

3rd ESCAPE WP2/DIOS workshop

21-23 March 2022

[Indico Event](#)

Participants

Agustin Bruzzese	Ian Bird	Patrick Fuhrmann
Alba Vendrell Moya	James Collinson	Paul Millar
Alessandra Doria	João Lopes	Paul-Niklas Kramp
Alex Clarke	Julia Andreeva	Raymond Oonk
Alexander Verkooijen	Kilian Schwarz	Riccardo Di Maria
Antoon Prins	Lucia Morganti	Rizart Dona
Debra Turley	Maisam M. Dadkan	Rob Barnsley
Diego Ciangottini	Marcelo Soares	Rohini Joshi
Elena Gazzarrini	Mario Lassnig	Ron Trompert
Enrique Garcia	Martin Barisits	Rosie Bolton
Fabio Hernandez	Matthias Fuessling	Stefan Schlenstedt
Fabrizio Lucarelli	Markus Elsing	Stefano Gallozzi
Federica Agostini	Mieke Bouwhuis	Steven Murray
Francesco Giacomini	Mihai Patrascioiu	Tommaso Boccali
Gareth Hughes	Muhammad Aleem Sarwar	V.N. Pandey
Gino Marchetti	Nadine NEYROUD	Xavier Espinal
Gonzalo Merino	Panos Paparrigopoulos	Yan Grange

Monday 21 March 2022 (Day 1)

Welcome and Introduction

Speaker: Xavi Espinal (CERN)

[Slides link](#)

Experiment Vision: DAC21 take-aways and forward look

Convener: Fabio Hernandez (CC-IN2P3)

Fabio: Welcome! This session will focus on the scientific projects and on how and why experiments would use an infrastructure like ESCAPE.

SKA

Speaker: James Collinson (SKAO), Rob Barnsley, Rohini Joshi (SKA Organization), Rosie Bolton (Square Kilometre Array Organisation)

[Slides link](#)

Q&A:

Pandey: transfer rate 90%, what does it include? Did it include retries? What could be the cause of failure - overload or file size / type?, are the failures for specific category, size of files or linked to network, load etc.

Rosie: transfers done/transfers submitted, FTS retries not included

Fabio Hernandez: Are people more comfortable with using tokens?

Rosie: The community is still not clear, X509 is a bit harder to obtain. Tokens are also not as straightforward, depending on who owns them.

MAGIC

Speakers: Agustin Bruzzese (PIC), Gonzalo Merino (PIC)

[Slides link](#)

Q&A:

Fabio Fernandez: Great complexity involving deployment (slide 16). Ready-to-use infrastructure was already provided, which difficulty are you referring to?

Agustin: difficult to understand all pieces composing the infrastructure, not deployment. We tried to deploy rucio using docker in k8s, and struggled to understand connections between all components. It's more important to understand what you're doing and how things interact with each other, this is the complicated part.

Rohini: embargoed data models and view data permissions need to be granted and not assumed. This is something crucial also in SKA. Is it something to be concerned about?

Agustin: yes, it is also a concern for MAGIC and CTA. Authorisation tokens could be useful to achieve this goal - to protect data and not allow users to see everything.

Stefan Schlenstedt: can you explain what do you mean with scope elimination in slide 16?

Agustin: a functionality where you could delete a scope on rucio DB could be implemented. This would be useful to delete test scopes which are now useless.

Martin Baristis: a scope is to partition data in pieces in order to identify data and its origins. The scope elimination is not possible since we want to make sure that you do not rename things in the same way due to HEP necessities. [This](#) is blocked now but could be easily changed, but for now we decided to keep it this way.

Xavier Espinal: in slide 18, you say that you would like to have a 'real case transfer replica' to other physical sites. Could you expand on this please?

Agustin: right now transfers were only in a test environment, it was not real data. What I wanted to say is to apply the workflow to the actual transferring machine.

CTA

Speakers: Gareth Hughes (CTAO), Nadine NEYROUD (LAPP)

[Slides link](#)

Q&A:

Martin Baristis: you mentioned that some of the data transfers did not go through because of the size of data, or because of the amount of data?

Gareth: I think we ran out of time.

Martin Baristis: with a CTAO instance, do you run it on the Dirac branch or do you have your own branch?

