



Project Title	European Science Cluster of Astronomy & Particle physics ESFRI research Infrastructure
Project Acronym	ESCAPE
Grant Agreement No	824064
Instrument	Research and Innovation Action (RIA)
Topic	Connecting ESFRI infrastructures through Cluster projects (INFRA-EOSC-4-2018)
Start Date of Project	01.02.2019
Duration of Project	48 Months
Project Website	<a href="http://www.projectescape.eu">www.projectescape.eu</a>

## D2.2 Assessment and analysis of performance of the first pilot data lake

Work Package	WP2, DIOS
Lead Author (Org)	Rosie Bolton (SKAO), Rohini Joshi (SKAO), Xavier Espinal (CERN)
Contributing Author(s) (Org)	On behalf of ESCAPE Work Package 2
Due Date	31.01.2021, M24
Date	26/1/2020
Version	1.0

### Dissemination Level

<input checked="" type="checkbox"/>	PU: Public
<input type="checkbox"/>	PP: Restricted to other programme participants (including the Commission)
<input type="checkbox"/>	RE: Restricted to a group specified by the consortium (including the Commission)
<input type="checkbox"/>	CO: Confidential, only for members of the consortium (including the Commission)

## 1 Versioning and contribution history

Version	Date	Authors	Notes
0.1	15/1/2021	Rosie Bolton	Internal draft for review by ESCAPE management team
1.0	26/01/2021	Rosie Bolton	Final version for submission

## 2 Disclaimer

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 the Grant Agreement n° 824064.

## 3 Table of Contents

1	Versioning and contribution history.....	2
2	Disclaimer .....	2
3	Table of Contents .....	2
4	Table of Figures .....	4
5	Executive Summary .....	6
6	Introduction.....	8
7	Composition of the pilot Data Lake .....	9
8	Assessing the Pilot Data Lake .....	13
8.1	Dashboards and continuous monitoring .....	13
8.1.1	The DIOS monitoring suite.....	13
8.1.2	PerfSONAR.....	15
8.1.3	GFAL-level tests and dashboard .....	16
8.1.4	FTS-level tests and dashboard .....	17
8.1.5	Rucio level testing and monitoring .....	18
8.1.6	The IAM service .....	19
8.2	Live tests with Experiment dataflows.....	20
8.2.1	Overview.....	20
8.2.2	Functional testing .....	37
8.3	General Comments / Performance of the technology stack .....	41
9	Improvements identified for future work and next steps .....	42

D2.2 Assessment and analysis of performance of the first pilot data lake

---

9.1	Feedback from the experiments and how it will be addressed.....	42
9.2	Next steps: Continuous improvement.....	43
9.3	Next steps: Introducing new elements.....	44
9.4	Next steps: Nice to haves .....	44
9.5	Next steps: Continued collaboration .....	45
10	Appendix A: Use-case context from selected ESCAPE Experiments .....	46
10.1	CTA.....	46
10.2	EGO / VIRGO .....	48
10.3	FAIR.....	49
10.4	LOFAR .....	50
10.5	MAGIC.....	51
10.6	SKA Observatory .....	52
10.7	Rubin Observatory / LSST .....	53
11	Appendix B: Additional diagrams .....	55
11.1	ESCAPE Data Lake Rucio Storage Element (RSE) details .....	55
11.2	Daily Health Check flow diagram .....	56
12	Abbreviations.....	56



## 4 Table of Figures

Figure 1: The location of storage elements active during the period Nov-Dec 2020, taken from our Grafana dashboard. Marker size indicates the number of transfers initiated at each location. ....	9
Figure 2: Image showing dashboard view of Rucio Storage Elements (RSE) usage in the period November-December 2020. Capacity figures and technical details for each RSE are provided in the WP2 wiki space. ....	10
Figure 3: Screenshot showing the IAM group structure in ESCAPE. ....	11
Figure 4: The components of our pilot ESCAPE Data Lake, showing central services, monitoring, identity management and the storage elements. ....	12
Figure 5: Diagram showing the testing and monitoring architecture. ....	14
Figure 6: Representation of the nested data transfer stack and the continuous tests running at each level. Note that at each level delegation of the task is passed down to the level below – Rucio instructs FTS which uses GFAL to make transfers. Testing at each layer in the stack is essential to understand failures. At the base level, perfSONAR tests the link health. ....	15
Figure 7: GFAL testing results dashboard for the 24hr period of the "FDR" (17/11/2020). Tests run every 2 minutes - results are aggregated here to cover the 24hr period. If a GFAL test fails, we can trace that in the dashboard and understand which protocol(s) and sites are involved, and which actions (upload, download, delete) are affected. The error messages are captured in the dashboard. ....	17
Figure 8: FTS dashboard view showing the transfer matrices for all three protocols. ....	18
Figure 9: Dashboard view of the automatic Rucio-level tests running from CERN for a period midnight to 0400hrs on December 17th 2020. During this time, intensive tests were running to try out different configurations - 31k transfers were submitted in 4 hours. ....	18
Figure 10: As for previous figure, but showing the transfer matrix. The matrix view clearly shows that INFN-NA-DPM-FED and INFN-ROMA1 were not available at this time. ....	19
Figure 11: Dashboard view from November 17th showing an overview of the Rucio activity. ....	20
Figure 12: Nov 17th testing window Rucio transfer matrix view. ....	21
Figure 13: Rucio event activity before, during and after the Rucio server re-configuration. ....	22
Figure 14: Rucio activity during the 2nd testing window, 15/12/2020 ....	23
Figure 15: Rucio activity from the second testing window, but with time range extended to show the completion of the "Million file test". ....	23
Figure 16: Progress of SKA's pulsar search test during the first testing window. ....	26
Figure 17: Progress of SKA's single pulse search test during the first testing window. ....	26
Figure 18: Progress of SKA's 24 hour QoS lifecycle test ....	27
Figure 19: Completed transfers of the SKA SDC1 dataset during and after the first testing window. ....	27
Figure 20: Completed transfers corresponding to 7 TB of CTA's data ingestion and replication during and slightly after the first testing window. ....	28
Figure 21: Summary of ATLAS open data transfers initiated by QoS based rules. ....	29
Figure 22: Increase in CMS storage used corresponding to open data ingestion test during the first testing window. ....	30
Figure 23: MAGIC's usage during the first testing window depicting data ingest and data clean-up. ....	31
Figure 24: Schematic of the relationship between the LOFAR Ingest (Stage 1) and Processing (Stage 2) tests. ....	32
Figure 25: LSST's (Rubin Observatory's) usage pattern during their first data ingest test (bottom) and a similar usage pattern reflected at IN2P3-CC-DCACHE, their primary ingest RSE (top). ....	33

D2.2 Assessment and analysis of performance of the first pilot data lake

Figure 26: LSST's usage pattern during their second data ingest test. Note the scale of the second test was larger than the first. .... 34

Figure 27: FAIR's first data ingest test began at around 12 pm following issues at GSI-ROOT. Transfers queued, submitted, done and failed are shown here. .... 35

Figure 28: FAIR's second data ingest ran smoothly over 24 hours with noticeable dips coinciding with stress tests running on the same day. .... 36

Figure 29: VIRGO's data ingest can be seen here in the form of increasing DIDs over the course of the 1st testing window. .... 36

Figure 30: Plot shows increase in the number of files during the preparation and execution of the million-file test. Note the horizontal axis covers 13 November to 17 December ..... 38

Figure 31: Progress of the transfers corresponding to the million-file test. Transfers queued and completed the following day..... 38

Figure 32: The high throughput test moved 1 TB of data from EULAKE-1 to FAIR-ROOT ..... 39

Figure 33: Output from FTS logs showing the achieved throughput for the transfer of the 1008 files in the High Throughput Test..... 39

Figure 34: Image showing (top) that just after creation, all but one rule (the RSE to which the file has been uploaded) is in replicating state. A few moments later (bottom), the rules have all moved to OK state. They can be seen to have varying expiration dates as required in the lifecycle. .... 40

Figure 35: Image showing example of changing data storage workflows from one day to the next in order to balance the data storage loaf across the Bulk Data Management system of CTA. .... 47

Figure 36: The current EGO/LIGO data movement flow. Image credit: Pierre Chanial ..... 48

Figure 37: A possible alternative EGO/LIGO data movement option, making use of datalake technologies. Raw data flows into the datalake for offline processing. Image credit Pierre Chanial. 49

Figure 38: The MAGIC telescope data movement, replication and deletion flow, showing deletion of data at source following successful replication..... 52

Figure 39: Rubin Observatory / LSST Data movement and processing schematic. .... 54

Figure 40: Rubin Observatory (LSST) data product flow..... 54

Figure 41: This diagram shows the current (Jan 2020) version of the action tree that we have developed collaboratively to guide newcomers to the monitoring team in their health checks. This means that people new to the datalake concept have been able to take part in the monitoring activities, greatly improving our ability to rapidly detect that components are under-performing. .56



## 5 Executive Summary

Since the start of the ESCAPE project, WP2 (Data Infrastructure for Open Science - DIOS) has come a long way in its goal to open up federated, distributed data management technologies to the broad science communities of the project.

We have built a pilot Data Lake and deployed all the supporting services required to make this work, with not just the three storage locations targeted in our plan, but with ten sites active across ESCAPE partner countries. At least as importantly as the physical and digital infrastructures, we have built a community - onboarding users from across ESCAPE's Astrophysics, Astro-particle and Particle Physics science experiments and participating institutes, engaging them not only in the assessment of the performance of the Data Lake but also in the day-to-day operational aspects. Some of these users are completely new to the elements of the ESCAPE DIOS technology stack, whilst others have significant experience in the specific context (especially in High Energy Physics). Some experiments have existing systems operational and are seeking ways to improve automation of their data management, whilst others are looking ahead at large-scale data challenges and asking if and how the open-source technologies presented in the ESCAPE project can answer their needs.

Our ESCAPE Data Lake is relatively modest in scale by comparison to the enormous data volumes managed in production in High-Energy Physics data catalogues (e.g. ATLAS and CMS). Precisely because we are not running a production instance, we have been able to try challenging tests (which do sometimes break things) and to reconfigure or rebuild our services as we learn.

In preparation for our assessment period (Nov-Dec 2020), we built up our continuous testing and monitoring work - we have working suites of tests running continuously that give us feedback on the health of most areas of the technology stack. We developed an operations team from across sites and experiments which shares the daily health check rota and meets weekly to check status of issues.

The ESCAPE Data Lake is available continuously to all users, but we decided that having a focused testing time would help to ensure good user support and put more strain on the central Data Lake services, so more closely mimicking real-life scenarios. Our first testing window took place on 17th November, and we held a work-package wide workshop on 9-10 December to present the results and identify lessons learned. Some experiments then continued their testing during the second testing window on 15th December. This saw over a dozen different tests developed by teams representing ESCAPE Experiments over the course of both testing windows.

Through the broad user community in WP2, we have been able to undertake an assessment of the Data Lake that has provided genuinely useful feedback - for our own benefit but also for the developers of the underlying technologies.

In this document we describe our pilot Data Lake and present both the continuous monitoring and the experiment-led tests in detail. The context for each experiment is interesting too - because it gives understanding of the Data Lake features of interest to each experiment - we have included several context examples in an appendix to this document (page 46 onwards).

## D2.2 Assessment and analysis of performance of the first pilot data lake

---

The process of hardening our pilot Data Lake into something more mature has certainly been accelerated by this testing work, and it will continue as we move to the "Prototype" phase of ESCAPE - the results of the assessment presented in this document are our starting point for that future work.

Overall, collaborators in this work package have been positively impressed by the performance of the Data Lake, and of the resilience shown by Rucio in particular. The interaction between ESCAPE WP2 colleagues has been excellent too - forming a coherent team, with members comfortable asking questions about technologies unfamiliar to them, has been a great success - especially when travel restrictions have made it impossible to meet in person.

For the first time, we have seen our flagship ESFRIs in Astro-particle Physics, Electromagnetic and Gravitational Wave Astronomy, Particle Physics, and Nuclear Physics working together in a common scientific data management infrastructure – the ESCAPE Data Lake – that seeks to implement the standards of data FAIRness and Open Access.

This is an achievement worth highlighting; we are breaking historic sociological barriers between these different scientific communities. We are both searching for commonalities and understanding differences in their computing models with a view to successfully addressing the plans for these challenging experiments in the next decade, especially in data management and computing.

The results of the assessment work that we present here give us confidence that the Data Lake concept, and the tools we are testing, should be able to scale to meet foreseen future needs of the Research Infrastructures represented in ESCAPE, in terms of number of sites, storage volumes, data rates and number of users.

## 6 Introduction

In the first 2 years of the ESCAPE project, WP2 has set about building a Data Lake using storage sites from partners across Europe, managed collectively through the data management service, "Rucio" (<https://rucio.cern.ch/>). The principal aim of the WP2 Data Lake is to provide an environment within which use cases from across Astronomy, High Energy Physics, and Astro-Particle Physics Sciences can be tested, with a view to understanding how the differing distributed data storage, access and manipulation use cases across these sciences might be supported.

Early in the ESCAPE project (July 2019) a joint workshop between WP2 and WP5 was held (see <https://indico.in2p3.fr/event/19214/>). This workshop introduced the WP2 participants to the technologies being considered for the ESCAPE Data Lake and to the differing use cases for each ESFRI or experiment, through a series of presentations. The major outcome of this workshop, for WP2, was the project plan which enabled work to start on the pilot phase of the Data Lake.

Initially, Data Lake deployment work focused on: building up the Data Lake components - identifying and onboarding storage elements; deploying a server, with appropriate back-end database support to run the Rucio service; ensuring that monitoring hardware was in place as needed (perfSONAR, see later); and building the software, databases and dashboards to enable the health of the Data Lake to be readily assessed.

In May 2020, the major building blocks were in place, but we were aware that the detailed knowledge of their implementation still lay in a small number of people in the project. Many of these technologies were developed in the High-Energy Physics (HEP) community, driven by the challenging but specific needs of the HEP experiments. To ensure that perspectives from ESCAPE participants outside HEP could also be captured, and to broaden the knowledge base, we set up a "Deployment and Operations" (DepOps) team. This team meets weekly, raising and fixing issues and enabling knowledge and confidence in the Data Lake technologies to flow from the experts in each area to include relative newcomers. This forum has enabled us to share feedback and future ideas across the whole Data Lake technology stack, from bug fixes required in storage technologies, to dashboard development and how to use Rucio. Though the pilot Data Lake is not a production instance, running the DepOps team gives us a flavour of the challenges faced in production. The DepOps team shared the burden of monitoring the Data Lake so that issues were resolved faster than previously, and we could be confident that the pilot Data Lake was ready for some more pilot users.

In September 2020 work began to onboard representatives from several experiments as test users of the pilot Data Lake - they were supported and encouraged to add their own (representative) data to the Data Lake and to become competent users of Rucio.

The combination of a working pilot Data Lake, a functional operations team, and growing user base meant that we were ready to formally assess the pilot Data Lake. In order to maximise the load on the elements of the Data Lake, and to more-closely mimic a production level Data Lake supporting multiple sciences, we decided to focus the experiment-based testing effort on two 24-hour periods - one main day of activity (on 17th November), followed by a second 24-hour period for any additional testing (15th December). During these periods the DepOps team and site representatives provided excellent user support - quickly working to identify issues and resolve them as fast as possible. In



between these two 24-hour periods, we held the second WP2 workshop (9-10 Dec 2020: <https://indico.in2p3.fr/event/22693/>).

The tests that were conducted in these two 24-hour periods and the opinions given by the experiment and site representatives (both positive and negative), and the experiences within the DepOps team have provided us with a thorough assessment of the functionality of the pilot Data Lake and ideas for directions to move in the next phase of the project.

In this document we describe the pilot Data Lake, then look at the tests used to assess the performance of the Data Lake and the overall activity during the testing windows. We use these tests to develop ideas for improvements and future work. We include details on the use cases for some experiments, since these give context to the tests performed, and for requirements for future work.

## 7 Composition of the pilot Data Lake

The pilot Data Lake has been the first implementation of the conceptual design of building a common storage and data management infrastructure for open science involving different scientific communities.

The core of the pilot Data Lake is composed of several and diverse storage endpoints provided by the ESCAPE partner institutes across Europe (see Figure 1), resulting in a total of ten storage endpoints active in the current implementation. Figure 2 shows the use made of the Data Lake storage in our assessment period Nov-Dec 2020.

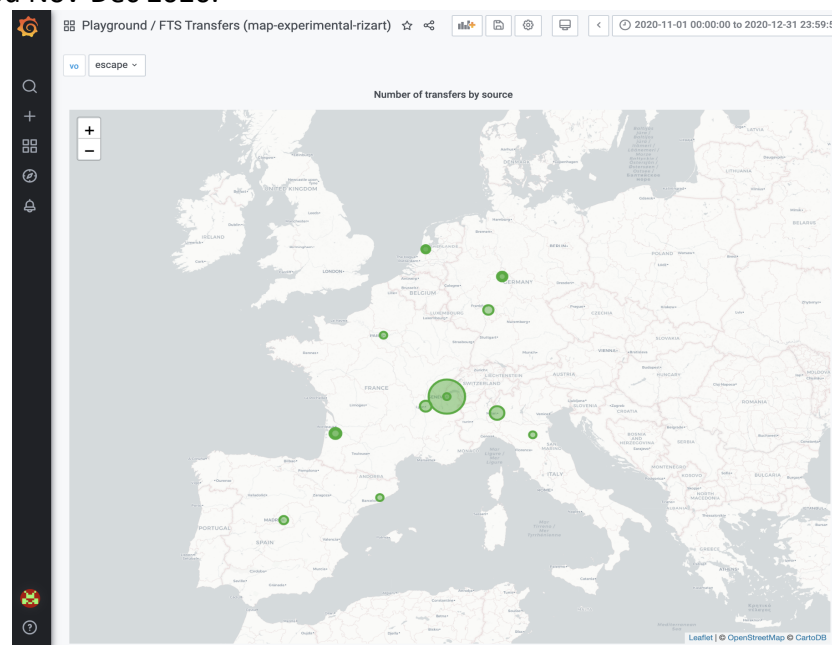


Figure 1: The location of storage elements active during the period Nov-Dec 2020, taken from our Grafana dashboard. Marker size indicates the number of transfers initiated at each location.

In preparation for the next stage of the project, several sites in the ESCAPE Data Lake have implemented [XCache](http://slateci.io/XCache/)<sup>1</sup>. This is disk storage managed by software providing a caching layer that both enables file reusability and provides latency hiding capabilities. This caching component is connected

<sup>1</sup> <http://slateci.io/XCache/>

## D2.2 Assessment and analysis of performance of the first pilot data lake

to remote storages and might help compute resources to be more efficient on CPU use, as the data is streamed from the local cache once the file either is fully replicated or starts to be consumed (streaming cache feature). This technology also opens several possibilities to enable data access to the Data Lake from non-standard resources (e.g., commercial clouds, HPCs, opportunistic resources, etc.) which may not have, or have a very limited, storage space. Hence having a caching layer at the sites or at the edge can iron out some of the classical issues (ingress/egress, security hurdles, etc.) when trying to make use of this type of resource. In ESCAPE WP2 we implemented a vanilla version of XCache which is ready to be deployed anywhere and it is already operational at several sites composing the ESCAPE Data Lake: CC-IN2P3, GSI, and CERN. We have started early assessment of this caching technology by running experiment specific workflows, but the full assessment of the effectiveness of caching will be addressed in the prototype phase.




Figure 2: Image showing dashboard view of Rucio Storage Elements (RSE) usage in the period November-December 2020. Capacity figures and technical details for each RSE are provided in the WP2 wiki space.

The different Data Lake endpoints cover a variety of different storage technologies (dCache, EOS, XRootD, StoRM, DPM, CEPH) with distinct types of disk backends in use (JBODs, Raid-based systems, Block Storage and Erasure Coding) - see the table on page 55 or the ESCAPE WP2 wiki site<sup>2</sup>.

These different storage endpoints are harnessed through an orchestration layer, composed by the data management tool ([Rucio](https://rucio.cern.ch/)<sup>3</sup>) and a service to coordinate file transfers ([FTS](https://fts.web.cern.ch/fts/)<sup>4</sup>) based on third party copy with several protocols enabled: http, gridftp and xroot.

