IAM@IJCLab: Status and Plans

IAM Workshop 2022

Michel Jouvin, CNRS/IJCLab

michel.jouvin@ijclab.in2p3.fr

IJCLab

- Probably some of you remember me from LAL…
- IJCLab is the name of a new lab which is a merger of 5 labs, one of them being LAL
  - Just a bigger lab! Covering HEP, Nuclear Physics, Astro&Cosmo, Accelerator research, Health & Environment Physics… (750 people)
- For the computing, trying to act as a bridge between those communities
  - Particular focus: make available to non HEP communities the solutions that have proven to work at scale

IAM@IJCLab: current situation…

- 3 production instances + 1 demo/test instance
- Demo/test instance (iam-generic.ijclab.in2p3.fr) used to demonstrate IAM to partners interested before starting a dedicated instance
  - Integrated into eduGAIN
  - No production service integrated in this instance as a client
- Iam-mesonet.ijclab.in2p3.fr: authentication service for a federation of HPC centers run by universities in France
  - https://www.mesonet.fr (unfortunately in French)
  - Authentication for web services like the user portal
  - Management of SSH keys at HPC centres using the SCIM API
  - Main identity provider: eduGAIN

… IAM@IJCLab: current situation

- iam-grandma.ijclab.in2p3.fr: authentication service for the GRANDMA (multi-messengers astronomy) project
  - https://grandma.ijclab.in2p3.fr
  - Main service is https://skyportal-icare.ijclab.in2p3.fr
  - Every GRANDMA participant **must** have a Slack account: Slack (OIDC) used as the main identity provider
- sso.ijclab.in2p3.fr: authentication + authorisation service for web application at IJCLab

IAM as the IJCLab SSO…

- Primary goal: allow an easy but controlled access to some IJCLab applications for entitled external users
  - Internal users all registered in the lab Active Directory: currently the core of the SSO at IJCLab (used by all services)
  - Some services (e.g. Indico) provides also an access to external users, using eduGAIN in the best cases or some accounts local the app in the others
- Idea: use IAM as the authentication+authorisation service for the applications that must be accessible to external users
  - Create an account for every internal user registered in AD using the SCIM API (Similar to synchronisation with HR database)
  - External users: eduGAIN as the identity provider

… IAM as the IJCLab SSO

- Authorization: allow consistent decisions across applications
  - Use groups to distinguish between local and external users: import group membership from AD for local users
- Additional benefit: only IAM is registered in eduGAIN
  - Change in apps may require an eduGAIN update that requires up to 24h
- Work in progress
  - IAM integrated with Active Directory using MS ADFS service
  - A couple of test applications integrated, like the (shared) datacentre management (openDCIM) and Indico test instance, with external users accessing it
  - Local user import not yet started: not completely clear how to create the link of an IAM user account with an OIDC/AD account

Deployment / Operation

- Currently all instances hosted by the same (CentOS Stream 8) VM
  - Current instances not heavily used: easy to split them on several VM if need be
  - VM services: Nginx, MariaDB, Podman (at IJCLab, VM deployed with Quattor)
- IAM deployed using the container distribution
  - Simplify a lot the deployment and the maintenance
  - 1.8.0 upgrade recently: < 5mn per instance including the small configuration update, easy/quick to revert in case of problems (as long as DB remains compatible)
  - systemd script written that support multiple instances whose parameters are defined in an environment file

SCIM / oidc-agent Experience

- SCIM API used by the MesoNET project to retrieve the SSH keys registered by IAM users
  - Goal: automatic configuration of SSH keys at HPC centres for Mesonet users
  - Groups used to select the users of a particular subproject if a centre provides resources only for some subprojects
  - A Python script written as a demonstrator for retrieving this information
- The existence and the completeness of this API is clearly an Indigo IAM asset
  - Was one thing that convinced some people in the MesoNET project!
- Interested by sharing experience and possibly some tools
  - We created a GitLab repository for this…

The eduGAIN problem

- Not strictly related to IAM but may be worth to mention…
- Using IAM with eduGAIN as the identity provider has the potential to open an IAM instance to the world in a controlled way!
- Unfortunately eduGAIN promise is not really delivered
  - IMO eduGAIN promise was to build a trusted federation of IdPs and SPs that would prevent the need for bilateral agreements
  - To increase the trust, Research & Scholarship Entity Category was added to tag the really trustable/academic sites
  - Unfortunately a large number of university IdPs are not comfortable with the concept of trust and they forbid the access to their IdPs if you don't contact them, explain the need and request the access…
  - It is almost a general problem with German universities, also with ½ of the French ones… and many others

Wishes - Feature requests…

- It should be possible to require Affiliation when a user create an account
  - Important to identify if a user is legitimate to request an account
  - Should be a visible attribute in the profile
  - https://github.com/indigo-iam/iam/issues/480
- Keep a fluid navigation when login with user/pwd is disabled
  - We typically disable the ability for a normal user to create/use a user/pwd account but still want to be able to use a user/pwd login for some specific accounts, e.g. admin
  - Works in 1.7/1.8 by adding '?sll=y' to the login URL but this means that you loose the redirect URL passed as a query string when login is triggered by another request/url (e.g. account validation)
  - Would be great to be able to have a link at the end of the login page allowing to switch to user/pwd login (when disabled) that will keep the redirect if any

… Wishes - Feature requests

- High availability – multisite resilience: a requirement for large-scale projects
  - At the same time, can be very complex: be pragmatic and try to address first the most likely use case: temporary unavailability of a server
  - Optimal reliability would be provided by the ability to run the service at 2 sites at least: would allow to be resilient to a site disaster (never say it will not happen…)
  - Most difficult part is the database resilience: SQL clusters are not easy and difficult if there is a WAN connection between two of the servers. MariaDB is also probably less appropriate than PostgresQL for this.
  - Look at caching strategies in the various instances that would allow to relax the need for a highly available central database?

EURO-LABS: a multi-tenant IAM?

- EURO-LABS: a European project to “glue together” the nuclear physics, HEP and accelerator communities around future R&Ds
  - https://web.infn.it/EURO-LABS/: Kick-off last week in Bologna
- WP5.2: Open Nuclear Physics
  - Federated data in nuclear physics to enable open science: extend to nuclear physics community the ESCAPE concepts
  - 1 key deliverable: a central/federated authentication and authorisation service
- IJCLab is part of WP5.2 for this AAI work: the plan is to deploy Indigo IAM for the EURO-LABS project
  - Act as a demonstrator for NP community of what would be possible at a larger scale
  - EURO-LABS is made of several (disconnected) subprojects and NP is made of many projects: are groups the appropriate concept? Do we need to think at federating several IAM instances? What impact on authorization?

Miscellaneous

- A GitLab repository created to share IAM-related developments
  - https://gitlab.in2p3.fr/indigo-iam/iam-tools: accessible through eduGAIN
  - Currently only the SCIM test script but plan to put more as soon developments progress
- Would some sort of forum be useful
  - Slack is not very efficient for discussing matters over a long period
  - iam-support email is good for reporting/discussing a problem with developers before may be opening an issue
  - Currently difficult to discuss an idea not only with developers but with the “experts” from the wider community
  - GitHub forum?