The Abingdon 59th EUGridPMA+ Meeting Summary

The 59th EUGridPMA+ meetings, in conjunction with GEANT's GN5-1 Enabling Communities activity, EGI Security, EOSC ISM, and the AARC Community Policy Area takes place on October 3-4 in Abingdon, kindly hosted by STFC RAL and organised by Tom Dack and Dave Kelsey. Thanks very much for their good care and efforts in making this meeting memorable!

In this summary, we will - jointly and collaboratively! - try to give an impression of the main discussions and results. As usual, much is also contained in the slides and ancillary materials that are attached to the agenda pages at https://eugridpma.org/agenda/59 (https://eugridpma.org/agenda/59) and linked therefrom.

These notes have been created collaboratively by those present - thanks for this work and for sharing your thoughts, questions, and answers for the benefit of those remote and for posterity.

> For the next meeting, we have (provisionally) agreed to have the 60th EUGridPMA +meeting in Copenhagen on the **Monday January 29th**, adjacent to the FIM4R and TIIME Unconference workshop for a full day.

We look forward to meeting you again shortly.

Kind regards,
DavidG.

Introduction and note takers

09.35
Present locally: Eisaku, Jens, Marcus, Liam, DavidG, DaveK, CasperD, Manikantan, Ian Collier.
Remote: AndersW, Adeel, Miroslav, Lidija, JanC, Baptiste

Developments in the Asia Pacific and the APGridPMA

09.49

APGridPMAupdate

Current chair and vice-chair are Eric Yen and Eisaku Sakane.
Next APGridPMA meeting at ISGC (planned March 24-29, 2024 in Taipei), and the 34th APGridPMA with APAN58 (August 2024) in Pakistan

- eMudhra CAs approved in APGridPMA meetings in July and August 2023. These are now included in the distribution
- HPCI is moving to a new (OAuth based, KeyCloak, OIDC Agent) based AAI model:
  - Considering cooperation with GakuNin.
  - There is also JWT-Agent in HPCI and now being deployed to the community as a token-handling software for users.
- ASGCCA and KEK are similarly deploying token-based AAI for the benefit of broader user communities (and also granting access to eduroam)
- The Malaysian Access Federation (SIFULAN) offers IdP-as-a-Service for country-wide user communities

GakuNin update

GakuNin (the Japanese Access Management Federation) has 296 IdPs, ~2M identities (with the majority being students), and 207 serice providers. The growth necessitates a new trust framework in Japan for academia. This is especially viable for those communities that rely on features not currenly present in the GakuNin IdPs - since these IdPs don't satisfy the community requirements. The same drive that pushed home organisation accounts elsewhere - the need for users to have just a single account - is very natural also as a push in GakuNin. Therefore, a new working group in GakuNin started in FY2021. The WG is also mindful of the international cooperation requirements. New components:

- GakuNin IAL/AAL Assurance definition)
  For assurance, GakuNin is targeting policies at IAL/AAL level 2, with draft policies (both in Japanese and in English) for polic and practice staements are available (https://meatwiki.nii.ac.jp/cnfluence/x/JoSfBQ (https://meatwiki.nii.ac.jp/cnfluence/x/JoSfBQ)). This is now under review by all stakeholders (RIKEN, NIMS, RCOS, HPCI) and checking interop with REFEDS and IGTF framewoks.
- Authenticator registry
  Evaluate authenticators basedon the GakuNin AAL, with in FY2022 formulating the criteria, and in FY2023 trial operations (MS Authenticator, RFC TOTP through Google Authenticator, PKIX client cert, FIDO2 credentials).
  Collaboration with FIDO2-Alliance in Japan
- Orthros proxy (assurance matching, bridging, attribute enrichtment)

implements also IAL/AAL managment and attribute assurance. In 2023 migrated away from 'OpenIdP' and moved to external IdP linkage and new home org bindings (including ORCID and mobile identity providers). In FY2023 there will also be enhancements in authorization attribute handling.
- IdP hosting service
  Since all orgs are in GakuNin, in FY23 this would be a potential next step. Planned in FY23 and FY24
- Ambitious timeline for Evolution already defined (also for group management and attribute schemas)

This basically comprises the complete AARC BPA stack :)

Questions and Answers

- SKA: none of the GakuNin IdP provide SirtFI, RAF: Can this be changed?
  - for the new trust framework, adding SIrtfi capabilities (also in the federation software, so that the university IdPs can assert Sirtfi, but also the federation can propagate it to eduGAIN. At the same time, pushing for expression of the assurance framework in terms of REFEDS Assurance Framework (REFEDS RAF v2). The capability is likely already there, but it is just not expressed?)
  - With GakuNin being a strong federation (and much better controlled that, e.g. the AMF Jens referred to) it could be relatively easy to take this step. "A shining example"
  - Best practice example from SWAMID: Shared their updated federation policy with the eduGAIN Steering Group for consultation. Usually Federations update their policies without informing other Federations in eduGAIN. Makes it difficult to assess whether Policies remain in line with Common Standards. Staff changes may not be reflected at all times.
  - Should eduGAIN also have a peer-review self-assessment process? Annual audit? But do call it *Assessment*, rather than Audit, to keep esp. the US side calm…
- Assurance ("the research communities require assurance"): do they require more than social, or does it vary betrween communities?
  - all communities in Japan seem to align on Kantara/NIST SP 800-63v2 LoA2
- Attribute Schema
  - currently consider commonly requested attributes by the commities, but that are typically scoped to the community.
  - Maybe look at the attribute profile that is being aligned in the AppInt/AEGIS group. (AARC, WLCG, or SciTokens).
  - Attributes in GakuNin are SAML for the moment, but the profile/attributes are technology agnostic
  - Grand Unified Token profile to be disucssed at TIIME in January 2024 in a joint session with Mischa Sallé. Agree amongst the technicians first, to avoid the politics.
  - Tokens might be scoped to projects within a defind set of relying parties/SPs.
- 'Complex Group Management' - what is intended here?
  - authorization of end-users for composite/complex conditions calculated based on groups? On-demand dynamic group created for users.
  - Attributes provided to SPs that can be used for authorization directly? Currently there is

