


Live Notes for the Open Technical Forum #1 on Wednesday, November 13, 2024

Indico: <https://indico.cern.ch/event/1471694/>

Welcome to the WLCG Open Technical Forum (OTF).



The WLCG Technical Coordination Board (TCB) is responsible for the technical evolution of WLCG services in line with the needs of the experiments and the capabilities of the infrastructure providers. The TCB defines a multi-year roadmap for such evolution and is responsible for its implementation. The TCB achieves these goals with a bottom-up approach through an Open Technical Forum (OTF) which welcomes the participation of all contributors to the technical evolution in the WLCG community.

Attending: 66 on Zoom, about 20 in the room.

[Introduction](#)

[Q&A](#)

[Tokens Session](#)

[Status Update from the Migration to Tokens Taskforce](#)

[Rucio Token Status Update: Dimitrios Christidis](#)

[CMS Token Status: Stephan Lammel](#)

[Tokens in LHCb/WLCG: Federico Stagni and Christophe Haen](#)

[Status and concerns from FTS: Mihai Patrascioiu](#)

[IAM: Hannah Short and Berk Balci](#)

[Summary](#)

[CernVM workshop summary and plans: Valentin Volk](#)

[DOMA General Meeting](#)

Introduction

Q&A

- Maarten would like us to utilize the meeting slots foreseen for the GDB in scheduling future OTF meetings. We can try to take advantage of those slots, if practical. However, scheduling is a challenge. Note that pre-GDB meetings took place on Tuesdays. Starting at 1pm is a challenge.

Tokens Session

Status Update from the Migration to Tokens Taskforce

Rucio Token Status Update: Dimitrios Christidis

Q&A:

- Example of a token that can do what it should not? Data deletion on some SE(s) without the storage.modify scope, for example. This has been fixed. Too much flexibility can lead to problems in the implementation. A validator tool might help. Access tokens have a fixed lifetime for a specific purpose.
- ALICE have been using tokens for 20 years. Some disagreement over the scope of token lifetimes, scope in Rucio. There is resistance to having long-duration access tokens.
- CMS uses tokens for bulk data transfers, so there is a desire to use them for datasets. There are pros and cons.

⇒ **Token lifetime and scope is a source of disagreement within the community.**

CMS Token Status: Stephan Lammel

Lots of progress. Expectation that IAM and vault will be 24x7 services.

Q&A:

- Support model for IAM and vault is best-effort. An analogue legacy service is myproxy. CMS recognizes this, however, imagine IAM going down over the Christmas period. There is no out-of-hours on call service. We need a robust service. Previous credentials lasted 8 days. This is a huge transition. The shorter credentials are the more we will be dependent on the robustness of the service. Note that one weekend of downtime is a 1% loss of Grid resources for the year.
- Tier-0 by MoU must have 24x7 service. This should be escalated to management.
- Will any experiment depend on IAM for online operations? CMS definitely. Then we need 24x7 support or long-lived tokens for online use cases. Why is there such a dependency for HL-LHC when there is no such dependency now? Data taking involves not just writing but access to services for the online systems.

- Support level might need to change as we move towards short-lived tokens. Not acceptable that the Grid goes down over a weekend.
- There are security issues as well. This topic needs to be discussed in the WLCG MB.

⇒ **Disagreement over service level needed for IAM and vault, impact on data taking during HL-LHC.**

Tokens in LHCb/WLCG: Federico Stagni and Christophe Haen

Little progress in workflow management during the past six months. Pilot submission with HTCondorCE works, situation less clear with ARC. For data management, Dirac can submit FTS transfers with tokens but not used in production yet. Concern about scalability of IAM, support level. Want long-lived tokens to mitigate issues seen in DC24, for example.

Q&A:

- What do you mean by “technical debt”? Concerning internal component modernization, what they would call maintenance.

Status and concerns from FTS: Mihai Patrascioiu

Many end uses. Needs to be simple, scalable, robust, pragmatic, and complete (tape and disk transfers). Many spikes in demand, 30 Hz average, spikes to 150-200 Hz.

⇒ **Needed functionality of a tape REST API needs further discussion (in February?).**

Will we be able to do bulk tape transfers fast enough? There is a bottleneck in the token exchange. Questions needing convergence between stakeholders on file ownership and token granularity.

Q&A:

- No questions or comments

IAM: Hannah Short and Berk Balci

Recent milestones: Moved to HA K8s deployment in October, continuity exercise completed successfully. Upcoming milestone to turn off OpenShift instances. Choosing options for DBOD driven by balancing complexity and stability.

What are the technical limitations? Recent load test results.

Q&A:

- Question about limitations, needed rates. Do we have a problem? To be discussed. Benchmarking tests should be continued in the future to understand the technical limitations.
- Were the tests run against production systems? No, against development instances with 1M access tokens in the database.

- Does the migration to tokens taskforce have regular meetings?
- BTW there are also periodic INDIGO IAM hackathon, next one in 2 weeks:
<https://indico.cern.ch/event/1443968/>

⇒ **There a scalability concern about access token handling IAM.**

⇒ **Concerns about effort level for IAM going forward.**

Summary

There is a lack of consensus about token lifetime and scope as well as the service model and impacts on data taking in HL-LHC. The tape REST API is a crucial functionality that needs more discussion, perhaps in February. There is a scalability concern for access token handling in IAM. Concerns about IAM effort level going forward.

Maarten will have a couple of slides summarizing the session as well. Not yet on Indico.

CernVM workshop summary and plans: Valentin Volkl

CernVM users workshop Sept. 16-18, 2024 [[Indico](#)]. Selected highlights: technical retrospective, FUSE developer presentation, EESSI (European HPC software distribution using CVMFS as a filesystem layer), snapshotter benchmarking, CMS experiment report (distribution of gridpacks and LUMI HPC).

Varnish is a possible replacement for Squid, for which there are maintainability concerns. CVMFS use case seems covered, but FRONTIER is less clear.

- Welcomed by ALICE, but 200 sites would have to change. Could Varnish be bundled with CVMFS itself?

⇒ **Need an evaluation of Varnish as a possible replacement for Squid including a plan for deployment**

There was participation by the UK fusion research community, who are interested in software distribution with CVMFS.

Challenges of scale for large-core-count (512 CPU) machines concerning memory needs. Container tools: the CVMFS snapshotter tool has potential to be a benefit.

Milestones: stop support of CentOS7, deprecation of x509-authenticated CVMFS (only one user, LIGO/VIRGO).

- LIGO has a timetable to phase out x509 user certificates. They may want token support, which is not on any roadmap.

Open question about the importance of hot patches for CVMFS without remounting and interrupting workloads, but this does increase complexity.

Next workshop after CHEP26.

⇒ **Investigations of scalability of CVMFS, snapshotter, question about importance of hot patching.**

More Q&A:

- On slide 11, improvements in Stratum-1 operations would be welcome.
- On slide 12, concern about unpacked.cern.ch, potential to cause high cache turn-over for one VO also to affect other VOs served by the same cache. Have we been hit by this in the past already? VO caches could be isolated, but increases complexity of the configuration.

DOMA General Meeting

Please see Indico: <https://indico.cern.ch/event/1470863/>