Gareth+Agustin: We have a CTAO Dirac branch.

Fabio Hernandez: I think google documents collect a lot of knowledge but are not very efficient, we should pay more attention to how we build documentation. Many people today referred to complexity, which can be mitigated by attention being paid to better documentation.

Martin Baristis: We have expanded our documentation, but if somebody wants to add to it, please submit it in markdown format to our [rucio docs git repository](#), even if in disorganized format!

KM3NeT

Speaker: Mieke Bouwhuis

[Slides link](#)

Note: "Possibility to define applicability period in metadata" -> calibration/positioning of the PMT strings in the water, need to be combined with data.

Q&A:

Xavier: interesting approach, Mieke was using infrastructure as it is, and in a short time KM3Net managed to achieve a lot. ! Also, some labs might be interested in sharing the infrastructure at the site, as some sites act as "Tier-0" for one or several experiments or they need Data management for local facilities. Hence having a consolidated Rucio instance shared by the different communities will help in providing core service to the community, saving manpower. Rucio multi-VO is a key component of this approach towards a common Dm system at the sites serving different communities.

LSST

Speaker: Lionel Schwarz (CC-IN2P3)

[Google docs link](#)

Q&A:

Rosie: FTS parallelism - ideas for what was happening?

Paul M: yes, could be database query frequency, if links are fast and files are very small.

<https://github.com/paulmillar/http-tpc-utils/blob/master/bin/concurrency.py>

A good practice was to use a scheduler inside FTS to see what is going on 

Xavier: Similar model for MAGIC/ CTAO re. Use Rucio/FTS to manage data recording (and deletion) of raw data from Telescope's buffer to SLAC?

Fabio: Data will be transferred from telescope site to SLAC, team looking at Rucio but need to transfer some alerts at high priority - 60s

EGO/VIRGO - absent, rescheduling for tomorrow (?)

Speakers: Pierre Chaniel (European Gravitational Observatory), Rhys Poulton (European Gravitational Observatory)

ATLAS (*) (excused)

[Slides link](#)

Arturo Sanchez Pineda (LAPP)

Thoughts on the future of Storage Orchestration and Quality of Service

[Slides link](#)

Speaker: Paul Millar (DESY)

Q&A:

Raymond Oonk: Some of the major compute centers also run their own FTS services, which could perhaps be shared. I don't see the OpenStack/K8 level becoming interesting for most users, with the absence of a few capable ones.

Paul: agreed, in the HEP community developing the grid out of k8s technologies is something that hasn't taken off yet.

Raymond Oonk: There should be some potential options in the HE 2023-2024 work programme to continue ESCAPE

Paul: that sounds promising.

Ian Bird: clusters are encouraged to continue at some level, but we are trying to make ESCAPE continue as a collaboration outside Horizon. There will be opportunities, but it won't be on the scale of ESCAPE. It will be a subset of partners from ESCAPE collaborations.

Fabio Hernandez: FTS as continental service is still a realistic idea? If I want to have Rucio for my needs, I also need to operate FTS, so it would be good to have a better understanding of it.

Paul: I believe there is work by EGI to provide FTS as a service in terms of backend. There has been work to provide a web user interface as well. People would rather buy something.

Decoupling is possible.

Xavier Espinal: there are ongoing discussions. Gareth mentioned it would be good to have a DL infrastructure to spawn. Another way is facilitating data management services for other communities, for simple use cases with less data volumes.

Pandey: I wanted to know your thoughts about storage media. When talking to vendors, the costs of SSDs will be lower than magnetic disks by 2026 if we compare \$/TB of storage. If that happens, is it possible we will not have spinning disks anymore?

Paul: yes, I see it happening. It depends on your faith in the human mind to overcome problems. Magnetic spinning disks have barriers, and each time you come across a barrier (like too large heads), this provides ability to move beyond limitations. Magnetic disks have a long history, and I think in 5 years time we will still not be at the limit. If the price of SSDs is the same, then we will start switching to SSDs for sure.

Tuesday 22 March 2022 (Day 2)

Experiment Vision: DAC21 take-aways and forward look

Convener: Mieke Bouwuis

LOFAR

Speakers: Vishambhar Nath Pandey (ASTRON) Yan Garange (ASTRON, the Netherlands Institute for Radio Astronomy)

[Slides link](#)

Q&A:

Riccardo Di Maria: What do you mean with code expected for DLaaS? The server itself or the integration with ESAP?