MAP Core (“Membership Attribute Provider” system)

- Is GakuNin a H&S federation? No … Orthros is there for the communities/RPs
    - Orthros can bind the user with ORCID and associate with a local ID in Orthros, and Orthros can bridge between credential providers (account linking).
    - GakuNin is a full-mesh

CA Update: TCS Personal Authentication and upgrade experiences

11.00

The TCS service (started as the SCS Server Certificate Service in ~2005) provides a range of public trust certificates for Server SSL, client SMIME, and code signing (OV and EV CS, e.g. for Windows kernel drivers). Recently, triggered by changes in the public S/MIME trust model initatiated by CABF, also private trust CAs specifically for client *authentication* purposes have been added. These include an (RSA and ECC) root for Research and Education, and a (for now single) subordinate issuing CA for GEANT TCS Authentication CA “4B”.

- Driven by Geant members: 45 NRENs
- eScience use-cases have been extended now - following the SMIME BR updates - to apply more broadly to authentication, and are increasingly becoming an important factor in the TCS portfolio.
- The trust structure is modelled through contracts, and flows via NRENS to GEANT.
- IGTF Classic <=> CABF OV are pretty much matched, and the recent updated to CABF SSL BR are still compatible with the IGTF namespaces.
- User and personal Robot certificates are supported by the underlying SAML authentication, leveraging eduGAIN and Seamless Access. The assurance is expressed as an explicit eduPersonEntitlement, which signals formal acceptance by the subscriber of the TCS CPS (and the issuing CA CP and CPS).

As of August 28th 2023, CABF Baseline Reqiurements for SMIME kicked in:

- Different profiels an validations available:
    - strict: 2yr, s/mime
    - multi-purpos: 2yr, …
    - legacy
- validation profiles
    - Sponsor validated
    - Mailbox validated
    - Individual validated
- Bottomline for personal email (SMIME) certificates: **s/mime certificates based the commonName on names are, and were never, unique!!!** and they must now be used for authentication purposes. The IGTF MICS *was* unique (since it includes ePPN in addition), which is why also the new “authentication” specific private trust client certs are OK (since they follow the iGTF MICS formatting and uniqueness requirements).

- What TCS made of it
    - s/mime personal certificates: OV
        - Sponsor-validated BR compliant
        - defined all TCS members as Enterprise RAs
        - require all orgs to use Gov-Info-Soure or LEI (3.2.3.2.1)
        - /clientgeant SAML endpoint, since it uses signed SAML statements, can upgrade individuals to "high" validation status based on SAML statement - and hence allow for sponsor-validated status for SMIME.
    - Move client auth to trust a private CA retaining DNs, but a differenct ICA issuerDN and root
- New names for user communication: Renamed IGTF proviles to "user understandable names"
- Practicals
    - OV Certificates available via Webportal, adn are not affected (if you do the LEI/GovSource updates in the right way!)
    - No hiccups on the IGTF side
    - Reevaluate all organisations based on govt. information source. Process 'rather slow' for now, but 'should' be fixable soon

Other CABF things of note

- Server certifiates will live only 90 days from end 2024 on. => ACME-OV (+client_id +client_secret) to fix this.
    - TCS offers endpoints for this
- Automate it NOW!
- `certbot` example commandline is in the slides of David:

```
[root@hekel ~]# certbot certonly \
--standalone --non-interactive --agree-tos --email davidg@nikhef.nl \
--server https://acme.sectigo.com/v2/GEANTOV \
--eab-kid DUniqueID_forthisclient --eab-hmac-key mv_v3ryl0n9s3cr3tK3y \
--domain hekel.nikhef.nl --cert-name OVGEANTcert
```

- Get client_id + client_secret from you CA :-O

CA Update II: eMudhra

11.42

- Private IGTF Roots are now in the distribution for eMudhra - both public trust and private trust since en dof August 2023.
- Service works without hickups and has a single consistent (REST & Web) interface.
- Looking forward to more users/subscribers!

SHA1 roots operational status

11.45

We know SHA-1 is no longer secure, but some projects and distros are deprecating SHA-1 (needlessly) also for self-signed certificates (where the trust flows from explciiexplicit installation, not from signature path validation)

- This affects (at least Rocky9, Alma9, and other RHEL9-following systems, ...)
- Workaround: `update-crypto-policies --set DEFAULT:SHA1` , but it is kind of a blunt tool to fix this.

Yes there are a lot of SHA-1 root CAs in EL9: these rely on non-standard bytes being added to the DER certificate blobs at the end, to express 'trust flags'. But since this is a proprietary format, it does not port between OpenSSL, BouncyCastle, CANL-Java, and other client libraries.

