

WLCG transition from X.509 to Tokens: Progress and Outlook

Tom Dack (STFC, UKRI)

on behalf of:

WLCG Authorisation Working Group

October 2024


Agenda

- WLCG IAM Token Issuers
- Token Evolution in the WLCG
 - Data Challenge 2024 - Lessons Learnt & Follow-ups
 - Token Profile Improvements
- Outlook & Next Steps
 - Token Usage in Grid Jobs
 - User Experience

WLCG IAM Token Issuers


WLCG IAM Token Issuers

- INDIGO IAM Deployed at CERN on Kubernetes
 - Believed to be the first **Highly Available** production deployment
 - GitOps workflow with ArgoCD hugely simplifies maintenance
- Experiment instances: ATLAS, ALICE, CMS, LHCb, FCC, ILC, AMBER, CALICE, COMPASS
- Operations: DTEAM, OPS
- Collaborating closely with INFN CNAF team for prompt updates



WLCG IAM Token Issuers

- VOMS(-Admin) services were fully retired by July 2024 (see next page)
 - IAM provides backwards compatible VOMS endpoint
- Kubernetes replaces OpenShift deployment
 - Greater control
 - Fewer dependencies
 - Easier to replicate between data centers for yet better HA



From VOMS(-Admin) to full dependence on IAM

- VOMS(-Admin) services were **not foreseen** to be supported beyond CentOS 7
- The **CentOS 7 EOL** thus became the VOMS(-Admin) service EOL: **30 June 2024**
- WLCG and other VOs have been **preparing for the phaseout** of the VOMS(-Admin) services for more than 2 years
 - **Critical VOMS-Admin use cases had to be moved to IAM**
- The transition **finished on schedule**, with only minor fallout
- This was made possible **thanks to** a **big collaborative effort** from **experts** of all parties concerned!

Token Evolution in the WLCG

Token usage in the 2024 Data Challenge

- **M.7 (Feb 2024):**
Rucio, DIRAC, and FTS have sufficient token support in released versions to perform DC24 using token authorization.
- **M.8 (Mar 2024):**
Sufficient storage endpoints support tokens to allow DC24 to be done using only tokens.
- DC24 was a **major** milestone in the WLCG transition to tokens
- DC24 enabled **scale tests** with tokens of services involved in data management
 - Rucio (ATLAS & CMS) and DIRAC (LHCb)
 - FTS
 - IAM
- Whilst M.8 was not wholly met, DC24 was a **big success** overall and allowed us to draw conclusions from **millions** of transfers completed with **tokens!**

Lessons from the 2024 Data Challenge

- Token Lifecycle Management:
 - Implement shorter token lifetimes to facilitate quicker cleanup processes during peak usage periods.
- Tokens & Data Management:
 - Revision of token orchestration to be more efficient for large scale data transfers
- Token Management Enhancements:
 - Stop storing access tokens in the DB to improve the performance. This needs a modification and/or replacement of the token management engine (MITREid). IAM developers are working on this.
- Performance Testing:
 - Enhance IAM performance tests to make them closer to the real use-cases and include closer examination of latency issues.

Tokens & Data Management (1)

Several ideas for more **sustainable** use of tokens in large scale data management have been discussed between experts of the services involved

- Current focus is on FTS workflows

A **new model** was proposed for testing purposes:

1. Tokens have scopes per individual file and long-ish lifetimes
 - A stolen token thus could be used for some time, but with only little potential damage
2. The FTS just uses those tokens without any exchanges or refreshing
 - Thus avoiding a big load on itself as well as IAM
3. If a token expires, its corresponding transfer will fail, passing the ball back to Rucio/DIRAC

The FTS codebase now supports this flow, alongside the process used in DC24 – which remains required by communities within and outside of WLCG

- Further enhancements were also discussed and will be considered later

Tokens & Data Management (2)

- ATLAS have started using this **in production** as of late August
 - 15 sites, all served by the CERN FTS, which has the new code
 - Starting with SCRATCH_DISK transfers, followed by DATA_DISK
 - No use of tokens during weekends for the time being
 - Typical token rates are 1-2 Hz, which occasional spikes of 5 Hz
 - Lifetimes currently are 2 weeks, to be reduced with more experience
 - The removal of tokens is left to the background cleanup job in IAM
 - Avoid Rucio complexity & interference that may affect IAM performance (see next slide)
 - The max number of concurrent tokens stored so far has been **3 M**
 - Max accurate as of Monday 21st October
 - Already more than the overall maximum seen in DC24, no problems so far

Tokens & Data Management (3)







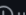













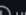





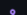





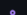
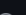
- The intended solution here is to have **IAM stop storing access tokens**
 - WLCG does not plan to use Token Introspection, instead relying on offline validation
 - Removing tokens from the IAM DB requires moving DB handling code out of the existing third-party framework, which first requires other changes
 - This change is planned to be completed before the end of 2024

Performance Testing

- **High-rate stress tests** are desirable, but currently not an option, as the same IAM instances are serving all other use cases
 - During DC24 downtimes of 1-2 days could have been tolerated, as IAM use cases were less time-critical then
- An **opportunity** has been identified: the migration of IAM instances from **OpenShift** to **Kubernetes**
 - Rather than decommissioning the old services immediately, these can first be reused for **stress tests** – when reconfigured with their own DB instances
- This would allow token usage to be ramped up in a steady manner, until instabilities encountered
 - Possibly due to **DB Limitations** encountered by the current IAM code

