Token transition state of affairs

Journées LCG-France, 20 June 2024

M. Litmaath

Background (1)

- Since 20 years, WLCG users have been using X509 certificates to identify themselves to grid services
  - Certificate lifetimes are up to 400 days
- It is not sufficient for users to able to identify themselves: there must also be a way for grid services to know to which *virtual organization* (VO) a user belongs
  - E.g. which LHC experiment
- For that purpose we have been using a Virtual Organization Membership Service (VOMS) for each VO that we need to manage
  - In particular the LHC experiments
- The membership of each VOMS instance can be queried by grid services, but most grid services do not need to do that anymore, as is explained on the next pages
  - Such queries are done against **VOMS-Admin** endpoints

Background (2)

- Only in browsers are user certificates used directly
- For most grid workflows, short-lived proxy certificates ("*proxies*") are used that are derived from long-lived (proxy or real) certificates
  - To limit the risk of theft of long-lived credentials and their potential abuse
- Per workflow, such proxies can be equipped with additional attributes indicating to which **VO** the user belongs, to which **groups** within the VO, and which of their allowed **roles** the user is assuming for the given workflow
- Such attributes are captured in an *attribute certificate* that is obtained from the VOMS service of the VO and then **embedded** in the proxy that the user will use: a *VOMS proxy*
- VOMS-aware services will normally bestow client requests with **privileges** based on the **contents** of those attribute certificates

Why change?

- While the previously described machinery has done the job for 20 years, there are several aspects that we do not like:
  - **User certificates are cumbersome**
    - *“I got a new certificate and nothing works anymore: help!”*
  - **Proxy certificates are used only in grid workflows, not in industry**
    - We thus need to maintain our own certificate handling stacks
  - **VOMS attributes are still rather coarse-grained**
    - A stolen VOMS proxy could still open the door to significant abuse

The solution: move to **tokens**

- **Standard practice** in industry and academia
  - Used **under the hood** → users do not have to know about the details!

- **Fine-grained access controls** → improved security
  - *Capabilities* a.k.a. *scopes*
    - Which operations are allowed
    - Can be set for individual files as needed
  - *Audiences*
    - Which services should accept a given token

- The **price** to pay: decide per operation **which** token to use!
  - Complicating middleware and workflows…

Token example

```
$ ls -l tmp-$$.dat
-rw------- 1 alicesgm alicesgm 881 May 25 15:57 tmp-3140407.dat
$ decode-token.sh tmp-$$.dat
{
  "wlcg.ver": "1.0",
  "sub": "a4f952ab-6e43-059c-c530-80df119a018b",
  "aud": [
    "ce01.some.site:9619",
    "ce02.some.site:9619"
  ],
  "nbf": 1716645428,
  "scope": "compute.create compute.read compute.cancel compute.modify",
  "iss": "https://alice-auth.web.cern.ch/",
  "exp": 1716991028,
  "iat": 1716645428,
  "jti": "852f012c-5bcb-4a9c-a2e6-ec2e25801560",
  "client_id": "ae76ab27-cc15-4082-a9bc-50ad587a73d6"
}
$
```