OSG created a "dual blob" mode and deployed it widely, but that broke:

- CANL-Java (BouncyCastle)
- dCache was affected, affecting many sites
- Upstream CANL bug opened by Paul Millar, will not be fixable overnight
- If the fix does not work across all systems and tools, it may lead to different certificates for Java and others (or use upstream Workaround)

For the CAs that can (even though it will not be a complete fix):

- Re-issue also your self-signed root CA key-hash with a recent SHA-2, IGTF distribution can then be updated
- Older public trust CA can not be fixed that way

New package signing key and packaging

IGTF packages will update to a new key (v4)

- Package ships new key 4
- Two variants of packages are available: Signed with v3 and v4
- Private key needs to be shared with Anders, as a backup person to care about IGTF packages
- The current repository should be migrated to drive the transparency of the distribution process
  - github.com/igtf (http://github.com/igtf) is is probably the best choice
  - (No disagreement among all participants)
- distribution managers will be invites to github.com/igtf/ (http://github.com/igtf/) where in the (near) future repositires will be created
- the Github location may also contain projects to build derived distribution for major relying parties

Token Trust and Traceability WG, WLCG AuthZ working group

14.15

New Trust & Tracability Working group for wLCG (TTT) in AUgust 2023, similar to the previous Tracability/Isolation WG. Working alongside the AuthZ WG and under the Security Group aegis in Indico, and is chaired by Matt Doidge.
The WG is targeting to work across more communities, including DUNE, EGI, SKA, and others, and produce both policy (best practice) and documentation.

There is scope to extend beyond just this group, and is worthwhile to consider also IGTF/PMA input for the trust in those tokens (and the issuing proxies).
Plenty of opportunities:

- AARC G071 is already there, and concerns the trust of the issuance side (and some of the content) of tokens
- AARC-TREE has scope for policy work in this area under WP2, and we should import the results/discussions there as well for the benefit of the other AARC communities

Meanwhile the WLCG infrastructure itself is migrating to tokens, although the milestone are slipping (and the infratructure thus keeps working) but work is ongoing on the migration. VOMS is still there, as more Indigo-IAM is evolving and with new develops shouldbe ready in the near future. As of now, VOMS is still being used (for data movement, predominantly)

The token-enabling upgrade for those EGI sites using HTCondor to upgrade to >= 9.0.x. Any necessary ARC-CE upgrade implementation should now be complete and done.

In wLCG, Atlas aims to have the major sites ready for tokens for their “Data Challenge 24” (DC24, March 2024) storage end-points. Rucio, DIRAC, and FTS have sufficient token support in released versions to support tokens for DC24.
Meanwhile, there are (unrecorded) discussions around storage access performance that prefer pre-signed URLs, coming from the LHCb DIRAc use. This is not widely disseminated and quite unknown outside.

End-user need for certificates to access WLCG should go away by March 2026.

Grant Unified Token profile

Within the wLCG use cases there are concens around the number of groups a user may be a member of, and since (under the AARC profile) the group names are long, the size of the assertions might be worrisome.
There are contentious issues aound introspection issues (due to the load on the endpoints), when used in the native flows. Whereas it may be needed for interoperability, so when proxying/interopping infrastructures. Mischa will take this forward, and set up a mailing list in a neutral and infra-agnostic manner with a strong community engagement rather than a single infra pushin for it.
Should be brought to FIM4R/TIIME as well.

Q&A

- till now GridFTP was used to transfer data. What has wLCG used to replace GridFTP transdfers with tokens?
    - most of the testing has been using FTS using WebDAV/http transfer. There is a third-party copy option via DAVX for this mode. The orchestration is usually done via Rucio.
    - SKA Authorization has worked on …
        - the SKA project is currently more in a development stage, with mostly UML diagrams (but not implementations) as well as an active Wiki.
        - open to a joint token profile
        - SKA pushed authorization attributes out to the proxies operated by the SRCned federations, as proposed by AAI working group
- for group naming the scope (community) has to be known. In wLCG, the scope is implicit from the issuer, how is that going to be clarified?
    - this needs to be solved in the GUT profile
    - the AARC token profile has that way in the URN structure, which makes tokens unnecessarily large

- Some implementations have a problem with too large http headers (whicis why wLCG is reluctant to use the AARC profile)
- But many proxies have a single issuer for all their communities hosted - actually almost all other proxies do that (but not wLCG, who defined the model some years go, but does have many groups :)
- This must be solved in the GUT during AARC-TREE

Updates from the Americas and from ACCESS-CI

15.00

Derek presents the status of TAGPMA (mostly a rerun of TechEx23, though). The monthly meetings are open to all of IGTF, and the 17 members ought to be joining them consistently. Some of the former XSEDE CAs have been retired, and some are retiring in the future. The NCSA-operated CILogon Cas will be retired in 2025. Meanwhile InCommon upgrade to "*InCommon RSA IGTF Server CA 3*".
The old IGTF InCommon CA was expiring by the end of December '23, and a roll-over was warranted to get a new trust chain as well. The new CA 3 is included as a new CA in the June 2023 IGTF distribution.

InCommon has also convened a stakeholder group of representatives to advise updates in the InCommon certificate service, and a survey is out to provide input and an expert group to advise on the future directions. Derek is part of that group. Importantly, also some non-edu orgs in the US (like the DoE national labs) can now participate.
Automation (with ACME or APIs) is a high-interest topic (given that we are moving to 90-day certs somwehre in likely 2024).

CILogon retirement:

- turn off the end-user facing /getcert endpoint by January 2024
- "create password-protected cert" by June 2024

GridCanada is doing a whole-sale replacement of their issuance infrastructure, and thus needs to go a refreshed audit with TAGPMA.

Google CA application to TAGPMA

Google in Pittsburg has been prodded by the PSC about Google cloud services in the context of the Google Cloud services in LHC's Atlas experiment. They would have liked joint-trust, but the Google CA services people were not initally responsive. But since Atlas insisted, they came back to info@igtf.net (mailto:info@igtf.net) and Derek/TAGPMA are now explaining the conditions and membership mechanisms.

