The Copenhagen 60th EUGridPMA+ Meeting Summary

The 60th 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 January 29 in Copenhagen, kindly hosted by the Niels Bohr Institute and organised by Anders Wäänänen.

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/60 (https://eugridpma.org/agenda/60) and linked therefrom.

These notes are being 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.

Introduction and note takers

09.30
Present locally: AndersW, DavidG, Jens Jensen, Maarten Kremers, Casper Dreef, Daniël Geerts, Mischa Sallé, DaveK, David Crooks, LiamA, Eisaku Sakane, Midde, Mads, Tom Dack
Afternoon: Ivan, Christos,
Remote: Jan Chvojka, Scott Rea, Miroslav, Mirvat, Derek Simmel

Round table updates

- Miroslav is still issuing certs, but now only for people outside of the Institute. Looking for other options for providing IGTF-accredited certs to Slovakia that would be a 'more natural' place.

Self-assessment review and discontinued authorities

(deferred until Cosmin can be present)

Developments in the Asia Pacific and the APGridPMA

09.57

APGridPMA updates

Now Eisaku-san is the Chair of the APGridPMA, with Sai Prasad being the vice chair. The next meeting will - as usual - be colocated with ISGC Taipei, with the summer meeting colocated with another e-Science or APAN meeting.
The mode of operations remains the same, with yearly self-audits and a catch-all CA service at ASGC.

Highlights:

- eMudhra CP/CPS approved private PKI services for IGTF
- new AAI provisioning to user communities and resource federation infrastructures - token-based infrastrictures with keyCloak and OIDC-agent and JWT-agent. The latter is developed by another team in the HPCI and allows updating the access token based on the provided refresh tokens - typically needed for storage access and gFarm middleware.
- closer cooperation of HPCI with GakuNin identity vetting sources and delegation of that to GakuNin IdPs that can provide identity assurance assertions.
- ASGCCA and KEK CA also deploying Indigo-IAM
- and SifuLAN offers IdP-as-a-service

33rd APGridPMA in Taipei March 26, 2024; 34th likely at APAN58, Pakistan.

GakuNin has the Othros authentication proxy service, based on the AARC BPA, and Eisaku has joined the AEGIS group to foster continued developments wih the AARC-Community (and linking to AARC-TREE). The credential policy and CPS for CPS are currently being discussed, and it's important to populate the policies and practice statement to provide certainties to the Japanese research community in terms of assurance level of the credentials issed by Othros. It's important to show that assurance level clearly to the research community in Japan. A more lengthy discussion of Othros at the ISGC conference in a few weeks.

Discussion between IGTF Assurance framework and the GakuNin levels. This is still being developed. This can be discussed in May.

Questions and Answers

- About the policies for GakuNin: is these also update of Sirtfi and R&S in GakuNin. This is still very low (with just one entity). Request is much more takeup of Sirtfi (https://refeds.org/sirtfi), to get update levels comparable.
- Typically security incident response is already taken care of quite well (implicitly, and maybe strengthened by existing federation agreements), but they just don't dare (or forget) to express it as an entity category. GakuNin would need an approval procedure for the IdPs to express Sirtfi.
- GakuNin assurance level effort for IAL2 LoA is progressing but there are no IdPs at that level yet. [this is independent from Sirtfi, actually - and Sirtfi might be easier]. But the expectation is that most *will* satisfy IAL2, but it's not easy to collect the evidence.
- What about REFEDS RAFv2 (https://refeds.org/assurance)? Might make it easier since now all the text is in the same doc.
- JWT-agent and OIDC agent are kind-of similar - discussion with Marcus ongoing. But OIDC agent does not satisfy the requirements of large-file-storage system services (at least not yet). The other WG in HPCI needed that function urgents: automaticaly get access to token update based on Refresh token. So design and implementation are similar to MyProxy - but not quite MyToken ;).
  However, JWT-agent does not support ssh integration. So for now in HPCI both are promoted. A review of the security model might be good to ascertain that the access token has not been 'implicitly' been turned into a quasi-refresh token now. To be discussed at FIM4R.

