Tokens, Trust & Traceability

GDB, 13 September 2023

M. Litmaath

v1.1

Computing (1)

- ALICE
  - Token deployment campaign done, for HTCondor CE sites only
  - VOMS proxies will continue being used with ARC CEs for now
    - And with HTCondor CEs *in addition*, because APEL currently expects that
    - ALICE jobs do not need that proxy anymore
- LHCb
  - Token deployment campaign mostly done, for HTCondor CE sites only
- First HTCondor CE token campaign on EGI done
  - HTCondor CE v5 + condor-9.0.x

Computing (2)

- Second EGI campaign to be launched in the next few weeks
  - HTCondor CE v5 + condor-9.0.19 ← when available in UMD-4 (CentOS 7 only)
  - It allows X509 / VOMS proxy *identities* (no FQANs) to be mapped via the SSL method
    - No wildcards needed
    - Clients relying on SSL mappings also need to use a recent condor version
  - When all customers of a CE can be mapped through tokens or SSL, i.e. no longer need GSI support, the CE can be upgraded to HTCondor CE v6 with condor >= 10.7.x
    - Those recent versions also support the plug-in call-out for EGI Check-in tokens
    - Mind this setting for APEL: `USE_VOMS_ATTRIBUTES = True`
    - Available for CentOS 7 and EL 9 (→ UMD-5)

- EGI Check-in token details were presented in the June GDB

AuthZ WG meetings

- 7 meetings (Mar 23 – Aug 24) since the March 22 GDB update
  - Next Sep 14
- Plus the CHEP talk, May 9
- And the IAM hackathon organized by RAL, July 25-26
  - 20 participants, a few via Zoom
  - Such events are appreciated also for community building!
  - Various matters were worked on
  - Of particular interest: stable access token rates achieved up to 800 Hz when 3 login pods are used
    - Thanks to NGINX + OpenShift resource tuning
    - 600 Hz with just 1 pod
    - Further increases still expected from DB handling improvements

AuthZ WG selected discussion topics (1)

- How to determine the VO for various kinds of tokens?
  - Needed e.g. by APEL
  - See page 7

- Use of `scope` vs. `wlcg.groups`
  - Groups are primarily foreseen to provide *context* information, but may also be used for authZ decisions by services that have been configured to use groups *instead of* capabilities, as agreed between a VO and those services
  - WLCG token profile to be updated accordingly

- Token rates
  - Avoid IAM being the bottleneck
  - High rates vs. transparent service downtimes
  - Mitigate through longer lifetimes and/or less fine-grained scopes
  - Impacts IAM, FTS, Rucio, DIRAC
  - Different treatments of read / create vs. modify (delete)
  - Today no such distinctions with VOMS proxies!
  - We hence do not need a perfect system right from the start

We will gain operational experience in **DC24** (Feb 12-23)

AuthZ WG selected discussion topics (2)

- Move CERN VOs from OpenShift to Kubernetes, as used at CNAF
  - Start with “ops”, “dteam”, “alice”
  - Current IAM LSC files will need to be replaced

- VOMS(-Admin) server vs. EL 9
  - VOMS server rpms are in EPEL 9, unclear how much testing was done
    - Could serve DUNE after June 2024, CentOS 7 EOL
  - VOMS-Admin: no intention to spend time on it
    - Not used by DUNE: VO managed through CILogon

- WLCG profile issues & improvements
  - To be followed up as of September, leading to v1.2

- Release of `httokensh` in OSG
  - A shell that manages the token for a child process

How to determine the VO of a token?

- WLCG tokens: 1 VO per issuer
- SciTokens: only big VOs have their own issuers (costly)
  - Fermilab *sub-VOs* have different `scope` paths and `wlcg.groups`
- EGI Check-in: 1 issuer for all VOs
  - VO encoded in verbose `eduperson_entitlement` claim values (AARC-G069)
    - `urn:mace:egi.eu:group:<vo_name>:role=member#aai.egi.eu`
- VO could be inferred from `issuer + subject`
  - Would require an external, potentially *fragile* mapfile
- Standard claims have been registered in IANA
  - `groups, roles, entitlements` may not be suitable for our **lean** (!) tokens
  - And even AARC has `eduperson_entitlement` vs. `entitlements`
- Each service could allow several ways, to be tried in a configurable order
  - Complexity…