The current Google CA idea is to just use their DCV-only non-namespaced certs (issued through ACME with HTTP-01 and DNS-01 validaiton only).

Now if all infra is in Google (even if used by Atlas, and thus used from the outside as storage endpoints). But to be integrated in the IGTF RP trust fabric, the CA subject *must* be namespaced with a unique-elements. The assurance needs to be OV (since they *can* be used as clients).

There is one other 'solution', which is to ensure that the host certs **never** have tlsWebClient eKU set. Never. This would prevent their use as client certs to the rest, and lower the requirements for OV org validation.
This spctacularl fails with LE, since it's *not* LetsAuthenticate. This restriction on *all* certs issued off that (Google) CA with always `tlsWebClient not in eKU` would also mitigte part of the issues, and

permit RPs to install this CA next to their other trusted CAs used for authentication in the same trust store. But that required ultimate trust in the policies of the (new) ICA(s).

And this does not address the generic trust issues that the wLCG Trust Evolution WG is looking at ("who said this was Atlas storage?!"). This is probably outside the authentication scope, but a worthwhile discussion for the WLCG trust evolution WG.

Trust evolution

Owen was born on Oct 2nd!

Self-audit review & status of suspended authorities

Cosmin Nistor periodically pokes the reviewers, and the on-line meeting was intended to speed up the process. Even if the CA manager is well-known.

- NorduGrid CA : the peer review meeting (Jens, DaveK) with Anders need to find a timeslot but there was no time to meet. Will be re-scheduled, and Anders should do a self-assessment before the meeting. No guidance is needed, but does need to complete the assessment form.
- SlovakGrid CA : also here the meeting still needs to happen. Jan and KJens to follow up.
- KENET CA : they have accidentally stopped issuing the CRL for the KENET ROOT CA. But since 22.05.2023 - Ronald Osure has not been response on this issue. Changes in the (governmental) IT governance have been changing. DaveK should organise a meeting and prod Ronald (just from a different mail domain than DavidG)

The others will be prodded by Cosmin Nistor.

From the APGridPMA:

- CNIC (e2023333) and SDG (77637f58) did not complete their self-audit review. Moveover, they have been disfunctional since Feb 10, 2023 due to a lack of up-to-dae CRL. They may be now suspended in the distribution.

Next meetings

For the 60th EUGridPMA, try to get the Monday 29th before TIIME, concurrently with the midpoint session, in Copenhagen. We might want to ask AndersW for a room if we cannot get a room at the KU location. Then people can and should stay on for FIM4R (Tuesday) and the Unconference.

29 January - 60th EUGridPMA+ combined meeting (PKIX & EnCo day)
30 January - FIM4R and TIIME (Copenhagen) (https://tiime-unconference.eu/)
31 - 1 Feb - TIIME Unconference

For the May meeting, we should strongly consider co-hosting with the AARC-TREE project meetings.
For the September '24 meetings, maybe prefer CERN again?

Day 2, Thursday Oct 5th

Enabling Communities: third steps in GN5

09.05

GEANT 5-1 has now started, and will run (for 2 years) in parallel alo to the upcoming AARC-TREE project. And there are EnCo activities that are not AARC-TREE, sinc they are ongoing activities, such as FIM4R and AEGIS. AARC-TREE is more like an 'accelerator' for the work in EnCo as well. So now is the time to add new work items or course changes to EnCo, since the next few months (until the end of the year) is the time to do it, since now GN5-2 is being defined. Organisationally, EnCo and AARC-TREE are independent, and they are formally entirely separate. The idea is that there is an extra push for the work, so there should also be more people working on these topics, adding through AARC-TREE. The work in AARC-TREE is well-defined and (it being a lump sum project) will - at least - deliver specifically those results. EnCo should supplement that, e.g. in relation to the work on Orthros and AEGIS, where AARC-TREE cannot provide support. For example this was incorporated via AEGIS, and EnCo is the mechanisms to provide support from the European end to collaborate, coordinate, and engage with the (in this case Orthros) documents. Since research engagement is global, this can even be expanded.

In working with communities, EnCo has support working *with* the communities, but the main implementatin effort should come *from* the communities, as they have specific targetted funding already to do that kind of work and it is any way good to ensure adoption within communities. "EnCo" is about *enabling* communities, and it is of course good to see that the output is picked up and consolidated, but specific communities should not be named in the wok programme for EnCo. It is all about coordination.
(details about the organisational structure and finances not recorded here - in general Brexit remains troublesome)

All that I can see, is just another AARC TREE

09.30

- AARC Success goes even to EC where they call for AARC compliance, yet show a series of nicely coloured diagrams that appear to be an AARC blueprint picture, but most defniitely are not :(We saw those in e.g. EOSC Procurement documents, where the BPA is re-interpreted in weird ways :)
- AARC TREE to address the gaps in policy and architecture space. Addressing the perceived excessive complexity
- We need to keep up the work on outreach, to make sure that people understand the BPA is not just a nice picture, but as the set of recommendations that come with it
- Assurance adoption at individual IdP level seems to be unrealistic (if a 100% coverage is aimed for). Government ID Systems get on the agenda
- Who needs to be reached out to? Essentially all Research Communities / EU and global initiatives, collaborations / research infrastructures and e-Infrastructures, Service/Resource providers, EOSC Ecosystem
- Evolving the architecture:
  - Evolution of the BPA, OIDC-Federation (in particular including a deployment profile to define the subset which is actually requited)
  - Harmonisation of attributes
  - SSI World: EU Digital Wallet, Verifiable credentials, decentralised storage
- Policy development
  - RI alignment and policy harmonisation: How can more proxies adhere to the community guidelines
  - Review what happens in the infrastructures, collect and harmonise (i.e. keep users from having to consent at every possible place?)