There is also a central information system ([CRIC](http://escape-cric.cern.ch/)<sup>5</sup> - the Compute Resources Information Catalogue) that keeps track of the different storage endpoint URLs, the protocols supported at the site, and the prioritization among them. When this information is updated it is picked up by the Data Lake

<sup>2</sup> [https://wiki.escape2020.de/index.php/WP2 - DIOS](https://wiki.escape2020.de/index.php/WP2_-_DIOS)

<sup>3</sup> <https://rucio.cern.ch/>

<sup>4</sup> <https://fts.web.cern.ch/fts/>

<sup>5</sup> <http://escape-cric.cern.ch/>

## D2.2 Assessment and analysis of performance of the first pilot data lake

components (e.g. Rucio server, FTS and storage endpoints) so that the clients used by the user community, pick up the correct configuration.

The [ESCAPE IAM instance](#)<sup>6</sup> is the central authentication and authorization service serving the ESCAPE Data Lake. IAM ESCAPE, deployed on a Kubernetes cluster at INFN CNAF, supports authentication through [EduGAIN](#)<sup>7</sup> and implements user registration and management as well as token and VOMS attribute provisioning services (representing both future and legacy solutions, respectively). IAM has currently 86 registered users organized in 12 groups, where each ESFRI has one (or more) dedicated groups, as shown in the screenshot in Figure 3.

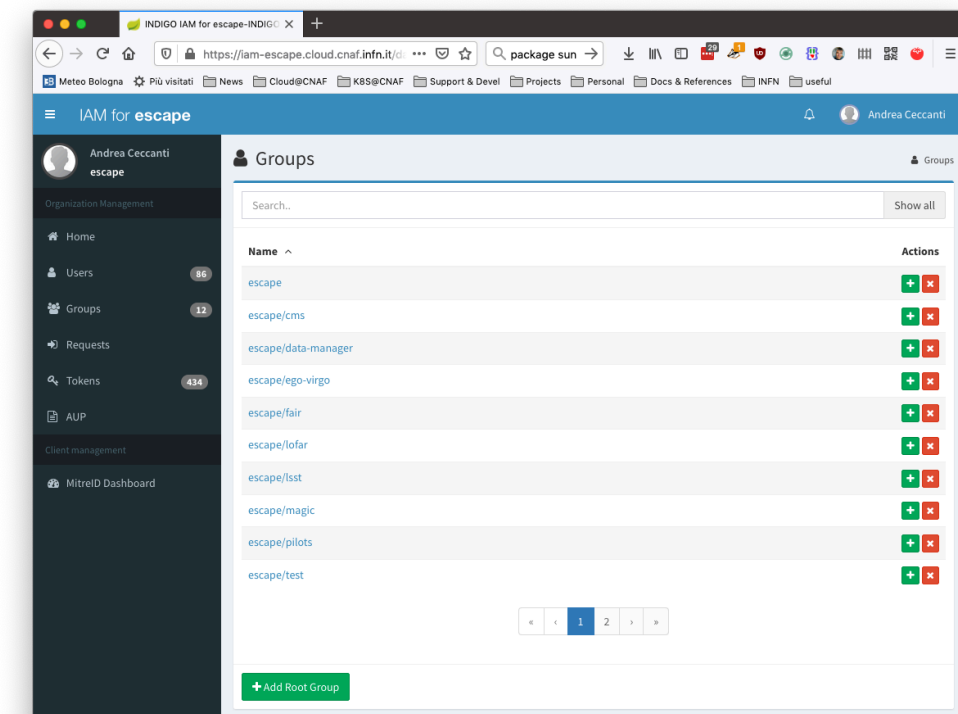


Figure 3: Screenshot showing the IAM group structure in ESCAPE.

<sup>6</sup> <https://iam-escape.cloud.cnaf.infn.it>

<sup>7</sup> <https://edugain.org/>

D2.2 Assessment and analysis of performance of the first pilot data lake

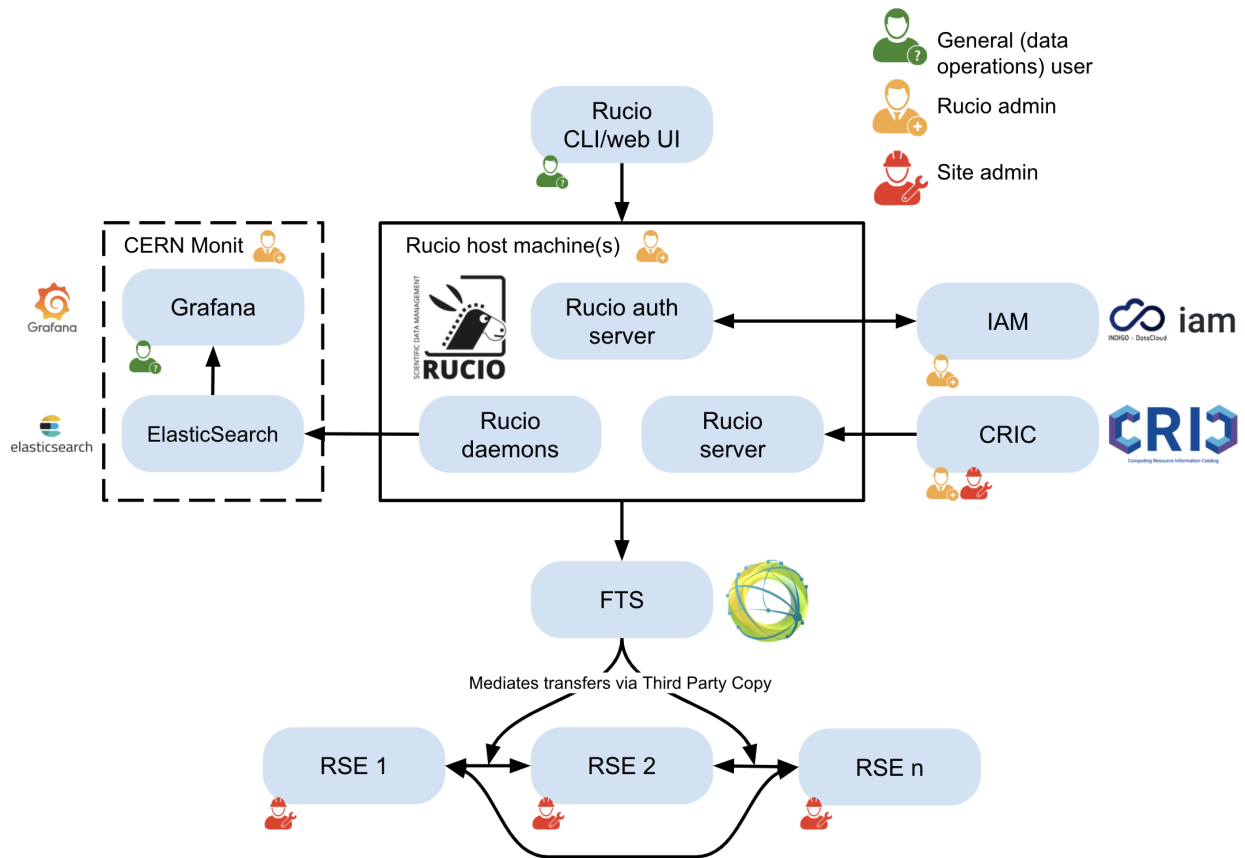


Figure 4: The components of our pilot ESCAPE Data Lake, showing central services, monitoring, identity management and the storage elements.

The interplay between the various technology elements is shown in Figure 4. The essential layer that translates the technical requirements of the storage elements and the data transfer mechanisms into the ability to manage data placement across a distributed, heterogeneous system is the Data Management Service, Rucio. Data operations users (i.e., Experiment-based personnel responsible for ensuring that their science users get access to relevant data products in the correct locations, for appropriate lifetimes) define "Rules" specifying how they want data objects (files, or collections of files) to be stored, replicated, or deleted. Data objects are addressed with Data IDentifiers (DIDs) in Rucio that are composed of a scope and a name (scope:name). Rucio interprets these rules and then performs operations, e.g., data replication, such that the rules are satisfied while transfers are kept to a minimum. As one possible example, if a storage site goes down or if data are deleted, provided that there is another copy of the data elsewhere within the Rucio ecosystem, Rucio (through a daemon named judge-evaluator) will detect that the rules for that data are now no longer satisfied, find the location of the file on another RSE and make duplications to bring the rules back into "OK" state.

Having different storage systems with different disk backend configurations opens the door to implement storage Quality of Service (QoS) intelligence, meaning that experiments can define data lifecycles and optimize the usage of the storage: data are usually more frequently accessed during the first days (hot) and gradually transitions to a state of rare access (cold). Hot storage requirements are usually more expensive than cold storage requirements, hence enabling QoS transitions may also

provide smarter use of storage resources, improving the user experience, and optimising the use of the overall storage. The QoS concept support is under development in the Rucio project, and it is possible to construct rules based on the QoS tag, as an alternative to specifying particular RSEs (e.g., "one copy is required in "SAFE" storage and one in "FAST" storage, rather than "one copy is expected in site A and one in site Y"). The usefulness of this aspect is something that can differ between experiments, and we were keen to test the concept from the perspectives of the different ESCAPE experiments. Currently the QoS classes used in the ESCAPE Data Lake are labels only - we use them to test the theoretical functionality of QoS-based rules, but e.g., "FAST" does not necessarily mean that the RSEs with this label actually have low latency - similarly, "SAFE" does not mean (yet) that RSEs with that label actually are more reliable.

In the first WP2 Deliverable, we stated the aims for the Pilot phase: *"In the first, pilot phase the focus will be demonstrating at the functional level the choices we made in terms of architecture and reference implementation. We will therefore deploy a small-scale pilot, spanning at least three of our partner institutes and demonstrate the core functionalities in terms of data transfer, data access and storage."*

Our Pilot Data Lake is in fact much more advanced than we had hoped, with ten sites providing storage elements (some sites offering more than one RSE) - giving a total offered capacity of almost 800 TB (about 300 TB used at the time of writing). We have close to 100 registered users from across 10 experiments, representing 6 of the 9 ESCAPE ESFRI projects and several associated institutes.

## 8 Assessing the Pilot Data Lake

### 8.1 Dashboards and continuous monitoring

#### 8.1.1 The DIOS monitoring suite

In the past year, considerable effort has been put into developing a suite of tests which automatically interact with the Data Lake components, and a corresponding database and dashboard environment that can be used to quickly assess the current and previous health of various aspects of the Data Lake. For this pilot phase our focus for this assessment has been on qualitative measures (i.e. "Is it working?"), rather than on quantitative measures (e.g. "How fast is it?").

Test results are stored in databases that are accessed by a dashboard suite - the primary one is a Grafana instance hosted by the CERN Monit team, but some sites also have local dashboard monitoring, and SKA has a database hosted at the STFC Cloud, which enables a more data-centric view than afforded by the information in Rucio events themselves. We have developed dashboard views that give information on the different areas of the Data Lake, and complementary tests to ensure these are populated.

Figure 5 shows our dashboard and database architecture. This has evolved over the past year, relying on the [influxdb](https://www.influxdata.com/products/influxdb/)<sup>8</sup> and Elasticsearch databases managed by CERN's Monit team (a pre-existing database suite built for monitoring the CERN IT data centre and the WLCG sites).

---

<sup>8</sup> <https://www.influxdata.com/products/influxdb/>



D2.2 Assessment and analysis of performance of the first pilot data lake

Users can access the Grafana suite with a lightweight (readily obtained) CERN account. Apart from very few exceptions taken from local monitoring dashboards, all the dashboard views shown in this document are from the ESCAPE Grafana suite.

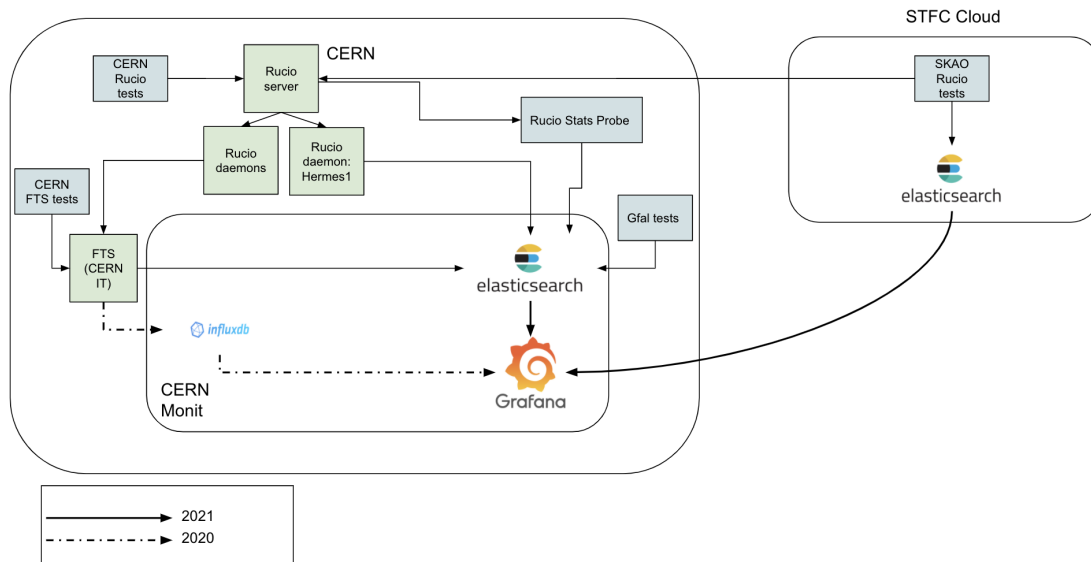


Figure 5: Diagram showing the testing and monitoring architecture.

As mentioned in the introduction (page 8) we have a built up a WP2 Deployment and Operations (DepOps) team to share monitoring and testing tasks and to provide a regular forum to discuss any issues. This team meets weekly to review issues identified in the preceding week and any ongoing issues in the JIRA issue management system. We developed a rota of people responsible for undertaking daily Data Lake health checks using the dashboards, with a flow chart to help newcomers learn how to identify errors and raise issues (Figure 41, page 56).

Our testing suite targets the various levels of the data transfer stack, which are described in Figure 6 below.

## ESCAPE WP2 Datalake Continuous Testing Framework and Network Stack

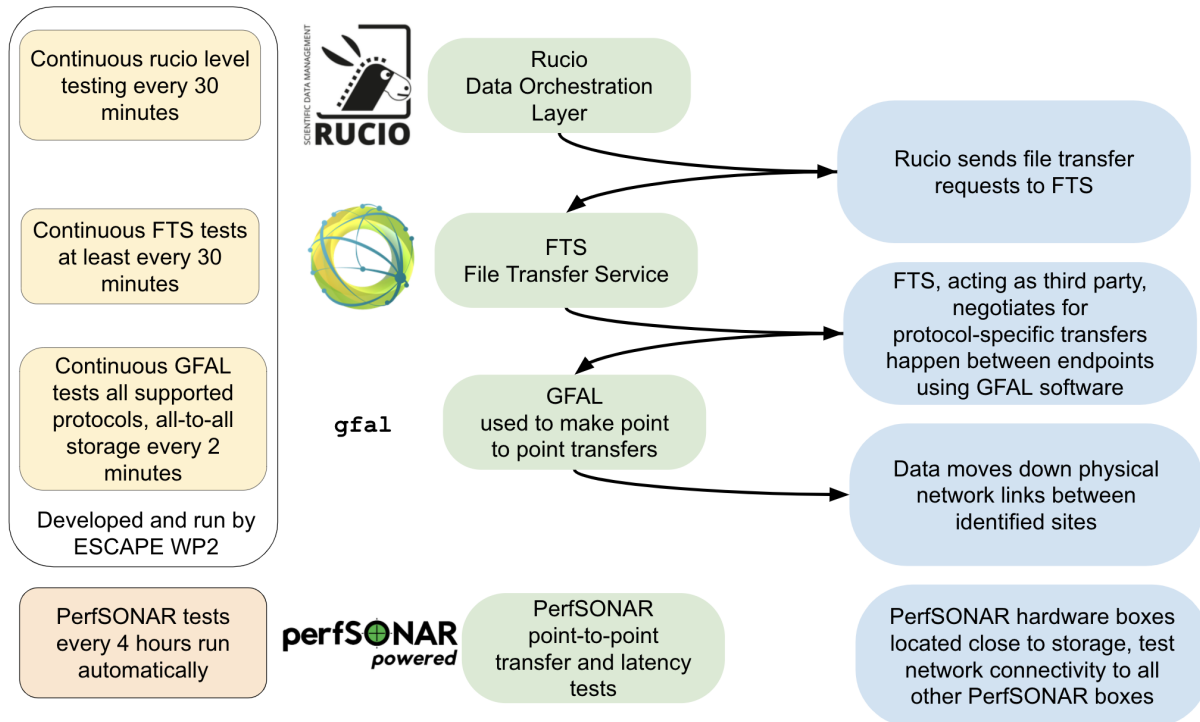


Figure 6: Representation of the nested data transfer stack and the continuous tests running at each level. Note that at each level delegation of the task is passed down to the level below – Rucio instructs FTS which uses GFAL to make transfers. Testing at each layer in the stack is essential to understand failures. At the base level, perfSONAR tests the link health.

### 8.1.2 PerfSONAR

At the most basic level, sites are connected to each other by network links, and it is essential to be able to see the health of these links. To do this, we made use of the existing perfSONAR infrastructure (<https://www.perfsonar.net/>). PerfSONAR requires a small hardware box to be installed within each site, which then performs automated transfer tests every 4 hours to all other perfSONAR sites. This all-to-all matrix is a powerful way of identifying which links are functioning well. We define an ESCAPE-Data Lake "mesh" to give us the (publicly accessible) [ESCAPE perfSONAR dashboard view](#)<sup>9</sup>.

The table below shows the average throughput figures for the ESCAPE perfSONAR boxes for dashboard-extracted data, during the period from 01/11/2020 to 31/12/2020 and thus covering both focused testing windows. Since there is no dedicated access to these links either for ESCAPE work, and neither for the perfSONAR tests, the throughput numbers are only indicative of the "spare" capacity on links rather than the actual line speed.

<sup>9</sup> <http://maddash.aglt2.org/maddash-webui/index.cgi?dashboard=ESCAPE%20Mesh%20Config>

D2.2 Assessment and analysis of performance of the first pilot data lake

source \ destination	PIC	SARA-MATRIX	INFN-T1 (CNAF)	INFN-ROMA1	INFN-NAPOLI-ATLAS	IN2P3-LAPP	IN2P3-CC	DESY-HH	CERN-PROD	GSI	mean as source, Gb/s
PIC	N/A	0.99	1.15		0.68	0.99	1.21	0.73	1.30	0.83	0.98
SARA-MATRIX	0.50	N/A	5.95	0.51	7.14	5.23	6.43	6.63	5.92	2.47	4.53
INFN-T1 (CNAF)	0.60	9.77	N/A	0.48	1.17	0.89	1.33	0.51	0.72	0.54	1.78
INFN-ROMA1	0.12	0.49	0.726	N/A	0.72	0.45	0.38	0.47	0.44	0.47	0.48
INFN-NAPOLI-ATLAS	0.19	2.98	6.48	0.82	N/A	5.08	5.89	5.95	6.36	1.91	3.96
IN2P3-LAPP	0.33	2.78	4.62	0.47	3.328	N/A	5.66	3.88	5.28	1.81	3.13
IN2P3-CC	0.36	4.66	6.89	0.52	5.88	6.76	N/A	6.45	8.66	2.22	4.71
DESY-HH	0.23	3.81	6.51	0.48	5.62	5.05	5.92	N/A	6.728	2.84	4.13
CERN-PROD	0.38	0.10	5.28	0.50	7.71	5.42	8.27		N/A	2.28	3.74
GSI	0.23	2.53	2.63	0.29	2.52	2.44	2.79	2.64	2.98	N/A	2.12
mean as destination, Gb/s	0.33	3.12	4.47	0.51	3.86	3.59	4.21	3.41	4.27	1.71	

Table 1: Results from the perfSONAR network link monitoring.