CA Update: RCauth.eu **(http://RCauth.eu)**

10.20

The RCauth.eu (http://RCauth.eu) service has its PMA that could benefit from more frequent meetings. Looking at the end of EOSC Future, that has supported the move to high-availability services. The HAProxy BGP-anycasted services is functional (Nikhef and GRNET), although the last few metres at RAL are still missing, and have been missing for the past years.

The anycast solution is not very well known yet - Jens sharing some slides.
It's a stateful service, shared between Nikhef, STFC, and GRNET, and you will always be able to connect to at least one of these, and using BGP anycast you will always be able to reach one of them. And the state is shared (Galera/MariaDB cluster) between the sites over an HA (VPN, originally VPC) back-end network. But each HAproxy connects by default to its closest delegation service (DS) which in turn connects to its issuing service back-end and talks to its local galera cluster node. Each site is conceptually identical (all have an HSM), but different vendors and throughput.

Randomness was used for key exchange with three different random sources.

Anycast explained using the LsIT course lectures. Failover is within seconds. And the HAproxy shields everything behind it, and you can use it (rather than terminating the anycast on the DS) on the HAproxy to take it out of the loop by the HAproxy server. And you only need one anycast address since the HAproxy can use SNI to use it for different back-end services. This basically is a 'shared webservice with a common backend state (database)' that can be used by any kind of

service. Not just CAs. For example this helps for WLCG where the push for extremely-long-lived-tokens comes from uncertainty about availability and uptime.

Finding the service performance indicators that would work at an organisation level to justify RCauth was a challenge (like adding 99.99999% availability, since performance/load was not the issue) was needed at RAL to justify this.

WAYF.dk (http://WAYF.dk) is doing distributed setups on the same subnet, but then managed redundancy by DEIC. Not using BGP for this, but with two different entrypoints for the same network on two dfferent locations.

The InAcademia DNS solution on Github maybe a bit outdated, but the software is used in production. The DNS caching at clients remains a problem though …?
WAYF is using a solution with weighing that can in real-time move users from one to another instance.

SSH Certs in a Federated World

11.30 - Mads

“Can you help scale ssh access to HPC with federated access” - and do that with SSH CA certificates. This then moved also to the GEANT T&I Incubator, and in the end the result was a set of ~6 solutions for federated SSH access.
Today, let’s discuss the DEIC/WAYF solution based on SSH certificates!

The SSH certificates are non-hierarchical, but in the end just a pubkey and some extra (expiration) information - more then just in the per-user pub keys. Default validity is 36 hours for the WAYF for now.
Now the WAYF solution does not need anything new on the client, and on the server side configure the extra certificates. Now, WebFed it!
The SSO login, now shows - after login - an ssh command line that will allow access to the chosen server.

It also works on the CLI only - and then you have to designate the intended CA and IdP. That will trigger opening the browser (platform specific).
There is a local “Principals mappings file” that maps global to local identity. This allows people to move between IdPs and remain the same user.

Constraints can be added in the issuing CA service based on e.g. assurance level and other IdP attributes. The CA itself can also be multi-tenant, and it being set up as a service in WAYF with a range of frontends and a back-end (HSM-based) services.

How does the server give access to the user? Based on certificate attributes and extensions. The user is typically pre-provisioned, and this is outside of the scope of this service. And in the map file the federated identifier should be already in (the principles_UID is a global config, per user).

There can be multiple CAs per server as well. All SSH-CAs are then created equal, and map all users.

For ensuring the users keys are on a token, the stronger FIDO modes needs to be used (so some Yubikeys and so). The token can prove that the key on the token was used by having an extra credential on the token which can be verified with Yubi tools. Can also be customised by organisational identity, eg NBI could ask for tokens with their identity for their own users. This may not be standard for all FIPS140 tokens. But that does not translate to the SSH CA domain. So that would ‘trust’ the SSH client - ahum …

And also the CA should issue ssh-server certificates, so that also the user has some trust. That makes it more reliable for the user in multi-server organisation.

Q&A

- the validity of the certificate is a policy issue. The MyACCESS-ID is 15min (a bit far too short). The WAYF default is now 36 hours (once a day) with discussions of moving to 11 days (once a week).
- it is possible to also do the provisioning over the SSH front channel
- Revocation is not supported - just keep the lifetime short!
- The Incubator work was completed a year ago, now working on their own - going to REFEDS in the afternoon. Could be a good solution for LUMI/MyACCESSid
- What about SSH forwarding? This also works with the SSH-CA, and the agent can also forward the keys here and talk back to the original privkey. Also for renewals.

