The Amsterdam 62nd EUGridPMA+ and AARC Policy Meeting Summary

Dear all,

The 62nd EUGridPMA+, AARC Policy and EnCo meeting is now over, and I would like to take this opportunity to thank again Maarten Kremers and SURF for hosting us in the Amsterdam offices. The next meeting will be at CERN in Geneva, from Wednesday 5th till Friday the 7th of February 2025, hosted by Hannah Short (the INDIGO IAM Hackathon will be on Monday the 10th also at CERN, so those attending both can stay over the week-end).

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 - thanks go especially to Marcus Hardt for singificant contributions to this summary.

I hope to see many of you in the continuous policy discussions and at the next EUGridPMA63 + meeting!

Kind regards,
DavidG.

Present: Maarten Kremers, Catharina Vaendel, Marcus Hardt, Eisaku Sakane, Liam Atherton, DavidG, Mischa Sallé
On-line: Adeel-ur-Rehman, Miroslav Dpbrucky, Hannah Short, Dave Kelsey, Jan Chvojka, Lidija Milosavlejvic, Nick Rossow, Nicolas Liampotis, Feyza Eryol, PeterK, Diana Gudu.

Next PMA & AARC meetings

The next EUGridPMA63+ meeting, in conjunction again with AARC TREE, GEANT5-1 EnCo, IGTF, and EGI will be held:

- **Wednesday February 5 noon till Friday 7th at ~16.00 2025**
  at CERN (Meyrin site, Geneva, CH) near the Computer Centre building (bldg 28).
- **Monday March 17th till Friday 21st: APGridPMA and ISGC**
  Taipei, TW hosted by ASGC
- **Monday March 31st till Thursday April 3rd: FIM4R, AARC and TIIME Unconference**
  Reading, UK, hosted by UKRI

In this summary

- The Amsterdam 62nd EUGridPMA+ and AARC Policy Meeting Summary
- Next PMA & AARC meetings
- IGTF Fabric updates and self-assessment status
  - TRGrid updates
- Developments in the Asia Pacific (Eisaku Sakane)
  - Federated credentials in HPCI (GakuNin + IAM)
  - Process for sign-up with federated Credentials
  - ID Proofing based on “genuine MICS”
  - Future Plans
- AARC Policy Development Kit: feedback from the Australian Access Federation (Nick Rossow)
  - T&I Framework in Australia
  - Challenges
  - Shear amount of raw AARC material: Where to start?
  - Terminology Challenge
- AARC and AARC TREE: policy development planning
  - Objective: support the diverse and different policies needed now
  - Overview of main policy activities
    - Proxies, AUPs, T&Cs, Privacy notices, … managing notice management
    - G082: Informational Document
  - Upcoming Deliverables and Milestones
- A new Assurance Profile: Elm and DCVOTA (Derek Simmel)
- Digital Wallets in GN51 (Maarten Kremers et al.)
- Token lifetimes and G081 (Nicolas Liampotis et al.)
  - Is low-medium-high life times enough of a differentiator?
- Validation and Adoption of policy guidelines
- Distributed proxy scenarios and ‘DAG’ AAI flows
  - Discussion about the user-path-discovery ambiguities
- Evolving notice management: G040 evolution and context
- SSH CAs by Marcus Hardt
- FIM4R at TechEx

IGTF Fabric updates and self-assessment status

The number of Authentication Authorities is slowly but steadily contracting, with move moving to GEANT’s TCS.
See updates in the slides - and move your root to SHA2+ if you can (but it will never be universal).

- All WaTTS instances have been silently turned off :(
- The 1.134 release will be withdrawing the HKU CA (HK)

TRGrid updates

The TR-Grid Root CA is expiring within one year, and it is time to re-think. The CA was accredited in 2005, and it has now been 20 years, used by the CERN Experiments as well as other use cases.

The CP/CPS was updated in 2021 (v2.4) and since then has been in operation. The number of users is getting smaller, with users using the CERN CA, with the number of hosts being mostly contant.

The self-assessment (GFD.169) for the off-line CA is done yearly, with consistent results. But there re now only 378 days left, the key is 2048 RSA bits, and the digest is SHA-1. This triggers the known issues on EL9, and this is also affectingthe Tier-2 WLCG and EGI centres in Turkey.

