FIM4R Summary

Overview

- Agenda: https://indico.cern.ch/event/605369/
- Attendees (52): https://eventr.geant.org/events/2580

COMMUNITY UPDATES

LIGO

Highlights

- COmanage registry at gw-astronomy.org, hosted and operated by University of Wisconsin Milwaulkee to allow broader purpose
- SPs individually registered in incommon
- **IoLR= Google, UnitedID, NCSA**

Future plans

- Move to CIlogon2 (marriage of CILogon and Comanage) – **outsource ID layer**
- Could maybe move to OIDC instead of SAML, however the cost is integrating with federations.
- Wanted to encourage attribute release and avoid use of a proxy, however this doesn't seem to work, may have to **move to proxy model**
- Hoping to move entirely to FIM, **remove LIGO IdP**

Challenges

- Budget constraints, pushed to work on visual aspects, e.g. GraceDB
- Sirtfi adoption stalled by incommon's requirement for C level approval
- Some eduGAIN partners not totally "in" eduGAIN e.g., Australia, Japan
- No role for **research communities in governance of federations**, perhaps the solution is to create an **IGTF federation**

ELIXIR

- **Proxy IdP**, SAML2 plus support for OIDC
- **ORCID** as an IdP, plus **social** options – **researcher LoA enriched separately**
- If the chosen IdP does not provide attribute bundle, helpful message is displayed – **user passed to local support group**, e.g. ELIXIR Germany, who will follow up with the IdP
- Group Management
  - Perun
  - User driven with custom application forms per group
  - Bona Fide management on top to grant additional access, e.g. check ORCIDID against publication, users can endorse other users
- Example. Beacon Network (query for DNA data sets), requires Bona Fide Research Status
- **VMs** created by users **cannot be trusted**, only allow mounting of data that was approved by committee

WLCG

- Existing certificate based federation
- New solution follows **proxy model** for authentication, all WLCG services behind CERN SSO.
- **Token translation on per service basis**, not classical blueprint architecture
- Some progress over last year, including one experiment moving monitoring portal behind SSO
- Difficulty is getting users to adopt new technologies when existing solution "works", albeit in a clunky fashion
- Requirements for FIM
  - Helpdesk essential
  - Sirtfi required now, restrict to known researchers registered with VOMS
  - Command line solution with minimal browser interaction
- Implementation uses STS, not maintained and suboptimal
- Reconsidering the **role of VOMS**, options inc. AA, token translator, etc

DARIAH

- 3774 users, identified by EPPN (possible weak point)
- Secondary authentication track that gets **OAuth2 authorisation** token
  - OAuth2 chosen for internal authorization rather than ECP, following a trial in which a number of problems emerged. OAuth2 much more simple (plus future-proof)
- Central Policy Decision Point, with **access rights centrally managed**

INAF

- Distributed communities, role based access, project for several decades so want **simplicity & sustainability**, use OTS components
- Fundamental constraint = open to all astronomy community (achieved by enabling eduGAIN)
- Member of AARC2
- CTA
  - **Enriching attributes** themselves, since IdPs insufficient. Using grouper for membership management
  - Internally adding **isMemberOf** attribute list & **entitlement** for access control
  - **eduPersonUniqueID** chosen
- Consent management
- 3 main experiments but having separate solutions for each experiment seemed simpler than creating single solution

Umbrella

- Used by photon and neutron facilities in Europe (14 partners + 2 pending) – all basically production status
- Integrating **ORCID** & pushing **umbrellaID** IdP in eduGAIN
  - Only 3 attributes -> no problem for data protection since it is all opaque
  - will join JISC instead of SWITCH due to registration requirements
- Member of AARC2
- Using **moonshot** at Diamond
- Using **eduTeams** for AA since users spread between multiple jurisdictions
- PR push, funded
- **Just IdP**, no SP – project called **eduGAIN bridge** to provide eduGAIN access to umbrella registered services

Some common themes?

- Proxy model
- Attribute enrichment, and per-attribute LoA as a consequence
- Preference for outsourcing and Off-The-Shelf components
- ORCID & Social Login
- Higher influence over (inter)federation governance

INFRASTRUCTURES

EGI

- **Diversity** of VOs raises complications
- Number of services & IdPs requires significant, **scalable policy work**
- Checkin, solution deployed in EGI in 2016
  - **Multiple IdP types** through single endpoint (inc. social & x509)
  - **Minimise overhead** for service providers
  - Not all services behind proxy but moving slowly
  - Central, unique, opaque, **persistent user ID created on first login**. Unique ID can be freely shared since opaque but remains useful for central logs
  - Previously had LoA Birch (IGTF), now LoA calculated based on user information
  - Stepped LoA requirements for different risk profiles, inc Sirtfi for PaaS
  - Checkin governs list of **trusted Attribute Authorities**, those trusted are harmonised & communicated with services
  - Unity connector to get LToS VO membership information
- Checkin integrated with RCAuth to provide x509
  - Users from trusted IdPs able to generate certificates
- Explicit account linking via COmanage