Discuss also how the instance of the SSH CA should be (self)assessed against AAC-G071 (AAOPS) or the IGTF Assurance Profiles. Probably AARC-G071 makes more sense given that it in the end translates FedID to SSH CA certs on a token-translation mode of a BPA proxy.

- also ready looking at G071.

IGTF fabric updates: roots, OpenSSL 3.2, hashing, and more

11.40

Deprecating SHA1 in self signed root certificates (insert emoji of choice). Set policy to accept SHA1 signatures will have unintended effects. Note that IGTF also has self signed SHA1 trust anchors.

BEGIN TRUSTED CERTIFICATE will not work with everything - as ever, Java crypto providers can break in interesting ways. CANL is maintained by a single person...

Discussion

- Commercial CAs will never change
- Some IGTF SHA1-signed trust anchors will go away
- Nudge CAs to move

New package signing key and packaging

IGTF packages sources have now migrated to GitHub and are built from there. There is also more supporting tooling put there now (including the monitoring script).

Still pending for 2024: update to a new key (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

CA Update II: TAGPMA updates and CILogon changes

14.00

- Paula Venoso has left as vice-chair from TAGPMA, and some other more minor re-shuffles in the TAGPMA member list.

- UNLP has retired their entire CA. Ale Stolk/ Venezuela at ULAGrid will also not continue as an RP (same from Vinod in UFF?Brazil)
- Google Trust Services *has* joined as new member. Several CAs are pushing for accreditation
- now 15 voting members (8 APs and 7 RPs)
- Grid Canada updates are under review right now
- NSF Cybersecurity Summit & WISE in Pittsburgh

Google Trust Services

Google Trust Services (GTS) has come back after the initial tests with ATLAS Data Distribution, and have now come back with a range of root CAs with various technology profiles. PKI.goog (https://pki.goog) has all the trust anchor fingerprints

- the challenge may be to ensure name uniqueness (probably DC naming), so they may be spinning up new ICAs under their existing roots. That would be fine.
- may they be doing keyUsage: TLSServerAuth only? That would be a nice one … but needs impact assessment on FTS transfers.

CILogon retirement

- turn off the end-user facing /getcert endpoint by January 2024 no done.
- “create password-protected cert” by June 2024
- continuing for LIGO and Fermilab
- EoL by May 2025

Jeff Gainer confirmed that the shared codebase between CILogon and RCauth will remain supported by the software (the certificate part in the CILogon code base), even if there is no CILogon service. The RCauth patches are to be back-ported into the main branch.

CILogon AAOPS self-assessment

Mischa should ask Jim Basney directly.

Token Trust and Traceability WG, Resource Evolution

14.15 // David Crooks

WLCG Resource Trust Evolution WG

Not too much evolution in the resource trust evolution woring group, so in real terms there is not much update on that specific WG. Risk assessment for cloud workflows needs to happen to understand what their model is, and do the workflows fulfil our trust needs. That might give more long-term answers. The use of LE is ‘just the outcome’ of the current limitations. But now with Google actually looking at their own cloud trust offer, that could be IGTF accredited, even this can change! This might also change the LE landscape.
And in the end, server-SSL is just a small question of ‘why do we trust that workload’.

- Where to do the risk assessment for the cloud provider workflow use cases? That needs a venue and definite push (maybe the WLCG Workshop).

Token, Trust, and Traceability TTTWG

Group works alongside (but not entirely identical) to WLCG AuthZ WG, looking more specifically at operational use of tokens and (re-)implementation of existing controls. For example: what does

central suspension look like? What is the guidance for developers in terms of logging requirements (?like the MWSG logging guidelines, used in e.g. Condor?).
Focus on traceability, rather than isolation at this point.

For the moment not yet a regular slot, but meets in the context of e.g. the EGI CSIRT F2F meeting. And there is an opportunity at TIIME in an unconference session.

- where are we now?
- document the issues
- how does the workflow use tokens?