For the migration, a complete refresh (with revocation of the old ones) is foreseen, since the number of subscribers it relatively small anyway. Another (or parallel) option is to move to TCS

Given tht most institutions in Turkiye get their public certs from the national public CA, the use case for TCS is reduced.
New new TR-Grid CA will be introduced. Lifetime for EEC on the old CAs can be gradually reduced.

Overlap of approx 6 months for introducing the new TR-Grid G2 CA.

Developments in the Asia Pacific (Eisaku Sakane)

Eisaku (NII) is the current Chair, with vicechair Sai Prasad from eMudhra.
The number of CAs is also contracting in the AP region, with ASGCCA acting as the catch-all now for ID, IN, MN, MY, LK, NZ, PH, TH, and VN. The HKU Hong Kong university CA also distinued, following issues with a failed update to a latest OpenCA distribution: the update was overly complex since OpenCA is not well-maintained and does not install readily on a modern OS distribution. They hence terminated their CAs, and their subscribers moved to ASGCCA (adding HK to the list).

The IHEP CA is building a new CA based on the new IGTF requirements, but they are based on the same OpenCA distribution. They are currently reviewing it, and there is no new update yet.

Moving toward new AAI according to federation infrastructure

- HPCI: using auth-ssh, KeyCloak, oidc-agent and jwt-agent and consider cooperation with GakuNin
- Token based AAI model deployed in ASGCCA, KEKCA (based on Indigo-IAM)
- Malaysian Federation (SIFULAN) offers IdP-as-a-Service nationally

Regional Collaboration on identity management

Federated credentials in HPCI (GakuNin + IAM)

- Addresses ID Proofing challenges
- Current face-to-face process + photoID process is too heavy-weight
- Delegate ID-Proofing to home-IdP
- Issues:
  - Current HPCI OpenID Provider does not use assurance
  - How to address industrial researchers
  - Do the IdPs meet the IAL (identity assurance level) imposed by HPCI IdM?

Process for sign-up with federated Credentials

- Traditional: f2f meeting + photo ID
- New: Sign-up with federated credential (self-service)
  - Scientific users come via GakuNin
- New (for industry):
  1. Use previously created account on Orthros
  2. Authenticate at Orthros with AAL2
     - Eg. using gBizID (a gov ID service)
     - Additional linkage (ORCID, other IdP) also available
- Both ways are accepted by HPCI IdM

ID Proofing based on “genuine MICS”

- HPCI CA is a MICS-based CA
  - Ext IdMs operated by Japanese HPC, connected via SAML to HPCI RA

Future Plans

- Design + Implement new sign-up to federated credential service
  - Prototype first
  - SAML profile for federated client
  - First results: Spring 25
- Otheros will be the primary entrypoint for industrial researchers, whereas GakuNin will be the main entry for academic users.
- ORCID will be at LoA level 1, without affiliation amnagement. The affiliation manager can usually preserve some attributes, such as home organisation of the researcher with high assurance. For companies, that will be using ORCID, this is less clear.
- for eduGAIN MICS integration, will leverage REFEDS Assurance Framework at LoA-2 “Cappucino”. This can constitute a push towards RAF adoption. Meanwhile, need to understand whether this assertion is based on the underlying actual assurance (based on ePAssurance and authenticating context class Reference ACR).

AARC Policy Development Kit: feedback from the Australian Access Federation (Nick Rossow)

T&I Framework in Australia

- Easy: BPA
- Hard: Policy (importance and understanding)

Challenges

- Where to start
  - How to present all the AARC information in a simple way
  - How can fresh people learn the output of five years of work
  - Best advice: “Just Start” (Dave)
- Terminology
- Governance
  - Change document names (to avoid lawyers from getting too excited)
    “If you use legalese-type language, you get a legalese-type response”
  - e.g. change name from “Policy” to “Best Practice”
  - => Should we rename “Policy” to “Procedures”?
- Implementation timeline: 1.5 years from “Desktop Mapping” to endorsing the PDK
- PWG Members:
  - 12 Organisations
  - 20 Participants
  - 8 research sectors

Shear amount of raw AARC material: Where to start?