The perfSONAR table above shows that the ESCAPE storage sites **typically** have *available* transfer capacity at a few Gb/s, though some sites have consistently less than this. For example, the upload speed to PIC is generally low (it has low values as a destination) because currently the site runs with a 20Gb/s link which is very highly utilised; an upgrade to 100 Gb/s is anticipated in Q1/Q2 2021. However, care must be taken not to over-interpret the perfSONAR numbers - the perfSONAR hardware is separate from the Rucio storage, and, though it may be physically close by, there can be additional bottlenecks getting data into and off the storage, so the site-to-site performance may not match the perfSONAR numbers.

Our main use of perfSONAR in our daily monitoring is to quickly see if a link is active or not - however, we have found that apparent failed perfSONAR tests are more often caused by failures in the perfSONAR boxes themselves (typically, archives getting overfull). We have never noticed an actual link failure to date.

### 8.1.3 GFAL-level tests and dashboard

**GFAL** (the "Grid File Access Library"<sup>10</sup>) is a software suite developed in the Worldwide LHC Computing Grid (WLCG) project to allow applications to access and move files stored in the distributed LHC compute grid. In ESCAPE we have developed an automatic gfal-based tool that runs every 2 minutes to test the ability for all ESCAPE RSEs to send and receive files to every other RSE through all the transfer protocols that are mutually supported by each pair of sites – results are visualised in the dashboard, as shown in Figure 7. These transfers are made directly from site to site without third-party management. The tests are run on a Kubernetes cluster at CERN.

Link to scripts: <https://github.com/ESCAPE-WP2/Utilities-and-Operations-Scripts>

<sup>10</sup> [https://www.gridpp.ac.uk/wiki/Grid\\_File\\_Access\\_Library](https://www.gridpp.ac.uk/wiki/Grid_File_Access_Library)





## D2.2 Assessment and analysis of performance of the first pilot data lake

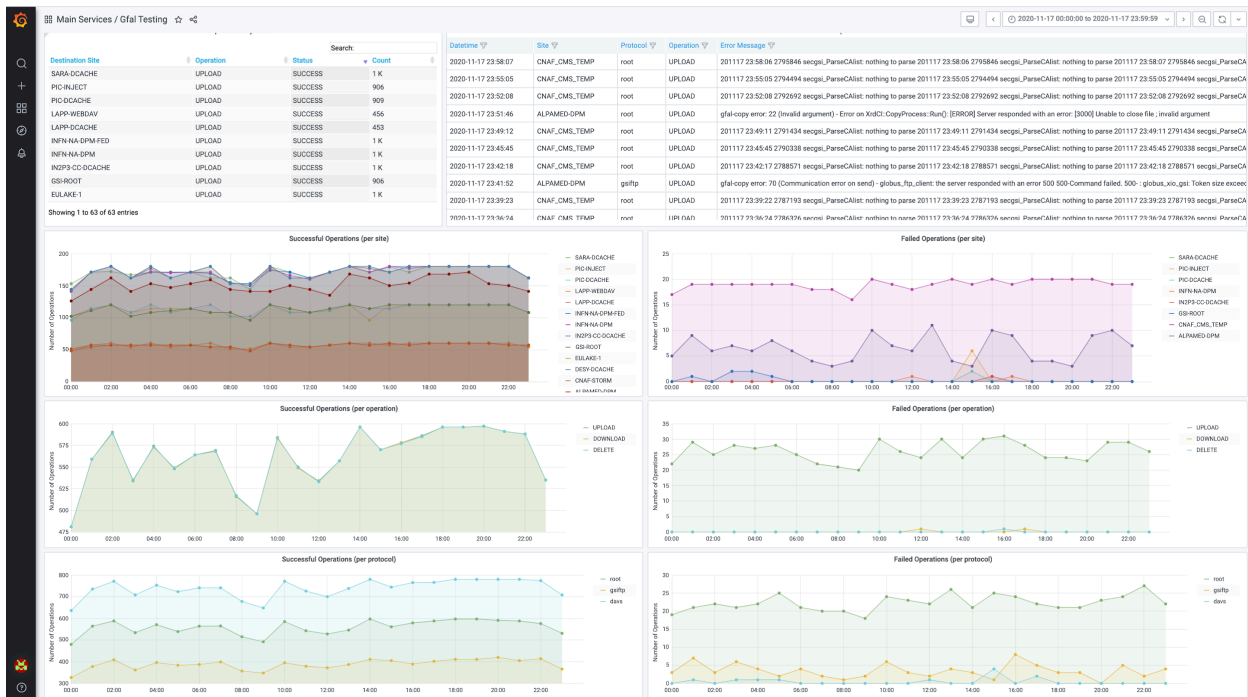


Figure 7: GFAL testing results dashboard for the 24hr period of the "FDR" (17/11/2020). Tests run every 2 minutes - results are aggregated here to cover the 24hr period. If a GFAL test fails, we can trace that in the dashboard and understand which protocol(s) and sites are involved, and which actions (upload, download, delete) are affected. The error messages are captured in the dashboard.

### 8.1.4 FTS-level tests and dashboard

The File Transfer Service (FTS - see <https://fts.web.cern.ch/fts/>) is a software suite developed at CERN that enables third-party copies between sites to take place. The FTS software is run on a server (currently using CERN's FTS-pilot server) and brokers transfers between sites, so that files can be moved from site A to site B without passing through the FTS server itself, provided that the two sites can use a common protocol. The Rucio service itself delegates file transfer management to the FTS service. Thus, a health-check on the FTS service is recommended. For this reason, a series of automated FTS-level tests runs continually, moving a range of different sized files (up to 5 GB) and repeating at least once per hour, to ensure good dashboard visibility. In addition, the transfers triggered by Rucio also appear in the FTS dashboard since they are mediated by the FTS service.

Link to testing software: <https://github.com/ESCAPE-WP2/fts-analysis-datalake>

Figure 8 shows an FTS dashboard view (taken in the morning of the 17th November 2020). The image shows transfer success rates between sites for the three different protocols our Data Lake supports. This particular view clearly shows (red) that GSI-ROOT is failing as a destination at this time, though still able to act as a source. In general, all individual red boxes (success <40%) are followed up, and any consistent orange boxes (success <75% also warrants investigation). By and large, when sites are working, we typically see FTS success rates at the 95% level or above - and of course, Rucio and the FTS both provide resilience to these failures, by re-attempting transfers, by using alternative protocols where possible or by (Rucio) choosing a different RSE if the replication rule allows. We have found the FTS service itself to be highly reliable.

## D.2 Assessment and analysis of performance of the first pilot data lake

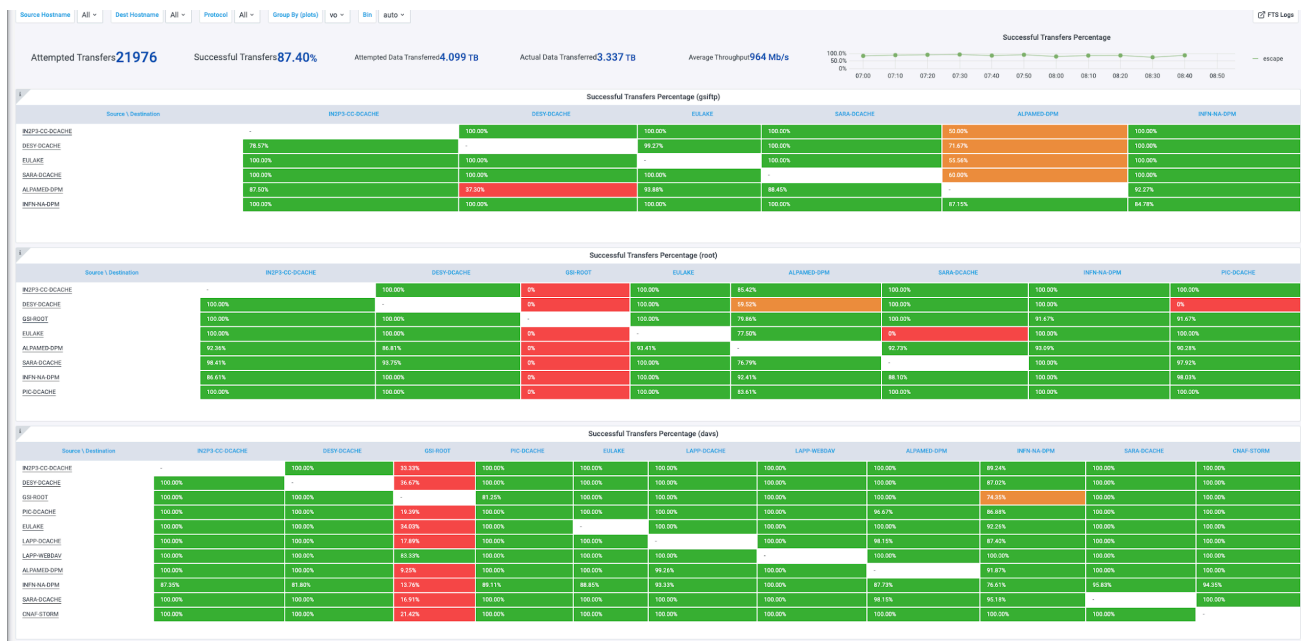


Figure 8: FTS dashboard view showing the transfer matrices for all three protocols.

### 8.1.5 Rucio level testing and monitoring

Automated Data Lake-wide Rucio tests are run continuously once per hour from Virtual Machines (VMs) hosted at the STFC Cloud facility in the UK. These resources are part of SKA's allocation on IRIS (<https://www.iris.ac.uk/>). These tests attempt upload and replication of randomly-generated files with a short (1 hour) lifetime, uploading to each site and then replicating from each site to every other site.

A second suite of broad Rucio tests is also run from the Kubernetes cluster at CERN. During the assessment period these were transfers of 100 MB every 30 minutes to all RSEs, including upload, download and replication. This has now been improved to give enhanced coverage, covering a range of file sizes from 1 MB up to 5 GB on cadences from once every 15 minutes down to once per 8 hours (the larger file transfer tests being run less frequently).

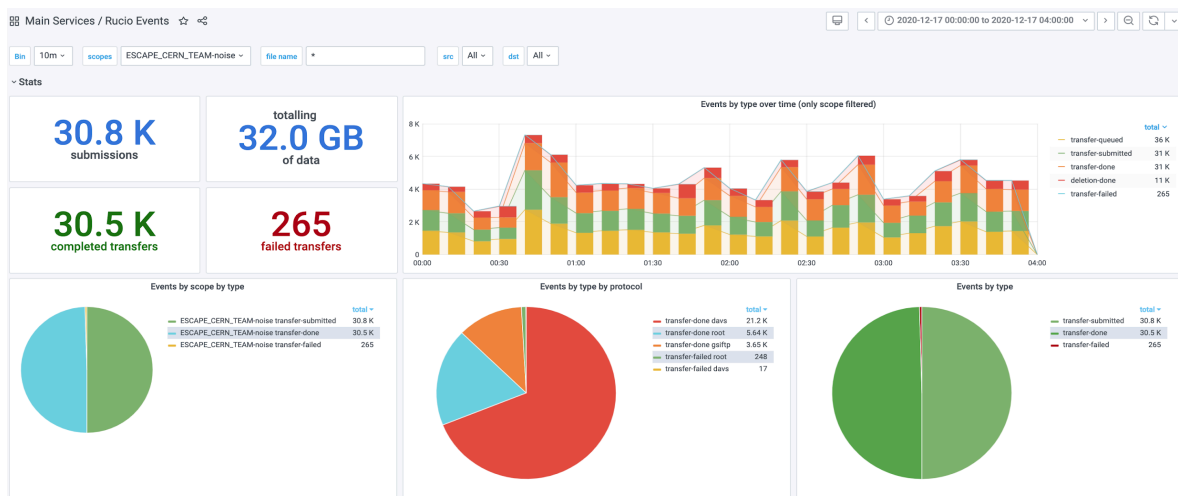


Figure 9: Dashboard view of the automatic Rucio-level tests running from CERN for a period midnight to 0400hrs on December 17th 2020. During this time, intensive tests were running to try out different configurations - 31k transfers were submitted in 4 hours.



## D2.2 Assessment and analysis of performance of the first pilot data lake

Transfer Matrix - Replica Creation

src \ dst	DESY-DCACHE	SARA-DCACHE	PIC-DCACHE	EULAKE-1	LAPP-DCACHE	IN2P3-CC-DCACHE	CNAF-STORM	ALPAMED-DPM	GS-ROOT	INFN-NA-DPM	LAPP-WEBDAV	INFN-NA-DPM-FED	INFN-ROMA1	FAIR-ROOT
DESY-DCACHE	NO DATA	101%	60%	100%	100%	100%	100%	92%	100%	99%	100%	NO DATA	NO DATA	100%
SARA-DCACHE	99%	NO DATA	99%	98%	98%	98%	98%	95%	98%	99%	98%	NO DATA	NO DATA	98%
PIC-DCACHE	100%	100%	NO DATA	100%	99%	100%	100%	95%	100%	100%	100%	NO DATA	NO DATA	100%
EULAKE-1	99%	100%	100%	NO DATA	100%	100%	100%	100%	100%	100%	100%	NO DATA	NO DATA	100%
LAPP-DCACHE	101%	100%	100%	101%	NO DATA	100%	100%	100%	100%	99%	100%	NO DATA	NO DATA	100%
IN2P3-CC-DCACHE	102%	100%	101%	101%	102%	NO DATA	101%	98%	101%	102%	102%	NO DATA	NO DATA	102%
CNAF-STORM	100%	100%	100%	98%	100%	100%	NO DATA	100%	97%	100%	99%	NO DATA	NO DATA	100%
ALPAMED-DPM	100%	99%	102%	100%	99%	98%	98%	NO DATA	99%	99%	100%	NO DATA	NO DATA	99%
GS-ROOT	98%	101%	100%	100%	100%	100%	100%	88%	NO DATA	100%	99%	NO DATA	NO DATA	99%
INFN-NA-DPM	101%	102%	101%	101%	101%	101%	101%	101%	101%	NO DATA	102%	NO DATA	NO DATA	101%
LAPP-WEBDAV	100%	100%	102%	101%	100%	100%	102%	100%	100%	100%	NO DATA	NO DATA	NO DATA	101%
INFN-NA-DPM-FED	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA
INFN-ROMA1	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA
FAIR-ROOT	100%	102%	100%	100%	100%	100%	100%	89%	100%	99%	99%	NO DATA	NO DATA	NO DATA

Figure 10: As for previous figure, but showing the transfer matrix. The matrix view clearly shows that INFN-NA-DPM-FED and INFN-ROMA1 were not available at this time.

Link to Rucio testing software: <https://github.com/ESCAPE-WP2/rucio-analysis>

### 8.1.6 The IAM service

Authentication and authorization in the Data Lake relies on assertions issued by the [ESCAPE IAM instance](#)<sup>11</sup>, either in the form of VOMS attribute certificates or JWT tokens. So, for the Data Lake to be functional, it is of crucial importance that the IAM instance is available and works as expected.

In order to continuously probe IAM availability and its ability to issue well-formed tokens and VOMS attribute certificates, an ESCAPE-focused [test suite](#)<sup>12</sup> was developed. A [Github actions workflow](#) was also defined to run availability and functionality tests against ESCAPE IAM every 20 minutes.

Looking at the results of those tests, and also considering the general Data Lake monitoring, ESCAPE IAM has been working very reliably with the exception of a single period of five days downtime (from November 25th to November 29th) caused by a serious storage incident at INFN CNAF, as summarized in this [post-mortem report](#)<sup>13</sup>.

In order to avoid instabilities in the future, work has started to realize a highly-available deployment of the ESCAPE IAM spanning multiple INFN data centres. It is anticipated that this setup will be ready by the second quarter of 2021 (see the “Continuous Improvement” section on page 43).

<sup>11</sup> <https://iam-escape.cloud.cnaf.infn.it/>

<sup>12</sup> <https://github.com/indigo-iam/escape-auth-tests>

<sup>13</sup> <https://jira.skatelescope.org/browse/EDLK-106>



## 8.2 Live tests with Experiment dataflows

As mentioned in the introduction, in November and December 2020 we had two focused windows of general testing from experiment representatives, in addition to the continuous technical tests that we run.

The goal was to undertake multiple experiment data workflows on a single day, testing data injection, data replication via rules, including QoS-based rules, and data access. We were keen to establish the perspective from the experiment representatives (e.g. "how usable is Rucio?"), from site maintainers (e.g. "what does supporting my site as a Rucio storage element entail?"), and to assess the ESCAPE Data Lake tools and services (Rucio, FTS, CRIC, IAM) themselves under pseudo-production conditions (e.g. "Is the service able to cope with tests at this scale?", "Are the dashboards giving me the information I need to see if the Data Lake is healthy?").

In between these two testing windows, we held our second WP2 workshop. This workshop was held virtually due to travel restrictions in place across Europe. We took the opportunity to record the sessions and identified a scientific writing team who took extensive minutes, helped capture actions and who have ensured that every presentation from the workshop is available publicly online, alongside recordings of the talks - please see workshop indico site to view these videos (<https://indico.in2p3.fr/event/22693/>.)

### 8.2.1 Overview

#### 8.2.1.1 17th November testing window

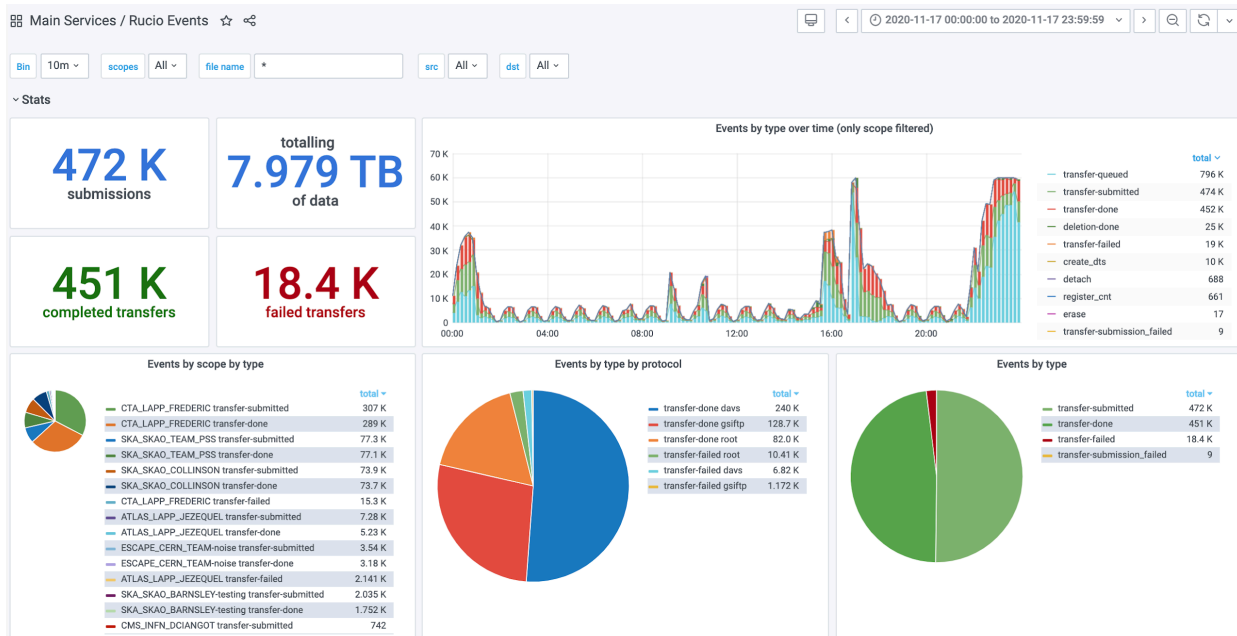


Figure 11: Dashboard view from November 17th showing an overview of the Rucio activity.

On Nov 17th, 9 experiments participated in our Data Lake tests, making use of 12 RSEs. In the 24 hours of Nov 17th, 8 TB of data were transferred (18 TB including the run-off into the next day). The dashboard view image above (Figure 11) shows the overall Rucio-level activity, with almost half-a-

D2.2 Assessment and analysis of performance of the first pilot data lake

million events submitted to the Rucio server. Of these, 96% were successful and 4% failed. Figure 12 shows the corresponding transfer matrix from the Rucio dashboard.