Yan: I already do queries on ESAP, but it would be great to have a 'make available' command already there, so basically integrating the DLaaS with ESAP.

Paul Millar: When you talk about the 'extra DLaaS instance' that you would like to build, we in DESY have already started this work, but ran out of time. However, we have another pilot demonstrator in the context of our photon sciences. You could look into duplicating it.

Yan: Cool, yes I basically would like to have a platform where I have the LOFAR software installed and the DLaaS functionalities.

Paul Millar: When you talk about storage, do you expect the services to come from the storage system or from Rucio? Because in dCache, this exists as a possibility already, if you don't care about it being dCache specific.

Yan: I don't really care about the services behind, as long as I can download data and have a place to store it!

Martin Baristis: You can use Rucio events for this as well. Secondly, you can represent sites not connected to each other in a directed graph (using the concept of "distance"), and Rucio can do a multi-hop transfer between different sites (??)

Xavier Espinal: you mentioned synergies between DLaaS and ESAP; this is also important for future work in view of VREs (=EOSC-Future). Current collaboration between ESCAPE and CS3MESH4EOSC, picked up as an EOSC success case story. Given this momentum, do you think it could be good to focus part of the WP5 work in consolidating DLaaS/ESAP as a "product" involving also the interested RIs/sites? (to be further discussed on Wednesday session)

Yan: I would agree with this,

Xavier Espinal: It is very good to deploy rucio on your side. In the context of what was discussed yesterday: saving resources/manpower and given the fact that NIKHEF is a common point for LOFAR and KM3Net, perhaps it would be also good to explore Rucio multi-VO capabilities? Also involving UK sites/people exploring the same direction (SKA/DUNE)?

CMS

Speaker: Diego Ciangottini (INFN, Perugia)

[Slides link](#)

Q&A:

Paul Millar: You mentioned the Hashicorp Vault (slide 7) as a way to distribute secrets, there's already a way of doing it in fermilab if you are interested.

Xavier Espinal: About the interoperability of the DataLake, you mentioned you have your jupyterhub installation there, and this is relevant for future directions in other projects like EOSC. How do you see mounting jupyterhub on a common platform involving these kinds of findings, and building an analysis reproducibility environment?

Diego: I think the missing part in CMS (?) is having a common layer at the data management layer. From the compute part(HTCondor, Jupyterlab, containerisation), everything is already flexible. Any problem might arise from the data management part. This week there is also a HSF meeting related to these kinds of activities, I can share the link.

Mieke: You say that you encountered no problem, and where did you find the support?

Diego: Both from the Rucio support channels and OIDC integration from ESCAPE community, for example. Cross interaction is quite useful to avoid duplicated work.

FAIR: R3B, CBM, PANDA

Speakers: Maisam M. Dadkan (University of Groningen), Paul-Niklas Kramp (GSI Darmstadt)

[RUG R3B slides](#) [GSI CBM-PANDA](#)

Q&A:

Xavier: within your research community, do you think there is something to use for analysis debugging, aggregating the user interface and then synchronizing the analyses with batch systems, etc?

Paul Kramp: It is a difficult question; it could work, but we would need a standardized way of performing analyses, as right now there are very different approaches.. It would be hard to apply the same concept to all the various different smaller experiments. .

Xavier: What are the tools for data access and analysis that people in your community are using nowadays?

Maisam: In our community, the researchers mainly work with command line interface, but recently there are requirements from 'older' scientists (i.e.supervisors) to see the analysis in form of notebooks, and I think the best place to achieve this is something like the DLaaS. This is also relevant for some journals in order to reproduce the results before publishing them.

Xavier: Yes, the scientific paradigm is changing, one follow-up in this context is EOSC-Future, which aims at providing a portal for researchers to link data management and software plug-ins, and interact with different kinds of resources.

Mieke: You were talking about querying the metadata, what do you mean by 'the metadata concept needs to change'?

Maisam: The first time I used the Rucio extension on DLaaS, I was expecting to search the data based on metadata, since it is really hard to find the DIDs of data. However, if for example I upload data on the DataLake, it gets assigned a DID; when I upload the same data, it is given a different DID, so I think this is a bit confusing for data querying. Moreover, the DID is not very useful outside the ESCAPE community, usually you need something like a DOI.