DOMA BDT WG meetings – selected topics (1)

- 10 meetings, Apr 5 – Sep 6
  - Next Sep 20

- Rucio & FTS token workflow designs
  - FTS clients are to provide source and destination access tokens
  - FTS will use token exchange to obtain its own access & refresh tokens
  - Cache & reuse tokens when possible
    - Depending on audiences, scopes & lifetimes
  - Optimizations beyond DC24: Rucio (and possibly DIRAC) *call-backs* to refresh tokens only when needed

- LHCb pre-signed URL proposal
  - Considered optimization R&D topic for after DC24

DOMA BDT WG meetings – selected topics (2)

- `storage.create` vs. potential *rename* abuse
  - Can be sufficiently mitigated with paths

- Use of tokens would benefit from consistent namespace across sites
  - Ideally the VO *base path* need not be specified by SE clients
    - dCache may insert it automatically as needed (not yet released)
  - Would simplify confining IAM clients to particular LFN paths
  - Supported at least by dCache, StoRM, XRootD
  - May require SE downtime for configuration change + adjusting the catalog of the VO
  - Symbolic links might be taken advantage of in some cases
  - Support for *wildcards* would necessitate extensive discussions

DOMA BDT WG meetings – selected topics (3)

- Proposals for DC24 transfers with tokens
  - Reminder: only production, no user workflows
  - Participating SEs need to run recent versions and configure token support according to experiment requirements
    - To be checked e.g. through SAM tests
  - When both TPC parties support tokens, Rucio & FTS & storage should rely on tokens instead of X509 VOMS proxies
  - Transfers with tokens should also generate realistic load on our token issuer infrastructure which allows us to better understand future requirements and necessary improvements
  - We expect successful transfers with tokens where the failure rates are not different compared to transfers done with X509
    - And no performance degradation in participating services
  - Monitoring enhancements are needed
    - Information about FTS authorization method per transfer
    - Performance data collected from IAM for different token flows and clients (rates, response times, ...)
  - Tests and mini challenges this autumn and January

Resource **Trust** Evolution TF meeting, June 29

- ATLAS procedures to use storage resources at Google and Amazon were presented
  - A load-balancer service offered by the cloud provider to sit in front of the actual storage needs to be given a CERN DNS alias in order to allow it to be equipped with an IGTF host certificate from the CERN Grid CA
- Felt to be a hack that also relies too much on individuals
- Turned out to be the proper way, encouraged by cloud providers!
- The procedures need to be polished and become standard
  - To be followed up e.g. in CERN IT
- The need for IGTF certificates may well go away in a few years

Token **Trust** and ***Traceability*** (TTT) WG

- Instantiated in August, intended to fill a role similar to that of the previous Traceability and Isolation WG.
  - And drawing from the findings of that group.
  - Working alongside members of the AuthZ WG, meetings within the Security Group in Indico.
- Meeting approximately once a month, no regular slot yet.
- Aiming to bring together collaborators from a range of communities
  - WLCG, EGI, DUNE, approaching SKA, others
- Intended to cover the Token side of the authZ coin - Federated Identity provision is important but mostly out of scope.
  - As is the user-side experience.
- Goal is to produce tangible outputs
  - Policy: Consider what is Best Practice. Risk Identification and Analysis. Building Trust in Tokens.
  - Documentation: Write down the above. And also produce “How-Tos”, guides and manuals.
    - E.g. “Understanding Token Flows for Admins”, “Token Job Tracing”, “Incident Response and Forensics in a Token-based environment.”
- Want to know more? Contact **Matt Doidge**
  - Or look up the CERN e-group token-trust-and-traceability-wg

Conclusions and outlook

- The main token *development*, *deployment* and *testing* objectives at this time are about:
  - Computing – HTCondor CE versions that no longer support GSI
  - Data transfers – preparations for DC24

- The use of cloud storage now is on a more secure footing thanks to findings in the Resource *Trust* Evolution TF

- The *trust* and *traceability* of token workflows are to be served by various kinds of documentation to be produced by the TTT WG

- To be continued…