The CERN 63rd EUGridPMA and AARC Policy Meeting Summary - HedgeDoc

IGTF fabric updates: status of authorities and trust fabric news

The 63rd EUGridPMA+, AARC Policy and EnCo meeting is now over, and I would like to take this opportunity to thank again Hannah Short and CERN for hosting us as CERN. And thanks to Pau for the underground CMS tour on the evening before.

In this summary, we give an impression of the main discussions, results, and resulting action items. As usual, much is also contained in the slides and ancillary materials that are attached to the agenda pages at and linked therefrom. These notes were created collaboratively by those present - special thanks also go to Marcus Hardt and Maarten for their extensive contributions to these notes. As usual, any mistakes are mine, and in the on-line version they will be corrected as we go.

Looking forward, start planning for the 64th EUGridPMA+ meeting, addressing more AARC guidelines and the 'AARC TREE' D2.1 trust framework due by the end of May 2025. Hence, the next meeting is set for Wed May 14th (atfer lunch) will Friday 18th until lunch (12.30), and is tentatively planned to be in Prague, CZ. But this location still needs to be confirmed.
In any case, looking forward to continue building that warm and fuzzy feeling of trust!

Regards,
DavidG.

Welcome

Present on-site: Tom Dack, DavidC, LiamA, DianaG, MarcusH, Petr Vokac, MaartenK, Mischa Salle, DaveK, SvenG, Derek Simmel, Pau Cutrina Vilalta, Ian Collier, DanielK, Eisaku-san, Hannah Short
Present remote: Scott Rea, Miroslav Dobrucky, Jan Jona Javorsek, Valeria Ardizzone, Nicolas Liampotis, Enrico Vianello

Next meetings

The next EUGridPMA+AARC meeting is scheduled from Wed May 14th (atfer lunch) will Friday 18th until lunch (12.30). Tentative location is Prague, kindly hosted by Daniel Kouril and CESNET.

T&I in GEANT 5-2

(Maarten and Marina - see the slides)

The GEANT 5 phase 2 project started on January 1st this year, continuing a tradition of service delivery of the GEANT org together with the NRENs in Europe. This includes networking, but also - and importantly - trust and identity services, like eduGAIN and eduroam. And InAcademia, a privacy-preserving way of flagging 'student-ness', which is now growing very rapidly.

While the project may be limited to a limitative list of beneficiaries, the actual work is done all in open forums so that we continue to

the ecosystem and can benefit from global input.

With eduGAIN being picked up by more services (like the EOSC EU Node), there is now much broader adoption of the 'additional features' of eduGAIN. Like the assertion of 'faculty' as affiliation. And on the CoreAAI platform, there is broader adoption as MyAcademicID is being used for EuroHPC access and for e.g. student mobility. The inner details of the CoreAAI are not discussed.

The T&I Incubator was very successful in previous round as well, and in GN5-2 can accomodate approx. 20 mini-projects.

New sub-task (in EnCo) on VC Wallets and the application to the R&E sector.

Contributions on wallets and interoperability are still needed and welcomed.

- MARGI suspension discussion
- Updates in 1.133
- for TCS discussion see dedicated agenda item
- SHA-1 issues will likely persistent due to RH misunderstading, but use the work-around for now. Then, push for CANL-Java updates t deal with the dual-cert format with the RH-proprietary trus bytes.
- Meanwhile, some CAs modes to SHA-256 or have been suspended

GEANT Trusted Certificate Service Gen 5

- See the slides - including the paraphrasing of Wittgenstein
- Generation of the private trust roots and ICA foreseen for Friday this week (Feb 7th)
  Distribution 1.133 will include these changes in as far as they are ready.

Token-based operations - security discussion

Concrete potential actions:

Develop list of token issuers

Need ~immediate solution. Currently rely on VO blocking their own service, could do this manually but currently no automation.

