future studies. In fact, the full assessment of the latency-hiding and caching layer is being addressed in the Prototype phase.

Synchronisation Services

A set of synchronisation tools has been designed and deployed to run every 2 hours in order for Rucio to pull necessary information from other ESCAPE services. These include IAM and CRIC, which provides Rucio with end-user authentication and authorisation profiles and RSE information and configuration, respectively. The IAM-Rucio synchronisation procedure is also used to produce authentication grid-map files for specific RSEs. In particular, the ESCAPE CERN endpoint profits from this method to ensure a correct mapping between ESCAPE users and EOS users for local access.

2.2 Deployment Model and Techniques

The necessity to prove resource awareness for infrastructure deployment is a fundamental point for the project sustainability and for the successful exportability of the proposed model and solutions to ESFRI partners. For this reason, it was crucial to demonstrate that either the entire infrastructure or specific functional elements willing to be adopted by ESCAPE experiments could cope with diverse computing models or resource availability at their disposal. A 24-month evaluation of the Data Lake demonstrated the feasibility of deploying the key components of the ESCAPE infrastructure in a Kubernetes cluster of solely 52 vCPUs, 104 GiB RAM, and 460 GB local storage.

The Kubernetes cluster runs on the CERN OpenStack [27] infrastructure, consisting of a master of 4 vCPUs, 8 GiB RAM, and 40 GB local storage, and 6 nodes each of 8 vCPUs, 16 GiB RAM, and 70 GB local storage. The choice of Kubernetes over bare metal or Puppet-managed implementation has been proven valid in terms of flexibility and reproducibility for the adoption by the ESFRIs. In fact, the cluster size can be adjust in case of a consistent variation of the load, and the base images of the key services are shared among ESCAPE experiments. The cluster hosts the entire testing suite, the synchronisation services, and the Rucio components described in the previous sections. Moreover, Filebeat logging agents send operational data of each node to ESCAPE monitoring via Logstash which is responsible for their aggregation and enrichment. Finally, the cluster runs a fully-flagged Rucio admin client with root privileges for ESCAPE Data Lake management purposes.

The aforementioned choice of technology allowed to deploy a manifold system utilising a lightweight but complete implementation. This contribution within EOSC aims to facilitate partnered sciences to deploy part or whole service stack according to their needs.

3 Full Dress Rehearsal

The Data Lake is available 24/7 to ESCAPE users, even though limitations and unforeseen down-times are expected by the non-production nature of the infrastructure. However, a specific testing-focused time-window has been planned and allocated to qualitatively mimic daily activity of ESCAPE sciences and to prove the robustness of the various Data Lake components, fulfilling the different requirements of the experiments' use cases.

A 24-hour assessment of the Data Lake, also know as Full Dress Rehearsal (FDR) exercise, has been performed involving activities from 9 ESCAPE sciences: ATLAS, CMS, CTA [28], EGO-Virgo [29], FAIR [30], LOFAR [31], LSST [32], MAGIC, and SKA [33]. As these experiments are at diverse phases of evolution with respect to computing model and data management, fortnightly assemblies for data injector demonstrators have been carried out with the goal to even out the differences in knowledge among them. The aim was both to familiarise with Rucio as an end-user and to raise awareness of the technology stack involved in the data movement, acquiring know-how on management and utilisation of the various

parts. The latter, in particular, was critical in view of the successful adoption and operation of such technologies by ESFRIs, following the term of the ESCAPE project mandate. During these sessions, a significant variety of specific use cases and challenges were addressed for tailored scenarios presented by each experiment. The joint initiative covered data injection and access, as well as the exploitation of QoS addressing requested life-cycles of data, both from the end-user and the experiment point of view. The planning of the activities involved also ESCAPE site-admins, responsible for the progressive improvement of the storage services' reliability and functionalities.

Furthermore, a deployment and operation team has been gathered in preparation for the assessment period to identify and solve problematic situations across the Data Lake by exploiting the ESCAPE dashboards. The common monitoring infrastructure was the keystone for the remote resolution of issues and for the support to both end-users and experiments.

The preparation period and the FDR itself exposed a few challenges with respect to the ESCAPE deployment and infrastructure. The identification and resolution of such were among the main successes and contributions of the community to the various technology stack that constitute the Data Lake.