New AAI architecture for WLCG

IAM = Identity & Access Management

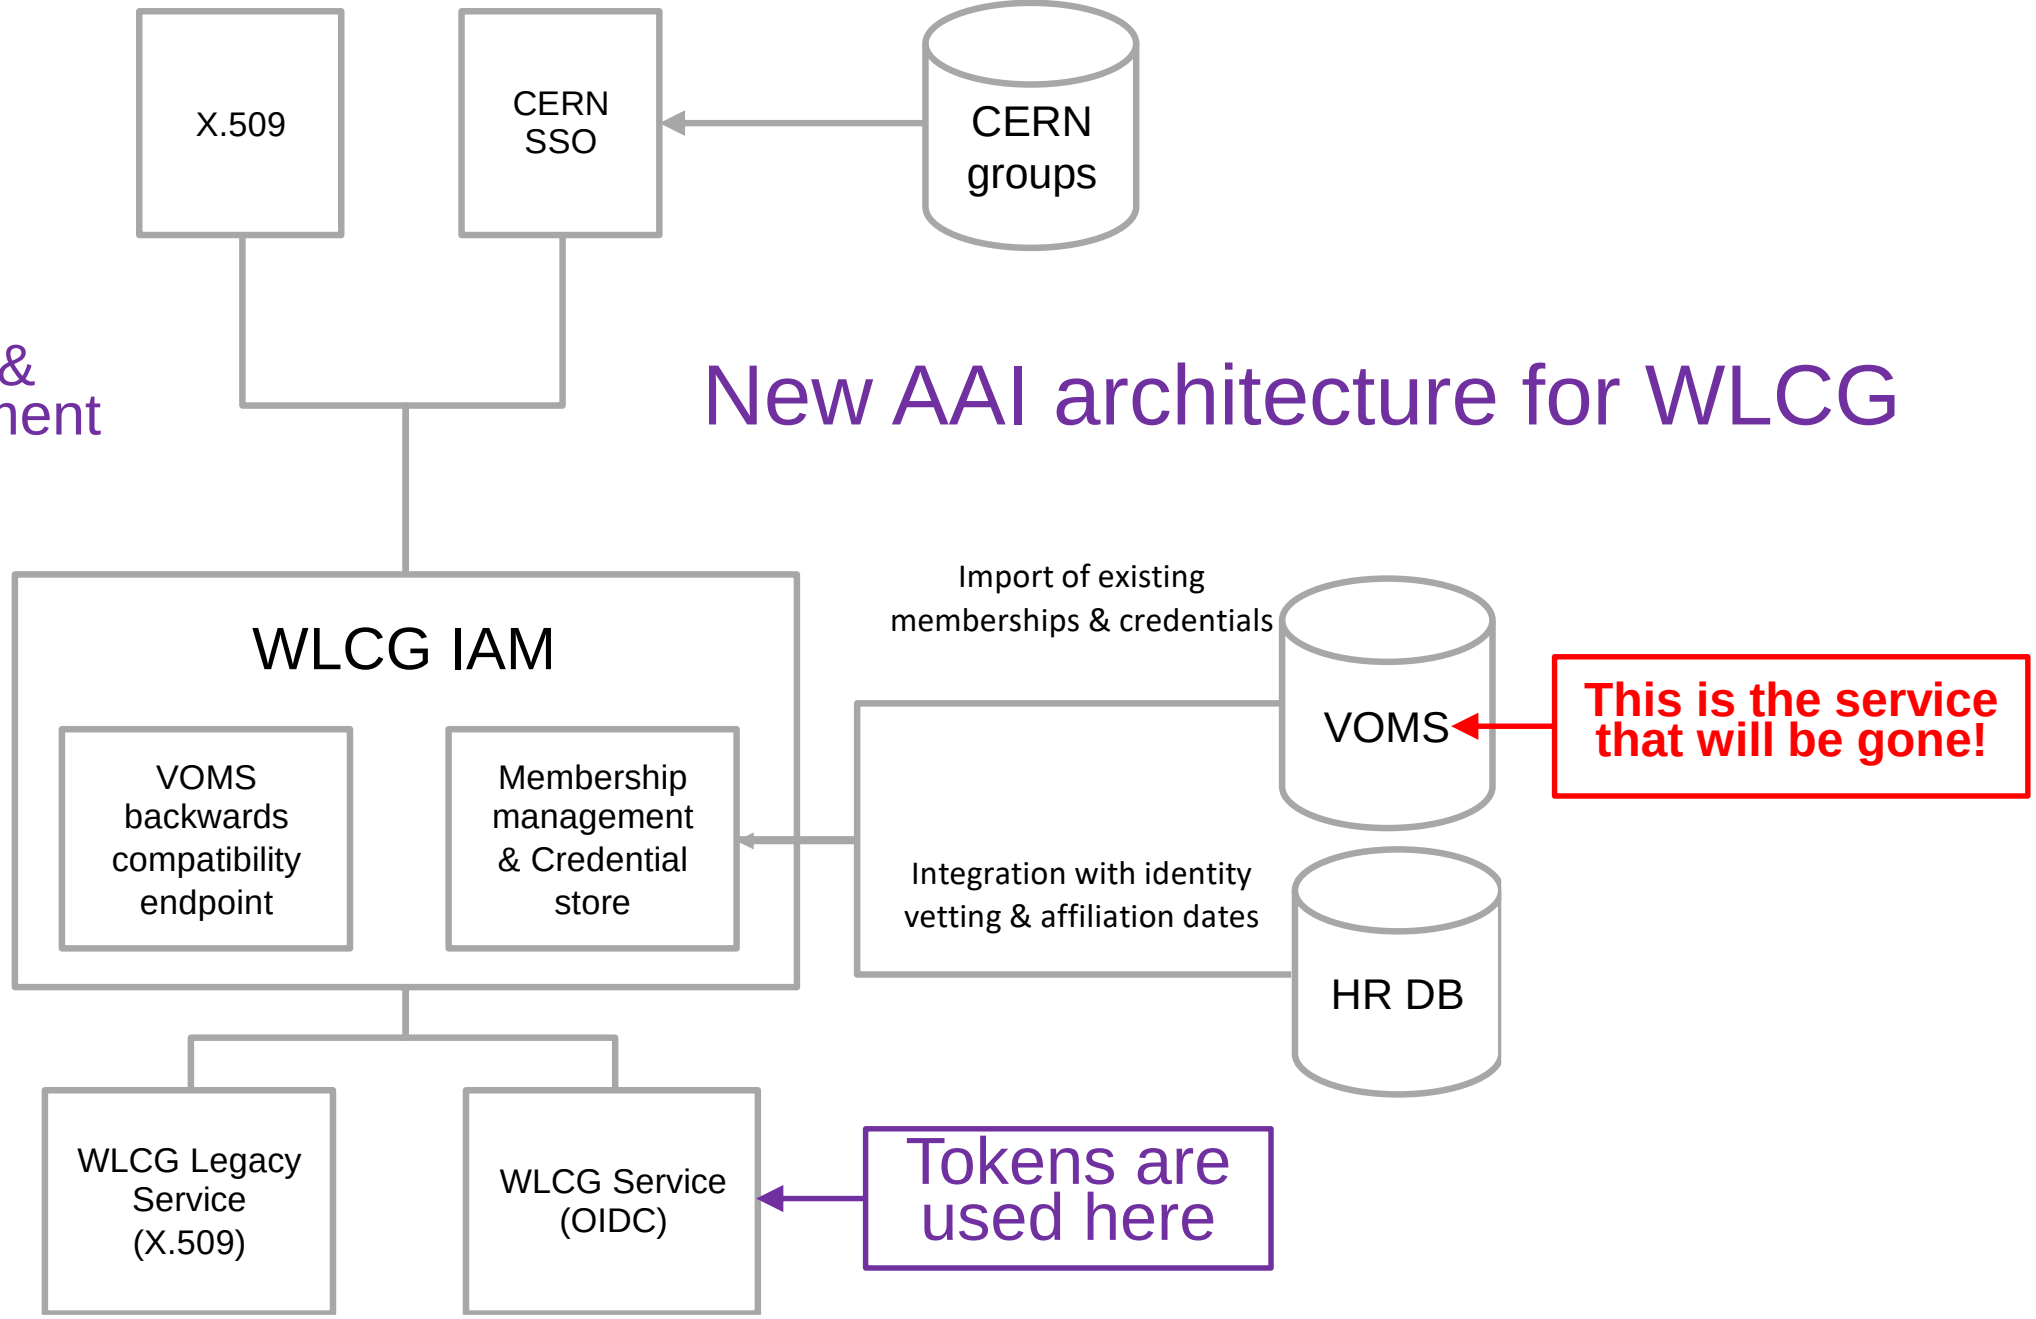

Crucial legacy functionality remains supported for now

WLCG Token Transition Timeline

- Since July 2017, the WLCG Authorization WG has been working on the design and implementation of the new AAI architecture for WLCG
- A token transition timeline with tentative milestones was published on August 22, 2022
- Though most milestones were postponed for several reasons, progress has been steady
- Details about various aspects are provided on the next pages

