GridPP and IAM

Daniela Bauer

Imperial College London

GridPP
UK Computing for Particle Physics

What is Indigo-IAM

In their own words:

- “The open source, self-contained Identity And Access Management (IAM) solution for Scientific computing!” (https://indigo-iam.github.io/v/current/)
- IAM is the AAI solution chosen to power the next generation WLCG Authentication and Authorization infrastructure.
- IAM is required to be able to use tokens.
- EGI uses EGI checkin for the same purpose.

What the plan was

WLCG Token Transition Timeline v1.0 (August 2022):
M.1 Sep 2022 IAM is also in production for ALICE and LHCb.
M.2 Dec 2022 DIRAC versions supporting job submission tokens deployed for concerned VOs (LHCb, Belle-II, …). [Note: Only because the DIRAC server is capable, this does not meant there are token issuers around, nor that sites will accept tokens, even if they exist.]
M.3 Feb 2023 VOMS-Admin is switched off for one or more experiments. (Moved to June.)
M.4 Mar 2023 HTCondor installations at EGI sites have been upgraded to supported versions > 9.0.x.
M.5 Mar 2023 End of HTCondor support for GSI Auth.
M.6 Mar 2023 Some storage endpoints provide support for tokens (at least one per service type).
(in 2021 this was: Mar 2022: All storage services provide support for tokens, Oct 2022 RUCIO uses token in production)
M.7 Feb 2024 Rucio / DIRAC / 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.
M.9 Mar 2025 Grid jobs use tokens for reading and stageout.
M.10 Mar 2026 Users no longer need X509 certificates. (Note: In 2021, this was still scheduled for for March 2024!)

Where we actually are

- Tier 1 has just (Mar 2024) been asked by CMS to enable tokens on their CEs: https://ggus.eu/?mode=ticket_info&ticket_id=165969
- Most UK Tier2s accept token pilot jobs.
- LHCb DIRAC: Upgraded to token compatible version in January 2023
- GridPP DIRAC: Pre-prod instance token capable sinde June 2023, production since July 2023:
  - Pilots only
  - Using an IAM instance co-hosted with the DIRAC server for most VOs
  - One CERN based VO (moedal): CERN just provided IAM server, three months after request: moedal somewhat confused on what to do with it
- Storage: Tier 1s and dCache sites
  - As noted in the CMS talk: There's a difference between SAM tests and production transfers
- Dec 2023: WLCG Token transition update mentions multiple IAM related issues

In summary: Best laid plans meet reality

Requirements are not written in stone:

- E.g. WLCG Token Profile: “The descriptions in v1.0 of the WLCG token profile are quite misleading if not outright unimplementable.”

GridPP Community support ≠ WLCG:

- More fragmented, less technical knowledge inside the communities, fewer people available for work
- But: It works. Let’s not break it:
  - Step 1: Work out what we actually need.
  - Step 2: Communicate our intrinsic requirements to the developers.

VOMS vs IAM

- VOMS: The current setup with three geographically distributed voms servers has close to 100% uptime over a decade:
  - We are about to make it worse for our users.
  - High availability IAM not a priority for CERN as they still ironing out more fundamental issues
    - Ongoing work at RAL for IRIS high availability IAM. No fully functioning prototype yet (project started in June): Clearly non-trivial.
    - Token transition milestones require us to have something on the ground now (and we do!)

What are we doing right now ?

Gaining experience: Can’t wait for a “production ready” service:

- IAM server co-located with DIRAC server at Imperial College to cover pilot submission with tokens and DIRAC (remember that milestone ?)
  - Using groups as stand-ins for VOs.
  - Co-location will hopefully minimise any downtime (we haven’t had one yet).
  - Lots of Tier2s accept our tokens, thank you.
- IAM server at Manchester:
  - For testing voms to IAM migration
  - Local submission using tokens

What are we doing right now ?

- VOMS can be ported to EL9: Robert Frank and Simon Fayer have basically done this:
  - This gives us breathing room until the nominal X509 optional date in 2026
- Multi-VO IAM:
  - One IAM server (x3 for HA) looks quite sensible if you have one VO with hundreds of users.
  - Less so with 20 VOs with < 100 (often less than 10) users each.
  - GridPP is not the only community interested in this.
  - There is no (technical) reason this would interfere with the way the WLCG uses IAM

What are we going to do next ?

- Continue gaining experience with running IAM
- Explore possibilities for HA deployment
- Discuss our use case with the developers with the aim of making IAM better for our communities
- Continue to support client side X509 until tokens are ready to justify user facing changes:
  - The goal is that our users only face change **once**.

Thanks for listening.

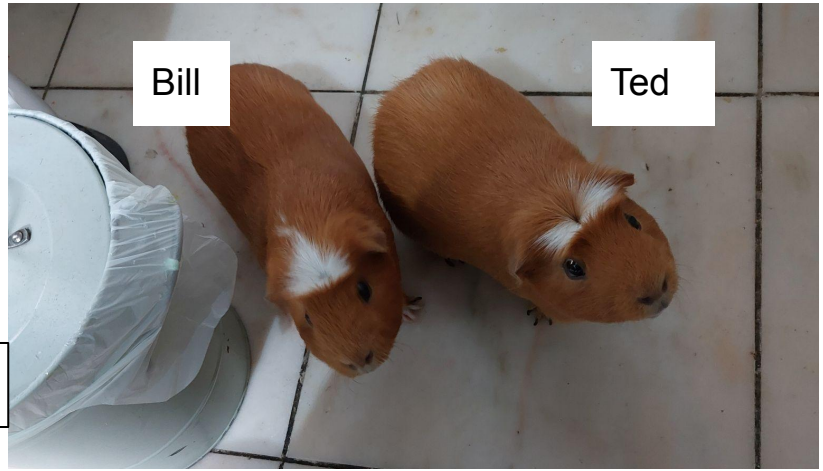

Imperial C
London

I see your cat and raise you two guinea pigs ->