The Oracle database used by Rucio is a critical pillar around which all the other services orbit, and for this reason, it is vital to ensure that its dependencies are properly addressed by all other components. The upgrading from a development to a production database instance, necessary to provide backup in the event of major services disruption during the FDR exercise, uncovered inconsistencies that were promptly taken into account by the Rucio community. Moreover, the Rucio daemon responsible to update the usage of collection replicas underwent extensive use only within the Data Lake, as other experiments utilise specific database procedures for this purpose. The maturity of this asynchronous agent was not sufficient for the scope, and provoked numerous issues in the database. In fact, it was designed to load the entire related table from the database during the cleaning procedure, saturating the available 16 GiB RAM memory of the underneath node if millions of rows were to be cleaned. However, ESCAPE significantly contributed to the resolution of this problematic as a result of the collaboration and synergy established with the Rucio developers.

Various use cases presented by ESCAPE sciences were extraneous to HEP experiments, and consequently to Rucio. As an example, the SKA experiment attempted a one-million files rule injection that has yielded an improvement of the standard Rucio deployment initially implemented in ESCAPE. In fact, the memory management of the daemons responsible for managing replication rules and for replicas deletion, as well as the resources pledged for these agents, were not designed to cope with this corner case. A joint effort between Rucio and ESCAPE developers has been carried out to implement the necessary changes on-the-fly, contributing to the success of these tests after an acceptable delay and without any failure.

The main achievement of the FDR was to have 9 different experiments working on and testing the shared ecosystem of services and tools with a variety of use cases and data workflows that greatly differ from one science to another. During this 24-hour assessment, about 1 million files totalling 20 TB and 10 TB of data were injected and transferred, respectively, by the various experiments activities in the Data Lake. The gauge for these numbers should take into account that the majority of ESCAPE partners have been building up their computing models and the knowledge of data management during the project mandate, consequently framing the FDR as a highly successful exercise.

SKA mimicked a 24-hour data-taking of observations of pulsar and single-pulse searches, pioneered a one-million files rule injection, and exploited QoS intelligence for data movement. EGO-Virgo reproduced 4 kHz-sampled data injection and access, whereas FAIR performed continuous data ingestion using several QoS classes. A real MAGIC observation

scenario, that expects a storage to be co-located to the telescope, was also mimicked by using an RSE as a buffer for subsequent data injection to the Data Lake. CTA uploaded over half-million files to simulate a night-data-captured from the telescope in Canary Island, and LOFAR reproduced an image of the astronomical radio source 3C196 by processing data already present in the Data Lake as shown in Figure 2.