EUDAT

- Central data centres across Europe - staging storage, sharing, etc – common layer is authentication
- **Multiple identity sources** (SAML, certificates, eduGAIN, CLARIN, ORCID (TBC)), pass through B2Access layer to internal services
  - Data enrichment performed if required
  - 6 attributes (from IdP or User)
- Investigating **LoA per attribute** since some come from users, external sources etc
- Challenges
  - Per federation **eduGAIN opt-in** policy is proving confusing for end users when IdP missing
  - **Attribute release**
  - Non-standard services, **desktop clients etc**
- Trialing **attribute based access** control due to flexibility but conscious of single point of failure

ONGOING PROJECTS

Security incident response (Sirtfi)

- Problems with security in federations
  - Highly distributed, e.g. logs are split
  - Bad guy doesn't sleep but IdP operators do
  - No mandate to investigate external organisations
- Sirtfi REFEFDS WG, ~2 years done, ~2 years left
- Workplan includes
  - Helping federations to adopt **procedures**
  - **Testing** Sirtfi process
  - Reaching out to **communities** e.g. TF-CSIRT, REFEDS, FOG
  - Targetting **SPs**, or highlighting how many Services behind one proxy
- Part of **Snctfi**

Hoping for slot in unconference to discuss 2017 Workplan – come along!

Snctfi

- Mechanism for building **scalable trust** for elements behind a proxy. Binds all participants in infrastructure (**proxy + innards**) together
  - Q: why should federations trust SP proxy?
  - A: because it **asserts R&S, Sirtfi, DP CoCo**
- Trust flows against the current of attribute flow
- Snctfi allows the R&S etc assertions to happen by setting **requirements on the infrastructure**
- No visible **assurance mark** in metadata but more a use for an infrastructure to cover its back -> feedback that a mark may be necessary
- Allows for different methods of internal infrastructure binding e.g. contracts, MOUs, policies
- SP Proxy's compliance could be peer assessed via IGTF?

Policies for Processing Personal Data

- Not legal advice but has been read by lawyers ☺ part of Snctfi
- Does not cover attribute release or personal data in research sets. **Scope is restricted to data collected on usage**
- New GDPR goes into force May 2018 – legally binding for member states
- Most research communities are **data controllers**, rather than processors, so must **define policies**
- Cloud Computing – gets complicated with multiple layers and jurisdictions
- Legitimate interest can support 3rd party sharing, but careful **balancing act** required
- BCRs are recommended framework to bind an organisation, though only applicable to legal organisation (many infrastructures are not)
  - Suggestion to create **a BCR-like policy**, which should prove sufficient
- Conclusions
  - In **EU legitimate interest & consent ok**
  - Outside EU, **BCR-like approach** might work. An **enforceable CoCo** might be alternative to getting specific authorisation
- Alternatives suggested
  - Could create legal entity for the community, with paid membership
  - Buy insurance

Data Protection CoCo

- Released 2013
  - 106 SPs support
  - 112 IdPs claim to release attributes to them
- Asked **WP29** for blessing. Results:
  - We can use it ☺
  - It cannot be endorsed by WP29 since doesn't provide **added value** (e.g. explain data minimisation in context of FIM) ☹
- V2 addresses WP29 requirements, GDPR changes, release outside EU (inc. international organisations)
  - **Longer** (4 -> 40 pages!)
  - 2 month consultation starting Wednesday at TIIME
  - Keep submitting to WP 29 for feedback
  - Aim to **submit for approval** in May 2018

AARC I & II

- Aims to build on **existing tools**, **avoid fragmentation** and bring FIM to Research Collaborations
- Many pilots produced, to show that the technology works, and then work to make them sustainable
  - E.g. CiLogon-like pilot, hide PKIX from users
  - Addresses non-web use cases & integrates policy elements
- Looking at many policy aspects and their interaction with existing groups
- AARC2
  - Support more research community use cases
  - Deploy results
  - Delivery platform includes community engagement -> continuously talk with research communities, help and identify new requirements
  - **Competence centre** for large r/e-infrastructures to co-develop new solutions
- **FIM4R** is a key community
- In addition, **create a forum** for infrastructures to exchange information (AAI, security, policy, pilots)