Be clear on what is in scope for AARC-TREE, and what is EnCo and long-term sustenance.

GN51 Enabling Communities

10.10 "What About the Forest"

Input for effective ideas is welcome. Maybe from the REFEDS community during the eduGAIN town hall meeting in Stockholm next week. Things that benefit from an open community process that would be engagement or small adoption ideas that would have a flywheel effect are potentially eligible for EnCo support under the GN51 partners.

Alignment between polic efforts of EnCo and the business development side (as identified in the eduGAIN Futures working group) since there is still limited adoption of new concepts in existing federations. More discussion in the informal space (bilateral pushing of new ideas) to get more adoption, or through EnCo funding get people involved in working groups in eduGAIN. With the new eduGAIN Constition approved (fully approved now, comes into effect on Jan 1st, 2024), the steering will now be more effective, and more direction setting rather that just informational meetings. The Assembly remains the current set of all members, but there is now also an executive Steering Commitee. On the latter, there are two at-large positions on that Committee from the wider community.
This will open the opportunity to initiate and chair working groups under eduGAIN, for example:

- branding strategy
- technical development
- looking 'over the fence' to import the beautiful policies and recommendations and how they can be imported in eduGAIN. This would strengthen eduGAIN as a service, but also the internal effort within the individual federations.

This will take shape early 2024, after kickstarting the new governance processes.

Does this open the possbility to use eduGAIN as a mechanism to ensure mandatory implementation of new standards (such as Sirtfi v2). For now we just have standards and documents, but there is no clear incentive to implement it, even in advanced federations. If there is no clear eduGAIN-level policy with deadline ('baseline requirement') you never get results. So this can be a mechanism to have an 'eduGAIN with Baseline' to adopt the REFEDS outcomes.

Outreach to emerging and smaller federations (through EnCo) for policy and engagement may create more of a drive for this, more than the technical outreach that is currently happening. Light-weight policies could be welcome for emerging federations, ad their IdPs and SPs.

On branding: the current guidance "Login with your Institution" is actually confusing, since we want in research cases for the uswer tologin throuh their community, and *not* pick their home org directly. An "eduGAIN" branding may not fully solve that, but is a least better?
The order of login options on a discovery page is important (and not start with username-password!). eduGAIN should be on top ;)

To make eduGAIN better adopted for reseach, also make it easier for service to ask for the right thing.

FIM4R evolution and assurance

11.20

FIM4R is an essential part of the AARC-TREE process, and there have been a few in-person

meetings again in the past year (Denver was the first time, then at CERN, and now the next one will be a full day 09.30 till 17.00 in Copenhagen in January 2024).
Assurance, and how to get the IdP to implement REFEDS Assurance Framework, has been a key element of FIM4R over the past years. The institute lawyers are often a blocking factor (since they are scared of asserting anything), and now AARC-TREE will be looking at alterntives through government eID and wallets.
But given the March 1st start of AARC-TREE, the number of FIM4R meetings in Europe actually happening during the project timespan is potentially limited (since the number of FIM4R meetings is limited).
And the smaller communities (10-100 sized) are usually less engaged since they don't have the people to send to these kind of meetings (or look at this topic). One option is through EnCo, or use their AAI providers to proxy for them?
Package up AAI with other collaboration tools (like Indico, docand meeting minute editing, basic wiki/file sharing) may increase the attractiveness, but an AAI 'on its own' is not by itself a useful proposition for getting mid-sized communities to move to a non-ad-hoc solution.

How can communities effectively engage over longer distances maybe (e.g. from Autralia, other time zones). At some point during AARC-TREE we should try to build a FIM4R event in a location that is more likely to get to (Taipei or so, at least in a similar time zone). The collaboration with InCommon is easier since there is more personal exchange.

The Agenda for the 18th FIM4R meeting (Copenhagen) is now being crafted:

- https://indico.cern.ch/e/fim4r18 (https://indico.cern.ch/e/fim4r18)

which will include registration and logistics information shortly. Topics could include:

- should a proxy be transparent to be trusted, or should it be opaque on purpose? Also if the proxy is for a commercial providers. And *should* the proxy expose it? Or is it just a fully-responsible entity and has the liability thereby shifted?
- The proxy could state adherence to all (legal) obligations and be done with it?
- Can opacity and technical proxying be separated?

To ensure user participation in the future, join with a related meeting. For the USA based participation, TechEx in December '24 would be fine. In Europe, adding it to an AARC-TREE meeting might be an option, and push in a (single) session or side meetigs at TNC, where there are half-day meetings an option (not full days any more). TNC sessions, as in the past, might be an option subject to the TNC PC working mode. May of the rest of the world is coming to TNC anyway, so this is a good global engagement point. That is less true for TechEx, for instance.
The TNC PC in previous editions insisted on making their own choice, so there was no assurance that any resulting session was complete or coherent. And you really don't want somthing in parallel to TNC!
And still, the research communities may not be present at TNC, but their providers mostly.