Rob Barnsley: Related to metadata searches, the new release of the Rucio code is being developed to include more complex metadata searches, inclusive of AND/OR operators, inequalities, etc.

Discussion on ESCAPE technologies

Convener: Mieke Bouwhuis

Why use FTS?

Speakers: Mihai Patrascioiu, Steve Murray (CERN)

[Slides link](#)

For FTS support:

non-CERN: fts-support@cern.ch or fts-devel@cern.ch

CERN: <https://mattermost.web.cern.ch/it-dep/channels/it-fts>

Q&A:

Paul Millar: ESCAPE dashboard related. We have a Grafana dashboard with some events coming from Rucio, some from FTS. The granularity of Rucio events, while in the FTS dashboard it seems like there is a limitation where data is binned once per hour.

Steve: FTS could give more information, but we have agreed that the smallest unit is 1 hour, which I agree could be too much time when debugging. It is more of a monitoring team issue, we can leverage.

Paul Millar: one of the experiments was LSST in yesterday's talk. They have a large number of small files (different from HEP community needs). We hit some problems with those small files. How much do you see FTS supporting other communities with different needs? Seems like FTS provides a bottleneck for data transfers, i.e. the transfers are so fast that they happen before the scheduler runs.

Steve: We know that we are going to modify FTS for Run4 and scale it up to make it more efficient.

Mihai: there is a feature in FTS which we didn't expose in slides which is useful for small files. There is a session re-use to send a big block of files and send them all together to the scheduler. We can also do it in bulk.

Rosie Bolton: Does the FTS team have advice for best-way in?

Steve: In the documentation we have instructions on how to install as a developer, I would suggest doing that. The software itself is simple to install, getting authentication and all credentials right is the hardest bit.

A Global Information System: what we used CRIC for and what else it can provide

Speaker: Panos Paparrigopoulos (CERN)

[Slides link](#)

Q&A:

Rohini Joshi: in terms of the CRIC development team and release cycle, what is in place?

Panos: We are three developers, we try to have 3 releases per year.

Rohini Joshi: If there was a VO specific CRIC development, would it have to maintain compatibility with a CRIC vanilla configuration?

Panos: It is hard to maintain compatibility with plug-ins, but we are happy to help develop plug-ins, we already did it in the past.

Rohini Joshi: Are you in contact with Swiss Dirac?

Panos: Yes, but Dirac is not available through RestAPI, so we have already discussed how to pull data from there, and we are trying to pull data with .json requests.

Xavier: What are the main functionalities that CRIC provide?

Panos: CRIC's main functionality is the topology for files and services (e.g. configure protocols, downtimes - good for new communities): they come out of the box without extra configuration.

Lessons learnt on satellite Rucio developments and FTS

Speakers: Agustin Bruzzese (PIC), Rohini Joshi (SKA Organisation)

[SKAO-Slides link](#)

Q&A:

Martin: ATLAS or WGL timeline of OIDC. About Multi-Rucio vo deployments I don't see why this should not work. It's probably better to bring this question to the multi vo rucio developers.

Agustin: What we did was a proof of concept. 2 separate instances or we can use multi vo instance?

Rizart: Refresh tokens are not supposed to last forever. You can do it but it is not a good practice. I suggest doing in this case what is called "device flow" where you trust the device and then use tools such as oidc agent.

Wednesday 23 March 2022 (Day 3)

ESCAPE/DIOS in the near future: service consolidation, cooperations, strengthening collaborations

Convener: Nadine NEYROUD (LAPP)

Rucio: vision on the potential long term ESCAPE collaboration agreement

Speaker: Martin Barisits (CERN), Rob Barnsley

[Slides link](#)

Q&A:

Mieke: I am surprised they don't have embargoed data in ATLAS. Why is it so difficult? Can it not be implemented on account level?

Martin: There is a disconnect between storage permissions and Rucio catalog permissions. We cannot stop users from going to the storage directly since we do not set permissions on the storage. Having consistent permissions over all storages is hard. Tokens could solve the problem. Right now we can hide the data from the catalog, but users still have access to the storage.

