The Abingdon 61st EUGridPMA+ and AARC Policy Meeting Summary

The 61st EUGridPMA+ meeting, a joint event with AARC (TREE) Policy, the IGTF, GN51 Enabling Communities, and EGI, is now over. Thanks to the kind hospitality of UKRI/RAL, Dave Kelsey and David Crooks, we had an intensive and interesting meeting with ample opportunity for trustbuilding - as usual. But of course the discussions are enabled by everyone engaging in the discussion, and I thank both the on-site and remote people for their stamina, and for making the workshop a success.

In this summary, we - jointly and collaboratively! - 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 https://eugridpma.org/agenda/61 and linked therefrom. These notes were 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.

The Abingdon 61st EUGridPMA+ and AARC Policy Meeting Summary
Next PMA & AARC meetings
EUGridPMA+ Welcome, agenda, minutes last meeting, note taker, introductions
IGTF trust fabric updates: status of authorities, fabric news, RHEL9/OSSL issues
- Self-assessment peer reviews and audits
- Signing key for the IGTF trust anchor distribution (PKIX technology)
- Operational challenges: automation
- Operational challenges: self-signed roots in Redhat/EL9
- TCS Updates

Developments in the Asia Pacific and the APGridPMA
AARC Policy Coordination and AARC-TREE: introduction to AARC TREE
GEANT Enabling Communities in GN5-2 and beyond
Planning for policy outreach
Table top exercises (TTX) for proxies and federations
AARC Community Survey: input for questions and context
FIM4R
WISE
- intermediate milestones and progress planning for AARC Policy

AARC PDK Feedback from the Australian Access Federation
Policy frameworks for PII 'as a result of Infrastructure use'
- collecting consent in proxies
- Recap and evolution of "G040" AUP and Privacy Notice model

Authorization and Tokens: updates from the GUT and the WLCG TTT
AARC Policy: token life time and revocation guidance
Eucalyptus or Elm, that's the Question (GTS) and TAGPMA update
Operational trust and Baseline - Snctfi v2
Jens' Soapbox

Next PMA & AARC meetings

The next EUGridPMA62+ meeting, in conjunction again with AARC TREE, GEANT5-1 EnCo, IGTF, and EGI will be held:

Monday 23 09.30-17.00 and Tuesday 24 09.00-17.00

at the Amsterdam SURF offices, Science Park 140, 1098 XG Amsterdam, The Netherlands. This will be close to the (parallel) European HTCondor Week then taking place at Nikhef on the same campus (Science Park 105). Logistic details will be circulated shortly, but for those prefering early arrangements, look for example at the Nikhef hotel list.

EUGridPMA+ Welcome, agenda, minutes last meeting, note taker, introductions

Round of introductions for local and remote participants. Present on-site were DanielK, JensJ, MarcusH, DaveK, MaartenK, CasperD, LiamA, JenC, DavidC, and DavidG. Remote joined BaptisteG, Adeel UR, EisakuS, MiroslavD, Nick Rossow, Jan Jona, LiciaF, as well as Cosmin Nistor and Derek Simmel for selected parts of th meeting. We received greetings from Ian Neilson.

IGTF trust fabric updates: status of authorities, fabric news, RHEL9/OSSL issues

The EUGridPMA sees continued evolution of the trust fabric, with some issuing authorities in the PKIX domain closing down, or transferring duties to new providers. We will be loosing the Portuguese LIP CA, and thank Nuno Dias for his long-lasting commitment to the trust fabric. Happily, we will continue meeting Nuno in other trust, identity, and other venues. The role has been taken over by the GEANT TCS as available in Portugal.
While some authorities are consistently engaged (like we heard from Dave and Feyza about Algeria in its self audit, others at times disappear from the radar (like KENET).

The Swiss users are slowly but steadily moving to eMudhra, although currently still using the legacy DutchGrid CA (being EGI members and affiliated to EGI.eu in Amsterdam). CERN is moving to TCS.