- Developed instructions:
  1. Define the Collaboration
  2. Assess Risk
  3. Manage Membership
  4. Privacy Management
  5. Security Management

Terminology Challenge

- Personal Data
  - Often mistaken as research data
  - Rename to Personal Information?

    - This “Personal Information” is the .au legal term for what in the EU is GDPR Personal Data. Once that was clear, the conversations became much clearer.
    - Post Hoc addition by Arnout (sorry): in research (and outside), the term PII is quite common (personally identifiable information)
    -
        - Post Post Hoc comment (from Nick) - PII is mostly used to refer to research data, particularly in Health research.
- Community
    - Community and Infrastructure are differently mentioned different policy documents. The vague defintion kept keep-up of the wording
    - => Use one term consistently
- Infrastructure / Service Providers / Participant

There is now a terminology document, merged with the top-level document, that states all definitions.

- -> Rename? Top Level Policy to ?

AARC and AARC TREE: policy development planning

Objective: support the diverse and different policies needed now

- Infrastructure alignment and policy harmonisation: helping out the proxy
- User-centric trust alignment and policy harmonization: helping out the community
- Anchored in the research user communities by co-creation with FIMAR, through policy workshops validating the restructured policy framework ... together with the new BPA

Overview of main policy activities

Proxies, AUPs, T&Cs, Privacy notices, ... managing notice management

For large ‘multi-tenant’ proxies:

- some subset users in some communities use a set of services— how to present their Terms and Conditions, and their privacy policies, so that the users
    - only see the T&Cs and notices for services they will access
    - this does not to need to be manually configured for each community
    - is automatically updated when services join
- as well as for community and dedicated proxies:
    - when new (sensitive) services join, who needs to see the new T&Cs? beyond AARC-G040
    - can we communicate acceptance of T&Cs to services even if ‘we’ are small and ‘they’ are large?

G082: Informational Document

- Frame the requirements for proxies, and which (elements of) the framework are relevant for each category of proxies (community proxy, multi-tenant community proxies, infrastructure proxies, &c). This both strengthens Snctfi, as well as lays the foundation for the Baseline and AARC TREE's "D2.1" (due in May 2025).

Upcoming Deliverables and Milestones

- M2.1 Guidance for privacy notice management by proxies (M10): lead needed
- D2.1 Trust framework Trust framework for proxies and SNCTFI research activities (M15)
- M2.2 eID assurance model suitability (M18)
- D2.2 AARC Policy Development Kit Revision (M24)
- Generally: Problems with availability of people, yet documents need to be written:
- Work items:
- Research Infra Alignment (Nikhef)
    - Operational Trust for Proxies: NIKHEF
    - 'Snctfi' R&E Baselining & Integration: Maarten + Dave (aligning it with OTfP)
    - Models for Cross-Infra AUP & Privacy Notices: Catharina, Dave, DavidG, Arnout
- User-centric Trust Alignment (RAL)
    - Lightweight Community Management Policy: Hannah kindly (much appreciated) volunteers for leading the team here
    - Guideline for Novel Federation Models: Marcus / Diana, Nikhef input (evolving G071)
    - Assurance in Research through eID: KIFÜ & Geant
    - FIM4R Policy Evolution: Continuoously ongoing (all)

A new Assurance Profile: Elm and DCVOTA (Derek Simmel)

Atlas and Google worked on getting forward with enaabling ATLAS to trust Google's certificates. To support this, the IGTF Trust Profiles are suggested to be extended with:

- DCVOTA: Domain Control Validation Only Trust Assurence
- ELM: Main difference being "The authority must describe how it can reasonably verify administrative control of Internet domains in the subject names by the subscriber at the time of credential issuance"

Two quick questions were also addressed

- the requirement for unique subject DN prefix naming (what we typically see reflected in e.g. '/DC=org/DC=incommon/...' and '/DC=org/DC=terena/DC=tcs/...' so that the accredited subordinate ICA can be included in the distribution?
  It is *not* needed for a HLCA self-signed root, but it *is* required for the end-entities (even if the end-entity-issuing CAs does not need it itself, it should have an expressible non-omnipitent RPDNC namespace)