Computing (1)

- Campaign to have HTCondor CEs upgraded to maintained versions
  - Intermediary version: v9.0.20
    - Supports tokens, SSL (no VOMS mapping) and GSI (with VOMS mapping)
  - Versions >= v23.x
    - Support tokens and SSL without VOMS mapping
  - Versions >= v23.5.2
    - Support tokens and SSL **with VOMS** mapping! (release notes)
  - To use SSL mappings with proxies, clients must also run **recent** versions!
  - All versions support *delegation* of VOMS proxies to be used by jobs and APEL
    - Mind this HTCondor (CE) setting for APEL: USE_VOMS_ATTRIBUTES = True
  - 53 tickets, >= 18 solved
  - Many sites prefer upgrading to **EL9** at the same time
    - APEL client, parsers and python-argo-ams-library available from the WLCG repository

Computing (2)

- APEL support for *tokens* is discussed separately between concerned parties
  - APEL, HTCondor, ARC, several sites, EGI Ops, WLCG Ops Coordination
  - Stopgap approaches are proposed for the time being
    - Map token issuers / subjects / … to **pseudo** VOMS FQANs
    - The rest of the machinery can stay unchanged
  - Medium-term solution expected from the GUT Profile WG (see later)

IAM service developments (1)

- All production instances at CERN are on a v1.9.0 pre-release since June 13 (to be updated next week)
  - Fixing various high-priority issues in the area of VO management
  - Other fixes are still expected in a few weeks

- The "dteam" instance is usable for service monitoring with tokens
  - Users are imported from the VOMS-Admin service until its retirement
  - VO membership is managed by EGI Operations and WLCG Ops Coordination

- A campaign has been launched on April 19 for sites to configure support for the instance for the "ops" VO by June 1st
  - 156 tickets, >= 140 solved

IAM service developments (2)

- New instances for the LHC experiments have been created on **Kubernetes**, sharing their DBs with the *OpenShift* instances
  - For better **load-balancing**, **logging**, **monitoring**, **GitOps** and **HA** options
  - They will eventually replace the current production instances on OpenShift
    - Dates to be decided per experiment
  - Sites have been ticketed to add support for the future VOMS endpoints and token issuers by May 31st
    - Many tickets are still open

- A timeline with tentative milestones for the transition from VOMS-Admin to **full dependence on IAM** was agreed for the LHC experiments
  - ATLAS and ALICE have migrated, LHCb look ready for June 24, CMS still WIP
  - **Supported use cases** for the time being are:
    - VO management
    - VOMS proxies
    - **Low-rate** token issuance for pilot jobs, SAM tests etc.

Why is the VOMS-Admin phaseout happening now?

- The VOMS(-Admin) service was expected to be **unsupported beyond CentOS 7**
  - Packages are available for EL8 and EL9, but not fully tested
    - And incompatible with MySQL 8
  - The developers at CNAF cannot afford to keep maintaining two VO management solutions

- The CentOS 7 EOL thus became the VOMS(-Admin) EOL
  - **June 30, 2024**

- However, **limited support beyond CentOS 7** has been obtained by EGI as announced through a broadcast on June 19
  - Further details in this section of the June 10 EGI ops meeting notes

VOMS-Admin phaseout snapshot

- April 29
  - Remove legacy VOMS servers from “vomses” – in production for Puppet at CERN as of **May 7**
    - *Many tickets from users who were disabled in IAM due to sync issues, all **fixed** manually*
    - Versions 2.0.0 of the wlcg-voms rpms only contain the LSC files, no “vomses” files
    - Broadcast sent to wlcg-operations list
    - Not critical: voms-proxy-init would fail over to a VOMS server that works
- May 06
  - VOMS-Admin switched off for first VO → delayed until after the WLCG workshop
    - No VOMS-Admin service is deleted yet
- May 31
  - Deadline for sites to have configured support for the Kubernetes instances, *including “ops”*
  - Start considering switching off OpenShift instances – unlikely to happen before July → September…
    - Possibly depending on HA situation on Kubernetes
    - Ultimately also the oidc-agent menu listing IAM instances should be updated accordingly
- June 03
  - VOMS-Admin switched off for last VO → *was in fact the **first** VO: ATLAS !*

Data Challenge 2024 (1)

- DC24 was a **major** milestone in the WLCG Token Transition Timeline

- It has allowed **scale tests** with tokens of services involved in data management
  - Rucio (ATLAS & CMS) and DIRAC (LHCb)
  - FTS
  - IAM

ALICE use "*access envelope*'' tokens with XRootD services since 20 years, but a **future switch to WLCG tokens** is an option being worked on!

