Astronomical data organization, management and access in Scientific Data Lakes

Y. G. Grange[1], V. N. Pandey[1], X. Espinal[2], R. Di Maria[2], and A. P. Millar[3]

[1]*ASTRON, the Netherlands Institute for Radio Astronomy, Oude Hoogeveensedijk 4, 7991 PD, Dwingeloo, The Netherlands;* grange@astron.nl

[2]*CERN European Organization for Nuclear Research, Esplanade des particules 1, 1211 Geneva, Switzerland*

[3]*DESY, Notkestrasse 85, 22607 Hamburg, Germany*
On behalf of ESCAPE WP2

Abstract.   The data volumes stored in telescope archives is constantly increasing due to the development and improvements in the instrumentation. Often the archives need to be stored over a distributed storage architecture, provided by independent compute centres. Such a distributed data archive requires overarching data management orchestration. Such orchestration comprises of tools which handle data storage and cataloguing, and steering transfers integrating different storage systems and protocols, while being aware of data policies and locality. In addition, it needs a common Authorisation and Authentication Infrastructure (AAI) layer which is perceived as a single entity by end users and provides transparent data access.

The scientific domain of particle physics also uses complex and distributed data management systems. The experiments at the Large Hadron Collider (LHC) accelerator at CERN generate several hundred petabytes of data per year. This data is globally distributed to partner sites and users using national compute facilities. Several innovative tools were developed to successfully address the distributed computing challenges in the context of the Worldwide LHC Computing Grid (WLCG).

The work being carried out in the ESCAPE project and in the Data Infrastructure for Open Science (DIOS) work package is to prototype a Scientific Data Lake using the tools developed in the context of the WLCG, harnessing different physics scientific disciplines addressing FAIR standards and Open Data. We present how the Scientific Data Lake prototype is applied to address astronomical data use cases. We introduce the software stack and also discuss some of the differences between the domains.

1.   Introduction

The ESCAPE project brings together the communities of astronomy, astroparticle and particle physics, as well as the relevant research infrastructures from the European Strategy Forum on Research Infrastructures (ESFRI). The project works towards the implementation of a functional link between the concerned projects and the European Open Science Cloud (EOSC).

Work Package 2 (WP2) of ESCAPE, called Data Infrastructures for Open Science (DIOS), aims to assess how the infrastructure, that has been in use for managing

the data for several LHC experiments, commonly known as the Scientific Data Lake, could be applicable to the broader scope of the facilities within ESCAPE. Apart from the use cases that produce large amounts of data which needs to be managed, this also includes projects producing more moderate amounts of data. Those are the projects that need to ensure that data is replicated at the appropriate level, and that perform processing in a distributed fashion, on different compute environments.

2.   Scientific Data Lake

The Scientific Data Lake infrastructure consists of several services that work together to manage the large data volumes from the LHC experiments as briefly illustrated in Fig. 1. It handles heterogeneity of storage systems by design, supporting endpoints with different file systems managed by different external parties including supporting different access protocols. This also allows for integration with heterogeneous computing resources, like grid-like sites, high-performance computers, commercial, private and hybrid clouds, as well as sporadic resources. The Scientific Data Lake helps in making experimental data *FAIR* (Wilkinson et al. 2016)[1]. *Findable* and *Accessible* through a common name space, *Interoperable* and *Reusable* by supporting easy transfers to various processing systems.
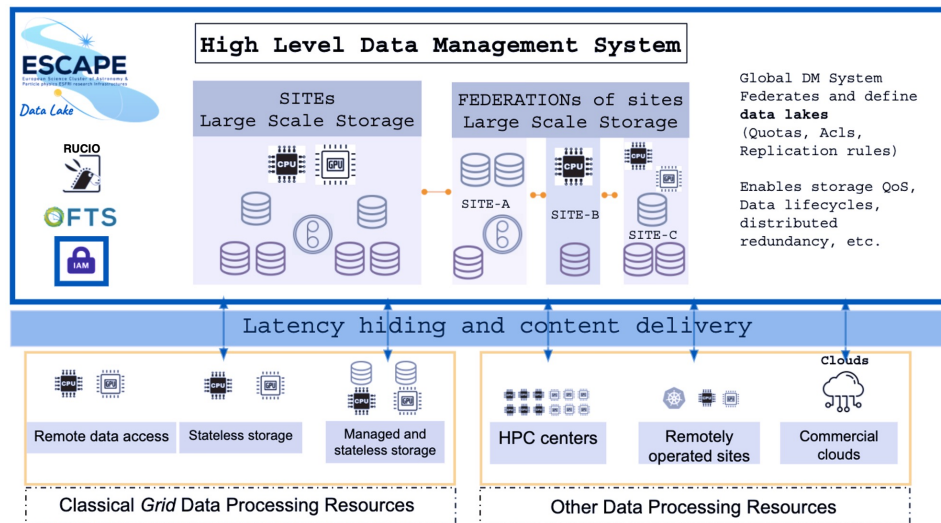

Figure 1.   A brief schematic overview of the Scientific Data Lake. In the top part, the three main pillars Rucio, FTS and IAM are shown, next to an illustration of a distributed set of storage end points, each with their own storage file systems (e.g. eos, dCache) and hardware (e.g. disk, tape). Some sites may also offer processing facilities. The bottom part shows different types of processing environment. The left half represents the classical Grid infrastructure, used to process the data from the LHC experiments. The right half lists other broadly adopted processing systems, like High-Performance Computing (HPC) or clouds. The content delivery and latency hiding layer of the infrastructure ensures that data can be accessed from any of those.