Self-assessment peer reviews and audits

(Cosmin Nistor)

- DZ (Algeria) has completed, updated the self-assessment and the new CP/CPS is ready and approved! Feyza and DaveK reviewed and got an updated CP/CPS. All changes are fine and all issues have been addressed. The name of the institute changed, but the people are the same and the OID has been updated.
- IRAN-GRID is waiting for funding to completely re-do the CA and get a new CA instance and do the self-audit. This should complet in 2-3 months from now (September)
- NorduGrid CA (Anders): Jens and DaveK will peer-review but are awaiting input.
- SlovakGrid: in progress, and Miroslav is almost ready. The next step is to meet with JanC and JensJ.
- KENET is still suspended, and they have not complained in the last 6+ months either. They will remain suspended, and on return they should be reviewed as a new CA. So cannot re-enter, and since the KENET Root also did not have an updated CRL for a long time, they were disfunctional anyway.
- KIFU is discontinued in an orderly fashion and review can be cancelled.

For BYGCA and TSU GRENA: their suspension is now imminent following lack of effective communications channels. DavidG will send a suspension notice with the suspension-review@ in CC. DaveK will inform WLCG on TSU GRENA.
Action for DavidG and DaveK.

Signing key for the IGTF trust anchor distribution (PKIX technology)

We will be changing the distribution signing key; the current one is "too short" for modern OS distributions to be installed by their standard package managers.
A new key has been out for about one year now, and will become the primary one after the end of CentOS7 (so July 2024).

Operational challenges: automation

The move to short-lived (90-day, maybe even shorter later) certificates is definitely coming, and automation is the only effective way to deal wih the issue in the logn term. Manual updates will fail. The use of ACME, esp. with external account binding for stability and for OV/EV, has been validated with several CAs, including TCS.

- Additional tools for TCS and ACME are available at https://software.nikhef.nl/experimental/tcstools/

The limitation to 90 days, pushed through CABF by (mainly) browsers, reflects the lack of adoption of CRL distribution and chechking in that domain, and the fact that - despite the possibility for OCSP stapling - the checking of OCSP status incurs a lot of round trips and hence delays in page rendering.

For the TCS and InCommon services, several automation options are possible. The long-standing one is the API, but also ACME is now supported well by the Sectigo issuance portal for automated issuance of TCS certificates

- some ACME endpoints are very powerful indeed (maybe overly powerful)
- GEANT EV and OV server URLs with EAB have be added and can be used

For clarity of discussion (also on e.g. the Elm profiles later), please *do not confuse* ACME, DCV, LetsEncrypt, 90-day certificates, and the ACME verification profiles (http, dns, eab kid/mhac). They are all different, and they can be combined in many ways. Conflating these *will* create confusion (sorry!).

Operational challenges: self-signed roots in Redhat/EL9

RedHat in its Enterprise Linux 9 (EL9) product deprecated SHA1 as an algorithm in almost all cases, *even* for self-signed roots where this makes no sense. This has been discussed extensively, and is not a fault of the roots or the IGTF distribution, but it is annoying for the maintainers of other products that get inappropriate (and unanswerable) issue reports (like https://github.com/dlgroep/fetch-crl/issues/4).

- Except for certain preset root certificates working with RH OpenSSL, to which proprietary 'trust bytes' have been added ("BEGIN TRUSTED CERTIFICATE")
- fetch-crl is one of the products collecting bug reports from users that do not identify root cause or remediations (where actually the ticket should be opened on the Enterprise Linux 9 with RedHat Inc.)

> But still, if you as a CA can re-sign with SHA-2 please do as it alleviates feth-crl support response load...

TCS Updates

The TCS communnity is currently doing a wide review of future possibilities, looking both at the scope of PKI services in the GEANT community at large, and at the implementation of the current TCS.

Developments in the Asia Pacific and the APGridPMA

The APGridPMA met as customary during ISGC24 in Spring 2024. eMudhra joined the APGridPMA in Sepember 2023 and is now providing also the co-chair, next to Eisaku as chair.
Also the AP region sees a storng move to new AAI models, with HPCI in Japan starting to operate a token-based AAI basedon OAuth-SSH, KeyCloak, OIDC-agent and JWT-agent - in coorpration with GakuNin. Also ASGC and KEK are moving in a smmiluar direction.
Meanwhile, SIFULAN (.MY) has started to odder IdP-as-a-service to encourage country-wide federated communities.

Next meeting:

- 34th APGridPMA will be virtual (APAN58 was not accessible enough)
- 35th will be again in conjunction with ISGC2025

Is the APGridPMA Slack channel open to members of other PMAs? Yes, it is!