SRAM community cross-service AAI, Infra proxies, and the WAYF explosion

11.59 [MaartenK]

SURF Research Access Management (SRAM) in an AARC compliant proxy, which takes a central role (next ot the SURFconext federation) to link research services to IdPs (that are linked through SURFconext):

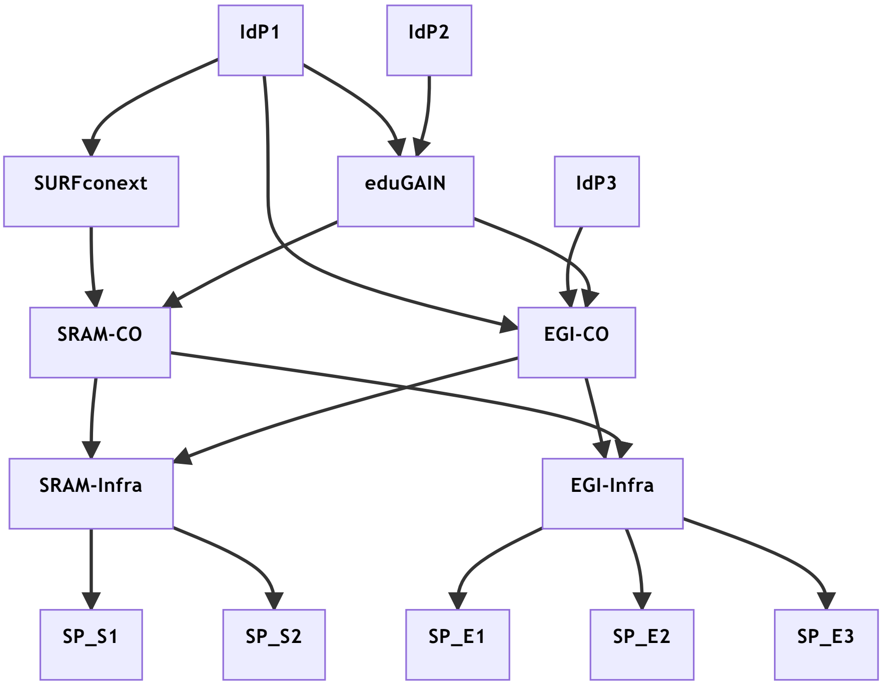

SRAM is meant to have service providers connected to it, and collaborations (COs) are configured in it, with the PI being)maybe the only) Dutch participants in a collabroative organsiation. The SPs (or OIDC RPs) are connected to SRAM. But … if we have nested proxies, where wee want to 'import' groups from e.g. EGI-CO into the SRAM Infra side, or vice versa. But if EGI Infra now also use services used by SRAM, should we stack the Infra proxies as well? Without too many WAYF!

Trying to prevent stacked WAYF that would confuse the users, and improve the user flow. The Germans NDFI/HIFIS federation has a similar issue.

Now AARC defind IdP hinting guidance, but how do we explain to the user that depending on your flow you should pick e.g. coming through the SRAM infra proxy , go through either EGI, or pick an SRAM-CO CO??
How do we automatethat choice? The knowledge may be at the SP, where the user is given the choice by the first component (SP) and then propagate the links through the system of proxies. The person giving the link to the new user can include the IdP hint parameters in the URL that is shared, and the IdP hints can be stacked and nested.
That does require all proxies support the hinting, and even then: the URL for the service should be 'simple' as well.

One option could be a *per-community masking/redirect URL*, that through a 403 redirect creates the actual Idp hiting URL with a G061 IdPhinting (chain) of IdP hints. This kind-of mimicks whatmany commercial services do with per-customer service login URLs. This requires the SP to have that knowledge (of their own customers, luckily), and build that redirection into their service. But the mechanism is straightforward.

Trust and information exchange between the proxies is another option (with *token exchange*) are two new guidelines that AARC AppInt is working on.

Cross-provisioning in proxies is another option, but to prevent data duplication you would need *"remote token introspection"*, and connect friends to eachother (behind the scenes). But then you have a name-mapping issue that needs to be cross-resolved. This is AARC-G052, which is de-facto approved.
But how does the user know to select the 'peer friendly' proxy? That would again present a double wayf, and that issue is not solved by the this back-channel comms. So the UX is still rather

complex. The service is unaware of where the token is coming from (for the non-web flow).
So this is not usable for a web flow, since the double wayf must be avoided.
But maybe only once, if the Infra proxy stores that choice persistently?

There might be adverse interactions with the impending loss of cookies across site (see W3Cs prvacy discussions). How hat impacts this may be limited, but it does need some consideration?

There are potenially similar considerationin the Earth Sytems Grid ESGF context, across the two parts of the community on both sides of the Atlantic. Use Case #3 therein is similar (on a private repo at https://github.com/ESGF/esgf-iam-uml/blob/main/docs/use-case-3.md (https://github.com/ESGF/esgf-iam-uml/blob/main/docs/use-case-3.md)).

At least EGI CheckIn kind-of support remote tokn introduction. And in ESGF, any opaque token is treated as 'from Globus', since only they are creating such opaque tokens. But just ask the users to login everywhere?

NDPF, CLI based access and HIFIS, OpenID Connect Fed

14.04 [MarcusH]

"Token solutions for the commandline … and more!" (https://cvs.data.kit.edu/talks/2310-token-based-solutions-and-more/)

OpenID Connect token management

One of the most well known producst to get token son the commandline is OIDC-Agent, which obtains the refresh tokens trough an authcode or device flow in OpenIDConnect. In version 5, security is improved, configuration and usability improvements, and it support 'mytokens'. It's available at https://repo.data.kit.edu/ (https://repo.data.kit.edu/). And several public clients are pre-registered in the software, such as *atlas* as well as other communtiies. It was security reviewed and the (minor) recommendations have been included in this version v5. Tis v5 is not in the public (EPEL, Debian, OpenSuSE, Ubuntu) repos yet. But the KIT repo has dev and release candidate versions (that have gone through an extensive 22-distro scanning) release procedure.