Discussion (notes from Indico)

- Assume that (at least physics) workflows all go through a workflow manager
  - Sven: make sure that direct submission is not possible in that case
    Having a refresh token on the WFMS is not ideal fpr several reasons
  - Petr: too many token requests to IAM (believe that it cannot handle the throughput)
  - Luna: only know if IAM is available at the point when you want to refresh (at that point it's too late)
    No list of trusted token issuers permitted on the grid. This is done by the experiment publishing which issuers it uses.
    What should an issuer need to do to get on this list? Several policies.
  - Luna: Such a list could also be used for revocation? Revocation of an entire issuer (would require a semi automatic way for software to read the list)
    Token issuer can effectively block itself by stopping publishing its JWKs
  - Mischa: Suggestion that JWKs be hosted separately to the issuer - benefit to security as otherwise JWKs would be compromised at the same time the issuer compromised
  - Maarten: Brian B has also suggested a CDN
  - Mischa: SciTokens puts them in github
    There should eventually be 1 token issuer per VO plus probably issuers for WFMS
    Several solutions proposed for issue of how to do tokens for aynchronous workflows
    Very long tokens (not ideal for security)
    Local token issuers for the WFMS
    Could you have a future token? No, can't know when a job will run
    FTS = 3rd party copy. Flow for writing also requires reading at the end to make sure it worked fine.
    Identity of DDM is used
    Some difference of opinion on whether an individual user could be blocked inside a DDM or WFMS
  - DavidC: for central suspension we need to be able to tell all token issuers not to issue for a particular user
  - Marcus: one way would to have long lived but revocable access tokens
    This can be expensive and increase load on the issuer (if the same process)
    WFMS issuers should be constrained on audience, scope etc to mitigate risk
    In WFMS we are talking 200K running jobs. Status sent every few minutes, need a valid access token.
    If the compute node is talking back to WFMS every 10 minutes why can't it do a refresh flow at the same time?
    Could refresh tokens not be stored in the data base? Perhaps with signing, but this is not commonly done
    Enrico: shouldn't this be a client credential rather than a refresh token? Possibly the wrong workflow. Client credential could be passed around the infrastructure and used wen required to get an access token.
    Whatever gets put onto the grid needs to have a finite lifetime (client credentials current don't)
    One of the initial aims was to use standard solutions - but perhaps we've done all we can and after a certain point within the infrastructure we will have to do our own thing
    If we are doing something proprietary, how does that affect downstream parties (e.g. sites or EOSC compute or cloud). If we are non standard we will potentially lose them.
    It is still not clear exactly what performance would be required of INDIGO IAM (and by when) to avoid experiments doing workarounds
    We should escalate things to the OIDC Federation where possible to make it part of the standard
    We have to GUT working group to try and get to a unified OIDC standard, which should help integration

Trying to extended the standard may actually help as additional expertise would join
What would be out of the standard spec? Spec is very flexible
Access token revocation not done at the issuer
How is middleware defining who to trust? Manually adding it to their config
Anyone creating a token issuer should be aware of what policies it should follow
Raise github issue to externalise JWK endpoint in the well-known endpoint? for several reasons:
- Don't rely on IAM to be up to validate the tokens
- Can block an issuer without the input of the issuer
- However, this does shift the risk to wherever the JWKs are hosted
or use RPMs for this?
- What performance needed for INDIGO IAM? Perhaps we can make this possible, collaborate with SKA ? Separate low throughput attribute authority from high throughput access token issuer
- This is fundamentally a cost benefit analysis, with experiments more likely (in some cases) to trust issuers that they run themselves

Actions

- DavidG DaveK Hannah (and others?): Write a page of what policies are required for a token issuer and share with TTT working group
- Petr/Maarten: provide required performance for IAM (potentially for a few models) and a timeframe for Berk/Enrico to analse faesibility

"Thank god X.509 is still there" - M L

AARC Architecture - token profiles and life times (G081 token lifetime)

From https://sharemd.nikhef.nl/h1cH5HhrQM6TnBTYYDVuDA :