Transfer Matrix - Replica Creation

src	dst	DESY-DCACHE	SARA-DCACHE	PIC-DCACHE	EULAKE-1	LAPP-DCACHE	IN2P3-CC-DCACHE	CNAF-STORM	ALPAMED-DPM	GSI-ROOT	INFN-NA-DPM	LAPP-WEBDAV	INFN-NA-DPM-FED
DESY-DCACHE		NO DATA	100%	51%	100%	104%	100%	100%	93%	35%	98%	100%	100%
SARA-DCACHE		100%	NO DATA	98%	100%	100%	100%	98%	88%	26%	98%	98%	96%
PIC-DCACHE		100%	100%	NO DATA	99%	100%	100%	100%	100%	23%	100%	100%	96%
EULAKE-1		100%	75%	47%	NO DATA	100%	100%	100%	100%	42%	100%	100%	100%
LAPP-DCACHE		100%	100%	99%	100%	NO DATA	98%	100%	98%	16%	98%	94%	96%
IN2P3-CC-DCACHE		100%	100%	89%	100%	100%	NO DATA	100%	91%	35%	98%	100%	100%
CNAF-STORM		100%	100%	98%	100%	100%	97%	NO DATA	100%	18%	100%	100%	100%
ALPAMED-DPM		28%	94%	100%	100%	100%	100%	100%	NO DATA	49%	93%	100%	100%
GSI-ROOT		100%	99%	94%	100%	99%	100%	100%	89%	NO DATA	100%	97%	95%
INFN-NA-DPM		100%	100%	100%	100%	99%	100%	99%	90%	45%	NO DATA	98%	NO DATA
LAPP-WEBDAV		100%	100%	100%	100%	100%	100%	98%	100%	100%	100%	NO DATA	100%
INFN-NA-DPM-FED		100%	100%	96%	100%	93%	100%	96%	81%	40%	NO DATA	96%	NO DATA

Figure 12: Nov 17th testing window Rucio transfer matrix view.

Notes from 17th November:

1. GSI-ROOT was added to the Data Lake as a small test instance, with 1 TB storage available on a single disk. GSI-ROOT's storage became overfull because the automated removal of testing files was not active, since a "Minimum free space" flag had not been set. Moreover, high concurrency of incoming data transfer saturated the storage input bandwidth. This meant that GSI-ROOT was unable to accept new data, though still successfully acted as a data source, hence the red column as a "destination" in the Rucio transfer matrix above. After noting the issue, we disabled GSI-ROOT's QoS flag, and removed GSI-ROOT from the automated tests whilst the data were deleted and a "Minimum free space" limit was set. GSI-ROOT came back online in the early afternoon and performed well. This failure led to a better understanding of the need for "Minimum Free Space" settings, and the rapid deployment of a new RSE at GSI called "FAIR-ROOT". This has a larger capacity (20 TB) and higher I/O performance and was ready in time for the second testing window.
2. We had advertised our testing window in advance to partners but in spite of this, we were unable to avoid maintenance on all our RSEs - on 17th November SURFsara underwent maintenance and was not scheduled to be online until after 17th November. This mainly affected the LOFAR tests, since the plan was to use data on the SURFsara RSE. As it happened, SURFsara actually came back up sooner than expected and performed well during the testing day, but this did still impact the LOFAR tests. We learned from this that we need to see a



## D2.2 Assessment and analysis of performance of the first pilot data lake

common database of RSE maintenance windows and to get downtime notifications - we will address this by using the GOCD<sup>14</sup> from EGI in future.

- The great number of transfers requested during the testing period on 17/11 triggered a first restart of the Rucio authentication server that was promptly spotted and monitored without any consequence for either the infrastructure or the end-users. Rather than allowing the situation to deteriorate, the configurations of the servers were updated and a restart of the services followed. This resulted in a Rucio downtime for the end-users of just a few minutes at 15:30. Afterwards, the Rucio event rate started increasing again and the backlog steadily cleared.




Figure 13: Rucio event activity before, during and after the Rucio server re-configuration.

### 8.2.1.2 15 December Testing Window

During the (optional) second testing window, 6 experiments undertook repeated or extended tests based on the outcomes of the first testing window. Throughout the 24-hour period, almost 30 thousand submissions were made to the Rucio server, totalling almost 9 TB of data. During this day the Data Lake performed smoothly as far as the tests were concerned and there were no issues to note.

<sup>14</sup> <https://wiki.egi.eu/wiki/GOCD>

D2.2 Assessment and analysis of performance of the first pilot data lake

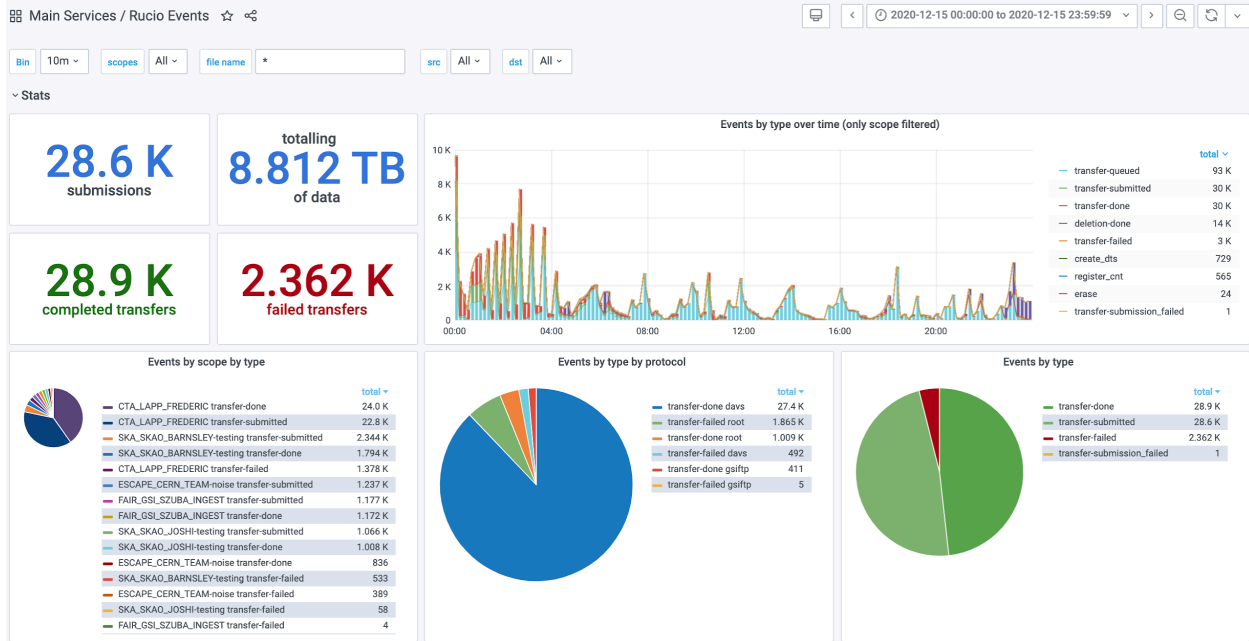


Figure 14: Rucio activity during the 2nd testing window, 15/12/2020

Activity during the second testing window (Figure 14), particularly the stress tests, resulted in a spill over into the following day (16 Dec) – as shown in Figure 15. Rucio event activity over both days shows more than 15 TB transferred, and **more than a million successful transfers**.

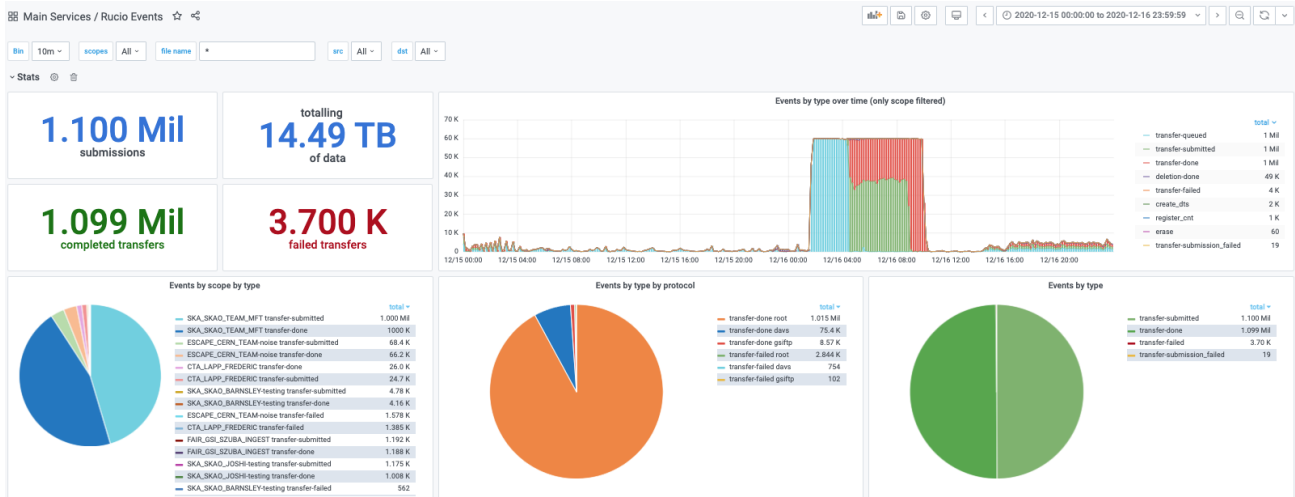


Figure 15: Rucio activity from the second testing window, but with time range extended to show the completion of the “Million file test”.



### 8.2.1.3 Tests performed

We now present the test performed in the two testing windows in detail. Table 2 below summarizes the tests performed by the experiments over the course of the testing windows for Data Lake assessment.

T E S T	Experi- ment	Test name	Brief description of the test  Outcome	Data volume uploaded/ transferred
1	SKA	Pulsar observations	<b>Description:</b> 24 hour test to simulate two sets of pulsar observations where pulsar search (PS) and single pulse search (SPS) observations are occurring commensally. A day's worth of data uploaded and replicated. Data corresponding to PS and SPS projects separated via scopes. <b>Outcome:</b> Partially successful	4.68 TB
2	SKA	QoS lifecycle	<b>Description:</b> 24 hour test uploading and moving data to simulate data lifecycle based on QoS classes. <b>Outcome:</b> Successful	1 GB
3	SKA	SDC1 data	<b>Description:</b> Science Data Challenge (SDC) 1 data replicated to FAST QoS storage to simulate making data products available to the end-user. <b>Outcome:</b> Successful	17 GB
4	CTA	Data Ingest	<b>Description:</b> Simulate ingestion of one night of observations by uploading and replicating data of varying file sizes. <b>Outcome:</b> Successful	7 TB
5	ATLAS	Open Data QoS test	<b>Description:</b> 24 hour test to demonstrate Data Lake reliability along with QoS functionality through continuous uploads and QoS based replications. Target data rate is that of the ATLAS open data workflow (2 GB/min). <b>Outcome:</b> Successful	3 TB
6	CMS	Open data ingest	<b>Description:</b> Import CMS open data, replicate based on QoS and use Rucio for data discovery. <b>Outcome:</b> Successful	300 GB
7	MAGIC	Data injection	<b>Description:</b> 24 hour test simulating data injection by moving data from a non-deterministic RSE to a long-lived copy at a deterministic RSE. Temporary copy deleted once replication complete. <b>Outcome:</b> Successful	338 GB



D2.2 Assessment and analysis of performance of the first pilot data lake

	Experiment	Test name	Brief description of the test  Outcome	Data volume uploaded/ transferred
8	LOFAR	Data ingest	<b>Description:</b> A LOFAR observation was ingested into the Data Lake <b>Outcome:</b> Successful	11 TB
9	LOFAR	Data Processing pipeline	<b>Description:</b> Data downloaded, processed further and outputs uploaded back to the Data Lake. This was carried out automatically in sequence for a list of datasets. The test was repeated on a few different days. Occasional failures mainly due to unscheduled Data Lake maintenance, but ultimately succeeded. <b>Outcome:</b> Successful	450 GB
10	LSST	Data ingest	<b>Description:</b> Data ingestion at LSST production rates by uploading data to an RSE. <b>Outcome:</b> Successful	2.4 TB
11	LSST	Batch data ingest	<b>Description:</b> Data ingestion at LSST production rates by uploading data to an RSE using batch jobs. <b>Outcome:</b> Successful	5.6 TB
12	FAIR	Data ingest I	<b>Description:</b> Data ingestion at a sustainable fraction of the production rates, and replication based on QoS. <b>Outcome:</b> Partially successful	2 TB
13	FAIR	Data ingest II	<b>Description:</b> Continued data ingestion and replication based on QoS over 24 hours. <b>Outcome:</b> Successful	3 TB
14	EGO/VI RGO	Data ingest	<b>Description:</b> Data ingestion (upload) and successful verification of data by downloading all the samples. <b>Outcome:</b> Successful	2 TB
15	Stress test	Million file test	<b>Description:</b> Replicate a million files via a single rule from a fixed source RSE to a destination RSE. <b>Outcome:</b> Successful	1 TB
16	Stress test	High throughput test	<b>Description:</b> Replicate 1 TB of data in the form of a container of 1 GB files from a fixed source RSE to destination to capture throughput achieved. <b>Outcome:</b> Successful	1 TB

Table 2: Outline of use-case-motivated tests performed in this assessment period.

## D2.2 Assessment and analysis of performance of the first pilot data lake

## 8.2.1.4 SKA

SKA will generate 600 PB of Observatory-level data products per annum, from 2028 onwards, and deliver these into a global collaboration of SKA Regional Centres (SRCs). If a global data management service can be used, large cost savings are possible as data will not need to be replicated so much. More detail on SKA's use case is given in the Appendix – see page 52.

In order to emulate data movement of Observatory-level data products to SRCs, pulsar observations tests were conducted (test 1 in Table 2). This test uploaded 50 MB files to two separate Rucio containers<sup>15</sup> in separate scopes at varying cadences to inject data into the Data Lake, followed by replication to another RSE. A DID's scope isolates the namespace into several sub-spaces. The test mimics a data intensive pulsar search (PS) observation with a more lightweight single pulse search (SPS) observation happening commensally. A summary of both the PS and SPS tests are shown below (Figure 16 and Figure 17). A dip in activity can be seen in case of the PSS test that corresponds to the Rucio auth server upgrade in the afternoon. Since each individual run of the SPS test is quite short, it was not impacted by the server upgrade.



Figure 16: Progress of SKA's pulsar search test during the first testing window.

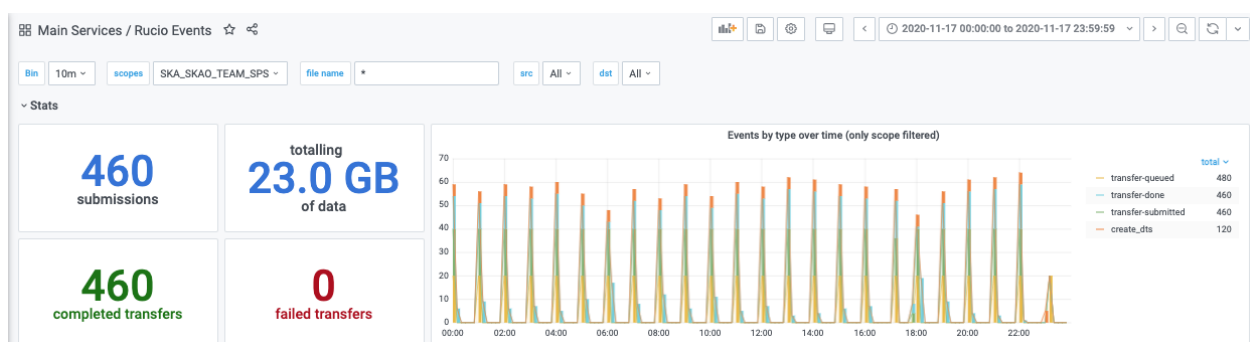


Figure 17: Progress of SKA's single pulse search test during the first testing window.

The second test (number 2 in Table 2) simulated the data lifecycle and movement from the Observatory to/through the SRC network based on Quality of Service (QoS) at the storage sites. The 24 hour test attempts to replicate a data lifecycle hourly based on QoS labels on RSEs in the Data

<sup>15</sup> Files in Rucio can be grouped into collections, namely, datasets and containers. Collections can further be organised freely.

## D2.2 Assessment and analysis of performance of the first pilot data lake

Lake. The test uploads a 100 MB test file to an RSE based on QoS (FAST) with a lifetime of 0.5 week and replicates this with increasing lifetimes to QoS CHEAP-ANALYSIS (1 week's lifetime), OPPORTUNISTIC (1.5 weeks' lifetime) and SAFE (2 weeks' lifetime), with all files for the day grouped in a single Rucio container.

Since replications were based on QoS, there was a good spread of RSEs used over the course of the day and the test was not impacted by issues at GSI-ROOT. Issues at GSI-ROOT can be seen in the form of blue 'transfer-failed' bars in the image below (Figure 18). QoS based selection of RSEs meant that after a failed transfer to an RSE, a different RSE satisfying the QoS condition was automatically selected. The issue was identified at around 8:50 am and no transfers failed in subsequent runs of the test.

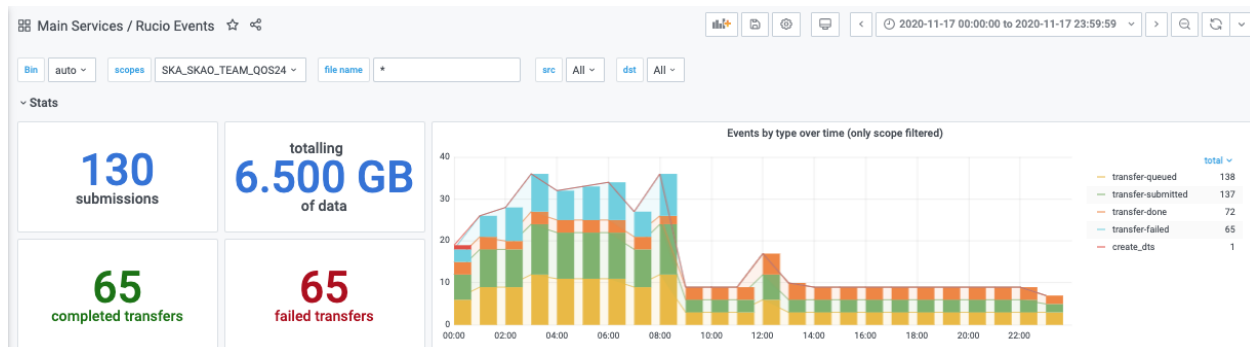


Figure 18: Progress of SKA's 24 hour QoS lifecycle test

The final SKA test (test 3 in Table 2) replicated data from the first Science Data Challenge (SDC1) to storage with FAST QoS to simulate making data available at a site ready for a user to process it efficiently or download it. The test creates a dataset-level rule for moving the files in the Rucio dataset "SKA\_SKAO\_TEAM:SDC1" to an RSE with FAST QoS with a lifetime of two weeks. Transfers were completed as expected on Nov 17<sup>th</sup> (see Figure 19). An interesting point to note is that when new files were added to the SDC1 dataset starting from the 20<sup>th</sup>, this initiated replication of the new files to FAST QoS as well since the test's rule was still active.

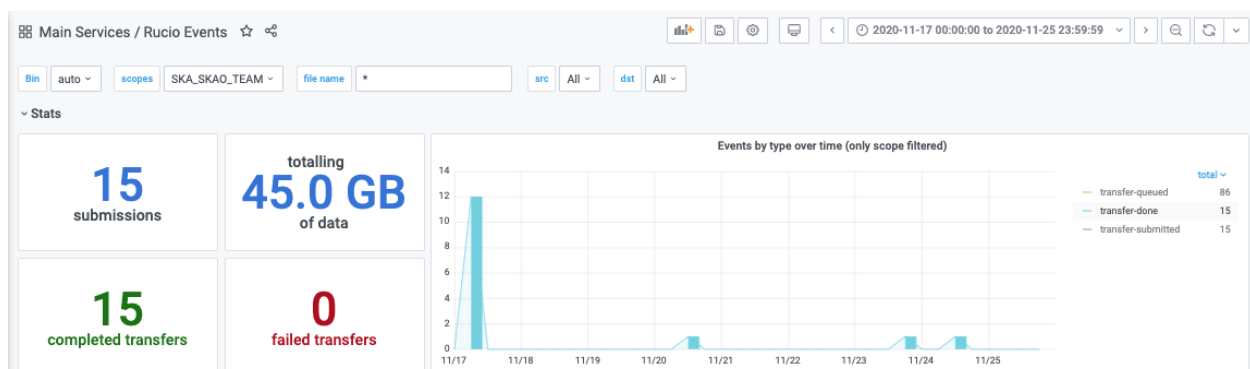


Figure 19: Completed transfers of the SKA SDC1 dataset during and after the first testing window.