Long lived jobs remain a (usability) challange, and there are plenty of anti-patterns around, including long-lived tokens, infinitre refresh tokens, or even worse things. MyToken provides a mitigation by providing limitation capabilities as well as geo-fencing, re-use limits, scoped, audiences, and ip address restrctions - also split over various periods and timeslots (so you compute first, and only later in the week retrieve the data). See e.g. https://mytoken.data.kit.edu/#mt (https://mytoken.data.kit.edu/#mt) – and the production service is at https://mytok.eu/ (https://mytok.eu/) compatible with the AAOPS G071 guidelines. This is a lot better then sending refresh tokens with the job!
But of course the user can still make silly mistakes or bypasses. ANd users like infinitely-long tokens to just 'not bother'. But then you can groups users by what they are allowed to do, so maybe 'users without training' get fewer capabilities and limited validity periods :)
There *are* alternative solutions, e.g. the vault tokens propmoted by HTCondor, but then those vault tokens are not known to be limitable.

In addition to MyToken, there is als an Account Linking service ("ALISE") that is now being developed with the Helmholtz AAI (linking Helmholtz AAI, EGI, and socialID).

Secure Shell with Federated Identity

SSH is the favourite method to access a large range of cmoputing resources, and the go-to solution for everything interactive & HPC. There are now several options, including SSH CAs

(DEIC), provisioning, and the auto-provisioning one from KIT. And within OIDC ssh clients, there are several options as well either using some of the OIDC flows, or using the access token directly. This includes KIT, STFC, and Masaryk - usually based on the PAM scheme.
The KIT solution maps to an account, which may be either auto-provisioned on the fy, or pre-provisioned with (human) intervention. Also pool-accounts are an option.

Mandatory is only the `pam-ssh-oidc` , the rest (feudal, RST motley_cue, ssh-certificate `oinit`) is optional. The last one is new, mimicking Krb's `kinit` .
There is on the client side not much special needed, and MacOS, Windows, and posix-like systems are supported. And this is interactive, there are some constraints in automated (rsync, cicd-pipelines) tasks through ssh.

And for a more gui-minded users, there is a WebSSH client that can be linked to cloud VM provisioning interfaces. The WebShell is otherwise mostly harmless.

`oinit` is a tool that configures a target host (updating .ssh/config), obtain a ssh certificate from a scoped ssh CA, and the tool will automatically retrive (and update/manage) the SSH certificate. This combines the KIT and DEIC solution. This one does not work with Windows yet (since it uses pipes), and git needs work.

OpenID Connect federation

This work, part of the GEANT T&I Incubator, provided the final report on how to push the federation spec to its final version, driven by Roland Hedberg and the Italian government federation. The incubator pushed the tools and stability:

- https://wiki.geant.org/display/GWP5/OIDCfed+support+on+SimpleSAMLphp (https://wiki.geant.org/display/GWP5/OIDCfed+support+on+SimpleSAMLphp)

This includes a gap analysis, and implemented OIDCfed support in SSPHP!

NFDI and the HIFIS infrastructures

The HIFIS AAI is based directly on the AARC BPA, including the 2019 evolution. It is implemented in Unity (also used for B2ACCESS in EUDAT), and it has self-service group membership services and prmarily targeting OIDC services.
There are three levels of access to the infrastructure, based on either *i* communities, *ii* home-organsiation based since it considers home org as 'automatic ' VOs, and *iii* REFEDS RAF Assurance based, where you get assurance assigned based on membership in the DFN-AAI.

In some cases ther was an interst in streetAddress for the user, but no assurance framework provides for that. Maybe because address it is so transient, with (students?) moving frequently?

In HIFIS, ~25% of the users come from non-Helmholtz IdPs (!), and the largest HH AAI IdP is DLR. As for scale, there are ~5000 users active per month.

Next to HIFIS there is now the NFDI, with 19 consortia awarded in two rounds of calls, assiging 90 MEur for the first 5 year period. NDPF now is based on Base4NFDI-IAM project to provide this community AAI for all of Germany, with 4 community AAI installations: AcademicID (GWDG) didmos (DAASI), RedAPP (LDAP Facade) by KIT, and Unity (FZJ). and the 19 consortia can pick their favourite fromthese four. And all AAIs are proxy based, and use a common attribute set, originally retreived from EOSC.
And NFDI has an Attribute Conformance Checker (NACo) - this is automatically run against the proxies (see https://cvs.data.kit.edu/~naco/ (https://cvs.data.kit.edu/~naco/)).

Also NFDI is faced with stacked/layered proxies, basically encountering the same challenge that

SRAM/EGI had before. This will be the common pattern also in the EOSC AAI and a work item for AARC-TREE.

The NFDI policies is implementing the ARC PDK and make it mandatory for the participants. For IdPs, Sirtfi is required, and for SPs and infrastructure Snctfi will be mandatory. Wolfgang is currently working on a legal *opinion* on the policies and the AARC PDK. DFNs role in Germany will help and has sufficient standing with the community. And since Germany has 17 Bunderlaender, and thus 17 different privacy regimes …

Jens' Soapbox

15.30 [Jens]

A less technical soapbox than intended, but it *is* full of soap. And of access control and delegated credentials!

On Delegation

But first: how can be deegate credentials to automate data transfers in SKA, through Rucio, and with 3rd party transfer, but without delegatable RFC3820 proxies.
"Everything should have the minimal permissions" (least privilege principle). So you don't want a Christmas-tree full of privileges, butjust enough to have the service do what it seends to do. Can that actually be achieved, when all people look for is "does it work", and tke functionality as the primary metric.

