WLCG Token Transition Timeline

WLCG Authorization WG

v0.91 - August 10, 2022

Introduction

Since July 2017, the WLCG Authorization WG has been looking into how authentication and authorization technologies used on WLCG may evolve towards taking advantage of federated identities and standards commonly used in industry and academia, in particular through the use of JSON Web Tokens (JWT). Notable results include the publication of the WLCG Common JWT Profiles and the WLCG Bearer Token Discovery specification.

This document aims to define milestones concerning technical aspects of the transition from X509 proxies with VOMS attributes toward the use of WLCG tokens in all relevant workflows. These milestones are tentative.

A number of prerequisites were tended to from 2020 through 2022 H1. In particular:

- To allow the VOMS-Admin legacy services to be decommissioned eventually, all its important use cases will need to be migrated to alternative services (IAM, CRIC), while the remaining use cases will be phased out. The list of use cases was obtained from analyzing the query logs and several were successfully migrated, e.g. the ATLAS AMI service, while others feature in the milestone table if they will not be phased out.

- Since the use of tokens is not supported by GridFTP operations, all GridFTP services will have to be replaced with alternatives making use of the HTTPS/WebDAV and/or Xroot protocols. This matter has mostly been taken care of under the aegis of the DOMA Third Party Copy WG and the few remaining cases will be followed up as part of the work of the DOMA Bulk Data Transfer WG.

- Storage service developers and the FTS team have implemented support for tokens to the extent possible at this time. Basic functionality and interoperability are checked daily by a token testbed comprising development endpoints for all relevant SE flavors. The testbed does not involve an FTS instance so far. The implementation and testing of advanced features (e.g. token exchange) will have to be driven by the data management frameworks Rucio and DIRAC, with corresponding implications for FTS functionality as well as testbed requirements.

- Because of the started phaseout of GSI authentication support in HTCondor, the first use case where tokens will be critical is job submission to HTCondor CEs, with ARC CEs to follow next. To help get job submission framework developers, grid experts in the experiments, site administrator representatives and other interested parties up to speed with the use of tokens, a CE and Pilot Factory Hackathon was held at CERN on June 3-4, 2021. Another such hackathon will be held on September 15-16, 2022, hosted by NIKHEF.

- IAM services were put into production for CMS and ATLAS in the course of 2021 and a campaign was launched for sites to equip their services with the *LSC* files of the VOMS endpoints provided by those IAM instances as well as the ones planned for ALICE and LHCb.

- OSG ran a campaign to have all CEs supporting ATLAS or CMS upgraded to versions no longer dependent on the Grid Community Toolkit (a.k.a. Globus) by May 1, 2022, thereby ending the support of GSI authentication for job submissions to those CEs.

- On the EGI side a campaign was launched on June 1, 2022, to have all HTCondor CEs upgraded to the 9.0.x series that supports GSI for authentication as well as tokens, and for all ARC CEs to have their REST interface enabled to allow HTCondor engines used in pilot frameworks and by the SAM ETF to continue submitting jobs to those CEs after the last HTCondor release with Grid Community Toolkit dependencies has reached its EOL by February 2023. As of August 2022, the majority of sites on EGI have already complied with those requests, allowing tokens to be used by ATLAS and CMS for job submission to HTCondor CE instances at those sites. ARC CE instances can already be configured to support tokens in *plain* job submissions, i.e. for which the ARC CE will not have to do any data management operations, but as GSI will continue working, it was deemed better to ask for just the REST interface to be enabled at this time.

- ALICE have successfully tested job submission to HTCondor CEs using tokens from the WLCG VO and can enable the use of tokens per site. LHCb depends on the support for tokens in the DIRAC framework, as do many other communities, e.g. Belle-II. The state of affairs was presented in the June 2022 GDB meeting: v8.0 should be able to use tokens for job submissions to HTCondor CEs and is expected to become available in the summer of 2022. The support for job submission tokens may be backported to v7.x.