### 8.2.1.5 CTA

CTA will generate 6 PB of Observatory-level data products per annum, in its initial phase and 1.7 PB from 2022 onwards for its first threshold phase. During this first phase, secondary data products to be archived will total more than 24 PB in the first year of operation with an additional 2 PB per annum the following years.

## D2.2 Assessment and analysis of performance of the first pilot data lake

To manage this data, two archives have been defined: a Bulk Data Management System and a Science Archive system. The two systems could be based on the same technology. At the moment, the WP2 Data Lake is foreseen as an implementation of the Bulk Data Management System of CTA. This system will serve as an authorised archive of all products from raw data (Data level 0) to science data (Data level 3). More information about the use case can be found in Appendix A (page 46).

In order to mimic the behaviour of a bulk data management system, data of volume equivalent to that of one night of observations were uploaded in the form of 500 datasets each with 10 files (test 4 in Table 2). File sizes varied between 0.01, 0.1 and 2 GB. Data were initially uploaded to IN2P3-CC-DCACHE and later to LAPP-DCACHE. Datasets were then replicated to PIC-DCACHE and ALPAMED-DPM. A small number of replications corresponding to tests run on 17th Nov spilled over into the next day, but the outcome was considered successful (see Figure 20).

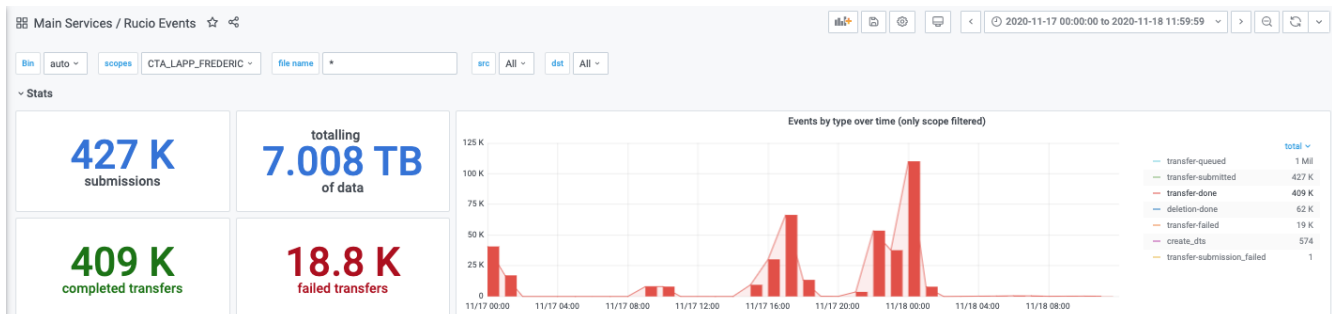


Figure 20: Completed transfers corresponding to 7 TB of CTA's data ingestion and replication during and slightly after the first testing window.

### 8.2.1.6 ATLAS

Both ATLAS<sup>16</sup> and CMS (below) use the Worldwide LHC Computing Grid, WLCG, to manage their data - this is a global project bringing together computational and storage resources offered by funding agencies around the world. Rucio itself was developed by the ATLAS experiment, where it is used in production to manage over 450 PB of data across almost 40 different countries. The ATLAS collaboration includes over 3000 scientific authors. The Rucio-centric Data Lake model is being considered as the data management model for the High Luminosity Large Hadron Collider, HL-LHC, due to be in operation post 2027.

Data reprocessing during HL-LHC (DAOD\_PHYSLITE) will result in a data rate of 1.5 GB/min. The dataset used for testing was a 2 GB reduced, public  $H \rightarrow \gamma\gamma$  analysis dataset and was uploaded once every minute to replicate these conditions. The main interest from ATLAS was to exercise the Storage QoS functionality. The test (test 5, Table 2) uploaded data from the LAPP cluster to two sites: ALPAMED-DPM (France) and INFN-NA-DPM (Italy). Following data upload to these two sites (alternating between them), QoS-based rules were applied, to have one copy of each dataset on storage with SAFE QoS and two copies on storage with CHEAP-ANALYSIS QoS. While there were some concerns about throughput out of ALPAMED-DPM (network saturation due to ATLAS data rebalancing happening concurrently), 7k files were transferred in 24 hours and the QoS transition successfully finished. This can be seen in Figure 21 below. Failed transfers in the morning (shown in

<sup>16</sup> <https://atlas.cern/>

## D2.2 Assessment and analysis of performance of the first pilot data lake

orange on the 'Events by time over time' plot) were largely due to issues at GSI-ROOT, and were resolved after 9 am.

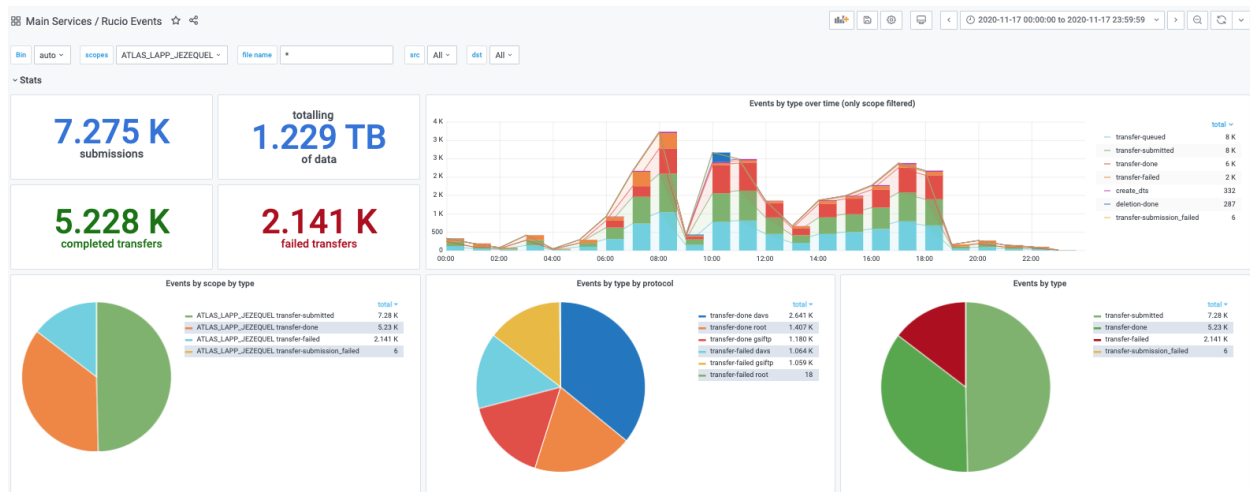


Figure 21: Summary of ATLAS open data transfers initiated by QoS based rules.

### 8.2.1.7 CMS

CMS<sup>17</sup> has been using Rucio for data management in production since late 2020, having successfully migrated from PhEDEx.

The test from CMS (test 6 in Table 2) was based on importing an open data dataset to Rucio from existing grid storage outside the Data Lake with one command line. Replication of the dataset was performed to OPPORTUNISTIC QoS storage and a simple data discovery workflow was executed asking for that dataset. CMS ran this dedicated test for a few hours - in the end achieving ~300 GB of open data correctly imported and analyzed, using Rucio for data discovery of the location. A corresponding increase in the storage used by CMS on Nov 17th can be seen in Figure 22 below.

Tests were performed using the tools found here: [https://github.com/dciangot/rucio\\_escape\\_inject](https://github.com/dciangot/rucio_escape_inject)

<sup>17</sup> <https://cms.cern/>

## D2.2 Assessment and analysis of performance of the first pilot data lake




Figure 22: Increase in CMS storage used corresponding to open data ingestion test during the first testing window.

### 8.2.1.8 MAGIC

The Major Atmospheric Gamma Imaging Cherenkov (MAGIC) telescopes are a system of two Imaging Cherenkov Telescopes located at Roque de los Muchachos Observatory (ORM) on the Canary Island La Palma, Spain. The 17 m telescopes, the largest Cherenkov telescopes in the world, have been designed to study the Universe and discover new Gamma-ray sources. The experiments will generate 127 TB of Observatory-level data products per annum. The MAGIC data centre is located at the Port d'Informació Científica (PIC) in Barcelona.

Data generated at ORM will be moved to the data centre in PIC in order to archive the data and delete it from the ORM data store to free up storage space at the observatory. Further explanation of the MAGIC Data Lake context is given in the Appendix (see page 51). The goal of the data injection exercise (test 7, Table 2) was to test the configuration required for the movement of these large datasets. Both RSEs for the test were located at PIC, a *non-deterministic* RSE emulating ORM storage named PIC-INJECT that allows users to register files with their existing path on storage and a deterministic RSE PIC-DCACHE. Injection was done using 24 files of 600 MB each, replicated every hour from PIC-INJECT to PIC-DCACHE. Once data was replicated to destination, source files at PIC-INJECT were removed.

The progress of the test over the course of the day can be seen in Figure 23, in the form of a steady increase in the overall used storage till approximately midday when a round of file deletion occurs as corroborated by a corresponding dip in storage used at PIC-INJECT. This is simulating freeing up storage at ORM as expected.

## D2.2 Assessment and analysis of performance of the first pilot data lake




Figure 23: MAGIC's usage during the first testing window depicting data ingest and data clean-up.

### 8.2.1.9 LOFAR

LOFAR is a low frequency radio interferometer having antenna fields located across several countries in Europe but with a dense core site in the Netherlands. LOFAR is an SKA pathfinder instrument and has been functional for several years undertaking a broad observational programme. Initially data are accessible only to the proposing scientists but after some time data becomes available for public use. The call for proposals is open to astronomers all around the world. Further description of LOFAR's data management context is given in Section 10.4.

LOFAR's Long Term Archive (LTA) is used to store both open and proprietary data of each LOFAR project and access is granted to the appropriate group of people. Data is made public after the embargo period is complete.

A data ingestion test was performed during the weeks leading up to the first testing window. A LOFAR observation dataset totalling 11 TB, was uploaded from LTA (Poznan) to SARA-DCACHE.

In order to prepare for the data processing test, a smaller separate 150 GB dataset comprising 23 frequency sub-bands (SBs) was uploaded to the SARA-DCACHE RSE with a replica made at EULAKE-1.

LOFAR's testing plans for both tests were affected by several issues just before and during the day of the first dress rehearsal (Nov 17th) including urgent maintenance at SURFsara, EULAKE-1 download failures, and upload timeouts on large files (due to time taken during Rucio's pre-upload checksum computation). This necessitated adapting the test plans.

## D2.2 Assessment and analysis of performance of the first pilot data lake

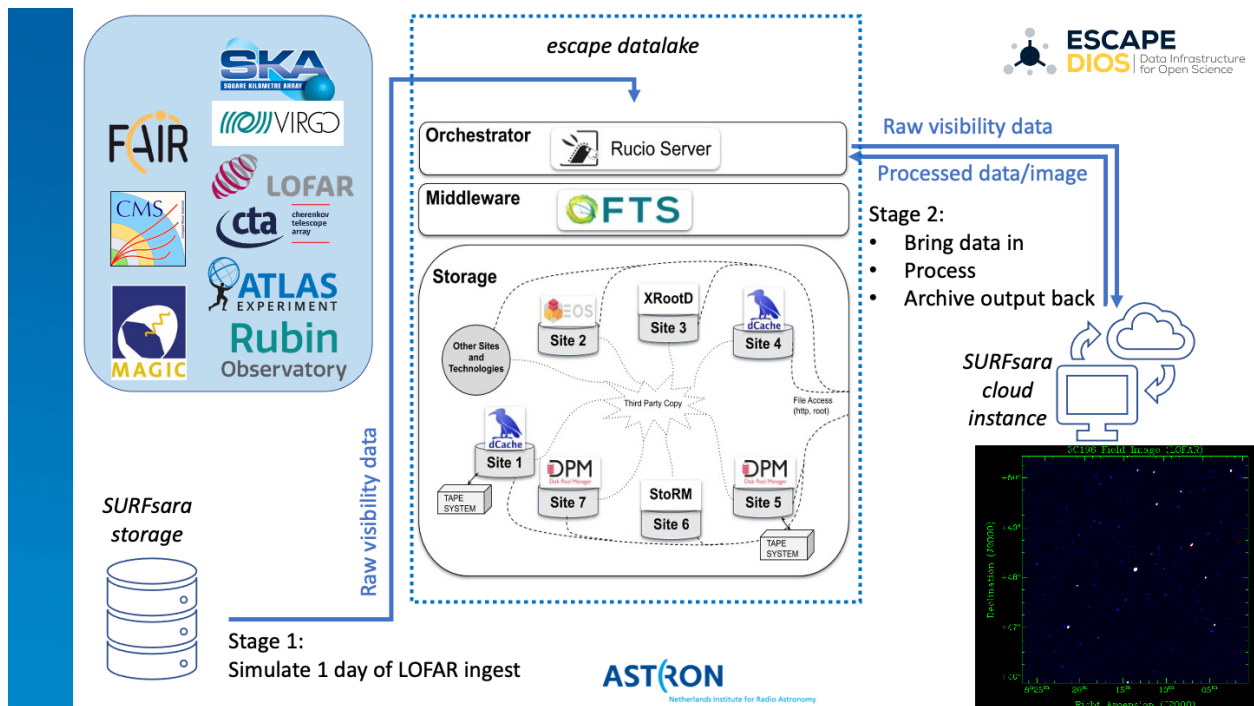


Figure 24: Schematic of the relationship between the LOFAR Ingest (Stage 1) and Processing (Stage 2) tests.

The Data ingest (Stage 1) and workflow Processing tests (Stage 2) using the Data Lake for LOFAR data are shown in Figure 24. The resulting image at the right-hand bottom corner is of the 3C196 field at 135 MHz using 23SBs in the Stage 2 tests.

The ingest test (test 8 in Table 2) was carried out using simulated datasets (still mimicking the complexities and size of a full LOFAR dataset) stored at the SURFsara cloud instance, which were uploaded to EULAKE-1 with different approaches (e.g. a single dataset at a time vs multiple datasets transferred together).

For the data processing test (test 9 in Table 2), to simulate a workflow data were downloaded from the Data Lake, processed on the SURFsara cloud instance, and the outputs (calibrated data set and images) were then uploaded to the Data Lake to be archived. The data processing test used, in sequence, all 23 sub-bands (SBs) at EULAKE-1 and the results were also archived back to EULAKE-1. The full test was repeated randomly on a number of days, and once the issues causing partial failures were understood and resolved, all the 23 sub-bands were (consistently) successfully processed, imaged and archived, indicating that there are no fundamental blockers to this test.

Going forward, a smoother production workflow and emulating an LTA endpoint on the Data Lake in combination with QoS functionality of the Data Lake could be quite useful. In addition, the scale and complexity for the data processing test will be increased towards a more realistic scenario.

### 8.2.1.10 LSST / Rubin Observatory

The LSST use case for a Data Lake is given in Appendix A – see page 53.



## D2.2 Assessment and analysis of performance of the first pilot data lake

The goal for the tests performed in the assessment (tests 10 and 11 in Table 2) was to ingest data at realistic LSST data rates. The expected data rate is  $\sim 20$  TB/day, which translates to  $\sim 230$  MBps. While the initial plan for the first test was to simulate data ingestion with data generated in a batch farm, this plan was pushed to the second test. Three runs were conducted in the first test, each run uploading a copy of the  $\sim 800$  GB HSC RC2 dataset (originally from the Subaru Telescope) from local IN2P3-CC machines to IN2P3-CC-DCACHE (with no replication). The image below (Figure 25) depicts LSST storage usage increasing throughout the day and a corresponding increase at IN2P3-CC-DCACHE. Each run uploaded  $\sim 800$  GB in 40-90 min with some runs having a longer tail than others resulting in a rate of  $\sim 313$  MBps at peak activity, which is substantially higher than the expected production rate of LSST.




Figure 25: LSST's (Rubin Observatory's) usage pattern during their first data ingest test (bottom) and a similar usage pattern reflected at IN2P3-CC-DCACHE, their primary ingest RSE (top).

The second test was run using a different ingestion strategy than the first with batch jobs running on standard workers instead of interactive scripts on IN2P3-CC Data Transfer Nodes. This resulted in a significantly higher sustained data rate of 3.8 Gbps over 2+ hours. In total, the HSC RC2 dataset was ingested 7 times over the course of the day,  $\sim 315$  K files and 5.6 TB as seen in Figure 26.

D2.2 Assessment and analysis of performance of the first pilot data lake




Figure 26: LSST's usage pattern during their second data ingest test. Note the scale of the second test was larger than the first.

### 8.2.1.11 FAIR

The Facility for Antiproton and Ion Research (FAIR) experiment is an upcoming international accelerator centre in Darmstadt, Germany. Of the four experimental pillars of the project - PANDA (antimatter studies), CBM (heavy-ion physics), NUSTAR (rare/exotic beams for astrophysics) and APPA (Atomic, Plasma Physics and Applications), tests were conducted for simulated data ingest of CBM (Compressed Baryonic Matter) data.

CBM, a modular detector, will entail raw-data ingestion at the rate of 200 TB/day with 3-month runs occurring every year. Once ingested, data will need to be archived and transferred based on QoS. The PANDA experiment will generate 1 PB/year raw data which has to be securely stored in a cold storage, mirrored on different computing centres, and, in addition, needs to be accessible with a low latency for reprocessing in a hot storage. In addition, up to 2 PB/year high-level data is generated with which users perform physics analysis. This data is complemented by 3 PB/year simulated data which is



## D2.2 Assessment and analysis of performance of the first pilot data lake

needed for the analysis. Both the high-level data and the simulated data must be available on a time scale of 10 years.

FAIR anticipates needing to store about 100 PB annually at between 5-8 centres for access by 900 science users. It is a good size fit to the current scale of the WP2 Data Lake.

The first test (test 12 in Table 2) demonstrated data ingest at a manageable fraction of the expected ingest rates in production. 1 GB files were uploaded to GSI-ROOT (failures in the first half of the day), and further replicated to storage with SAFE and CHEAP-ANALYSIS QoS. Since GSI-ROOT has the QoS label CHEAP-ANALYSIS, only replication into SAFE QoS was required as Rucio attempts to minimise the number of transfers.

Although the test commenced late on Nov 16th, issues at GSI-ROOT (see point 1, page 21) caused all uploads to be stuck and no activity was seen until the issue was worked around at about midday Nov 17th. Uploads then began at a frequency of every 10 minutes. This rate was doubled to once every 5 mins after 13:30. The Rucio server upgrade caused a dip in transfers at about 15:30 as can be seen in Figure 27. The test proceeded as expected for the rest of the day.




Figure 27: FAIR's first data ingest test began at around 12 pm following issues at GSI-ROOT. Transfers queued, submitted, done and failed are shown here.

The second data ingest test (test 13 in Table 2) was a 24-hour test conducted on Dec 15th. Uploads were performed to FAIR-ROOT at an increased frequency, a new file every 30-60 sec. Files were then replicated to storage with QoS SAFE and CHEAP-ANALYSIS as in the previous test. The second round of data ingest went smoothly as can be seen in Figure 28 below, and the new FAIR-ROOT RSE performed well. There was a noticeable sustained drop in the activity rate in the afternoon due to a Rucio server stress test that began around this time (see the Million file dataset stress test, page 37). A smaller dip can also be seen around 10:30 am that coincides with the high throughput stress test that replicated data to FAIR-ROOT (see page 39).

## D2.2 Assessment and analysis of performance of the first pilot data lake




Figure 28: FAIR's second data ingest ran smoothly over 24 hours with noticeable dips coinciding with stress tests running on the same day.

### 8.2.1.12 EGO/Virgo

The data management context for the Gravitational Wave experiment, EGO/VIRGO, is given in Appendix A, see page 48.

The goal of the test performed (number 14 in Table 2) was to ingest public VIRGO data into the Data Lake and download it for verification. A 4-hour dataset was uploaded to the EULAKE-1 RSE containing 14,400 one second data chunks. The test's progress on Nov 17th can be seen in the form of a corresponding 14.4 k increase in the number of DIDs associated with the VIRGO experiment, shown in Figure 29.