2. A consensus about MUST vs SHOULD in our recommendations.
    - Recap of discussion:
        - MUST: Might be too restrictive, infrastructures may choose they can't support the doc
        - SHOULD: Too open for "evasive maneuvers"
    - Potential way out:
        - MUST, unless "appropriate" mitigation is in place?
        - MUST, unless risk is well understood, deliberately taken, and mitigated to a reasonable extent
        - There is also an IGTF SHOULD: If you deviate, you must be transparent, explain the reasoning and get the endorsement by the involved relying parties
            - Original IGTF Text: "If a 'should' or 'should not' is not followed, the reasoning for this exception must be explained to relevant [accrediting] bodies to make an informed decision about accepting the exception, or the applicant must demonstrate [to the accrediting bodies] that an equivalent or better solution is in place."
3. A close-to-final draft for the lifetimes, for document "1"
    - Currently "doc 1"
    Objections considered only before Feb 19 2025
        - AT (verifiable offline, de facto non-revocable)
            - Default same as above (could be aligned with typical session lifetimes at a service)
            - Min same as above
            - Max 6h (In line with foreseen incident response times)
        - AT (verified online, de facto revocable)
            - Default 1h (Same as SAML, enough for reaching protected resource and doing something with it, like login)
            - Min 15min (enough for reaching the protected resource)
            - Max 25h (Taken into account: Revocability, Enough time to run a short job, checking results of previous day)
        - RT
            - Default
            - Min
            - Max
                - Either an infinite number of 90d (maybe client_credentials is a more applicable flow)
                    - Rotation is a MUST in this case
                - Or 13months
            - Deprovisioning is the actual underlying question
    - Future "doc 2"
        - List of use cases paired with **mitigation strategies**
        (e.g. how to turn offline verified, non-revocable to revocable, or how to limit the power)

Developments in the Asia Pacific and the APGridPMA

(see slides by Eisaku Sakane)

Updates supplementary information:

- IHEP is updating its CA systems, but since it was using OpenCA this requires a lot of additional work (it has not been developed further). At the moment there are no further updates, but at the next APGridPMA meeting they will provide a status update.
- The KEK and ASGCCA providers are now also provding token-based AAI services based on IndigoIAM for communities (and HPCI has also developed a JWT-agent complementing OIDC-agent and OAuth-SSH)

Communications for the APGridPMA is now also using an 'apgridpma' Slack channel.

The next 34th meeting will be March 18, during ISGC2025 - and the next 35th meeting in summer or autumn. Note also APAN59 in early March.

HPCI and GakuNin

Medium-scale Demonstration Experiment FY2024:
Demonstration experiment to realize the use of academic e-resources based on the ID operation management of IAL2 and AAL2, which new GakuNin trust framework plans to newly establish in the future.

- Period: January 17, 2025 to March 31, 2025.
- 20 participants (University, Institute, Company)

Online self-service for identity proofing

Sign up with federated credentials.
Traditional identity proofing based on a face-to-face meeting with photo-ID presentation. Being authenticated by

Plans in FY2024 (until March)

- GakuNin
  - We will summarize the result of the medium-scale demonstration experiment FY2024.
- HPCI
  - We will organize the required attributes for online self-service for identity proofing.

DCVOTA and Elm - TAGPMA update

(see Derek's slides)

CILogon certificate services are going to be retired by May 2025 - affecting those users in the US from communities that are non-US based and relied on CILogon certificates.

Elm and DCVOTA

Derek presents Elm (update v1.5 2024 assurance level documents)

Having circulated this exact 'Elm' v1.5 version of the text in November 2024, and again on Janury 17, 2025 for final comments. That having been two weeks without comments, and no further comments or concerns having been raised at this EUGridPMA meeting, the EUGridPMA hereby approved the version 1.5 of the IGTF Assurance Profiles.

For the DCVOTA profile, this having been two weeks without comments, and no further comments or concerns having been raised at this EUGridPMA meeting, the EUGridPMA hereby approved the version 1.1 as presented by Derek (allowing for a typographical correction of "an TAGPMA" to "a TAGPMA").

WLCG and EGI *will* update the acceptable assurance to include DCVOTA next to ASPEN, BIRCH, and CEDAR

Following this approval, **DavidG will**:

- assign and enter the OIDs for Elm and DCVOTA in the objectID registry
- update Alvestrand.no and https://oid-rep.orange-labs.fr/
- upload the documents to igtf.net and the versioned pages

DCVOTA CA: Google Trust Services

Derek/TAGPMA will get the *relevant subset* of the GTS CAs so that the number should be limited (for namespacing and CRL retrieval). Not all of https://pki.goog/

The namespace prefix - for the ICA(s) - needs to be specified (probably in the CPS).