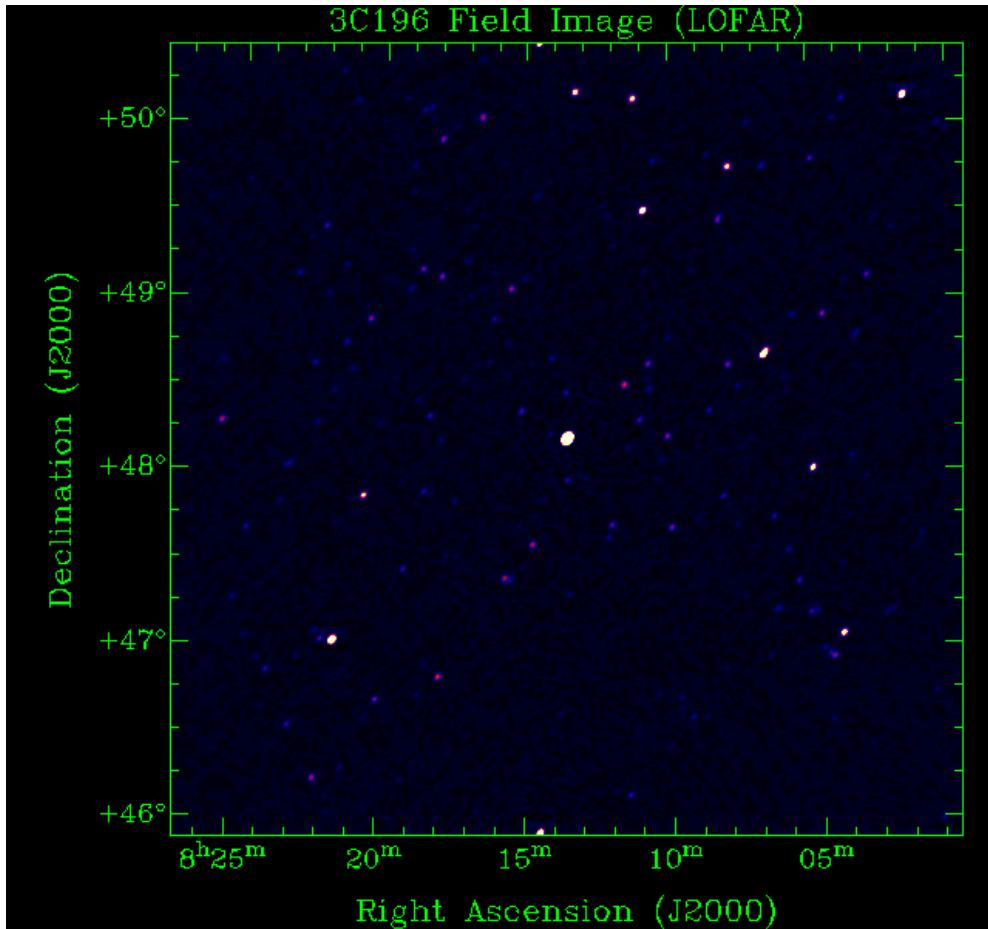


Figure 2. The image of the astronomical radio source 3C196 reproduced by using LOFAR data present in the Data Lake.

LSST injected more than 150k files to simulate production conditions, i.e. 20 TB per day, cross-checking integrity and accessibility of data via JupyterLab-Notebook. For the latter part, the integration of Rucio in JupyterLab, developed within ESCAPE, has been used, providing the accessibility of the Data Lake also for scenarios involving notebooks. Finally, ATLAS and CMS focused on the usage of open data and their life-cycles supported by the QoS intelligence. Moreover, the flagship ATLAS Higgs-to-photons analysis has been reproduced using open data on a notebook, profiting from the JupyterLab-Rucio integration.

Flagship experiments from Astro-particle and Particle Physics, Electromagnetic and Gravitational-Wave Astronomy, and Nuclear Physics found in the ESCAPE Data Lake a successful playground to address data management challenges and pursue FAIR data principles.

4 Prototype Phase and Future Challenges

As the Pilot stage successfully engaged an unforeseen number of partner sites, the Prototype phase aims to consolidate the Data Lake functionalities and scale up the performance of the different services involved. The goal of such ecosystem is to ensure the needs of the ESFRI partner projects in terms of functionality, performance, and usability.

The orchestration of the cloud-like storage federation should involve data injection and distribution at GB/s scale, where the net flow targets the order of Terabytes per day. The provisioning of different storage QoS should ensure the transition across several layouts according to the data life-cycle expectations of the experiments. The infrastructure should mostly rely on token-based authentication and authorisation, whose integration is expected for all services constituting the Data Lake. A common set of monitoring tools should provide a detailed picture of the status, usage, and activity within ESCAPE, including new experiment-specific functionalities. Stress testing based on representative workloads from experiments should assess data access and processing, also involving heterogeneous facilities such as HPC, and a content delivery and caching layer based on XCache should prove the validity of latency-hiding approach at scale. Moreover, event-based data processing should be demonstrated through a per-storage messaging system for file-level information.

5 Conclusions

The modular ecosystem of services and tools constituting the Data Lake concluded its Pilot phase beyond expectations, and has entered the Prototype stage. The FDR demonstrated the robustness of the choice made with respect to the reference implementations of the various technologies. Moreover, science projects in ESCAPE at different phases of evolution, regarding data management and distributed computing knowledge, underwent a profound even-out of their inequalities. The on-going phase of the project aim at consolidating the services functionalities and at improving the infrastructure, such as integrating token-based authentication and authorisation or deploying a tailored content delivery and caching layer.

The ESCAPE Data Lake has been proven a successful scientific data management infrastructure where flagship ESFRIs in Astro-particle Physics, Electromagnetic and Gravitational-Wave Astronomy, Particle Physics, and Nuclear Physics pursue together the FAIR and open-access data principles. This represents a fundamental achievement under both sociological and computing model aspects for different scientific communities that should address upcoming data management and computing challenges in the next decade.

Acknowledgements

Authors acknowledge support from the ESCAPE project. This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 824064 (ESCAPE, the European Science Cluster of Astronomy & Particle physics ESFRI research infrastructures).

References

- [1] R. Bolton, S. Campana, A. Ceccanti, X. Espinal, A. Fkiaras, P. Fuhrmann, Y. Grange, EPJ Web Conf. **245**, 04019 (2020)
- [2] *ESCAPE Project*, <https://projectescape.eu> (2021), accessed: 2021-02-27
- [3] *ESFRI*, <https://www.esfri.eu> (2021), accessed: 2021-02-27