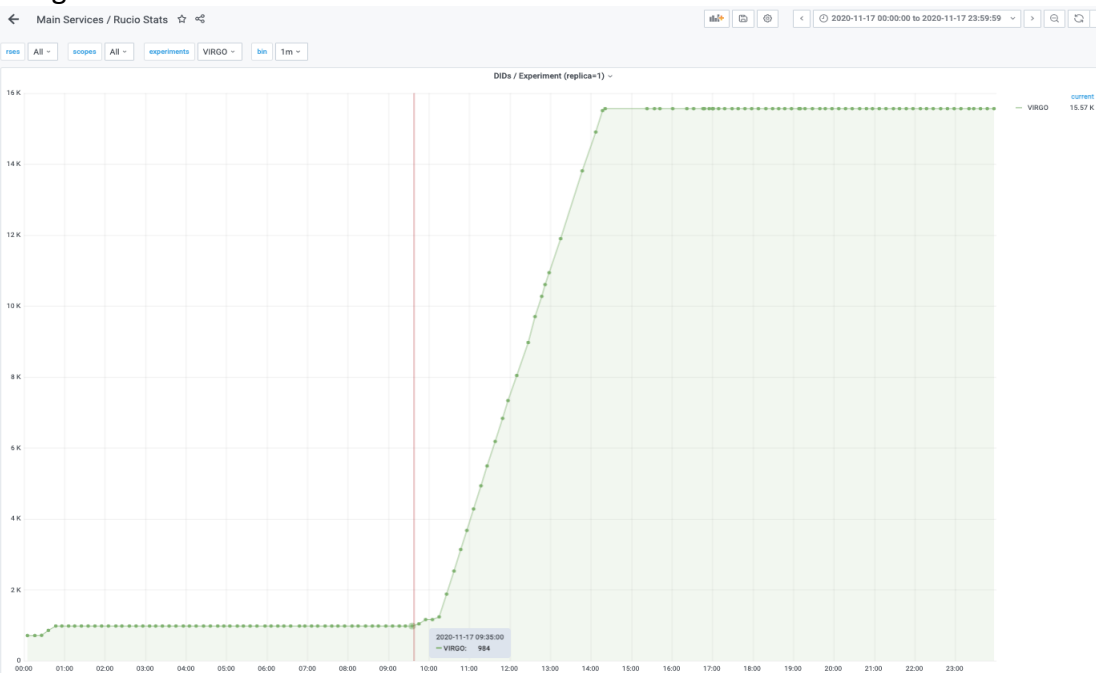


Figure 29: VIRGO's data ingest can be seen here in the form of increasing DIDs over the course of the 1st testing window.



## 8.2.2 Functional testing

### 8.2.2.1 *Million file dataset stress test*

The goal of this test (test 15 in Table 2) was to test Rucio and the surrounding architecture with a heavy payload in terms of number of files. Whilst large volumes of data (and large file sizes) are useful for analysing the behaviour of the Data Lake's storage endpoints, network connectivity, packet loss, etc., a test that involves a large number of individual (small) files tests the resilience and performance of the middleware infrastructure i.e., how well Rucio and the surrounding technologies manage a large number of individual entities (files).

#### 8.2.2.1.1 Preparation

Since upload times, overheads, and timeouts can become troublesome, the test aims to transfer pre-uploaded 1 million files. The files were uploaded over the course of 9 days (13th-21st Nov) to several RSEs for redundancy – see Figure 30. Files were 1 MB in size and grouped together in a Rucio container SKA\_SKAO\_TEAM\_MFT:million with a long-term copy of the container at EULAKE-1.

During this preparatory phase, it was discovered that Rucio's injector daemon was not able to natively cope with handling a replication rule with a million files even if it were given large amounts of resources. Through excellent collaboration from the Rucio development team we were able to deploy a new configuration that more efficiently handled the daemon's memory usage. The daemon attempted to process all the files in a rule all at once thus running into memory issues for rules with large file numbers. The new config works well because it processes the files in a rule in batches, thus allowing the injector daemon to, ideally, deal with any number of files in a rule. This new configuration was already built but had not been tested by the Rucio team - in our Million-file test it was found to be very effective and as a result, this new configuration has been retained.

D2.2 Assessment and analysis of performance of the first pilot data lake




Figure 30: Plot shows increase in the number of files during the preparation and execution of the million-file test. Note the horizontal axis covers 13 November to 17 December

8.2.2.1.2 Execution

A single container-level rule was created to move the files from EULAKE-1 to IN2P3-CC-DCACHE, and the test kicked off at 12:30 pm 15th Dec (Figure 31, (below) shows the progress of the test). The rule was in INJECT state for about 12 hours before beginning to queue transfers in the early hours of the 16th, as seen by the blue ‘transfers-queued’ in the image. The transfers were then submitted to FTS (seen in green, concluding by ~9 am) and the transfers were completed (and the test finished) at about 10:30 am 16th Dec (transfers-done can be seen in red). The image above reflects the corresponding increase in the number of files at IN2P3-CC-DCACHE.

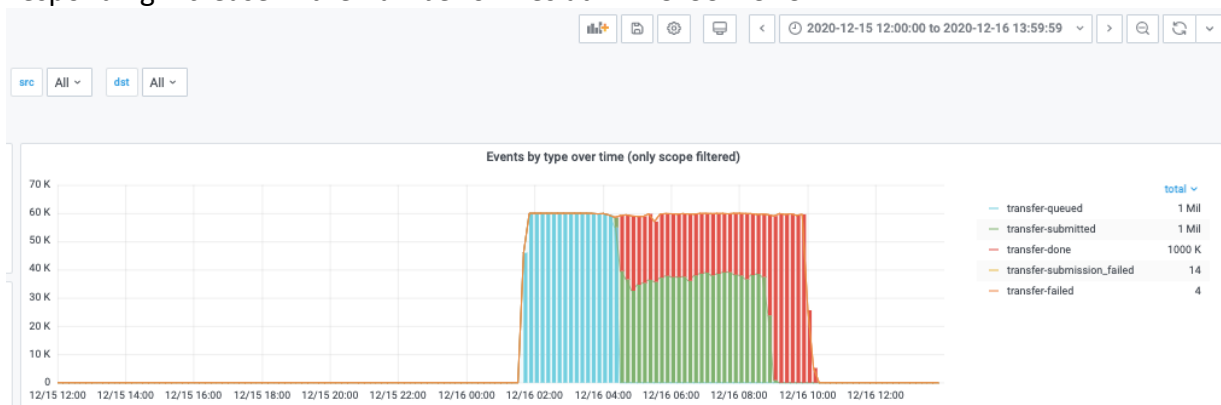


Figure 31: Progress of the transfers corresponding to the million-file test. Transfers queued and completed the following day.



### 8.2.2.1.3 Results

The test successfully completed on Dec 16th. The Rucio ecosystem was able to adequately cope with the payload of the test. Once the rule had been created and the test began at 12:30 pm on Dec 15<sup>th</sup> it drastically increased the load on the Rucio server and specifically the injector daemon. While Rucio handled the simultaneous payload of the million-file test in combination with other tests, the former did have an impact on other activities that occurred during the afternoon/evening that day. Noticeable dips were seen in other tests' activity due to the million-file test, for example as seen in FAIR's test in Figure 28.

#### 8.2.2.2 High-throughput test - measure overheads for Rucio vs FTS

The goal of this test was to move a significant amount of data (1 TB) from one RSE to another to capture the throughput as seen by FTS, Rucio, and perfSONAR. Since we wanted to focus on the throughput aspect and not burden Rucio with a large number of files, the individual file sizes used for this was 1 GB with 1008 files grouped in a ~1 TB Rucio container, SKA\_SKAO\_JOSHI-testing:1G-files-stress-test.

On the day of the test, the Rucio container was replicated from EULAKE-1 to FAIR-ROOT and the test completed in about 20 min – see Figure 32.




Figure 32: The high throughput test moved 1 TB of data from EULAKE-1 to FAIR-ROOT

The theoretical maximum throughput that can be achieved is limited by a single 10 GE interface for TCP/IP connectivity at FAIR-ROOT and thus is **10 Gbps**. The average throughput from the perfSONAR dashboard captured right before running the test was **2.7 Gbps** (337 MBps). However, the test demonstrated a significantly better FTS throughput of **7.8 Gbps** (978 MBps) as shown in the screenshot below (Figure 33). Lastly, based on the timestamps of the rule creation and completion (21 minutes apart), the Rucio perceived throughput was **6.4 Gbps (798 MBps)**. As expected, the Rucio perceived throughput is a bit lower due to overheads both in Rucio and in the FTS queue.

Source	Destination	VO	Submitted	Active	Staging	S.Active	Archiving	Finished	Failed	Cancel	Rate (last 1h)	Thr.
davs://eoseulake.cern.ch	davs://lxwp2d1d1.gsi.de	escape	1008	6	0	0	0	1008	6	0	99.41 %	978.63 MB/s
			0	0	0	0	0	1008	6	0	99.41 %	-

Figure 33: Output from FTS logs showing the achieved throughput for the transfer of the 1008 files in the High Throughput Test.

Future tests should measure the Rucio overhead for multiple file sizes and run consistently to understand the distribution of values.

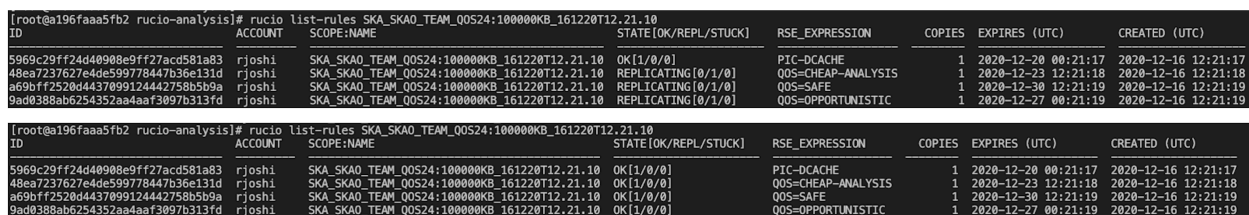
### 8.2.2.3 Testing QoS rules and data lifecycle simulation

Upon consultation with experiments and discussing their use-cases, an initial set of experiment-agnostic QoS classes were established: FAST, CHEAP-ANALYSIS, OPPORTUNISTIC, and SAFE. The classes largely cover the different stages of the data's lifecycle as applicable to an experiment. Loosely speaking, the FAST and CHEAP-ANALYSIS labels refer to "hot" and "warm" storage respectively where it is likely that data on the storage is being interactively used either by an end-user or a process. The OPPORTUNISTIC label refers to "cooler" storage and the SAFE label was created for the general use case of permanent data storage and archival.

For most experiments, a data product can be seen to move through storages with all these QoS labels as it progresses through its lifecycle, from ingestion to archival. **Thus, a generic test to simulate data lifecycle using QoS labels was conducted.** It involves a 100MB file uploaded to an RSE based on QoS with rules created for replicating it to storage with varying QoS labels and varying lifetimes. For simplicity an order is assumed (which is sufficient for this exercise but can also be trivially changed): upload to FAST QoS and retain for 0.5 week, replicate to CHEAP-ANALYSIS for 1 week, replicate to OPPORTUNISTIC for 1.5 weeks, and finally replicate to SAFE storage for 2 weeks. This order aims to show data moving from "hot" storage early in its lifecycle where it is likely to be heavily used, up to "cold" storage where it is archived.

At present rules cannot be scheduled for creation in the future<sup>18</sup>, so all rules were instantiated when the test ran, and file replicas placed on all QoS classes, simply with varying expiration dates (shorter for the "CHEAP ANALYSIS", longest lifetime for the "SAFE" QoS).

The QoS requirements were successfully interpreted and file copies placed according to these, with rule lifetimes as specified (see Figure 34). The missing functionality we identified in developing and running this test was the ability to have a future rule to initiate a file replication to a particular QoS class at some future date - for example, to plan ahead to create a copy of a file on tape archive only after it is (say) 1 year old.



ID	ACCOUNT	SCOPE:NAME	STATE[OK/REPL/STUCK]	RSE_EXPRESSION	COPIES	EXPIRES (UTC)	CREATED (UTC)
5969c29ff24d40908e9ff27acd581a83	rjoshi	SKA_SKAO_TEAM_Q0524:100000KB_161220T12.21.10	OK[1/0/0]	PTC-DCACHE	1	2020-12-20 00:21:17	2020-12-16 12:21:17
48ea7237627e4de599778447b36e131d	rjoshi	SKA_SKAO_TEAM_Q0524:100000KB_161220T12.21.10	REPLICATING[0/1/0]	QOS-CHEAP-ANALYSIS	1	2020-12-23 12:21:18	2020-12-16 12:21:18
a69bfff2520d443709912442758b5b9a	rjoshi	SKA_SKAO_TEAM_Q0524:100000KB_161220T12.21.10	REPLICATING[0/1/0]	QOS-SAFE	1	2020-12-30 12:21:19	2020-12-16 12:21:19
9ad0388ab6254352aa4af3097b313fd	rjoshi	SKA_SKAO_TEAM_Q0524:100000KB_161220T12.21.10	REPLICATING[0/1/0]	QOS-OPPORTUNISTIC	1	2020-12-27 00:21:19	2020-12-16 12:21:19

ID	ACCOUNT	SCOPE:NAME	STATE[OK/REPL/STUCK]	RSE_EXPRESSION	COPIES	EXPIRES (UTC)	CREATED (UTC)
5969c29ff24d40908e9ff27acd581a83	rjoshi	SKA_SKAO_TEAM_Q0524:100000KB_161220T12.21.10	OK[1/0/0]	PTC-DCACHE	1	2020-12-20 00:21:17	2020-12-16 12:21:17
48ea7237627e4de599778447b36e131d	rjoshi	SKA_SKAO_TEAM_Q0524:100000KB_161220T12.21.10	OK[1/0/0]	QOS-CHEAP-ANALYSIS	1	2020-12-23 12:21:18	2020-12-16 12:21:18
a69bfff2520d443709912442758b5b9a	rjoshi	SKA_SKAO_TEAM_Q0524:100000KB_161220T12.21.10	OK[1/0/0]	QOS-SAFE	1	2020-12-30 12:21:19	2020-12-16 12:21:19
9ad0388ab6254352aa4af3097b313fd	rjoshi	SKA_SKAO_TEAM_Q0524:100000KB_161220T12.21.10	OK[1/0/0]	QOS-OPPORTUNISTIC	1	2020-12-27 00:21:19	2020-12-16 12:21:19

Figure 34: Image showing (top) that just after creation, all but one rule (the RSE to which the file has been uploaded) is in replicating state. A few moments later (bottom), the rules have all moved to OK state. They can be seen to have varying expiration dates as required in the lifecycle.

More generally, the QoS classes have helped shape several tests performed by the experiments. SKA tests involved moving data to FAST QoS storage to simulate making data "available" to the end-user. ATLAS open data was moved to SAFE storage and two copies on CHEAP-ANALYSIS storage. CMS data was replicated to OPPORTUNISTIC storage followed by a workflow that discovers the data. FAIR tests replicated data to SAFE storage to simulate data archival. Lastly, the LOFAR data ingest will also incorporate QoS based rules to simulate data lifecycle in the future.

<sup>18</sup> This has been identified in a later section as an action item.





RSEs advertising particular QoS classes could also remove this designation and reduce traffic in - we took this approach with the GSI-ROOT RSE which became overloaded. By removing the QoS flag from this RSE all tests with rules based just on QoS continued fine using alternative RSEs with the same QoS class, highlighting the power of QoS-based rules as opposed to RSE-specific rules.

### 8.3 General Comments / Performance of the technology stack

For this asynchronous system, overall reliability is much more important than speed - data placement is planned to provide good user experience and it is not anticipated that Rucio would be used to initiate urgent transfer requests. Instead, a system that does, eventually, get into the required state (where every Rucio rule is "OK") is crucial. Many of the Rucio features have been developed to achieve this, in the heterogeneous and very complex environment of the WLCG.

It is very advantageous to have **multiple protocols supported** by as many sites as possible. We have observed occasions where some protocols fail but Rucio still succeeds because an alternative protocol is automatically attempted. Sites that support only one protocol are at more risk to absolute failure at software or hardware upgrade points.

We have observed some failures associated with a particular storage technology (in this case DCache) and the way that the file transfer service used a particular protocol, to interact with the RSEs using that technology. If our Data Lake was fully homogeneous, we would have seen failure across the board, our heterogeneous system was able to flag this issue. However, it required the combination of Rucio and FTS-level monitoring to spot that even though Rucio transfers did succeed, there was an issue to act on.

Collection level rules also proved to be very useful - one nice example being the (accidental), automatic management of replicas when adding new files to an existing dataset covered by a rule (SKA added new files to a dataset by uploading to a specific RSE). This automatically replicated the file across the other RSEs used by that dataset - exactly as it should, but without requiring any specific action to initiate that replication process.

The million-file test was a very interesting learning exercise, and truly behaved as a Rucio stress test. While Rucio is being used in production by several communities, a single rule initiating the transfer of a million files, had never been done before and heavily pushed the Rucio deployment.

When preparing for this exercise, it was found that if a single rule has a large number of files to process, the injector daemon runs out of memory and crashes before it has a chance to process the rule. This issue was taken to the Rucio development team, they were aware of this possibility and had previously built a more memory-efficient configuration for this daemon. However, it had never been deployed or tested on any Rucio deployment. This new and improved configuration was deployed and tested on the ESCAPE Data Lake and was found to work very well. This configuration has since been retained, strengthening our deployment while giving back to the Rucio community and demonstrating the benefits of the ESCAPE Data Lake as an excellent testbed.

The GSI-ROOT RSE was deployed as a small test instance only with intention to upgrade when need arose. As a result of the inundation of this RSE during the first testing window and from the increased

load of the automatic testing framework, GSI rapidly and successfully configured a new, larger RSE, called "FAIR-ROOT" and were able to include it in the Data Lake before the second testing window. As has been previously mentioned, this new RSE performed very well during the second testing window. This highlights the speed at which a skilled team was able to deploy new RSEs, and Rucio's flexibility to cope with such changes.

## 9 Improvements identified for future work and next steps

Here we gather the feedback from our Science Experiment partners, Operations team and Site representatives, and group it into common themes. After general feedback we list numbered action points to pick up in our future work.

### 9.1 Feedback from the experiments and how it will be addressed

Experiment tests and stress tests conducted over the course of the testing windows were largely successful and served as an extensive learning exercise for everyone involved. Several common themes emerged in the feedback received from the experiments. They are listed below as well as ways to address them that have been implemented/will be implemented in the near future.

- **Need for Embargoed data:** A large number of experiments (CTA, SKA, MAGIC, EGO/VIRGO, LOFAR) have/will have proprietary data that needs access control. This will become an increasingly important requirement to proceed with Rucio testing. The specific nature of data protection varies from experiment to experiment depending on where Rucio will sit in their use case. For example, CTA's bulk data management system will extend to CTA observatory users only and not end science users. This is different to SKA, where the data lifecycle in the SRCs extends from data visible only to the PI and their collaborators, up to the data becoming public and visible to all. Several experiments are looking into deploying their own Rucio deployments, in order to learn more about Rucio and have more control over their data. These needs will be addressed going forward and steps to do so have been outlined in the following section (action point 4).
- **Rucio availability:** The Rucio server was taken down momentarily during the first testing window as described earlier (Note 3, page 22). This impacted the tests run by SKA, CTA, MAGIC and FAIR. This highlighted the need for improving redundancy in the Data Lake infrastructure. This has been addressed by deploying two replicas of the Rucio server and Rucio auth server in the Kubernetes cluster. One of our goals in ESCAPE is to understand the minimum required Rucio server configuration, so exercises that stretch the capabilities of the existing configuration are valuable.
- **Site maintenance and downtime communication:** As described earlier (note 2, page 21), SURFsara was undergoing maintenance during the first testing window. This primarily impacted the LOFAR tests, but also required a change of plans for SKA tests. This will be addressed in the future by utilising the WLCG notification and ticketing system. This is described further in action point 16.
- **File deletion:** There are several data ingest use cases across experiments that involve moving data from a telescope site where storage is limited, to a data centre for post-processing, user

interaction and archival. This flow of data necessitates data deletion from the telescope site once it has been safely replicated elsewhere. This is applicable to the CTA, MAGIC, LOFAR, and SKA experiments and the MAGIC data ingestion test explicitly evaluated data deletion in the Data Lake. Data clean-up doesn't always happen as expected and this will be addressed in the future as described in the action point 2.