- in order to meet currently-deployed RP software, we discussed putting some constraints on the extendedKeyUsage extension, so that the certificates cannot be used for client authentication in a transaction. Not being able to limit that was one (the other) reason why LetsEncrypt was consistently troublesome.
  For client auth, robots (or other automated agent credentials, or token-based auth) should be used. It would mean that in the PKIX rendering of Elm, i.e. in the DCVOTA AP profile, in section 2, it would be good to add
  "
  In clarification of the IGTF PKI Technology Guidelines, the issued end-entity certificates MUST include the extendedKeyUsage extension. In list of values in the extendedKeyUsage extension MUST NOT include "clientAuth" (SSL/TLS WWW Client Authentication)
  "
  This essentially clarifies the context of section 3.4.3 in GFD.225, and the text in section 4.4 of the PKIX Technology Profile, where it now says "For end-entity certificates issued to network and service entities with an extended validity period due to the organisational sub-domain name ownership also having been validated, the ability to act as a client (i.e., extendedKeyUsage 'TLSWebClient') must not be asserted." DCVOTA is in that role ("organisational sub-domain name ownership ... having been validated"), and must therefore NOT have TLSWebClient.

  By explicitly making sure that these are not clients, the RP software will automatically enforce the 'correct' behaviour (i.e. refuse those cases attempting to *authenticate as a client* with a server certificate), while retaining full compatibility with public-trust browser access and automated ("grid") use cases. But it will prevent that suddenly we see clients whose actual owner cannot be traced, and you 'just' end up with a domain admin contact that knows nothing about the user involved - which would lead to a security incident response breakdown.

Digital Wallets in GN51 (Maarten Kremers et al.)

Please join the webinar at the URL provided (recording available).
https://connect.geant.org/2024/09/04/trust-and-identity-infoshare-wallets-23rd-september-1400-cest (https://connect.geant.org/2024/09/04/trust-and-identity-infoshare-wallets-23rd-september-1400-cest)

Token lifetimes and G081 (Nicolas Liampotis et al.)

Points for discussion:

- Which tokens to consider
    - More commonly used types or focus on OAuth2 tokens
        - Include “token protection tokens”
    - Focus on OAuth, put others to Appendix
        - Move ssh and similar stuff elsewhere
- Refine certain poinnts:
    - Bound
    - Revocable / Offline use
    - Accept out of band revocation
    - e.g. time limit on emergency suspension
- Support flexibility
    - If you need a longer lived token, allow for reducing the risk-profile of the use-case, e.g. by scoping down the token
- Risk based
    - Do you cross admin domain or not?
    - Can another site use the token? (Offsite needs more restrictions)
- OP can / cannot know what the token is going to be used for
    - THis is only partially correct. Ex WLCG knows exactly what happens
    - Delegated Proxy scenario can be addressed by out-of-band assurances between Proxy-A and Proxy-B (see slide 10)
- We should have guidance on the relation between CAAI and InfraProxy lifetimes
    - Avoid situation in which InfraProxy will be in place to extend lifetimes beyond sanity
    - B should probably not ask A for more scopes and info than necessary
    - Differentiate scope and audience

        - If a token can be used across sites, responsibility is higher
        - site requirements GFT32 (2004)
        - An RT of 1 y is not a risk for GFT32
        - 13months should be the limit

- Updating information from upstream
    - Proxy-A -> home-idP : Probably more static, but some communities (ELIXIR) require more frequent updates
    - Proxy-B -> Proxy-A: Information is more dynamic and might need more frequent updates
- Q1:
    - Q1: Should Proxy B dynamically adjust refresh token lifetimes based on the specific requirements of each client?
    - yes
- Q2:
    - Q2: How can Proxy B ensure that the RT lifetime RT with RT with issued by Proxy A does not conflict with the lifetimes MIN MAX required by Client 1 and Client 2? lifetime ® lifetime
        - A can always give max lifetime tokens to B
- Q3:
    - Q3: What mechanisms can be implemented to manage different lifetimes while maintaining consistency across the chain?
        - One could encode the lifetime in a scope
        - In federations trust marks might be a proper way to use it

Is low-medium-high life times enough of a differentiator?

