

ESCAPE DIOS Final Meeting Meeting Notes Amsterdam 7-8th June 2022

Participants: Federica Agostini, Martin Barisits, Mieke Bouwhuis, Simone Campana, Diego Ciangottini, Riccardo Di Maria, Rizart Dona, Xavier Espinal, Enrique Garcia, Elena Gazzarrini, Frederic Gillardo, Yan Grange, Fabio Hernandez, Haili Hu, Gareth Hughes, Rohini Joshi, Maisam M. Dadkan, Gonzalo Merino, Roberta Miccoli, Dave Morris, Raymond Onk, Ghita Rahal, Muhammad Aleem Sarwar, Ron Tromper, Alba Vendrell, David Groep

7th June afternoon session (14:00) introduction and thematic discussion session

The ESCAPE DIOS experience: Introduction talk, ESCAPE goals, achievements and future prospects.

Discussion Forum #1: ESCAPE impact on Scientific Computing in the Research Community

Pre-discussion organised discussion in 4 working groups, provide answers to the following questions, and then share the outcome.

- *Did ESCAPE achieve the expected impact? Elaborate answers to the following questions:*
 1. *What is the impact for the ESCAPE sciences?*
 2. *What is the impact for other HEP/Astro communities?*
 3. *What is the impact for other sciences?*
 4. *What is the impact for the facilities providing the services?*
- For all the above, discuss what we achieved and what we should focus on in the future*

===== Report from the 4 Working groups: Discussion =====

Group1:

1. From the CTAO point of view: overall lack of expertise and computing, ESCAPE gave access to the community and shared development and some resources too (hw)
2. Helped in taking away complexity, worked on a common platform and subset of tools
3. Fostered potential inclusion of small experiments and potential for a multi-messenger approach. Still Rucio is perceived as difficult as an entry point to the model.
4. Data lake as a service, good as a facilitator, good entry gate to ESCAPE model hiding some of the complexity behind.
5. From DESY Pow, pushed the development of dcache (exposed to different use cases)
Storage and ESFRIs talking to one another.

Group 2:

1. ESCAPE fostered a natural alignment across RIs on the common infrastructure and the Data Lake model
2. DIOS offered the possibility to evaluate performance of the Data Lake model under real conditions; experts could provide advice to the relevant parts of the RIs with first hand experience on data management solutions.

3. Good inspiration and demonstration is that CMS and ATLAS using the same technologies for their Data Lakes
4. (DOMA experience and link to ESCAPE) ESCAPE facilitated the implementation and testing of the DL model (tools and infrastructure).
5. Knowledge transfer possibilities, e.g. PIC and PaN sciences, propose to look at the ESCAPE DIOS model, bottom up approach, from use cases to demonstrators.
6. ESCAPE is split in different layers, modular tools (RUCIO, FTS, IAM) that allow the integration of several underlying technologies (ie. storage: dpm, dcache, eos....)

Group 3:

1. ESCAPE provided a platform to share knowledge, tools and expertise and provide direct contact with developer in all the stack of services
2. Manifested evidence that it is important to have the right involvement, experts that understand the needs of the RIs and understand the services, perhaps we were lacking a bit the user facing infrastructures (besides APs/DLaaS)
3. SP? Which are the synergies? Help development(coding), reporting issues/bugs
4. Adding more levels of abstraction is OK but end-users should be shielded from this. Need to hide complexity to the final users and to address simple data management needs.
5. Metadata and IVOA needs. Should we put it into Rucio? Perhaps these services should be science agnostic and built as plugins?
6. Tokens in WLCG and HEP the activities in DIOS during the project makes the infrastructure more accessible to other communities
7. WLCG sites are not surprised by technologies (they know) but have a lot of value to see and get exposed to the workflows by other scientific communities, understand their needs, etc.
8. Foster stackoverflow-like forums, we have them separate them, perhaps missing one acting as a community gathering forum (What Rucio does, what is FTS...)

Group 4:

1. ESCAPE helped consolidate some of the tools in scientific communities. This knowledge was shared among them.
2. Astro communities main impact was the exposure to grid-like technologies, and a possible long term solution for DM, ie. CTAO, SKAO, KM3NeT
3. DIOS provided the ground to advance in the development and set the scene for further deployment on some of the hot topics in Scientific Computing: Tokens, Data Lake model, etc. With some synergies with WLCG/DOMA sometimes
4. Generalised some functionalities in a transversal way, eg. metadata support contribution by SKAO
5. Considerable dev effort, facilities receiving feedback about usage, expose to new use cases, workflows, helped the facilities to evolve.