- **Uploading large files, and large volume Data Sets to the Data Lake:** Individual files with large file sizes are quite common in radio astronomy as a result of the data processing pipelines. Uploading large files via the Rucio client involves a preliminary step that performs a checksum on the files, which can result in a timeout if the file is large. This impacted the LOFAR tests (using ~100 GB files). Also, SKA's pulsar observation tests had to be scaled down due to the issue of upload timeouts and delays. This needs to be understood further and will be addressed in the future by including non-deterministic RSEs for use cases involving ingestion of large data volumes.
- **Flat namespaces:** Rucio identifies data with Data Identifiers (DIDs) composed of a scope and name thus resulting in a flat namespace. DIDs are unique over all time. Grouping can be done with datasets and containers in Rucio but data are not managed in a hierarchical structure. The lack of POSIX-like file systems in combination with DID uniqueness can be tricky to work with, as has been identified by the LOFAR and SKA experiments. This can be addressed by establishing a naming convention, but this is simply a new step for some users.

## 9.2 Next steps: Continuous improvement

Smooth Data Lake operations depends on the continuous improvement of existing machinery including IAM, testing, monitoring, and better understanding of file deletion. These are described below.

1. Over the course of the Data Lake's young life, there has been one significant failure of the IAM instance. While service remained stable during the testing windows, **better resilience of IAM** is required, for which a high availability service is planned along with an IAM monitoring dashboard.
2. **File deletion:** Data deletion is an important aspect of the Data Lake orchestration that has been flagged by experiments as well as several sites, as storage footprints continue to increase. Rucio does not automatically force immediate deletion of files - they simply cease to have a rule requiring their retention. We need to understand better the implications of this, and how to ensure that the daemons responsible for file clean-up are working properly.
3. Testing improvements
  - Our automated testing suite must continue to develop as our storage system grows in scale and as we add more functionality to the Data Lake (each time a new feature is added we should design a test to exercise it).
  - Intuitive dashboards  
A more data-centric view option (showing where copies of data are) would give confidence that Rucio is applying rules correctly.
  - Monitoring of the storage space occupied at each RSE, with alerts if this stays above the acceptable "high water mark" for each one.



- Monitoring of the Rucio Kubernetes cluster itself will alert us to potential problems.

### 9.3 Next steps: Introducing new elements

Our Rucio instance is young and we are not yet taking advantage of all that the Rucio ecosystem can offer. A few points for future improvement pertain to deploying and testing newly available features in Rucio:

#### 4. Support for **embargoed data**

- At present, all experiments share a common deployment such that all experiments can see each other's data. For experiments to be more comfortable introducing proprietary data, we would look to **deploy a multi-VO Rucio instance**. This allows isolation of data, users, config, etc between experiments/VOs (Virtual Organizations) on a single Rucio instance.
- The ability to add **experiment specific policies** have been recently introduced in Rucio in the form of a python "policy package" that is installed globally. This package should contain a permission module implementing experiment specific custom permissions. Multi-VO Rucio support for this feature is in progress.

#### 5. **QoS policy mapping** to QoS classes

Currently QoS support in Rucio is quite new. Our tests have shown that being able to use QoS whilst writing rules is popular and likely to be useful. We would also like to explore the definition of QoS "Policies" (which can be experiment-specific), which then map to particular QoS "Classes" (e.g. "SKA Public User" might be a QoS policy for storage elements suitable for public data access - this might include multiple QoS classes (FAST, OPPORTUNISTIC etc) or the mapping might evolve over time).

#### 6. **Rucio subscriptions** offer a potentially very powerful way to plan ahead for data placement before files are generated - based on metadata matches checked at the point of registration of files or datasets, new rules are generated from a subscription that automatically manages data placement.

### 9.4 Next steps: Nice to haves

#### 7. **Token-based authentication and authorization** and moving away from the traditional x509-based approach has not only been flagged by sites and users during our experiment tests, but has been an aspiration in the wider WLCG community. This effort has been led in part by the DOMA project (<https://twiki.cern.ch/twiki/bin/view/LCG/DomaActivities>). While we do not have ownership of this technology shift, we plan to closely follow the work happening in this area, by maintaining strong ties to the DOMA project. Also, we plan to support this transition to token-based AAI by utilising the ESCAPE Data Lake as a rich testbed.

#### 8. **Rucio development**

- Currently, rules kick in and are in effect from the moment they are created. It would be useful for several workflows to be able to **create rules with a user-defined start timestamp**. Subscriptions, as mentioned above, would also assist these workflows.



- Rules can be created with a lifetime (default expiration is no expiration). At present, the lifetimes are defined as the number of seconds from rule creation. We would like to be able to **create rules with an expiration timestamp** instead of lifetime in seconds. Both of these points above help the planning of data placement using Rucio in absolute time, instead of relative time, and are widely applicable to several communities within and outside ESCAPE.
9. **Data corruption testing**  
We would like to test the presumed resilience that Rucio gives for file loss for this by always having (at least) 2 copies of files, and deleting one without Rucio's knowledge in order to determine if the loss is detected and rectified.
  10. **Real-life QoS classes, including Tape.** The current QoS classes are labels only - they do not formally reflect the specific QoS of each RSE. We would like to include some Tape storage in our Data Lake, labelled with an appropriate QoS class, to give us long-term backups of all experiment data.
  11. **Use of metadata** for data discovery and building metadata-based workflows is of interest to several communities, including SKA, CTA, LOFAR, FAIR, and EGO. Metadata can provide an alternate namespace for data discovery and can allow for data flows to be built based on metadata. While Rucio has support for a native metadata catalogue, an external catalogue can also be supported which will be of interest to communities that want to build a large, custom metadata structure.

## 9.5 Next steps: Continued collaboration

### 12. Complete Experiment Onboarding

Some experiments have not yet been onboarded (KM3NET, JIVE, ELT). We plan to ensure the onboarding process is made available to these experiments and will encourage them to join.

### 13. Increase Rucio experience across ESCAPE partners

Interested experiments are considering deploying their own Rucio instances. This will increase the knowledge base and make it possible to develop guidelines for how a Rucio server should be configured in non-HEP scenarios. In SKA's case we intend to deploy Rucio and use it to share experience both within WP2 but also with partners outside ESCAPE, ideally including sites outside Europe to test SKA data movement.

### 14. Challenging data transfer scenarios including long distance, low bandwidth connections, and handling 'shaky' long-distance connections where the transfer could drop and/or produce failures).

### 15. Rucio Collaboration workshop

We would like to strengthen the collaboration with the Rucio development team so that (1) our ESCAPE WP2 users can understand the features already available in Rucio, and (2) to better enable us to contribute to the Rucio code base to help make the improvements that would benefit the ESCAPE science community.

### 16. Operations improvements

- We need to see a common database of RSE maintenance windows and to get downtime notifications - we will address this by using the GOCDDB from EGI (<https://wiki.egi.eu/wiki/GOCDDB>) in future. This would allow us to draw on the experiences gained in WLCG and enable better communication of downtime plans.



- Similarly, some sites have reported that it would be easier to make use of the GGUS helpdesk / ticketing system since this is already used by most RSEs within the WLCG. We are considering this.

### 17. Documentation

We would like to help develop clear guidance on the requirements for RSEs, and how new RSEs can be onboarded and configured. Also, we would like to expand on the documentation for deploying Rucio, including deploying a stand-alone Rucio. This will assist new communities in experimenting and testing Rucio, which of course will lead in turn to improved documentation.

## 10 Appendix A: Use-case context from selected ESCAPE Experiments

We gathered some useful context from colleagues across several experiments, which helps to understand the motivation for the tests undertaken in the assessment period and for involvement in ESCAPE WP2 in general.

### 10.1 CTA

The CTA (Cherenkov Telescope Array<sup>19</sup>) Observatory will generate 6 PB of Observatory-level data products per annum when it reaches its production phase (2023-2027). But even from as early as 2022, in the pre-production phase, it will produce 1.7 PB per year. Adding the Monte Carlo simulation data (22.5 PB) and all resulting products gives more than 24 PB to be archived the first year of operation and additional 2 PB per annum the following years.

For this first phase, two archives are defined as separate entities and different usages:

- The **Bulk Data Management system**, *accessed by CTA Observatory users only* and dedicated to the archival of all products from raw data (Data Level (DL) 0) to science data (Data Level 3). File metadata could be updated by authorized users.
- The **Science Archive system**, *accessed by science end-users* (physicists, astronomers, etc) and dedicated to the archival of high-level data from science data (DL3), Intermediate quick-look data (DL4), science quick-look data (DL5) and Observatory Science data (DL6).

For the Bulk Data Management system, a planned computing model with 2 to 4 data centres is envisaged with a minimum of 2 replicas per file on tapes. This archive is mainly used for long-term archive of raw data and of upper levels to science data (DL2, DL3) as the result of event reconstruction productions. An annual reprocessing of all data is planned and automatic deletion of old versions (n-2, n-3,...). The science data are then automatically sent to the Science Archive system.

Raw data are produced in two telescope array sites, one in Canarian island (Spain), the other one in Chile. The current planned network bandwidth will be limited to 10 Gbps for the two sites on shared links.

For the Bulk Data Management, data could be accessed by authorised users only at CTAO data centre sites.

---

<sup>19</sup> <https://www.cta-observatory.org/>



D2.2 Assessment and analysis of performance of the first pilot data lake

For the Science Archive, data are public except proprietary data that are readable only by authorized users (Principal Investigator and delegates only) and for a defined period of time (period must be flexible). Academic or non-academic users can perform analyses and further data combining workflows on data products and data collections. Users will generate data products of their own which we will want to include in the global CTA Science archive. Public data must be Virtual Observatory accessible. The data volume is more in the TB scale.

The two systems (Bulk data Management and Science Archive) could be based on the same technology.

As a first objective and current study, the WP2 Data Lake is foreseen as an implementation of the Bulk Data Management System of CTA. Therefore, the Data Lake must meet the requirements of the CTA Data workflow. As part of this workflow, the plan is to have one copy of raw data generated on a telescope array site (CTA-South or CTA-North) in a given day stored in a primary data centre. Therefore, *in order to balance the load across 3 data centres*, one copy of raw data generated the next day will be stored in another primary data centre (Figure 34), showing Day 1 replication left (first into site A, the copying to site B with site C not involved), and Day 2 (first into site B and copying to Site C with site A not involved).

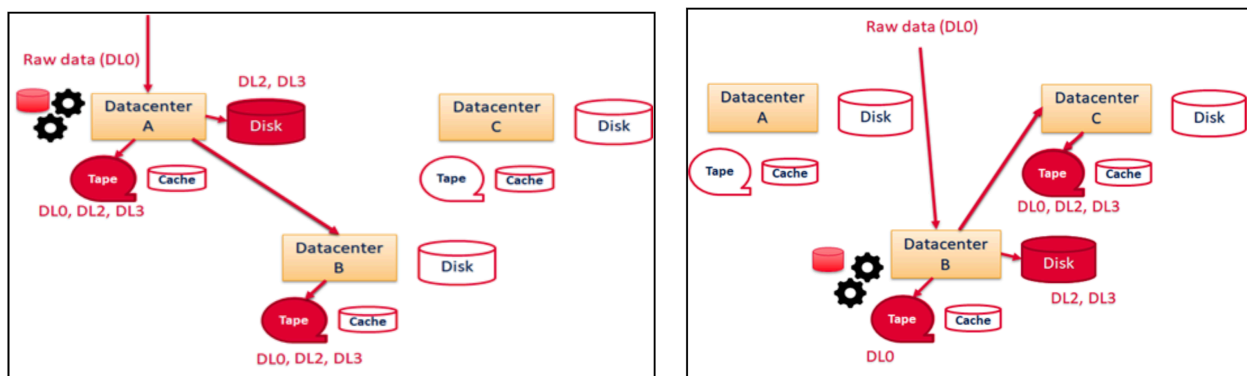


Figure 35: Image showing example of changing data storage workflows from one day to the next in order to balance the data storage load across the Bulk Data Management system of CTA.

In the basis phase more than 6 PB per year (for 30 years) that will be produced by the telescopes plus a fixed value of 20 PB for simulation. This is estimated to number tens of millions of files per year and millions of datasets. A maximum of 4 data centres and 2 telescope array sites is envisaged. For the Bulk Data Management System, the number of users are very limited (tens of users) when it could be very high (several 1000 users) for the Science Archive.

Therefore the features that are the most interesting for CTA are the following: Storage Quality Of Service Management; Lightweight version of a Rucio Storage Element to be deployed on Telescope array sites to help Bulk Data Management; OpenId management in order to be able to manage any academic or non-academic users; Experiment-wise metadata management with ability to have complex searches based on the metadata (Bulk Data Management and Science Archive); and Fine-grained proprietary data management (Science Archive) giving access to data only to authorised users.

## 10.2 EGO / VIRGO

VIRGO is a Michelson Interferometer with 2km-long arms, built to detect gravitational waves, located at the European Gravitational Observatory (EGO) in Italy. To date VIRGO has completed three observing campaigns alongside other gravitational wave detectors, over several years. The next observing run will not be until at least June 2022. VIRGO's data moving use case is interesting because it requires some very low latencies, in particular the need to perform online processing quickly and share any events with the LIGO detector in the USA - if both detectors see evidence for a gravitational wave at the same time (appearing in the Event Aggregation Database), this will lead to an astronomical alert being triggered and subscribing (electromagnetic) telescopes can move to observe the region of interest.

This need for low latency data movement has led to a complex data distribution diagram (Figure 36), which has some inefficiencies in data transfer since scientists at some European sites would still access EGO/VIRGO data through Caltech.

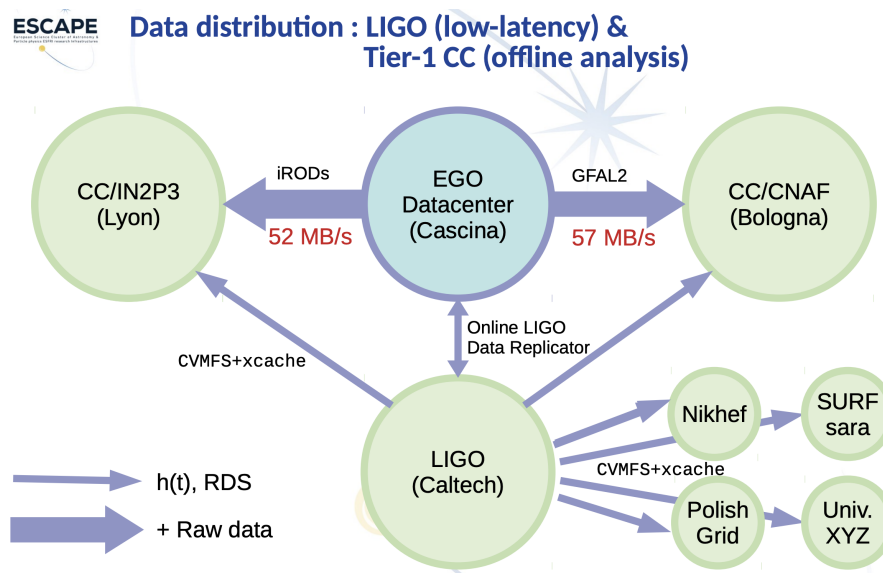


Figure 36: The current EGO/LIGO data movement flow. Image credit: Pierre Chanial

An alternative option could be as shown in Figure 37, where dedicated transfers still manage low-latency data movement if required, but a Data Lake concept makes data products available for offline computing at the compute centres at CNAF and IN2P3 and also to users across Europe and in the LIGO consortium.



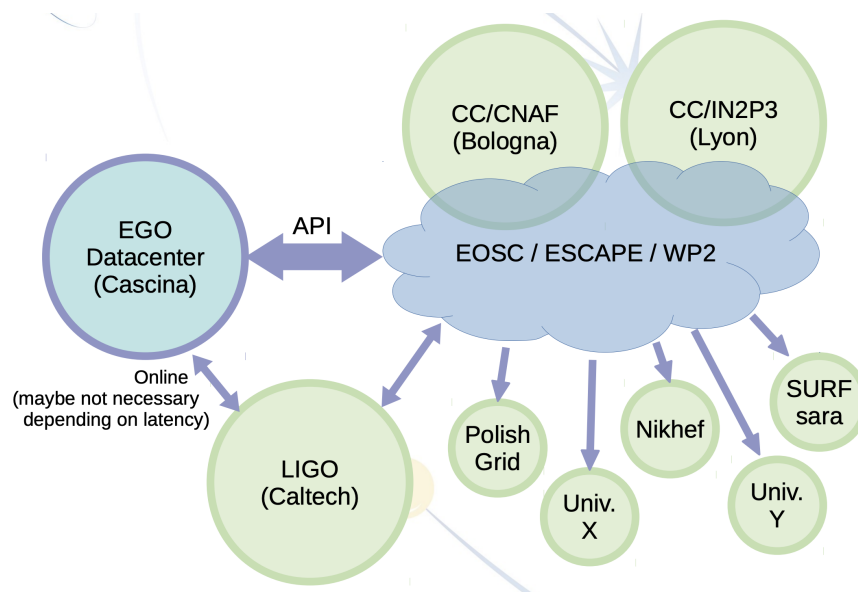


Figure 37: A possible alternative EGO/LIGO data movement option, making use of datalake technologies. Raw data flows into the datalake for offline processing. Image credit Pierre Chaniai.

EGO Raw data has a lifecycle - initially "Hot" and needed for low-latency processing and online analysis (at EGO), it then would move to "Cold" when it is used about once per month for detector characterisation. This would be data younger than  $\sim 2$  years. After that, data would be moved into the archive, from where access might be needed only 2-3 times per year to help calibrate the signals. Thus, data placement on different QoS classes (e.g. onto slow Tape or fast disk) could be determined by rules depending on the access frequency for each dataset - this is handled by Rucio using the rucio-c3po daemon.

### 10.3 FAIR

The mission of the international FAIR particle accelerator facility in Darmstadt is to study the structure of matter and the evolution of the Universe. FAIR consists of four scientific pillars: (1) NUSTAR is investigating nuclear reactions under stellar conditions; (2) APPA is investigating atomic interactions and plasma physics as well as corresponding applications; (3) PANDA intends to do research on topics around the strong force, exotic states of matter and the structure of hadrons; The goal of the CBM (4) research program is to explore the QCD phase diagram in the region of high baryon densities using high-energy nucleus-nucleus collisions.

In total the scientific FAIR collaborations consist of about 900 scientists.

The FAIR accelerator facility will have the unique ability to provide particle beams of all the chemical elements, as well as antiprotons at a previously unparalleled intensity and quality. At the heart of the facility is the underground ring accelerator SIS100 with a circumference of 1,100 meters. Connected to the SIS100 ring accelerator is a complex system of storage rings and experimental stations with 3.2 kilometers of beam lines in total. The existing accelerator facility of the GSI Helmholtzzentrum für Schwerionenforschung will serve as the injector.

The PANDA experiment at FAIR is using antiproton beams of up to 15 GeV/c leading to antiproton-proton collisions of maximal 5.5 GeV in the centre-of-mass-system and a data stream of 200 GB/s into the Green IT Cube. The average interaction rate of 800 kHz is reduced by an online filter by a factor 100 to an 8 kHz event rate which is stored for offline reconstruction.

PANDA will generate 1 PB/year of raw data which has to be securely stored in a cold storage, mirrored on different computing centres, and needs to be accessible with a low latency for reprocessing in a hot storage. In addition, up to 2 PB/year of high-level data is generated which is used by the users to perform physics analysis. This data is complemented by 3 PB/year of simulated data which is needed for the analysis. Both the high-level data as well as the simulated data has to be available on a time scale of 10 years.

CBM will run at an interaction rate of up to 10 MHz with typically Au+Au collisions of up to 11 AGeV leading to a peak raw data flow of 1 TB/sec into GSI's Green IT Cube. The First-level Event Selector (FLES) processor farm corresponding to about 60,000 cores will do full event reconstruction and data reduction by event selection at a rate of  $10^7$  events per seconds.