Timeline

Milestone ID	Date	Description	Dependencies	Teams
M.1	Sep 2022	IAM is also in production for ALICE and LHCb.		CERN IT, IAM devs
M.2	Dec 2022	DIRAC versions supporting job submission tokens deployed for concerned VOs (LHCb, Belle-II, …).		DIRAC, LHCb, Belle-II, …
M.3	Feb 2023	VOMS-Admin is switched off for some experiments. Prerequisites: ● Significant VO admin functionality issues in IAM have been resolved ● User registration procedures have been switched to IAM ● IAM services are sufficiently HA ● CERN IAM team is sufficiently staffed ● Remaining VOMS-Admin use cases have been moved or will be dropped		IAM devs, CERN IT, experiments
M.4	Mar 2023	HTCondor installations at EGI sites have been upgraded to supported versions > 9.0.x. Prerequisites: ● DIRAC versions supporting job submission tokens have been deployed for the concerned VOs (LHCb, Belle-II, EGI catch-all, …) ● HTCondor CE supports (adjusted) EGI Check-in tokens ● IAM or equivalent in production for ALICE, LHCb, Belle-II, …	M.1 M.2	HTCondor Dev Team, WLCG ops, EGI ops, sites
M.5	Mar 2023	End of HTCondor support for GSI Auth (link).		
M.6	Mar 2023	Some storage endpoints provide support for tokens (at least one per service type).		WLCG ops, storage devs, sites
M.7	Feb 2024	Rucio / DIRAC / FTS have sufficient token support in released versions to perform DC24 using token authorization.	M.6	Rucio, DIRAC, FTS, experiments

M.8	Mar 2024	Sufficient storage endpoints support tokens to allow DC24 to be done using only tokens.		Storage devs, WLCG ops, sites
M.9	Mar 2025	Grid jobs use tokens for reading and stageout.	M.8	Experiments, WLCG ops, sites
M.10	Mar 2026	Users no longer need X509 certificates.	M.9	Experiments, MW devs, CERN IT

Notes on selected milestones

M.3: VOMS-Admin is switched off for some experiments

Milestone Date: Feb 2023

While VOMS-Admin has not been switched off for an experiment, it should typically keep being used for user registration and removal. A VOMS importer utility is run multiple times per day to propagate user changes to the IAM instance of the experiment, which in turn is used to obtain tokens for the workflows that need them and acts as an additional VOMS service endpoint.

One significant impact of the IAM transition is the replacement of the VOMS-Admin API with the incompatible SCIM API provided by IAM. From the VOMS-Admin logs it appears the relevant VOMS-Admin queries concern 2 use cases:

1. Experiment middleware that needs to enumerate all users and their privileges.
2. Services that have not converted or cannot convert to using VOMS.

Examples of the first use case are **Rucio** and the ATLAS **AMI** service, both of which have been able to switch to IAM for their needs. Examples of the second use case are:

- **EOS**
  - Can take the necessary information from IAM or directly from the VO. For example, EOSCMS takes the mappings directly from CMSWEB.
- **GGUS**
  - Needs the lists of DNs that are allowed to open Team and Alarm tickets.
  - An automatic procedure is not needed per se: privileged VO members can open a ticket whenever a DN needs to be added or removed, which happens rarely.
- **EGI Operations Portal**
  - It needs the *number* of users per VO, for funding reasons.
  - That could be obtained from IAM through a simple, possibly public API.

M.4: HTCondor installations at EGI sites have been upgraded to supported versions > 9.0.x.

Milestone Date: Mar 2023

At the time of writing, EGI Check-in tokens cannot be used for job submission to HTCondor CEs. This mismatch has already been looked into and a practical solution is expected to be agreed e.g. during the CE token support hackathon of 15-16 Sep 2022. Check-in tokens are foreseen to be used by some of the EGI VOs as well as for the CE availability test jobs submitted by EGI ARGO framework and hence the support of those tokens is crucial for EGI sites, VOs and operations.

M.6 - M.10: data and job management milestones

Milestone Dates: Mar 2023 - Mar 2026

At the time of writing, workflow details involving Rucio and/or FTS vs. SEs have mostly been identified, but may need to be re-discussed if major implementation or operational hurdles are encountered. The token testbed covers basic functionality and interoperability. Most endpoints pass most of the tests. Rucio and DIRAC should drive these developments to allow FTS, IAM and SE development teams to know what functionality is expected of their service instances. An example of a potential hurdle would be the configuration of an SE to allow the concurrent, *consistent* use of X509 / VOMS and tokens, which is very desirable, if not crucial, for a smooth transition toward tokens. DC24 only involves third-party copies, no user workflows.

For grid jobs to be able to make use of tokens, the experiment job frameworks will have to be extended in ways to allow the necessary tokens to be made available at the right times, ideally shortly before they are going to be used, to allow their lifetimes to be kept short, while in practice it remains to be seen per experiment what is feasible and deemed sufficiently secure.

For end users to be able to switch to tokens, convenient mechanisms have to be put in place. Data and job management client tools need to hide the use of tokens as much as feasible, possibly making use of auxiliary services that manage the tokens for the workflows that need them. Examples of such services are HashiCorp Vault + *htgettoken* (in use at FNAL) and MyToken (used by EOSC Synergy). It is envisaged that users will have to go through grid workflow authentication dialogues much less frequently and more conveniently than today.