Open floor discussion:

- How is the contact with other fields (did it happen, will it happen)?
 - There is discussion with the other domain projects. With ESCAPE we have matured to a level that we could extend to other fields. PIC: there are

- discussions starting (ie. Bioluminescence). Fields with a need and knowledge but without legacy tooling, so that makes adoption easier.
- Of course the most relevant question is whether a community needs distributed computing or not. Collaborating on the supporting services (e.g. AAI) puts them on equal footing on which we could build with the data management technology when that will be needed. That is where the dialogue could be started in the future.
 - Photon-neutron sciences as an example, we had a workshop ESCAPE and PANOSC/EXPANDS their actual needs can be served by EUDAT-like tools but if they significantly grow they may actually need a more complex system, or the ability to do more and hide complexity to the end-user.
 - Some contact exists with earth observations in the intertwin project. There is a work package on storage federation and it has a good representation from the ESCAPE project. Maybe also engage with other relevant projects.
 - Distributed computing and open cloud are opposites. Should we do something about that perception?
 - Caching and CDN was not explored thoroughly, astro/radio perhaps not ready and also no HEP difficult to justify the overhead, but have a role in the future.
 - It is important that looking towards the future the caching needs to be more investigated. Also to find a way to define when a cache adds and when not.
 - One value of the instrumental values that DIOS project brought in is to have a testbed. It is important for people to test the tools, because that makes it much easier to gain experience and adopt them. However, setting up your own testbed is cumbersome. In the ideal world one should be able to set up the whole system on one's laptop. In future projects we could try to address being able to deploy your own Rucio, dCache, FTS, etc.
 - There is already a minikube instance that deploys a large fraction of that.
 - dCache and IAM missing
 - The main use case here is the data management expert who wants to know how this helps them with their problem.
 - This may be even too advanced for the "John Doe" user who just wants an installed demonstrator to play with (preferably even through a browser).

===== Day 2, Session 1 =====

8th June morning session: three thematic discussion sessions

(1/3) Rucio and the ESCAPE community, activities, collaboration and synergies

- Token: Moving to token based auth scenario, bringing in the option to enable embargo da. Strong need from some communities, and potentially community specific implementation. **Input from community is key.**
- Metadata: metadata functionality was historically limited. New communities bring more requirements on metadata. Metadata plugins were introduced. Some active support from the community (ie. SKAO). **Need to understand what metadata capabilities are missing, but also where we conceptually draw the line of supported metadata functionality.**
- Deployment and documentation. Need a dedicated effort to make substantial improvement on the documentation. This would be an effort from both newcomers

and the experts. Deployment has quite some complexity. Documentation is key, but also good demonstrators could be useful. **What can we do to make the deployment easier?**

- Collaboration. Rucio needs engagement of the community for its sustainability and evolution. Several ways to contribute: test/debug, active sw dev, help in documentation. Envisaging creation of a steering board to gather community requirements/needs. **How can Rucio better engage with the community?**
- **Metadata discussion**
 - Dave Morris (IVOA): Got standard metadata and APIs, sitting next to the storage system (with a RDDDB). How Rucio can adapt/plug into IVOA structure.
 - Users scan metadata via web services and return tables of metadata about the query, potentially data can be hosted in different Rucio instances. String returned with the SFN needs also the knowledge of the Rucio instance.
 - Plan with Yan Grange to explore and test a potential approach to this.
 - The recently developed filter engine is a step forward but still lacking functionalities (i.e. extended queries, etc.)
 - Rucio ability to plug-in external metadata sources: possible different ways, need consistency on PIDs. **Possible good PoC with IVOA**, using the current Postgres DB talking to Rucio in a single data lake scope.
 - Sync between metadata catalogue and actual Rucio catalogue can be engineered async or synchronous. Is a strategic structural decision. Synch delayed too much can make things complex when disagreement (high loads, long times with inconsistencies, etc.)
- **Deployment and documentation discussion:**
 - Starting with Rucio can be complicated also because needed adjacent services (ie. storage endpoint) perhaps need work to have an easier hello world start. Perception is that it takes time: go to forest step, succeed, then stuck, find out a solution and repeat.
 - Potential to consolidate documentation spread around
 - Infrastructure perceived as one of the main barriers, make use of K8s to ease deployment, etc.
 - Possibility to set up a **working group** to address this. Few people gather together and contact the expert to cross check.
 - Possibility to organise a **Hackathon** on Rucio documentation (3 days). Contribution from the community, especially the ones got Rucio working in their RIs.
 - Consultancy is key for large scale deployments after testing phase, data schemas, organisation, etc.
- **Collaboration discussion:**
 - Developers and community are the two main references to keep the community. Advisory board would leverage the two of them and could oversee the project's guidelines, identify common interests, etc.
 - Forking Rucio developments not very sustainable, get the forks aligned after some time is hard (examples exist)
 - Steering committee creation.