- [4] *EU Research Infrastructures*, https://ec.europa.eu/info/research-and-innovation/strategy/european-research-infrastructures_en (2021), accessed: 2021-02-27
- [5] *The FAIR Guiding Principles for scientific data management and stewardship*, <https://www.nature.com/articles/sdata201618> (2021), accessed: 2021-02-27
- [6] *EOSC Portal*, <https://eosc-portal.eu> (2021), accessed: 2021-02-27
- [7] *INDIGO-DataCloud: a Platform to Facilitate Seamless Access to E-Infrastructures*, <https://link.springer.com/article/10.1007/s10723-018-9453-3> (2021), accessed: 2021-02-27
- [8] *AARC/AARC2 Project*, <https://aarc-project.eu> (2021), accessed: 2021-02-27
- [9] *FIM4R*, <https://fim4r.org> (2021), accessed: 2021-02-27
- [10] P. Fuhrmann, V. Gülzow, *dCache, Storage System for the Future*, in *Euro-Par 2006 Parallel Processing*, edited by W.E. Nagel, W.V. Walter, W. Lehner (Springer Berlin Heidelberg, Berlin, Heidelberg, 2006), pp. 1106–1113
- [11] A. Alvarez, A. Beche, F. Furano, M. Hellmich, O. Keeble, R. Rocha, *Journal of Physics: Conference Series* **396**, 032015 (2012)
- [12] A.J. Peters, L. Janyst, *Journal of Physics: Conference Series* **331**, 052015 (2011)
- [13] A. Carbone, L. dell’Agnello, A. Forti, A. Ghiselli, E. Lanciotti, L. Magnoni, M. Mazzucato, R. Santinelli, V. Sapunenko, V. Vagnoni et al., *Performance studies of the StoRM Storage Resource Manager*, in *Third International Conference on e-Science and Grid Computing, e-Science 2007, 10-13 December 2007, Bangalore, India* (IEEE Computer Society, 2007), pp. 423–430
- [14] *XRootD*, <https://xrootd.slac.stanford.edu> (2021), accessed: 2021-02-27
- [15] *WLCG*, <https://wlcg.web.cern.ch> (2021), accessed: 2021-02-27
- [16] E. Lorenz, M. Martinez, *Astronomy & Geophysics* **46**, 6.21 (2005)
- [17] M. Barisits, T. Beermann, F. Berghaus, B. Bockelman, J. Bogado, D. Cameron, D. Christidis, D. Ciangottini, G. Dimitrov, M. Elsing et al., *Computing and Software for Big Science* **3**, 11 (2019)
- [18] A.A. Ayllon, M. Salichos, M.K. Simon, O. Keeble, *Journal of Physics: Conference Series* **513**, 032081 (2014)
- [19] *GFAL2*, <https://dmc-docs.web.cern.ch/dmc-docs/gfal2/gfal2.html> (2021), accessed: 2021-02-27
- [20] B. Bockelman, A. Ceccanti, I. Collier, L. Cornwall, T. Dack, J. Guenther, M. Lassnig, M. Litmaath, P. Millar, M. Sallé et al., *EPJ Web Conf.* **245**, 03001 (2020)
- [21] *perfSONAR*, <https://www.perfsonar.net> (2021), accessed: 2021-02-27
- [22] A. Anisenkov, J. Andreeva, A. Di Girolamo, P. Paparrigopoulos, A. Vedaae, *EPJ Web Conf.* **214**, 03003 (2019)
- [23] *VOMS*, <https://doi.org/10.5281/zenodo.1875371> (2021), accessed: 2021-02-27
- [24] *Kubernetes*, <https://kubernetes.io> (2021), accessed: 2021-02-27
- [25] ATLAS Collaboration, *JINST* **3**, S08003 (2008)
- [26] CMS Collaboration, *JINST* **3**, S08004 (2008)
- [27] *OpenStack*, <https://www.openstack.org> (2021), accessed: 2021-02-27
- [28] M. Actis, G. Agnetta, F. Aharonian, A. Akhperjanian, J. Aleksić, E. Aliu, D. Allan, I. Allekotte, F. Antico, et al., *Experimental Astronomy* **32**, 193–316 (2011)
- [29] F. Acernese et al. (VIRGO), *Class. Quant. Grav.* **32**, 024001 (2015), 1408.3978
- [30] P. Spiller, G. Franchetti, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **561**, 305 (2006), proceedings of the Workshop on High Intensity Beam Dynamics

- [31] H.J.A. Rottgering, R. Braun, P.D. Barthel, M.P. van Haarlem, G.K. Miley, R. Morganti, I. Snellen, *LOFAR - Opening up a new window on the Universe*, in *Conference on Cosmology, Galaxy Formation and Astro-Particle Physics on the Pathway to the SKA* (2006), astro-ph/0610596
- [32] J.A. Tyson (LSST), Proc. SPIE Int. Soc. Opt. Eng. **4836**, 10 (2002), astro-ph/0302102
- [33] *Square Kilometre Array*, <https://www.skatelescope.org> (2021), accessed: 2021-02-27