Anyway, with GSI and RFC 3820 proxies this worked, but does this translate to tokens? Client needs to authenticate, get a bearer token, and the client can do wht it needs to do. This also kind-of works, but the (OIDC access) token has no restrictions, and is bearer only (RFC 6750). So we might be wanting token exchange (so rather RFC 8693). So can we get token exchange at the time you delegate raher than copying a bearer token, in a verifiable way?

If the access tokens are signed JWTs, they will need a (shortish) life time, but can be used independently and at high frequency. Opaque tokens, such as used by Globus, will require a callback every time, so validit is less of a worry, but the OP needs to be 100% available – and incidentally the OP will have a nice stat and control to count usage, licensing, and of course a charging scheme for RPs.

To make that work, tokens have som recognisable features. In OIDCfed, they are max_path_length, naming_constraints, decorated with trust anchors, intermediates, leaf nodes, &c). This is basically PKIX and PKI bridging. Like we saw in the OICfed trust bridging scenario before [basically, this is abstract federation struture].

A bit further back, there was Kerberos, with a TGT (Ticket Granting Ticket). The TGT is the login and a long-lived toke that can be used to retrieve service-scoped tickets (service and host). And a service can forward tickets, but also a user can forward tickets on their own behalf. Destination then needs an identity to forward *to*, and the dest use the tiket to fetch data from the source.

As a 'new old school', there are macaroons, they being fancy cookies, and are now being used (not only dCache, but also in some Google services, apparently). They can be constraint on delegation, like a maximum usage count, or a time window. They might be woth a fresh look.

And of course also SAML assertions *can* be revoked, even after they have been used for example after a job has started, when used with software that would support it, and some XACML magic.
"Old new old school" solution :)

Authorisation is in the end also delegatable, but then you just give over the entire authorisation, without any delegation. Like FTS retaining its own identity, and in the golden potal scenario. K5 is a bit like that as well.

On Assurance

Many of the proxies had no revocation capability, since they were (supposed) to be shortish-lived. If they lived longer, suspension of the superior credential (or the community!) was about the only option. For OAuth2 tokens, RFC 7009 does define revocation, and normal tokens can hence be revoked. But still you have to check credentils: check on initiation only, or check during e.g. a long-running transfer or job? That influences the assurance level.

But the real questions to ask is: what can the credetial do? Will it ‘be’ the user and be a Christmas tree, or just have the necessary rights? Till now only the coarse grained stuff has been used (e.g. the policy constaints in proxies were never actively used, as that was ‘too complicated’, beyond the simple ‘can you run a new job’ and the ‘limited’ proxies). But when the service you try to access gets all your rights, it can worm along. Restrictions however can be really fine grained. How do you do that with the technologies above? Or with pre-signed URLs (where enforcement should still be needed)?

> Q: we could do this in 3820 but users neve did. Any better now?
> Any complexity should be moved away from the user and be autoamted in the services. That would help adoption for constraining user rights ...

Anyway, finegrained authZ should be defined, and users are not too fussed about it. Unless you have really valuale daa, like what is protected by REMS. And then you need an audit trail as well!

How to express constraints?

This can be anything from n-out-f-m multiperson control, to time constraints, the phase of the moon, or validity period. How can you express those requirements?

There are several language options: XACML, OIDCfed language, and before the user interactions starts (RFC7521). You need to remember the tokens that are doing legitimate work, even if the tokens areinvalid under a new policy (but were acceptable under the old policy before you updated).

This is again a proxy problem:

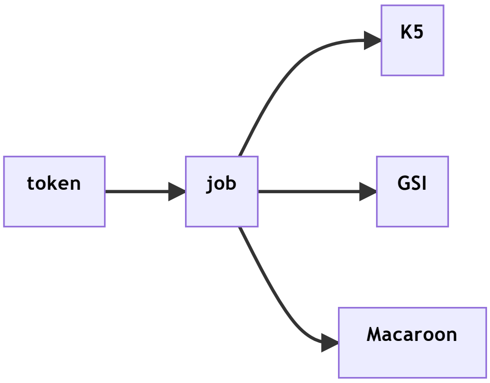

and you should be able to translate between the variants. This used to be Secure Token Services. OP (or in general an AS), MP. and TGTs are all effectively instnce of that:

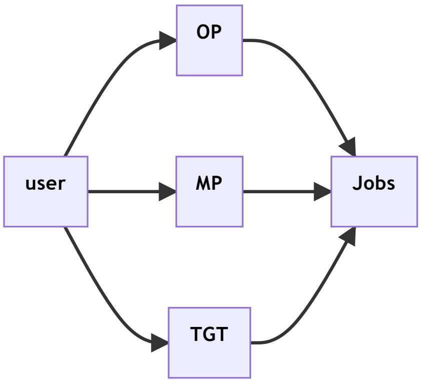

And then there are Robots ...

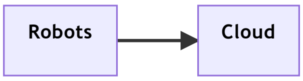

We have RCauth.eu (https://rcauth.eu/) for that. Which is technologically a good solution, but then you will need to add the authorization on top of it.
And there is [myproxy], mytoken (https://mytok.eu), for extended lifetimes.

Future meetings, FIM4R, and TIIME

- EUGridPMA+60: Monday January 29, 2024 Copenhagen, DK
- FIM4R: Tuesday January 30th, 2024 Copenhagen, DK
- TIIME Unconference: Wed January 31 - Thu Feb 1st Copenhagen, DK
- ISGC and APGridPMA: March 24-29, 2024 Taipei, TW