What is the uptake for Sirtfi in GakuNin/JP. How can be enourage further uptake? Kento and Eisaku are working on this, but there are no updates yet. In general federations are slow to change - this is also seen in the UK with R&S to gain access to cluster computing in Bristol for instance.
There is always the IGTF-to-eduGAIN bridge.

AARC Policy Coordination and AARC-TREE: introduction to AARC TREE

Licia introduced the background of the AARC Trees. The introduction of AEGIS at the end of AARC2 enabled both sustainability and deployability of the BPA and Guideines.

David gives an overview of the challenges that still needs be tackled on the policy side of interoperablilty (ie simpliyfing the PDK, need for SNCTFI) and new challenges that have risen (stacking of proxies, how to tackle security incident respones with SIRTFI - preferably a simulated incident)

- Need to revisit SNCTFI?
  - IdP only knows it has authenticated the user to the proxy
  - Freshness can only come from the home IdP
    - Which processes are applied by home org when a person leaves - is there conversely a process for their keeping their account?
    - Proxy needs to check freshness every so often and not just rely on the linked accounts
  - IdP afraid of realising attributes to proxy instead of (end)SP?

- EuroHPC tried to push for assurance
  - Not much happened so assurance is now provided by external service
  - HPC as a "community" (UK's Isambard cluster (Bristol), Japan's HPCI, EuroHPC)
- Easing the acceptance of AUP
  - If a new service is added requiring the same AUP as everything else and the user has already accepted it, there is no need to ask the user again
  - This would be eased by (re)using the WISE baseline AUP - without additions
  - Maybe AARC TREE needs to work on this: can the service communicate its AUP to the proxy, and the proxy tracks the user's compliance
  - HPC historically unwilling to use WISE AUP

GEANT Enabling Communities in GN5-2 and beyond

The activities in GN5-2 for EnCo will do

- Policy & Technology
- Business development
- Wallets
- Mentorship programme (TIM)
- Communications
- TIME conference coordination

There is also - new! yeah! - a comms expert from Norway engagend in the GN52 project.

Maarten presents the disjunction of the GN52 and AARC TREE projects tn ensure there is both continuity but no overlap in activities, i.e. they are fully complementary.

GN EnCo provides the coordination for FIM4R/WISE/&c since they can provide the sustainability. AARC TREE contributes a boost, and uses FIM4R as a mechanism to get consistent requirements and therefore contributes effort there, but can never provide continuity since it is an output driven 2-year project.

Meanwhile, organisations continue to make sure that there is no double counting.

For GN5-1 there is an November deliverable coming up, which is an overview of all services, incubator and Enco, and combines all those outcomes. The EC accepted the previous instance as well, so should be OK. Also the GN5-1 review went smoothly, but there we no EnCo-specific recommendations.

Planning for policy outreach

The plan to disseminate AARC TREE, including both architecture and policy, at the TNC24 event is working out, and there is a 4-person-90-minute session included for Rennes. This will be akin to the outreach talk as we had in the AARC TREE kick-off (but of course with fresh content).
But if we want FIM4R requirements, we need to reach out first. Then get input.

TechEx with TTX exercises - but this one should be developed. A roleplay with exeprienced operators might be more instructionalto the audience. So that the participants are experienced and then can e.g. comment on what should be happening. This might be 'easier' to organise and smooth out the rough edged that were still present in Taipei.
But then how to do that outreach in Europe? SWITCH ran a maptable exercise in their federation (talked about it at the JISC NetWorkshop).

How to we get AARC across to more? Danielk submitted the TTX abstract also to the EGI Conference (Lecce, in September) - still pending acceptance.

Initially push the TTX (and baseline security guidelines), and then the new DPK.
Organise outreach to NFDI in Karlsruhe?

JISC's Networkshop 2023 had a presentation from SWITCH with security training using games scenarios. Unfortunately the programme has disappeared but we can find the name of the activity.

International Conference on Research Infrastructures? (note, just before TechEx and in Australia) https://icri2024.au/

The challenge is to talk policy in e.g. plenaries or in *other* sessions, and not have specific tracks since then you keep talking to the same people. Talking in other tracks is more effective. Was done earlier in the EGI conference bringing security to the people'.

A more generic issue is how to find the meetings that the target communities meet. What is 'in it for them'?