Token Profile Improvements

- Version 2.0 of the **WLCG Token Profile** is under preparation
 - This fixes various issues identified within the v1.0 profile, others still under work
 - Several need to be agreed in AuthZ or DOMA BDT WG meetings
- WLCG is engaging with the Grand Unified Token Profile WG
 - Definition of an agreed VO attribute – for accounting, etc
 - Various details about the GUT WG can be found at: <https://edu.nl/atvkvx>

<input type="checkbox"/>	 Use consistent style for acronym and RFC hyperlinks in token glossary item #53 by DrDaveD was merged on Apr 1  updated on Apr 1
<input type="checkbox"/>	 Extend the token cache lifetime #17 by bbockelm was merged on Mar 30  updated on Apr 1
<input type="checkbox"/>	  Document edits not changing the semantic #49 by mambelli was closed on Mar 31  updated on Mar 31
<input type="checkbox"/>	 Correction of typos, grammar, etc. #51 by maarten-litmaath was merged on Mar 31  updated on Mar 31
<input type="checkbox"/>	 Improve the description of path handling in storage scopes. #48 by maarten-litmaath was merged on Mar 30  updated on Mar 30
<input type="checkbox"/>	 remove non-AuthZ, protocol-specific behaviour constraint #42 by paulmillar was merged on Mar 30  updated on Mar 30
<input type="checkbox"/>	  add definition of a VO #41 by paulmillar was closed on Mar 30  updated on Mar 30
<input type="checkbox"/>	 add alternative VO definition and description #47 by maarten-litmaath was merged on Mar 30  updated on Mar 30
<input type="checkbox"/>	  Propose a semantics for the wlcg.group claim #2 by bbockelm was closed on Mar 30  updated on Mar 30
<input type="checkbox"/>	  Provide a definition of wlcg.groups that is independent of VOMS. #15 by bbockelm was closed on Mar 30  updated on Mar 30
<input type="checkbox"/>	 Consolidation of several changes concerning wlcg.groups #46 by maarten-litmaath was merged on Mar 30  updated on Mar 30
<input type="checkbox"/>	 Change storage.write in wlcg.capabilityset examples to correct storage.create #16 by DrDaveD was merged on Dec 7, 2022  updated on Jan 11
<input type="checkbox"/>	 fix bad example of storage authz in scope #40 by paulmillar was merged on Jan 3  updated on Jan 3
<input type="checkbox"/>	 CE Scope definition #11 by hannahshort was merged on Jan 2  updated on Jan 2
<input type="checkbox"/>	 Update description of access tokens with capability and group assertions #23 by paulmillar was merged on Jan 2  updated on Jan 2

Outlook and Next Steps

Next milestones

- **M.9 (Mar 2025): Grid jobs** use tokens for reading and stageout.
 - Implies **significant changes** in workload management systems
 - Tokens to be provided just in time?
 - Scopes? Audiences? Lifetimes?
 - Scalability concerns?
 - Fallback on X509 + VOMS during transition period?
- **M.10 (Mar 2026): Users** no longer need X.509 certificates
 - **Tools** should be sufficiently smart to obtain the correct tokens for specific operations
 - **Auxiliary** services such as *Vault + htgettoken* or *MyToken* may be needed to simplify the user experience, used under the hood by tools for job and/or data management
 - Investigations in this space are already underway within some experiments

Other developments

- Various **IAM improvements** are still desirable in the short term
 - Fixes for the last of the current & new high-priority issues
 - The next IAM hackathon will be held Nov 27-28 at **IJCLab**, Orsay
 - Another IAM hackathon is planned for Feb 10-11 2025 at CERN
- Engagement with **Token Trust & Traceability (TTT) WG**
 - Aiming to equip site admins, VO experts, developers, ... with best practices for token usage, which will also provide input for policy documents
 - See the next talk!
- **Accounting (APEL)** adjustments for Tokens – short vs medium term needs
 - Some VOs have stopped equipping jobs with X509/VOMS proxies, not yet on the horizon for LHC experiments
 - Stop-gap solutions to be investigated by APEL team, medium-term solution will likely take input from GUT WG

Conclusions and outlook

- **Collaborative** efforts – especially with parallel WGs – will keep many stakeholders involved
- We look forward to increased **reliance** on the **benefits** of tokens
 - Aiming to reach the next levels in **data management**
 - Equipping **jobs** for reading data and uploading results
 - Making the **user experience** both simpler and more secure
- Stay tuned – ***tokens will return...***
 - ... at the next CHEP!

Some CHEP talks with token content

- Evolving INDIGO IAM towards the next challenges
- Preliminary findings and recommendations from the Token Trust and Traceability Working
- A Lightweight Analysis & Grid Facility for the DARWIN Experiment
- Fermilab's Transition to Token Authentication
- CMS Token Transition
- Supporting medium/small-sized experiments in the transition from X.509 to JWTs
- FTS3 Token Support for a Proxy-less WLCG world
- Latest developments of the PUNCH4NFDI compute and storage infrastructures
- Addressing tokens dynamic generation, propagation, storage and renewal to secure the GlideinWMS pilot based jobs and system.
- dCache project status & update
- Data Challenge 2024 - CMS activities
- Evolving StoRM WebDAV: delegation of file transfers to NGINX and support for SciTags
- Posters
 - A Managed Tokens Service for Securely Keeping and Distributing Grid Tokens