Nadine: Rucio is a global product, are the new functionalities universal or provided as plug-ins?

Martin: It is a component based architecture. Some workflow community plug-ins are very useful. Rucio is a community development, so we have to identify common collaborations before acting. For example, the data embargo is of low interest for the HEP community right now.

Xavi: More and more communities want to use HEP resources, commercial clouds, etc.. could the QoS be merged with content delivery? As there is a lot of momentum of getting data close to many resources.

Martin: We have to be careful, there are a lot of caching technologies that are transparent to Rucio. We would have to age the interest in the rucio community; right now, there is enough to make our own interest group, but we would need to look into it a bit more.

Xavi: Some projects like the Fenix project aim at building a common endpoint where users can ask data from rucio, fts, etc.

Martin: We were always discussing these temporary/volatile RSEs (for example in google notebook), we should look into it.

Xavi: how open are people to actively use (/collaborate) Rucio? Comments from the experiments are welcome.

Rosie: could be very interesting, for rucio team scaling up would be a real challenge, a longer lived escape collaboration could act as an envelope to share requirements and collaborate.

Matthias Fuessling: from CTAO point of view having a collaboration ongoing would be extremely useful. Our dev team would be interested. We would certainly engage in such a forum or even more if the technology is finally picked at CTAO.

Rucio extended services: DLaaS and further integration possibilities with Analysis Platforms/Facility

Speaker: Riccardo Di Maria (CERN)

[Slides link](#)

Q&A:

Nadine: DLaaS outside CERN? What have you planned for that, do you have some ideas?

Riccardo: Deployment and distribution model is all public and accessible for others. Storage backend and distribution in the current service is EOS - and it is really storage dependent, if you run other backends some more development needs to be taken care of.

Rosie: It is very interesting, broadening it out would be brilliant, but it depends on other collaboration. SKA is interested.

Riccardo: Jupyterlab extension is available (Hosted in [Rucio github](#)) and can be used for any purposes. On the other side, the DLaaS is an always-available service and has the feature of having the data storage as a backend (which is beyond the extension). Nevertheless, the extension is a standalone.

Gareth Hughes: What about a MyBinder deployment of DLaaS?

Riccardo: Some integration work should be done.

Markus Elsing: Binder plugin for rucio already exists. Will follow this offline.

Xavier: the DLaaS should give a final backbone between the data, local facilities and users. Merging analysis platforms with DLaaS could be pursued in order to have a good entry point for fostering sciences. Interest group?

Riccardo: I have been contacted by many people (SKA, MAGIC, CTA, ATLAS, KM3NET, LOFAR, FAIR), so it would be good to open the floor to discussion. Everybody who has already had a look at this is a 'useful' person who could keep the momentum going. I would keep in touch once in a while to share the knowledge with the newcomers. Let's keep it alive!

Xavier: It depends on local resources, and we should keep the flexibility.

Rosie: Raymond is making a good comment, 'It would probably also be useful for the storage providers to be part of the SIG (to ensure that the connection between users and providers is kept alive)'

Future perspectives for token-based authentication

Speaker: Rizart Dona (CERN)

[Slides link](#)

Q&A:

Raymond Oonk: I see two potential issues: (i) it is still unclear to me how (IAM) tokens will integrate with compute (local access and also global access e.g., via DIRAC or ARC/HTCONDOR-CE) and (ii) interoperability of IAM with other Token initiatives (EduTeams, Check-in, community AAI's etc).

Rizart: (i) not sure about DIRAC and HTCondor specifics. The idea here is more general: the services will implement openID specification. If the question is about what the tokens will contain, this is already my question. This is a discussion that needs to happen in the community based on their needs. On the technical side integration will happen with any service who supports openID connect.

(ii) interoperability: if they don't speak the same language, a translation layer is needed. Maybe people from CNAF can comment on this.

Francesco Giacomini: Federation in EOSC is in my presentation.

Mieke: Are X509 going to disappear?

Rizart: No, they are still gonna be around. Indigo IAM can support both.

Gareth: Comment on DIRAC: almost half of things to do are integration with tokens. It will be the focus of the next DIRAC workshop.

Xavier: Do you think it would be useful to send a survey to the ESCAPE community (Collecting the needs from our communities)?

Rizart: I think it is a good idea. A problem might be that the topic is not very mature yet, people would need technical expertise on this.