Christos: the EOSC setup with many proxies and chains is more complex than in WLCG, where it is unclear who 'the token issuer' is. Who should be issuing the tokens? Community AAIs or the infrastructure proxies? There may be many infrastructure proxies, and tokens must be traceable along the re-issuance and translation chain.
But now the user viewpoint becomes important as well: "how can I revoke the refresh tokens, if I don't know which ones have been issued as a result of my initial token??" !

- a bit like CT logs, but then scoped per user and access controlled to the user
- but you also need a mechanism to even look at that
- is there a delegated responsibility for proxies to propagate revocation? Then the proxies must have mandatory capability of doing that (remember SLO? This is even worse!)
- even if WLCG is much simpler, the interest from an incident responder view is that you need to understand all: also EOSC and all the other complex infrastructures.

Needs to be addressed at the AARC abstract level, and then the implementers can go ahead based on guidelines.

- separating the technical bridge and the trust bridge role of the proxy
- this is basically the OIDCfed proxy scenarios as discussed at ACAMP at TechEx23

Session based on the Community AAI, and all need to maintain the session and associate it with the session (during the session life time). But maintaining the session and make it visible to the user throughout the proxies is far from trivial.

- If revocation at the top, then the cascading refresh will in the end fail.
- Is the proxy at the infrastructure level not actually stateless. Should the infra proxy have session (like a refresh token)? It **has** to go upstream to validate the refresh token with its upstream (like like MasterPortal does, but nothing else does this!)
- When you start a new token life time it is now mostly an independent token, which is wrong™.
- On the infra proxy we connect services, and the responsibility of the infra proxy is to ensure that the right services connect to it. It should *also* go up to the community AAI and forward the requests to verify the user and get the claims and only then go back down to the end-service (client). That would make it much less stateful.
- Client authentication is at the infra proxy, User authentication at the Community AAI. Weird since OIDC mixes those two modes :(
- Token lifetimes may help make this less stateful? Controls are more important because of what the AuthZ flow is from an incident response perspective.

There is an AARC proposed guideline on token controls, which should be joint AARC Architecture and Policy, but is now orphaned. It is relevant for EOSC as well as WLCG, and there should be policy involvement as well. And should be done under AARC-TREE.

As a specific challenge on the TTT WG: also share this with IGTF/AARC Policy.

Inspired by `eduPersonAnalyticsTag` attribute: could not a compliant proxy look at a URL communicate in the incoming session, and 'log' source and destination of a token to the user-specified logging/transparency endpoint.

- this is only going to work on proxies are well behaved/controlled
- is in the end a **policy** issue on target/audience and scope (not have an exploding scope downstream)
- link with the traceability WG since this also provides traceability
- multi-step delegation is also similar, so it *could* even be relevant for WLCG. Otherwise, the complexity of EOSC is far bigger

The Wallets will have the same challenges (including the revocation capabilities) and there are analogies. In the central registry, wallets hold the proofs of what you did and act as a central hub. And there's always the B-word …

Enabling Communities: third steps in GN5

15.00 // Casper Dreef

There is still effort available, with approx. effort for two persons available in the last ~ 11 months of the project! This can include e.g.:

- REFEDS WGs
- Adoption of REFEDS policies and frameworks
- standardisation bodies
- effort to support federated SSH-CA … but that is already in for the CoreAAI for MyAccessID :)
- more people (and keep in mind the GEP)!

What is the status of OIDF RANDE working group? There is plenty of effort, but not enough physical warm bodies going there. How can we re-vitalise the R&E WG?
When we have something ready we can push, only then mobilise OIDF RANDE and then there is actually an operational issue and move things forward. The work can be in the AARC Architecture WG, and then push to RANDE. But you cannot quite go in with a fully speced document, because that also would not fly.

- e.g. Remote Token Introspection: if we really need this, and then we will need KeyCloak to support it, then it should urgently go to RANDE, since only after OIDF endorsement would KeyClock implement it.
- Remember what we did for RFC3820 …
- Is Remote Token Introspection compatible enough with the existing specs?

And we need mentorship programmes to get fresh people in!

