Token Transition update

GDB, 13 December 2023

M. Litmaath

v1.1

Computing (1)

More details in Sep GDB update

- Second EGI campaign was launched on Nov 3
  - HTCondor CE v5 + condor v9.0.19 ← using the WLCG repository
  - It allows X509 / VOMS proxy *identities* (no FQANs) to be mapped via the SSL method
    - Clients relying on SSL mappings also need to use a recent condor version
  - A fatal flaw was then discovered at a few sites: the desired *fallback* from SSL to GSI did not work
    - Somehow this had not been spotted in the many tests…
  - The fix turned out to be quite non-trivial and took a few days to design and test – thanks very much to Jaime Frey from the HTCondor Team!
  - The campaign then was restarted with **v9.0.20** on Nov 17
  - Out of 53 tickets, 10 are solved, many in progress, several on hold
    - Most sites will work on this as of early next year – reminders will be sent
    - Some sites prefer upgrading to EL9 at the same time, but APEL client & parsers not yet available for that platform → to be followed up with priority

Computing (2)

- Third EGI campaign to be launched in the spring of 2024
  - 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 >= v23 with condor >= v23
    - To get all HTCondor CEs on fully supported versions
    - Those recent versions also support the plug-in call-out for EGI Check-in tokens
    - Mind this setting for APEL: `USE_VOMS_ATTRIBUTES = True`

IAM service developments

- Upgrade tests for v1.8.3 ran into some setbacks
  - Issues encountered at CERN that were not seen at CNAF
  - More fixes were then deemed important enough still to be included
  - Latest release candidate was tagged on Monday and tested OK on Tuesday!
  - We still intend to do the upgrades this year

- An instance for the “dteam” VO became available on Dec 7
  - The first instance on Kubernetes at CERN
    - Other instances are still on OpenShift
  - EGI broadcast with LSC file details
  - VOMS client details will also be published soon
  - Users are imported from the VOMS-Admin service until its retirement
    - Expected to happen in a few weeks!

- Other small VOs at CERN will also get their IAM instances
  - As of early 2024, to allow them to get their token clients into production ASAP
  - VOMS(-Admin) service EOL == CentOS 7 EOL, or earlier if possible

AuthZ WG items (1)

- IAM v1.8.1 security update broke ALICE & ATLAS token renewal
  - The clients turned out to have been making use of a behavior that was only allowed due to the bug that got fixed
    - Addition of scopes by the client owner, but not consented to in the registration
  - The clients had to be re-established
    - Easy to do, but with unwanted operational repercussions
  - The matter was debated, but currently moot, because enforced by the underlying third-party framework

- A lot of issues and PRs have piled up for the WLCG token profile
  - Implementors know where to deviate from what v1.0 specifies
  - An updated version is expected early next year: v1.x or even v2.0

- The CHEP paper has been written
  - Some minor comments were received, still to be addressed

AuthZ WG items (2)

- A new WG was proposed and has started in the meantime: the Grand Unified Token profile WG
  - To try to establish a common base profile for AARC / EGI Check-in, SciTokens and WLCG tokens
    - In particular to determine the VO, groups and roles of tokens
  - Each profile can then have its own extensions

- The Token Trust & Traceability WG has discussed its scope and a first “test case” was proposed: XRootD token how-to
  - Configuration
  - Logs
  - Debugging
  - Testing → the “dteam” VO should also serve token use cases!
    - IAM tokens could be used in parallel with EGI Check-in tokens for “dteam”
    - The “wlcg” VO hosted at CNAF is for developers

DC24 workshop & DOMA (BDT) meetings

- Tokens featured in several workshop presentations – in particular:
  - FTS & Tokens
  - Storage & Tokens

- The DOMA meeting of Dec 6 had the Rucio presentation
  - ATLAS and CMS token behavior in DC24 should be the same by design, as there will not be much flexibility yet in the Rucio token configuration

- Though tokens are still deemed an optional ingredient in DC24, the plan is to have them used where possible right from the start
  - When source & destination of a transfer both support tokens
  - If a given SE fails transfers due to its handling of tokens, it can be switched to VOMS proxies instead
  - We should thus gain operational experience at a big scale and discover which things need to be improved where for the required reliability etc.
    - There are **no** concerns about token rates at this time

Conclusions and outlook

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

- There are 2 new WGs to try and help simplify what developers and sysadmins are confronted with to support tokens
  - The **GUT profile** WG aims to give relevant profiles a common base
  - The **TTT** WG aims to provide recipes for dealing with token matters

- To be continued…