AARC Blueprint Architecture

- Not trying to address generic use cases, but the **specific difficulties** that RCs have when operating internationally
- Blueprint architecture 1 was over simplified, AARCBA 2 adds realism
  - **Authorisation** layer added, further work expected on scalable methods for this
  - Focus on **pragmatic guidelines** for e.g. non-web access, token translation, authorization etc
- Non-web guidelines, realistic, production-ready suggestions that don't require 10 years of deployment history
- Many token translation combinations explained, including mapping

Guidelines open for comment until end of February
https://aarc-project.eu/architecture-guidelines-recommendations-for-comments/

Publishers & FIM

- At last STM meeting, information on RA21 (meeting of many key publishers) determined
  - **IP address authorisation** is bad
  - Need to make the experience easy and consistent for user
  - Battling open data platforms (legal or otherwise)
- **Survey for campuses** – please share is you can
- **Pilots** planned
- They know about REFEDS discovery, unclear how much they know about other schemes
- Hopefully more interaction with this community in future
- In AARC discovered that this is not really a technical problem for libraries, difficulty is that it is not in contract
- Publishers now seem more welcome to this topic

Kantara OTTO WG

- Open Trust Taxonomy for Federation Operators (OTTO)
- WG within Kantara that came to be when having a **holistic look at federations post OIDC**
  - 1/3 members from R&E
  - Clarifying assumptions made when defining the federation standard for OAuth2
- Defines actor roles in federation
- JSON-LD (JSON linked data) model is convenient for describing federation relationships
- 4 APIs that OTTO intends to develop
- Need to address **registration, revocation, metadata distribution, key management burden**
- Policies to be defined according to GTRI schema

FIM4R V 2

2012 Requirements

- **User friendliness (high)**
- **Browser & non-browser federated access (high).**
- Bridging communities (medium).
- Multiple technologies with translators including dynamic issue of credentials (medium).
- **Implementations based on open standards and sustainable with compatible licenses (high).**
- **Different Levels of Assurance with provenance (high).**
- **Authorisation under community and/or facility control (high).**
- Well defined semantically harmonised attributes(medium).
- Flexible and scalable IdP attribute release policy(medium).
- **Attributes must be able to cross national borders(high).**
- Attribute aggregation for authorisation(medium).
- Privacy and data protection to be addressed with community-wide individual identities(medium)

2012 Recommendations

- Recommendations to the research communities
  - Conduct Risk Analysis
  - Run Pilot Studies coordinated by experts
- Recommendations to the technology providers
  - Separation of Authorization and Authentication
  - Credentials revocation
  - Attribute delegation to the research community
  - Standardise efforts in Levels of Security/Assurance
- Recommendations to funding agencies
  - Fund FIM technologies that are focused on solving the described needs of the research communities

Progress Discussion

- Significant progress made
  - AARC I & II
  - We are here having rational discussions between RCs, Fed Ops, eduGAIN etc!
  - Many successes
- Some requirements remain, for others we have found work-arounds, some are new

Recommendations 2017

- Address command line and non-web use cases
- Integrate FIM with existing Community Membership Management Tools (AAs)
- Build operational security
- Support GDPR adequacy certification for intergovernmental organisations
- Commercial IaaS interaction
- Make FIM a production service and a corner stone of the European Open Science Cloud
- **eID**
- **Proxy model standardisation**
- **Analyse existing, available components in FIM marketplace**
- **Greater collaboration with non-EU partners (e.g. US)**

NEXT STEPS

FIM4R Document Plans – Proposal

Editorial board

- Rep from each community/infrastructure (you are probably here!)
- Define survey Qs
- Write summary of progress since FIM4Rv1
- Combine contributions from communities/infrastructures

Community/Infrastructure input

- Statement on own progress and challenges
- Complete survey

FIM4R Document Plans - Proposal

- Output?
  - Whitepaper
  - Include targeted recommendations to players, e.g.
    - Funding agencies
    - SPs
    - IdPs
    - Federation Operators
    - …
- When?
  - Don’t want to be too slow and need to be aware of calls for funding
- Where?
  - Previous published by CERN (and others?) could repeat
  - Proceedings TNC 18?

Coming up

- Run an 11th FIM4R Workshop to start (Chicago offered)
- When else to meet?
  - Feedback @ RDA Session, April 6, Barcelona https://www.rd-alliance.org/plenaries/rda-ninth-plenary-meeting-barcelona/rda-9th-plenary-registration
  - Where else?
- Who? Representatives of Research Communities and Infrastructures
- Timeline?
- Website – Scott owns domain. Agreed that could be hosted at CERN

Thanks

- Rainer for the organisation & dinner logistics!
- All speakers & attendees