IGTF certificates will be identitied by the policy oid for DCVOTA.

Risk assessment for tokens

David Crooks:

- one of the outcomes from yesterday was the structural and technical context of the work, and after dinner Sven, Marcus,

Daniel, and DavidC discussed the risk study and how to practically move forward
- with the experiments and WLCG management, sites and security there should be a stakeholder and scope inventory, where we want the outcome to be both detailed enough and timely enough to be actionable (and not be shelved)
- we can do that in stages, but the first thing to have is **we need something, and when/what is that going to happen**

SvenG:

- target audience, and who will be accepting any residual risks?
  - not be answered here and now, but outcome should be timely
- the term 'risk' shows up several times, but was not quantified (and rather vague). Risks that come with a certain technology are not considered as part of an ongoing process, rather than an architectural approach, taking into account the risk scenarios and how this risk is perceived by the major stakeholders (including relying parties, sites, opsec team).
  - It also depends on the perspective of the resource provider. If that is a commercial one, just providing it to a community but not bothering what is inside, many things otheerewise considered 'bad' (like cryptomining) are not that bad.
  - but as a community, such resource abuse may 'not be appreciated' - unless they are a cryptomining research community.
- But as a site you are still responsible for due diligence. And may jeopardize the status of the site in the (funding) community. So this does not quite hold.

This discussion is a lot of work, but we need quicker answers now. So in addition to a comprehensive review we may need to reach down the stack to answer more specific questions.

- rough skeleton of what the study would look like
- then flesh out with scenarios
  And ... what is the potential benefit that makes a site accept a risk? (i.e. a typical primary asset list).

Did other orgs (like FNAL) have already done such an assessment we can borrow from? And EGI has relevant documents as well. And there is the existing LCG and EGI risk assessment (DaveK).

What & when?

- two goals
  - broad risk assessment - can take approx. a year
  - identify focussed risk studies to answer short-term needs (for IAM service token engineering and related items, primarily WLCG) - within the next 6 months, i.e. **by the end of July 2025**
- subgroup (of this group) to participate in scoping:
  - DavidC: focussed assessment for WLCG as above;
  - for the broader assessment: *undefined group* but at least EGI (Valeria) is interested!
    SvenG and Mischa to provide input to the process in terms of 'questions'.
- initiate cultural change in the community
- making an architecture checklist
- architectural view of current solutions and their vulnerabilities (and there are as many solutions as there are people/ experiments)
  - what is the full picture?
  - what is the risk environment and show the result to the software developers and experiments. And make suggestions make on that ("all experiments should work in the same way", fat tokens vs. fine-grained tokens?, cross-community impact of incidents/vulnerabilities)
- scope: at least WLCG and EGI mngt as a community, 'site view', 'community (e.g. LHC experiment)', RIs (like WLCG) – predominantly for the broad risk assessment goal
- consumer(s): *undefined* (to be defined by the initial assessment groups?)

desired result: to build the case for suffcient resources to make it work and achieve the desired objectives (the word 'risk' may actually confuse part of the target audience?)

AARC Policy Coordination and AARC-TREE: policy development planning

Objective: support the diverse and different policies needed now AARC
Infrastructure alignment and policy harmonisation: helping out the proxy (M1-M18, 21PM)

- Operational Trust for Community and Infrastructure BPA Proxies Generalised consolidated Security Operational
- Increase acceptance of proxy identity providers through common baselines Baseline from PDK+EOSC
- Review infrastructure m coordinated AUP, TC, and privacy notices
- D2.1 Trust framework for proxies and Snctfi research services

User-centric trust alignment and policy harmonization: helping out the community (M6-M24, 26PM)

- Lightweight community management policy template .
- Guideline on cross-sectoral trust in novel federated access models
- Assurance in research services through (eIDAS) public ID assertion

AARC I082 trust frameworks and Policy Development Kit evolution

We reviewed the "Trust framework for proxies and Snctfi research services" (AARC-I082 / AARC-TREE D2.1) document (https://edu.nl/uu3qt), clarifying the roles that proxies have in the different phases of the interaction, and pointing out the structural

differences betwene infrastructure proxies that connect services across domain vs. 'site-local proxies' (a term introduced specifically to distinguish them from the proxies discussed in the existing AARC BPA) althouh both have a role in enabling research workflows across resources (within a single administrative domain, or across domains). This has been clarified in the introduction section to I082.