---

[1]https://www.go-fair.org/fair-principles/

The architecture is built on three main pillars consisting of :

Rucio Rule-based data management (Barisits et al. 2019). Rucio brings files and data sets that are stored in multiple locations under the same name space, and makes it possible to apply fine-grained data management concepts like polices, replication rules, data life-cycles, etc.

FTS File Transfer Service (Ayllon et al. 2014) is currently responsible for globally distributing the majority of Large Hadron Collider (LHC) data. It takes care of reliably transferring data between two endpoints, taking care of lower-level issues like matching protocols or retries.

Indigo IAM Identity and Access Management (Ceccanti et al. 2018) provides the AAI capabilities. Users can be authenticated via their institutional credentials (e.g. through EduGAIN) and user groups can be managed which are provided to the underlying applications (like Rucio and FTS) so that access rights can be appropriately applied to the group members.

3.   Science Data Lake and astronomy

One of the main goals of ESCAPE WP2 is to assess the suitability of the Scientific Data Lake infrastructure to support work flows from the astronomical use cases, suggest possible improvements and required new functionalities to the developers of the tooling. We give a brief overview of some of the main differences between the data work flows in both fields including ongoing relevant work towards addressing them.

Different metadata for files. Rucio currently offers full support for file metadata (size, checksum, etc.) and metadata keys that are accelerator physics-specific. Custom metadata can be provided through a plugin or JSON file. It is currently not possible to automatically apply rules based on this custom metadata. Also the rule engine presently only supports (partial) equality comparison of metadata keys. Work is currently being done to extend the metadata-based rule execution to custom metadata, and extending the rule engine to support inequalities as well in queries (e.g. 01h42m30s<RA<01h45m30s).

In particle physics, access to data is generally granted based on X509 certificates. Obtaining a certificate for access can however be cumbersome, as it requires a trusted party to confirm ones identity when requesting, and the certificate needs to be regularly refreshed (typically every year). During the project, support for OpenID Connect (OIDC)[2] tokens for AAI is being implemented and activated. This makes it possible for user authenticcation using home institute credentials resulting in a more accessible system for non-expert users as well. Currently, the ESCAPE WP2 Rucio instance now supports usage of OIDC tokens. However, some of the the storage system back-ends still require X509 certificates for data access (i.e. downloads).

Between astronomy and particle physics, the data access models significantly differ. While in particle physics user access is centrally managed by a project, in astronomy not only data access must be provided by the observatory, but additionally it should also be possible for the PI of the project to be able to decide on granting access to the

---

[2]https://openid.net/

data. In addition the concepts of proprietary and public data, which are very common in astronomy but not so in particle physics has implications on AAI.

Non-expert users represent a large fraction of public data users in astronomy, while in particle physics the data is generally accessed by experts. To improve the data access by non-expert users, a Jupyter Lab kernel that provides easy access to the Scientific Data Lake (called DataLake as a Service; DLaaS) has been developed within the project.

4. Conclusion and Future work

In the first phase of the ESCAPE project, we have successfully assessed how the Scientific Data Lake architecture can be used to manage data from a wide range of experiments from the fields of fields of (astro)particle physics and astronomy.

In the next phase, further integration between the Scientific Data Lake and the other components developed in ESCAPE will be one of the focal points. Further work is ongoing towards integration the DLaaS and the Science Analysis Platform developed in WP5 of ESCAPE (see Swinbank 2021, for an overview of this platform). Querying the Scientific Data Lake as one of the data sources in the Science Analysis Platform makes it possible to combine data sets from different collections or to customise query functionality for a specific field, for instance by offering astronomy-specific query parameters to astronomers. Also, there is an opportunity to investigate how the Scientific Data Lake could be integrated with the VOSpace standard (Graham et al. 2018), easing the integration of the Scientific Data Lake in the Virtual Observatory. Also, the support for both public and proprietary data and the granular group management, as both being mentioned in Sect. 3, needs further investigation. Initial work to demonstrate ability to group scientifically related astronomical observations (e.g. calibrators and target observations etc.) using the concepts of datasets, containers,and hierarchy on the scientific data lake and utilize them in processing work flows has shown promising results and thus needs further investigation.

Several project partners (MAGIC+CTA, SKAO and LSST) have set up their own instance for further exploration and ASTRON is planning to do so as well. This makes it possible to assess the architecture from the operational level, and also reap the gains made past the time frame of the ESCAPE project.

References

Ayllon, A. A., et al. 2014, Journal of Physics: Conference Series, 513, 032081. URL https://doi.org/10.1088/1742-6596/513/3/032081

Barisits, M., et al. 2019, Computing and Software for Big Science, 3, 11. URL https://doi.org/10.1007/s41781-019-0026-3

Ceccanti, A., et al. 2018, Indigo identity and access management (iam). URL https://doi.org/10.5281/zenodo.1874791

Graham, M., et al. 2018, VOSpace Version 2.1, IVOA Recommendation 21 June 2018

Swinbank, J. 2021, in ADASS XXXI, edited by B. Hugo, et al. (San Francisco: ASP), ASP Conf. Ser.

Wilkinson, M. D., et al. 2016, Scientific Data, 3, 160018