AA Operations Guidelines (AARC-G071) updates

16.30

Update GEANT CoreAAI (ChristosK)

Secrets management discussion ongoing - the same NetHSM infrastructure to protect secrets for the CoreAAI platfom in general, both the proxy and e.g. the SSH-CA that DEIC/WAYF is now

operating. Can we use the NetHSM for real-time signing of the tokens? Can they support that use case? They are in a fixed location, but the CoreAAI platform is in AWS and maybe in other places in the future as well. They are a few hops away. Can we do that securely as well as sufficiently performant (to keep up with the signing rates needed for the tokens). This is work for 2024 ;)

Lifetime of tokens also still needs consensus. That was flagged in the G071 discussions but is not resolved. For the users this inconsistency is very confusing. Also for developers by the way when implementing workflows. This will be more prominent in the future.

- The CoreAAI may also be the first service to go through auditing at GEANT in the next couple of years.
- Used as the basis for many (also EC designated) services, so 'general purpose' eduTEAMS, MyAcademicID, MyAccessID, and now also the shift by RIs to using shared services.
- This has some impact on G071, where the community-centric bits are managed centrally by the AA Operator as a common element for all tenants. "They don't care" - just make it work. And that layer should 'just exist'.

Next meetings

For the May meeting: Wed 29 and Thu 30th all day. Location Abingdon/Coseners or Amsterdam/Nikhef.

For the September '24 meetings, maybe prefer CERN again?

Jens' Saebekasse

16.37 [Jens]

What identity - and what is the context in which you use your identity: work, family/ies, hobby/ies, … So the physical entity remains the same, but all of the biometrics used, as well as name and email can be contextual.
The same applied to 'work' context: going into "organisation" -> "payroll", but then also "project submission system" -> "projects", or ORCID, … and these are all different representations. And "research" with linked infrastructures and resources, library accounts, etc.
And … an infinite number of passwords.

But in a Brave New World™:

- Some people use organizational home pages
  - But organisations can change
  - Difficult to keep up to date
- ORCID integrated view of things
  - Who manages permissions of what it sees
  - How we manage what is in there
  - How we provide assurance of what is there
- Also a technology question
  - Delegated credentials
  - Languages for communicating

How do you associate the right elements, and disassociate from the wrong ones. What will provide a proper consistent view? You can get a person to do that for you, or an AI, or … get a single identity associated with all of that? That was what ORCID was intended to provide, but –

that is also incomplete.

What is missing?

- Single transferable identity, satisfying
  - Identifier uniqueness
- Identity assurance customisable
  - Customisable `acr`, `amr`

and then the LoA should be commensurate to the value of the resource you're trying to access, and not release more attributes than absolutely needed (e.g. only know you're over 18, you don't need a DoB, but just that one single bit).
Kind-of like SSI, but not decentralised - and the blockchain would be RFC6962-ish.

Yet *actually* what is missing could be assurance:

1. User accesses storage service
2. Storage service says "you must authenticate'
3. User is directed to community AAI
4. Community AAI directs user to an IdP
5. User authenticates to the IdP
6. User is directed back to the storage service
7. The storage service says "your LOA is not high enough"
8. Go to 2
   And you do the login dance twice...

What to Fix?

Server and Issuer alignment points

- The `acr`
- Required attributes (*pace* CoCo)
- Revocable (delegated) credential
- Renewable (delegated) credential
- How long the credential is needed
- What it will be used for

In the setup for EUDAT all access to all attributes has to be released, since the proxy cannot predict what is needed, and is not clairevoyant. The service could need all of them. The attribute request statement was not communicated (which would have been SAML AttributeRequests or OIDC scopes).
But doing the long dance again and again is also a UX issue. Including signing of the AUP and other properties of the sign-in signing process.
Then finally there is *attribute assurance*. And the attribute release policy should preferably be constant (in some way or form). And per-attribute assurance (and freshness) is needed in some scenarios? REFEDS RAF only has a global freshness quality.

All these considerations go into building attribute federations. Can you put stuff in, but not take out? How can a user actually manage all that complexity? For the user to fulfil the request to actually gain access and release the right attributes will need new technical support!

We can do better!
(and keep in mind that service operators always prefer simplicity over completeness :)

... and now Wallets will make this both simpler as well as more complex at the same time (since you cannot concatenate VCs since they have to be verified against each issuer, unless the bundler itself becomes an issuer again in a proxy scenario). This still needs further analysis :)