- The Data Challenge sessions of the WLCG workshop also feature observations and discussions about the use of tokens

Data Challenge 2024 (2)

- DC24 has allowed us to draw conclusions from **millions** of transfers done with tokens!

- It is clear that some ways in which tokens were used are not advisable for the long term
  - FTS and IAM instances got overloaded in various ways, causing failures and requiring interventions

- Several ideas for more **sustainable** use of tokens will be discussed in the next months between experts of the services involved
  - Concerning token audiences, scopes, lifetimes, exchanges, refreshing, ...

- For the time being, we keep relying on VOMS proxies for most of our data management

AuthZ WG items (1)

- Various IAM code changes were desirable in the short term
  - In particular to fix VO management issues in view of VOMS-Admin EOL
    - They were the main focus of an IAM Hackathon at CNAF, May 29-30
  - Lessons learned from DC24 will be taken into account later
    - In particular, stop storing access tokens in the DB
  - **No high-rate usage** is foreseen for the time being
    - First, **sustainable** token usage patterns have to be agreed and tested between the parties involved in data management

- Version 2.0 of the **WLCG token profile** is under preparation
  - Fixing a number of issues encountered with v1.0
    - In the description and/or implementation
  - Most PRs have been merged and the corresponding issues closed
    - A few open cases need to be discussed in AuthZ or DOMA BDT WG meetings
    - More difficult cases will be postponed for future revisions

AuthZ WG items (2)

- The **Grand Unified Token (GUT) Profile WG** has met 5 times already (agendas)
  - Good progress with its current main challenge: how to determine the VO for the token profiles and the various use cases we need to handle
    - WLCG tokens, SciTokens, EGI Check-in tokens
  - A **new, common attribute** will be defined with practical semantics
    - Details TBD in upcoming meetings

- The **Token Trust & Traceability WG** will meet again June 25
  - Aiming to equip site admins, VO experts, … with best practices for tokens, which will also provide **input for policy documents**
    - Recipes, tools, log mining, testing, debugging, monitoring, banning, ...
  - For example, to prevent exposure of tokens through logs!
  - Or how to use the “dteam” VO for monitoring with tokens (see next page)

Auxiliary services

- The March 7 Ops Coordination meeting had a presentation on **MyToken**
  - Used at KIT e.g. to monitor dCache services with “dteam” tokens
  - Further details are available here
- The May 2 Ops Coordination meeting had a presentation on ***htgettoken*** + ***HashiCorp Vault*** as a Service for Managing Grid Tokens
  - In production at FNAL for various communities since >1 year
- Such auxiliary services are expected to facilitate various use cases
  - Production workflows
  - Monitoring
  - User workflows
    - To help avoid that users need to know anything about tokens!

Conclusions and outlook

- **Collaborative** efforts will involve many of us in the next months
  - VOMS-Admin EOL – planned deadline for CERN VOs June 30
  - IAM usability for VO administration by LHC experiments and others
    - Work on high-priority issues continuing for a **1.9.x release** in July
  - HA options for LHC experiment IAM instances – advancing
  - **Data management:** lessons learned from DC24
    - Aiming to reach the **next level** of token usage in the second half of this year
  - HTCondor CE versions that no longer support GSI – along with moving to **EL9**
  - APEL adjustments for tokens – short vs. medium term
  - GUT Profile WG progress toward a new VO attribute – for accounting and more
  - Version 2.0 of the WLCG token profile – to signal where we intend to go
  - More **deployment and operations** know-how – also providing input for policies
  - More use of auxiliary services – gradually benefiting more use cases

Thanks for your kind attention!