- RTs (Refresh Tokens) can live quite long, since they can be revoked anyway
- ATs and ID Tokens are the issue
- Out of band revocation should be added
- Do we make different recommendations based on
    - How well do I protect it?
    - Do I get longer tokens if I rotate them? => Yes
    - Scopes, do they make a difference? => Yes
- Default values: SHOULD
- Min / Max: MUST

Validation and Adoption of policy guidelines

Validation and adoption planning with Peter Balcirak (AARC TREE WP4)

Deciding now on the framework for the validator framework, so it is now timely to review what the policy group will be planning for validation.

- Operational trust security baseline: can we combine operational trust and security baseline. What we prepared for EOSC. We want to expand that one to a baseline model; and adopt the baseline. Which is input for a potential validator.
- User centric trust: Excelent feedback from Australia; we really have to do some work. Should be tested, validated and will be a new version of the Policy Development Kit.
- Peter started to document the use cases, they could do a pilot there (in the life sciences).
- G040 is outdated: We need to work with use cases to make the document suitable again (Dave/David/Catharina/Arnout)
- SRAM is also interested in this; for SRAM it is important that is in the final stage, but we (SURF)can play around with it.
- Security officers are an important target group, too

How much would go to Aegis directly? How much in the WP?
Needs to be split up; not direct numbers now.

“this is from the DoA: On M6 the task will start with a pilot, in collaboration with WP2, for the WISE Acceptable Usage Policy and Privacy Notices interoperability across RIs and the impact to the user experience. In addition, focused piloting activities will be executed in support of the work performed in WP1 for OpenID Connect Federations and the Authorisation for Federated Resources.”

Start with an existing guideline, and move on from there. The OIDC Frederation and Cross AUP are ready continue …
The exact framework is still to be decided: automated validators (which is more suitable for architecture and technical?) which checks compliance, and the other one being more of a self-assessment and a (peer) review thereof. Use G071 for the compliance through the existing review sheep.
The spreadsheet would be sued to create the model for an on-line questionnaire and a self-assessment tool. Any checks that can be performed on-lne, those can be verified direcly. Other questions can be answered with documentation, and check whether the AAI instance complies or not.
Re-doing the WLCG IAM one with the tool and then comparing it would be a good test case [Hannah].

G056 > you can test attributes, this can be done; there is some basics tools for.
G071 > the review shee[pt]

Pilot for the privacy notices and AUPs (G040+) and now focus on he expected output and how to measure success. WHat do we watn to know from the profiles.
DoD (definition of done)

DOD: Consistent experience across different proxies (e.g. SRAM/Checkin/Helmholtz/Checkin)

This may require:

- for users: a questionnaire
- for the service providers: how many exceptions (or explicit interstitial screens) do they still addto the workflow?

Are we going to 'have something' by the end of the year?

- A peer-reviewed G071 self-assessment should be sufficient to trust other proxies
  - There will always be people asking for contracts
  - Maybe attaching a leightweight contract to the self-assessment is all it takes
  - the Baseline work (Snctfi) can be folded in here as an additional mark

The 'opaqueness' of the proxy is a matter of debate, with Snctfi peer reviewed self-assessed proxies having a role. Example in G082 my indicate where mutual trust is needed.

Distributed proxy scenarios and 'DAG' AAI flows

As in G082, the mesh of services and proxies makes it confusing ot the user it there is no directed hinting at the service level.
This could take several forms, as long as it prevents loops?

See the use case of e.g. TCS ("/customer/surfnet") and combined with Seamless Acces that works to direct the user to the proper home IdP via the community ('surfnet' in this case) without too many clicks. And SeamlessAccess is widely adopted, stable, and globally available.

And the SP can have other sources as to determie the 'favourite' IdP, maybe geolocated IP addresses, or commiunity-specific URLs, or ...

"All proxies are equal, but some proxies are more equal than others"

At the infra layer, the WAYF SHOULD NOT present (eduGAIN) IdPs alongside community proxies.