This framework also defines the *structure* (not the specific content) of the PDK policy development kit. A new section has been introduced, taking into account the feed-back from the AAF on the position of the current PDK templates in the purpose-audience matrix. Also https://docs.google.com/spreadsheets/d/1EWBJaUrJPatWzs2_jdWnFJIqeZYMHvULTucLxQPaC9Y has additional audience information.

Each policy/procedure will be an individual AARC (informational) guideline, and each added to Zenodo for permanent reference. The PDK itself is then an overview document that refers to each of these (versioned) DOIs.

There are currently 9 documents in the PDK. In the refactoring:

- the risk assessment one should be removed, with a (single) line in the PDK overview document pointing to existing (external) risk assessment guidelines
- the Personal Data document is very regional (GDPR biased), but it was used by WLCG and EGI, but used to bind non-EEA countries that are thereby bound to the same framework. This is a basis for the pretty-binding-not-quite-corporate-rules (noting that today we also have the REFEDS DP Best Practice). As for the Template data privacy notice - keep the document, but reduce the content
- The Service Ops Security Agreement becomes the Baseline Operational Requirements (from EOSC and UK-IRIS). Discuss traceability and logging if not already there - there should be a procedure/process attached to the policy as well. This is the FAQ to the EOSC Baseline Op Security.
- top-level policy can become the webpage and explanatory document. Then there should be a document that gives authority to specific roles and functions. "Collaboration definition" - basically the question "what are we?"
- Incident response: the current one is a procedure, and is also very lengthy. The new EGI SEC01 is more check-list like and much more concise, but misses out a bit.
    - have a PDK document describing that an incident response procedure is needed to implement the Baseline, and give examples based on existing infras. Pro-active reporting is missing in EGI and similar
    - we do now have Sirtfiv2, so let's include that.
    - Guidance is "read this document and examples, and make sure that your procedure meets these requirements."
    - Refer to federated incidents and the need to inform eduGAIN.
- Acceptable use: WISE Baseline AUP. AARC-I044 has the guidance for that, plus AARC-G083 on combining notices. And start with the WISE Baseline AUP.
- Membership management -> community management policy with 5 bullets.

The AAF matrix also identifies the 'missing elements' (the empty boxes) - this shows tha the PDK has been 'resource protection' oriented, since it lacks other policy elements such as service level agreements and in general service management/service portfolio discussion typically found in ISO20k, FitSM, and ITILv3, including the 'promises' to users. While these are likely out of scope for a template (and there are many templates already!), pointing these out in the overview is useful.
Similarly, the PDK does not have any discussion on 'access guidelines' - incuding resource allocation and funding. Also here, there is no need to include actual templates in the kit, but it should identify that policies in this area will be needed. Now, converging on these policies typically takes decades and is more 'political' than technical ...

We should add a glossary in the top-level policy that defines the terms, in particular 'community'. "We use the word community in the templates and guidelines, and invite the users to the PDK to replace this by appropriate user-specific language".
There was discussion on the use of the word 'policy', but by keeping the term 'policy' it ensures explicit approval, whereas 'guidelines' never get approved by management bodies.

Wallet ecosystem - new federation models in AARC TREE

(Maarten Kremers et al.)

Planning for Wallets and VCs is ongoing in GN5-2 (and was already there in GN5-1). The current ecosystem (including the European Block Chain Initiative EBCI) and the whitepaper that is now being written following a risk analysis in the GEANT project.
Link to whitepaper

The idea was to investigate the 'threat' of wallets to the existing multi-lateral federations in eduGAIN, but the result indicates that the main changes would be above the community-proxy layer, where it seems to be largely contained.

- Should the community AAI also return 'stuff' as VCs in a wallet? Niels van Dijk's demo encoded group entitlements in VCs and put those in wallets, with the community AAIs acting as issuers:
    - Selective Disclosure JWTs (SD-JWT) allows the user to release a subset of the claims without breaking the signature on the whole blob.
    - Linked claims are less relevant to 'our' community.