CBM will generate 18 PB/year of raw data. The current computing model foresees that RAW data are delivered to on-site storage media (RAW\_HOT) and copied to long-term archive (RAW\_COLD). Two copies of RAW data have to be archived in two distinct data centres, one being FAIR/GSI. From the RAW data about 4 PB of AOD data are created annually which are transferred to the participating data centres, serving a regional CBM user community. Additionally, 9 PB of SIM data are generated in correspondence with the experiment settings, conditions and statistics. The AOD data will be kept for 5 years while the SIM data will be available up to 3 years after production.

## 10.4 LOFAR

LOFAR has ~50PB data in the LTA (Long Term Archive), a major part of which needs to be processed. The LOFAR Data Valorisation (LDV) is an action which aims to compress and process this previously archived data. We expect this to be done in phases and completed by 2023.

The science-ready products (e.g., images, spectra, source lists etc) will be made available to the project teams and to the general astronomer after any embargo period if applicable.

Prior to 2020, we used to have ~7 PB of visibility data archived on the LTA per year. Recent implementation of data compression combined with a correlator upgrade, give a LOFAR visibility data to be archived at LTA to be ~4 PB per year, contained in about a million datasets per year in the coming years.

Currently most LOFAR data is processed by a few "power users" in Key Science Projects. In the near future (2021/2022), the visibility data will be in general processed using various pipelines by the LOFAR observatory. This means that both the processed data and products (calibration solutions, images, image cubes, spectra, source lists etc) will be more easily available to the general astronomer. The aim is to move towards public data, as well as a more prominent role for a general astronomer/user - the need to support up to 400 users is foreseen. At present users can access (and

download) their own (and publicly available) data from the LTA. We aim to improve this access towards a more efficient and user-friendly approach.

The LOFAR LTA is currently distributed over three locations over Europe (Poznan, Amsterdam, Jülich). The high-speed network bandwidth as well its reliability does vary from one location to another. This is even more relevant when the LTA sites are accessed by a general astronomer (via public internet). This aspect has to be taken into account when new storage locations are added to the LTA.

One of the main data management challenges lies with ingesting data from our Central Processing (Offline) cluster, just after the correlator, into the LTA. This is based on several parameters (e.g. project requirements to have all their data at the same or preferred site, storage claims, telescope duty cycle) and available resources. Data transfers are currently managed using self-built tooling.

We aim to have a Science Data Centre which will cater the needs for LOFAR but importantly act as an SKA Regional Centre once SKA data starts flowing in. Thus, some of our interests mentioned below also arise due to our long-term ambition, and Rucio is a promising option for the following aspects:

- Geographically distributed storage
- Centrally managed and orchestrated
- Single name space for easy access irrespective of data locality
- Independent of backend storage technology
- Independent of specific transfer protocols
- Scalability of the Data Lake to several hundreds of PBs and more.
- Ability to add the existing LTA locations to the data lake
- Separation of concerns between the storage provider and the experimental operation
- Rule-based ingest from a location outside the Data Lake (e.g. Central Processing of LOFAR)
- Possibility that small data centres (storage facilities) can be included.

Central data management and automating the transfers without (significant) manual intervention, is another interesting feature of Rucio. The possibility of decision making on locality and duplication based on defined rules coupled to metadata is also of interest. Being able to link together observations that have some relation (e.g. observations for the target source and its calibrator source) would be helpful.

The amount of storage space at the central processing is limited so we need to know at the ingest time what remote location is available, and how much space is available. Rule based transfer do have a scope here as the decision making of the site where the data is ingested can be offloaded to such a service and transfers can be scheduled.

## 10.5 MAGIC

The WP2 Data Lake is envisaged as an implementation of the automatic management of the MAGIC experiments. The Data Lake must therefore meet the requirements of the MAGIC Data Workflow. As part of this workflow, we plan to have one copy of the data generated at the Observatorio del Roque

D2.2 Assessment and analysis of performance of the first pilot data lake

de los Muchachos (ORM), deleting the original files once they have been replicated at the PIC (Figure 38). The main focus is on these data creation, replication and (original file) deletion, but there may also be additional sites in future requiring a more distributed replication scenario.

The amount of data anticipated is about 130 TB per year, in approximately a quarter of a million files - this is not huge by Data Lake standards, but the automation of the replication and deletion would enable a smaller operations team to manage the data.

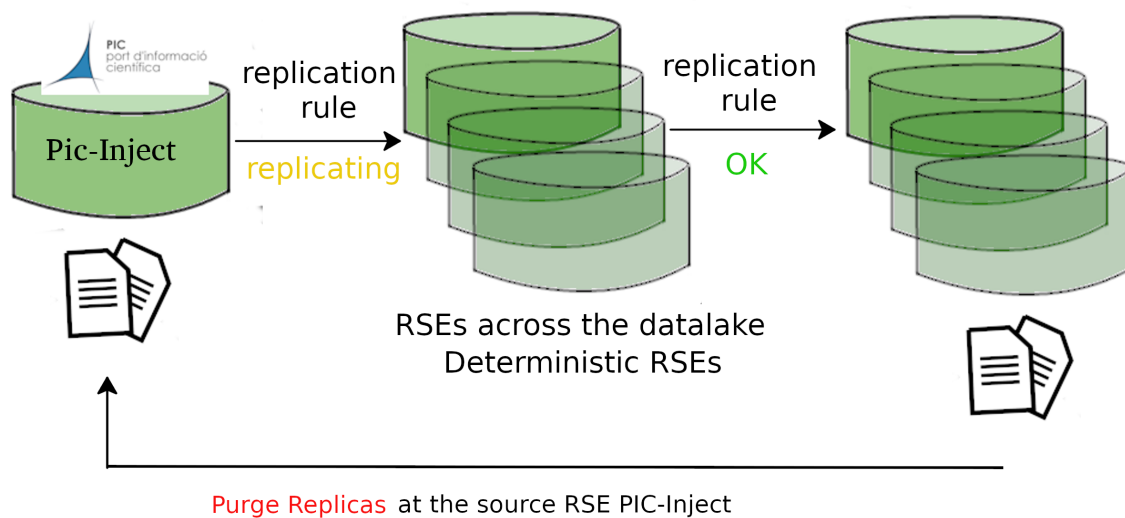


Figure 38: The MAGIC telescope data movement, replication and deletion flow, showing deletion of data at source following successful replication.

## 10.6 SKA Observatory

SKA will generate 600 PB of Observatory-level data products per annum, from 2028 onwards. These need to be accessed by authorised users (taking into account proprietary periods) at SKA Regional Centre (SRC) sites around the world so that astronomer users can perform analyses and further data combining workflows on data products and data collections. Users will generate data products of their own which we will want to include in the global SKA Science archive.

SKA will have a collaborative arrangement with the SKA Regional Centres and can see that if a global data management service can be used, large cost savings are possible as data will not need to be replicated so much. However, success here will require SRCs to interoperate - meaning that access to compute will need to be provided to connect users to sites holding copies of datasets unless transfers can be easily achieved - so we need a system that can intelligently balance moving files vs. "taking compute to the data".

SKA's data products will have different usage patterns, some of which will be predictable a priori - following a projected data lifecycle, but others will depend on the quality of data and the level of scientific interest - some data products will wax and wane in popularity and the data placement should be able to respond to this. Rucio is able to provide features to help with this - through the definition of QoS classes and policies (under development) and through the rucio-c3po daemon that can apply rules based on popularity (we have not yet tested this).

SKA's large data rate into the SKA Regional Centre network may require specialist transfers to manage the link, but it would be good to understand the degree to which data transfers into the SRC Data Lake could be handled asynchronously and delegated to a service such as Rucio - it would be best if this can be done whilst retaining some control over the network link (to enable urgent transfers to override Rucio ones).

SKA's data products will themselves be large - of order 1 PB for an image cube. This is far too large to be handled well in any data transfer activities, and instead SKA intends to break these cubes into individual files of between about 500 MB and 10 GB in size. This means that data collections with around 100k to 1 Million files will be common, and whatever data management system is ultimately used will need to handle data collections of this size readily.

## 10.7 Rubin Observatory / LSST

The Rubin Observatory<sup>20</sup> is being built on a remote mountain site in Chile. Data are collected at the Summit Site. A first copy is stored at the data access centre at the Base Site. A second replica is immediately sent to the LSST Data Facility (recently decided to be at SLAC National Accelerator Laboratory). At the data facility, the data are processed rapidly to generate alerts, and on multiple slower cadences (combining and comparing previous images) to produce a suite of calibrated data products for use by collaborators. Without significant delay (as quickly as practicable) a third copy of the data is sent from the LSST Data Facility to the Data Access Centre (DAC) in Lyon, France (Figure 39). Before the telescope enters full operations, commercial cloud will serve as the data facility.

---

<sup>20</sup> The Rubin Observatory was previously called the "Large Synoptic Survey Telescope", LSST.

D2.2 Assessment and analysis of performance of the first pilot data lake



Figure 39: Rubin Observatory / LSST Data movement and processing schematic.

There will be further DACs in other participating countries (including the UK), but the details of which data products will go to which DACs may change, some sites may only hold copies of data products temporarily to enable processing. It is anticipated that users will interact with the data products of their collaboration within the DAC, using the JupyterLab-based LSST Science Platform (see <https://ldm-542.lsst.io/>).

In terms of the assessment of Rucio in the duration of the ESCAPE project, our Rubin Observatory colleagues are interested in the ability of Rucio to manage data transfers and replication between a relatively small number of large sites (the DACs), and whether and how the scientists using the LSST Science Platform can call data using the Rucio API within their Jupyter notebooks, and how a data lifecycle with periodic re-analysis of files at compute-only sites can be implemented.

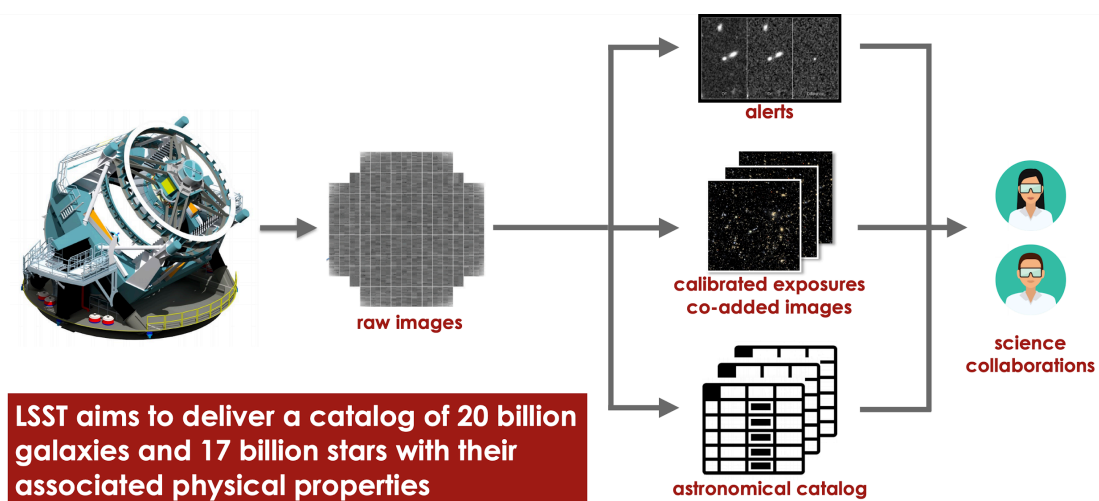


Figure 40: Rubin Observatory (LSST) data product flow.



## D2.2 Assessment and analysis of performance of the first pilot data lake

The expected annual raw data volume is ~7PB. Including data products increases this amount by a factor of 10. The number of files will be about 1 billion per year, so a total of 10 billion for the operational phase of the instrument. This figure of 1B per year also includes temporary files that need to be created for the duration of the data release processing time. Temporary files can be deleted for each new data release campaign (i.e. every year).

The number of users fully using the capability of the Data Lake would be small (likely a few per site): those are the users/teams in charge of producing and publishing the annual data releases. Once the releases are published, it is anticipated that hundreds of individual researchers would access (i.e. reading, not modifying) the released data at every site where a copy is located (Figure 40).

## 11 Appendix B: Additional diagrams

### 11.1 ESCAPE Data Lake Rucio Storage Element (RSE) details

This is the Data Lake RSE information that WP2 collaborators have made available on the WP2 wiki site - [https://wiki.escape2020.de/index.php/WP2\\_-\\_DIOS](https://wiki.escape2020.de/index.php/WP2_-_DIOS)

Institute	STORAGE Technology	Version	Quota (inode limit)	Min Free Space	XROOT	HTTPS	GSIFTP	ESCAPE VO (X509)	Token Based Auth/Z	perfSONAR Hosts <sup>Ⓔ</sup> (bandwidth & latency)
CERN	EOS	v4.8.20	300 TB	30 TB	YES	YES	YES	YES	NO	psb01-gva.cern.ch psl01-gva.cern.ch
LAPP	Federated DPM (ALPAMED)	1.14.2	100 TB	10 TB	YES	YES	YES	YES	YES	lapp-ps01.in2p3.fr lapp-ps02.in2p3.fr
LAPP	dCache	5.2	10 TB	1 TB	NO	YES	NO	YES	WIP	lapp-dcache01.in2p3.fr
LAPP	webdav		100 GB	90 GB	NO	YES	NO	YES	NO	lapp-esc02.in2p3.fr
SURFsara	dCache	6.0.9	98 TB	140 GB	YES	YES	YES	YES	YES	perfonar-bandwidth.grid.surfsara.nl perfonar-latency.grid.surfsara.nl
GSI	xRootD	4.12	20 TB (1M inodes)	1 TB	YES	YES	NO	YES	YES	dclxdlperfonar1.gsi.de dclxdlperfonar2.gsi.de
GSI	xRootD	4.11	1 TB	10 GB	YES	YES	NO	YES	YES	dclxdlperfonar1.gsi.de dclxdlperfonar2.gsi.de
INFN-CNAF	StoRM	1.11.19	10 TB (100k inodes)	1 TB	NO	YES	NO	YES	YES	perfonar-ps.cnaf.infn.it perfonar-ow.cnaf.infn.it
INFN-ROMA1	DPM	?	2 TB	200 GB	YES	YES	YES	NO	?	perfonar1.roma1.infn.it perfonar2.roma1.infn.it
INFN-Napoli	DPM local storage	1.14.2	68 TB	5 TB	YES	YES	YES	YES	?	perfonar.na.infn.it perfonar2.na.infn.it
INFN-Napoli	DPM Federated storage	1.14.2	46 TB	5 TB	YES	YES	YES	YES	?	perfonar.na.infn.it perfonar2.na.infn.it
DESY	dCache	6.2.2	40 TB	4 TB	YES	YES	YES	YES	YES	perfonar-ps-04.desy.de perfonar-ps-03.desy.de
CC-IN2P3	dCache	6.2.10	60 TB	1 TB	YES	YES	YES	YES	YES	ccperfonar1.in2p3.fr ccperfonar2.in2p3.fr
PIC	dCache	5.2.20	28 TB	27.99 TB	YES	YES	YES	YES	YES	psb01.pic.es psl01.pic.es

Table 3: The datalake RSE information that WP2 collaborators have made available on the WP@ wiki site - [https://wiki.escape2020.de/index.php/WP2\\_-\\_DIOS](https://wiki.escape2020.de/index.php/WP2_-_DIOS)



## 11.2 Daily Health Check flow diagram

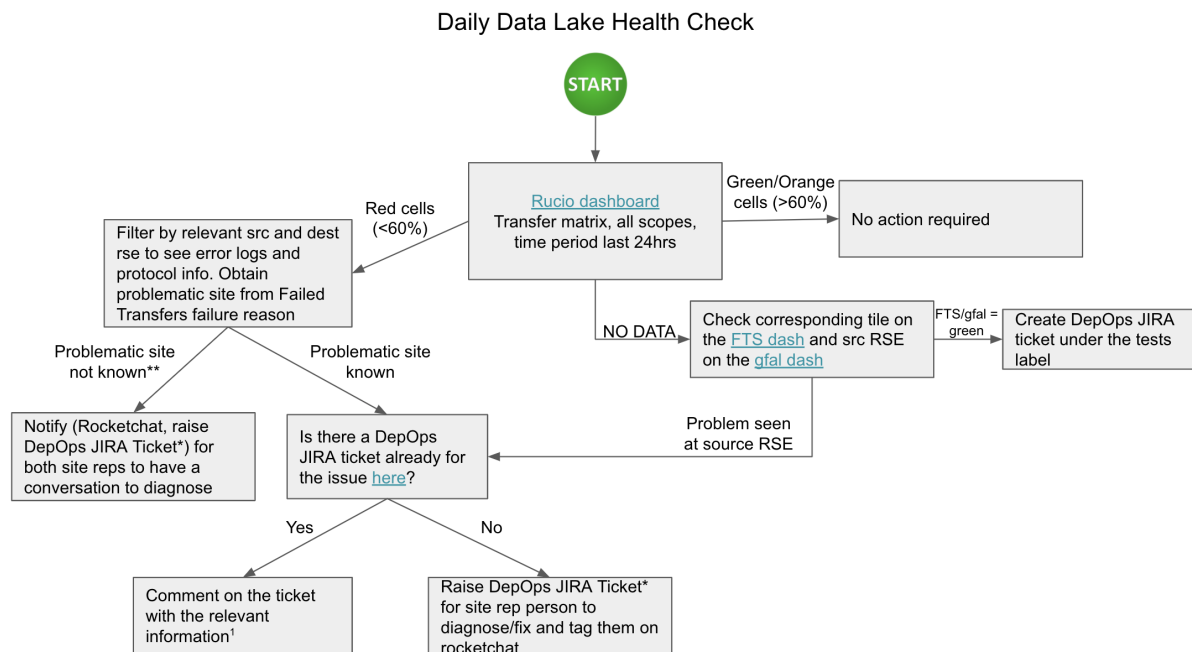


Figure 41: This diagram shows the current (Jan 2020) version of the action tree that we have developed collaboratively to guide newcomers to the monitoring team in their health checks. This means that people new to the datalake concept have been able to take part in the monitoring activities, greatly improving our ability to rapidly detect that components are under-performing.

## 12 Abbreviations

We include our most commonly used abbreviation here for reference.

CERN: Conseil Européen pour la Recherche Nucléaire", or European Council for Nuclear Research

CPU: Central Processing Unit

CRIC: Compute Resources Information Catalogue

CTA: Cherenkov Telescope Array

DID: Data Identifier

EGI: European Grid Infrastructure

EGO: European Gravitational Observatory

FAIR: Facility for Antiproton and Ion Research

FTS: File Transfer Service

GB: Giga Byte, Gb: Giga bit

GFAL: Grid File Access Library

GGUS: Grid Global User Support (helpdesk system)

GOCDB: Grid Configuration Database

GSI: Gesellschaft Für Schwerionenforschung (Society for Heavy Ion Research)

HPC: High Performance Computer

IAM: Identity and Access Management

INFN: Italian National Institute for Nuclear Physics

IT: Information Technology

JWT: JSON Web Token

LHC: The Large Hadron Collider

LOFAR: Low-Frequency Array (Radio Telescope in Netherlands)



## D2.2 Assessment and analysis of performance of the first pilot data lake

---

LTA: Long Term Archive (of LOFAR)

LSST: former name for Large Synoptic Survey Telescope (now called the Rubin Observatory).

MAGIC: Major Atmospheric Gamma Imaging Cherenkov – telescopes located at ORM

ORM: Roque de los Muchachos Observatory – host to MAGIC

PB: Peta Byte

PIC: Port d'Informació Científica, Computing Centre in Barcelona.

QoS: Quality of Service, here relating to storage ingress and egress performance and reliability

RSE: Rucio Storage Element (a storage entity within the Rucio system)

RO: Rubin Observatory (formerly known as LSST)

STFC: The UK Science and Technology Facilities Council, a national funding agency with a cloud compute facility which SKA makes use of for development purposes.

SKA: Square Kilometre Array

TB: Tera Byte, Tb: Tera bit

VM: Virtual Machine

VOMS: Virtual Organization Membership Service

WP2: ESCAPE Project Work Package 2 – Data Infrastructure for Open Science (DIOS)

WLCG: The Worldwide LHC Computing Grid