The trust demarcation is at the community proxy

Community proxy -> Opague
Infra proxy -> Transparant (but was does this mean?)

SP_E1 shoudl become splitted into multiple endpoints -> SP_E1_SRAM which is than directly connected to the SRAM community.

Discussion about the user-path-discovery ambiguities

- A service such as seamlessaccess.org (http://seamlessaccess.org) might be the key to help guiding the user through the proxy-maze
- Depends on users always using the same browser, which is probably always the case

Evolving notice management: G040 evolution and context

In preparation for the guidelines on notice management and the Operational Baseline (G071), a 'context framing' informational guideline (G082) is foreseen. This analyses proxy operator requirements and sets the scene for the notice management and baseline guidelines. It (not entirely by accident) also will be the basis for an AARC TREE deliverable (D.1) in May 2025.

- Discussed possible machine readible policies so multiple AUP's can be shown as one > could be in this iteration already as best practice
- G082 will feed in the deliverable
- Users should not click the same policy at many different places
- Services should not require the user clicking policy acceptance, since this will interrupt workflows
- Leaving the AUP solely to the Community may be heavyweight to them. A careful balance of Proxy-AUP and Community-AUP may be useful
- There is the `voPersonPolicyAgreement` attribute that carries a list of URLs to policies that the user has already agreed upon. This should be consulted by underlying services to spare the user clicking too many policies.

Update meetings: every Friday at 09.30 and 10.00 at SURF Utrecht and on-line. Send to policy@aarc-community.org (mailto:policy@aarc-community.org) the announcement of that meeting. (Catharina, Arnout, DavidG)

The items that do not make it in G08x notice manamenet, but should go in to G082 (and D2.1 for AARC-TREE) can be discussed biweekly just after the G08x+ update slot. With Maarten, DaveK added. (10.00 till ~10.30 every 2nd Friday).

DavidG to send to the mailing list.

SSH CAs by Marcus Hardt

Talk about ssh certificates work, and then how the KIT solution uses them. The tool 'oinit', as shown in the live demo, is inspired by 'kinit' for Kerberos. Both client and host certificates exist, and one can use the same CA for both purposes. From the 'host key' you create the certificate with `ssh-keygen`, giving the hostCA with the `-s` option.

In the host certificates, there are 'principals' defined, both 'ssh-server' as well as the host name. It uses TOFU (Trust On First Use). Compared to regular ssh keys, there is a validity period assigned to each certificates. For user certificates, the username is also bound in it as a principal, the extensions support specific ssh capabilities (X11 forwarding, agent forwarding, pty, port-forwading). The certificate can be retristed to the username, but it cannot be restricted to a specific host (but of course the host has to trust the CA in order to get in in the first place).

Critical options adds `forced-command`, so e.g. oinit-switch.

Marcus has a containerised demonstrator where the (EL8+) systems can use SSH CAs are shown working in practice. With both a client and server-side tool (using `sh oinit-demo.vm.fedcloud.eu` and a `@ca-certificate` entry in the ssh known_hosts file). Provided the OIDC tokens is available :)

This is also relatd to the DEIC solution, and a lot of information was mutually exchanged.

The SSH server has MotleyQ next to it, and it suggests the available OPs without authentication (this is public information). If you provision the accounts out-of-band, then MotleyQ is not needed (but then the CA has to present the proper user name as a principal).
The just in time provisioning is also there, using a variaty of attributes to contruct the username.

Token-based SSH has two major drawbacks:

- it needs a specific PAM module
- it needs a wrapper tool

The certifidate-based option also enables non-interactive use like rsync-over-ssh.

FIM4R at TechEx

Planning and presentations.

For TechEx Boston, 19th FIM4R on Sunday 8th, 2024:

- AARC3 outreach and presentations (DavidG, Christos, DaveK, Hannah)
- Identity Python stack (LIGO and NIH)
- Future planning (TomD, ...)

Question: “what is the role of the community?”, i.e. community first, or does FIM4R recommend evolution? What is the role of the proxy in hosting small (ephemeral) communities in a hosted solution?

FIM4R meetings after that:

- TIIME in Reading (UK) (in conjunction with AARC, maybe with overlap) on the 31st of March to the 3rd of April