Table top exercises (TTX) for proxies and federations

The role of exercises for federated IdM has been clear since the first realistics ones for Sirffi v1. WIth more proxies now part of the scheme, there are many more parties involved, and hence more complex.
So can we design or enhance exercises, togehter with eduGAIN CSIRT, and get input both on Sirtfi v2 as well as a handle on proxy complexity.

What kind of exercises will 'exercise' (composite) proxies?
What we can improve on the Taipei TTX to get more federation operators involved?

- learning goals included making the role of the proxy *operator* clearer. We want those surprises to hit home. Anticipating who is the audience will be helpful.
- fully -scripted roleplay (like at the EGI Lisbon conference) is more helpful for the audience
- Cornelia Puhze did something similar for SWITCH, designed to key federation participants up to speed in case of an incident: https://www.switch.ch/en//piece-cake
- Would an OP detect misuse of an account (as it would see all the clients)

The duration of the TTX, which was 2 hours in Taipei, is probably the sweet spot for these types of awareness and comms trainings. Keeping people busy for four+ hours is really difficult, so 2 hours is probably a good duration to target.
Maybe depends on the event.
A dedicated "CLAW" like event (2 days) is probaby too large for extending this scenario, s you don't get the right people to attend such an event.

What the proxy operator adds here, is that the proxy operator can check with its own SPs, but conveying the incident upstream is commensurately complex. So the proxy helps in scoping and getting (rapid) reaction from peer SPs.

The timeliness of information propagation is also an issue that needs to be addressed. How much delay do proxies introduce?
Can VOPersonEternalIdentity help by adding any intermediate identities in the assertion blob. The hierarchy will then not be known, but at least you have a list. Those can be targetted for communication in parallel. But there are privacy and tracability issues. Hashing would make it

loose all value to inciden response, even if it is privacy preserving and would still allpw identification of login loops &c.

This would apply for identity *as used for access to the infrastructure* (i.e. 'access' in CoCov2). There you don't ask for attributes you don't need, but here the tracability is for incident response, which is a legitimate interest. This might also be true for account de-duplication. This will probably help with the HPC use case.

Next slots:

- Tech-Ex: Dec 9-13
- EGI Conference Sep 30 - Oct 4

When this converges on a cossistent set of cards, turn this into a kit that eduGAIN can also use in emerging and wider federations. But also in those ccountries and regions some SPs and IdP, next and proxy operators, should be involved next to the federation operator.
With EnCo/eduGAIN to run a 'train the trainers' programme globally? Just like TRANSITS

Also TRANSITS-II is being started again, and it was full within 15 min. So there is demand for those as well. Maybe part of a specialisation progamme suite with also federation trainings, crisis-comms training, &c?

As for the presentation of the TTX:

- there should be a more solid introduction, so that the audience will understand what is intended, and what will be required of them. "A common view of what is intended to be achieved"?
- feed-back form for participants - for ISGC this is not available

AARC Community Survey: input for questions and context

The AARC TREE project provides for effort for an in-depth survey of Research Infrastructure requirements (supported by the Use Cases activity WP3: "*This work will use as the starting point the FIM4Rv2 paper together with requirements that AARC TREE partners may have collected via other activities. In addition, it will engage with relevant forums and stakeholders (such as FIM4R, AEGIS, EOSC AAI Task Force, National RIs and European initiatives such as the EU dataspaces) to gather the initial set of requirements and use cases. Based on this, an initial set of the requirements and use cases will be captured, to drive further work.*"

An updated questionnaire (accessible to AARC TREE participants) is available at https://docs.google.com/spreadsheets/d/1NXj0T0u2JkC_-htaDE7_whVcgTcZKL8amB42NuqTYds

Action:

- Policy input by beginning of June (~TNC time)
- Interviews will be over July and August.

FIM4R

Maarten summarised the FIM4R development and meeting cadence, as presented in the attached slided.

- FIM4R at TechEx will be on-site only (FIM4R19 on December 8, with mixed capacilty at TechEx not yet available).
- FIM4R at TIIME (early 2025, in the UK - Reading, Manchester, or Jodrell Bank to be decided).

How to we ge the right people in the room (from the communities, not only the proxy operators) and how to get the answers out?

Can we actually make the communities (at least the funded ones) to actually prioritise the (policy) requirements that come out of the Use Case questionnaire by AARC TREE WP3?
Analyse the outcome of the survey and takl 1-on-1 wth people/communities for further information. This is also "FIM4R", but in a specific setting to get to the long tail communities that usually do not attend the meetings.