Fabio: Since the beginning of ESCAPE, I wanted to try tokens and access all storages, and that proved to be difficult. Getting a token was not that easy. For this you have to install some tools, and to install these tools you need root privileges (Can't be easily installed in the equivalent of lxplus machine). There are many barriers for end users to even understand the questions you posed in the presentation. We should improve ways to get people familiar with the token. We would need the server to recognise you as an authenticated user.

Rizart: I agree on the maturity, but I think it is not very hard to enter in that domain. There is a lack of documentation, that's for sure. Hackathons from WLCG were very useful. I think they can find even more problems on requesting a x509. Kinds of engagement are needed (Like hackathons to educate).

Xavier: We should make the infrastructure work with tokens, also from chats with Reana. Andrea Ceccanti ran a Hackaton 1'5y ago, it would be useful to do something similar soon.

ESCAPE/IAM: an IAM service for your ESFRI/RI and towards a common AAI layer for EOSC ?

Speaker: Francesco Giacomini

[Slides link](#)

Q&A:

Martin: We have to do a high amount of token exchange, did anyone ever benchmark this?

Francesco: I know there were some load tests, but do not have the number. Someone from ATLAS mentioned 500 Hz frequency of uploads. No idea if it is a target already achieved by IAM, it needs to be at least the target. If you have numbers in the sense of targets you expect to reach that would be useful for us.

Martin: Not clearly yet but maybe soon. For downloads, not sure if we need a token per download.

Francesco: need to consider that you have multiple exchanges with IAM when downloading.

Rohini: ESCAPE IAM instance specifically, are there any forums/channels of communication? Proposes of specific ESCAPE IAM mailing list / common chat depending on how much flux it will be.

Francesco: ESCAPE IAM will not disappear. We have our support email list. There is also a Slack channel. If there is any other channel of communication, such as Rocket Chat, I am happy to follow it.

Synergies, collaborations, common initiatives

Infrastructure/services "inventory", explore options for the future: synergies, collaborations, common initiatives, projects, setting up the right forums discuss/follow-up: EOSC-Future, CS3Mesh4EOSC/ScienceMesh, FENIX-HPC, interTwin, etc.

Speaker: Xavier Espinal (CERN)

[Slides link](#)

Q&A:

Enrique: great initiatives, LAPP side will support them and VRE team is heavily relying on DLaaS!

Mieke: We should keep the DLaaS going, isn't there European funding on this?

Xavier: yes, in the next calls we should prepare proposals.

Mieke: How do we keep in touch?

Rosie: We should find common names to refer to when writing proposals. Like some 'ESCAPE products'.

Martin: we should keep in touch as well.

Xavier: European Union agreed in taking ESCAPE as an example.

Pandey: ESCAPE has helped sciences a lot, so we should find a way to be concretely put it into action.

Fabio: we can see what member of community commit to do, and we should work towards funding opportunities that are coming.

Elena: for whoever is interested in EOSC-Future meetings, please fill in the bottom of this document with your names and emails!

Preparation of the final deliverable

Speaker: Xavier Espinal (CERN)

[Slides link](#)

Q&A:

Yan Grange: I like the idea, happy doing that.

Matthias: Datalake for dummies document, I am keen to contribute.

Xavi: We have a draft but basically done. Will share it as soon as possible.

Rosie: What do you propose that people do? What incentives do you have available?

Xavier: editorial and writing team will be very useful. All names will be there. Final presidential event, will buy you a coca cola (or at least a beer? eheh). Don't be shy.

Fabio: What is the title?

Xavier: ESCAPE project with an outlook on the future.

~~~~~  
*Common themes noted: (list here suggestions which are recurring) [keep at end of document]*

1. Documentation/ knowledge - lots but scattered
2. FTS...
3. DLaaS deployments, global interest from many community.
4. RUCIO-DIRAC integration work mentioned by CTAO and SKAO
5. (let's take an action to make a channel for FTS onboarding then, as a long-lived ESCAPE collaboration)

~~~~~  
EOSC-Future common meetings interest: names, emails and institutes: we will contact you!

V.N. Pandey (pandey@astron.nl)

Raymond Oonk (raymond.oonk@surf.nl , SURF)

Yan Grange ?