As a result of the whitepaper there was a webinar for outreach (https://wiki.geant.org/display/GWP5/Trust+and+Identity+Infoshare+on+Wallets) in September last year, with Davide explaining the relation to eduGAIN and Christos on the relation with GEANT's CoreAAI product.

For the next (GN5-2) project phase will review use cases for the wallet with eIDAS and the EUDI wallets. And OID Federation adds to the wallter system since there are mny more providers, and will those be the ones to which the wallets authenticate?

The scope of the AARC TREE work is on analysis of whether and if so how this can help het govID assurance in the ecosystem in w way that fits the research collaboration proxies. That may work better than relying on home IdPs doing REFEDS RAF. The REFEDS white paper working document lists.

Germany (DFN-AAI) has implemented REFEDS RAF by virtue of Wolfgang. It may be however that every IdP still defaults to 'low', athough at KIT (of course) everyone with a cert does get 'medium'. Otherwise not well known.

The least we can expect is that in the identity proxy layer there will be a way to step-up the assurance by means of EU Identity Wallets.
Where the community proxy steps-up the assurance, with or via account linking, this will actually change the assurance level as defined under G009 (login with eID will be cumbersome). Here the proxy can step up

But even if there is no linking involved ... these docs should be merged (identity step-up and account linking)

Account linking and ID step-up are slightly differently defined

- G009 (linking) requires to always pass the assurance of the last method used for authentication
- ID step up, however, will allow an OP to assert the assurance level of the stronger method, even if a lower-assurance identification was used to log-in
- There are EIDAS (2014) regulations that requires accepting qualified signatures just as well as "blue ink" signatures.

Step-up for non-EU27 should be incorporated as well e.g. through the use of remote vetting cmpanies like Sweden is doing, or the e-stepup on Germany for non-EU citizens.

What is the current status of the work on Wallet in the working group mentioned by Markus?

- The bi-weekly calls were moved to Monday to review the whitepaper. The schem can be offloaded to a (new) AARC ARCH working team.

What are the plans for the activity dedicated to this topic in AARC Tree (is there a study we need to do with some use cases)?

- There is M2.2 on the step-up of assurance through government e-ID that will rely (to some/a large) extent on Wallets. Niels van Dijk (SURF) will link both wallets and the eID stepup based on earlier work on wallets in the T&I Incubator. - with an emphasisis on the policy/good practice aspects.
  This is a new guideline (to be submitted to AEGIS in AARC TREE M18).

Looking at the EBSI work, there is useful things in there besides the "B" thing.

https://edubadges.nl

AARC and eduGAIN federation TTX challenge and security day ISGC2025

The ISGC Security Day comprises two events: the eduGAIN Security TTX "table top exercises", based on the eduGAIN scenario and supported by the open source tool, reviewed by Daniel Kouril. A new addition this year are token-based elements of the sscenario, developed by MarcusH, where we also invlve authentication flows through an IdP-SP proxy where the result is a token that eds up in the service provider (RP) log files. Analysing these log files is a new element, and not yet widely understood - which is an additional learning element compared to the Boston scenario (which took 2-4 hours to run with the variou role players).

The exercises highlights a couple of risks for the participants. To make you understand the risks in the eduGAIN context, we go through building a risk inventory, identify the threats, and the foundational controls that help mitigate the most common risks in federated authentication. In this part of the course, we learn how to apply the federation terminology, how to relate that to concepts in the common risk management standards such as ISO27k5 and ITSRM2. How to use the risk managment process in this exercises builds on a presentation of common risk management process and - skipping the contextual part - continue to focus on identification of risks.
In an interactive workshop setting, we discuss in small groups how they will apply the process and stay at a 'risk management' level rather then being defeated by complexity in detail.

Who are the threat actors? What is the impact of popular eduGAIN multi-lateral scenarios? What is the intended outcome of the exercise and what are the lerning outcomes?

Implementation for ISGC2025

- will this work in the few weeks left? MaartenK and SvenG to sit together F2F in '(Zevenbergen|Utrecht|Amsterdam)' within X days [maybe colocated with the AARC ARCH meeting Feb Tue 18 at Nikhef, at 10AM, with a joint lunch with the WG meeting **Action DavidG: add eggless lunch participants**]