- Outreach and publicising of Rucio tools in other communities (astronomy) could be enhanced.
- Potential issue with a single person's knowledge/rights/authority to merge requests (make technical decisions). Rucio is community based and the general direction is to have a pool of experts.
- **Outcome/Summary:**
 - **Continue metadata collaborations, explore plugging-in external data sources: possible IVOA PoC** (Yan Grange and Dave Morris)
 - **Setup a Working Group or organise a hackathon (users+experts) to consolidate documentation, enhance consultancy for large deployments after testing phase.**
 - **Open Steering Board creation, identify synergies among communities, setup collaborations, and funding options. Martin will keep ESCAPE posted about its composition and mission.**

(2/3) AAI/IAM and token-based auth, future perspectives, Collaborations and fora

- ESCAPE IAM being integrated with other AAI/IdP providers: integration with EOSC AAI federation and FENIX-HPC project.
 - Need to make use of a proxy to translate OIDC to SAML (this proxy -Satosai is provided at the moment by Geant within EOSC-Future)
- Allow users from ESCAPE who do not have their own community AAI to access EOSC resources with the ESCAPE AAI to be the entry point for this.
- **Proposal:** Ask each experiment/RI to appoint a **Virtual Organisation admin** to approve membership, validate accounts, etc. Requirement on users to use their institutional e-mail and VO admin can accept (or reject).
 - IVOA as a community can have literally hundreds of VO admins (small communities), this might be an issue.
- General question about which infrastructure scale is needed for AAI activities? a single IAM instance and multi-IAM instance vs. a test/playground instance a la ESCAPE. Example is the instance at CERN for LHC, which is integrated in the CERN DDBB for IdM (Identity Management)
- After the testing/PoC phase, the **RIs need to evaluate how they foresee the needs of AAI** to set the service at the right scale for their purposes.
- **Need to explore further and set strategy on how to disentangle (or connect) Identity Provider (experiment/local user databases) and Auth/Authz (IAM)** higher level capabilities.
- How Fine grained IAM can be? groups/subgroups... not really a limit and each subgroup can have an admin, up to RIs where to set up the limits.
- Token exchange might be huge from the Rucio side O(1kHz), how is the IAM infrastructure deployed? Is it scalable enough? Is it a sustainable model? Also other services in the stack (storage) will be also exchanching (renewal, etc.) Is there a real limitation?
 - **Requirements from the experiment need to be in line with what they are ready to invest in AAI**, high performance does not come for free and need to be provisioned according to the expectations (high performance => \$\$\$)

- Token lifetime in LHC is now extended (less exchange) to be more voms-proxy to gain experience and learn in the transition.
- Explore the possibility to organise a Hackathon before the end of the project (~autumn)?
- **RIs to reflect and understand how to structure IdM/AAI**, what are their needs for groups/roles, VO admins, etc. Good to track and share the plans within the community to increase synergies and knowledge transfer among us. Perhaps a good option to strengthen a new proposal for community gathering: working groups, virtual communication channels, etc...
- **Strengthening the IAM awareness** across projects and working groups:
 - Setting up a common chat for IAM can be interesting - Collaborate, check how communities are merging.
- **Super simple instructions are needed**, digested, filtered and with no acronyms, copy paste like instruction. I need to download a file: here is the recipe. Felt tokens are complex, difficult to start.
- Mixing Web workflow and CLI workflow to get tokens can be confusing, perhaps disentangle both clearly in the instructions/recipes?
- Would a community development model work for IAM? several people contributing (similar to the Rucio approach). Plans need to be understood for the mid and long term dev plans, but might be a good idea.

(3/3) Analysis Platforms, Analysis Facilities. Linking data access and analysis with computing resources.. Scoping future work and identifying interested communities and related fora.

Discussion topics

- New technologies (besides the ones in DIOS) to be onboarded? CD/CI, containers, REANA, repos?
 - Suggest to not get stuck in techs, system flexibility should be preserved. Specially adding an abstraction layer for the user, so the underlying infra is *hidden* for the user.
 - Integration of technologies being used by RIs
 - Felt experiment specific software integration/repos is still weak.
- Variability of the model, adapt to user needs: multi-threading, caching, etc.
 - CDN/latency hiding not completely exploited in ESCAPE besides HL-LHC context, need to leverage while exploring integration of heterogeneous computing facilities
 - Provide knobs/controls to the user to decide what are the needs, ie. test job vs large submission (=> enable caching for large datasets)
 - Leveraging test jobs vs large scale job submission
- How to scale multi-user/multi-tenant?
 - Shareability of notebooks?
 - Integration of heterogeneous resources
 - Support and maintenance of the platform and the underlying services, standard tools will be attracting more persons, ie. K8s, Git repos, etc.

Meeting finished: 13:30