"We need to go to their venues". And then: via booths and exihibition at conferences? This is also a way to share the compendium at the end of the AARC TREE project. And include stuff from EGI, GEANT, ...

SRAM stakeholdr group as a source of community input.
Maarten will call for a meeting with usual suspects and discuss further.

WISE

See Dave's slides on the agenda.

The Wise Steering Committee will be conveived by Dave to discuss the status and future of the working groups (and which groups should take on which work items).

intermediate milestones and progress planning for AARC Policy

- Start with a ToC of the AARC documents end of June/early July
- Milestone easrly september
- writing completion during the 23+24 workshop /PMA62

AARC PDK Feedback from the Australian Access Federation

Opportunity for input and feed-back from our Australian colleagues. Do we need to revise and 'template' terminology? What should the new PDK structure look like, and what is the role of the 'top-level' policy document?
And are all things actually policies, where some are more like procedures, and some information guidance or a glossary?

Dave Kelsey et al. summarized the feedback from the Australian Access Federation as they reviewed adoption of the AARC PDK and the challenges and new ideas they encountered.

The've created a security policy group with membership from all important stakeholders in Research Australia. And they presetned the work of the PDK there, and asked for feedback. This is also how to engage stakeholders from the research communities.

Difference between policies, procedures, and concepts across the PDK.

- they used the 2019 version, so not including the service operations baseline. The updates were not publicised enough and thus were not found.
- VO membership management policy was full of procedures
- incident response was also more a procedure

Terminology definition

- PII vs personal data. In CoCo that is 'access' data.

- "Community" is not clearly enough defined. It in the end it' a 'grouping of users', but the word does not resonate with the .au community, where "VO" was more useful.
  For a globally applied PDK, it is unlikely that a single glossary will be possible. You need a template to full this in. "Whatever you call the community, put it in [here]". For .au, "Collaboration" was the word that 'worked'.

Also NFDI is also diversing from the original terms and definitions [MarcusH]. Something similar happend in UK-IRIS, where the communties could not agree that they were a community (in terms of terminology). So "Put in Your Work Here"? Community, Collective, VO, Project, Group, or just "Users". One distinction here is that the communities does not include the infrastructure resources as an essential element. (although the community may have services, like the Community Proxy as a component) Which is different from the original US/Globus concept of "VO" which *included* the resources.

Also the SKA (both Telescope and SRCNET) have yet to agree what the proper term is (and they have different assurance levels, even if the data moves from Telescope to SRCNet once observation is complete). Since it is partially controlled in a different way, it gets to be more complex that, say, WLCG.

Even if you do not participate in the work, you should still be kept up to date on the outputs and upgrades.
The Policy Group in Australia was about to endorse the use of the PDK as a basis for further evolution. And the engagement and the feedback are excellent!

For the small communities, they are 'self managed' and they do not 'see' the concept of a community, since they are more meshed together in a 'web of users'. That also depends on the workflows. This is more like a web of trust and FooF, but then thre is a Principle Investigator (PI) who may hve a stronger role. And these tends to become more formal, e.g. once a grant is involved. Then the organisations (universities) will be a the formal participants. The same holds for access to HPC, where this is linked to the PI. However, the PI does not usually do any 'real work'.

Even in the large communities, there are subgroups that have their own scope and authorization grants, even if they are subject to oversight (and the ability to fix in case of urgency) and access approvals. This is again a policy question, and the question of 'who owns the data?' [There is experience with authZ for small communities, like Jens provided during the meeting]

What about sensitive data? There a formal structure (and/or trusted third party that is managing the community) required. For special cases (emergencies, leave) there should be 'break the glass' procedures that are again subject to authorization.

The Australian Access Federation research communities are (probably) also the larger ones, which may have substructurebut they are at least large enough to talk policy to the experts. It would be interesting to identify smaller communties and how they 'react' to the PDK structure.

For 'new' infrastructures, deploying the whole of the PDK in one go will never work. Starting with a few (security baseline, privacy) and then gradually augment the set. And: do not read the policies, but just trust they are the 'good things' you expect it to do.

So what about the top-level policy? What about 'princples'?
WIth just a few foundational principles (like the Nikhef ones for digitisation) where they are the "Princples of Collaboration".

The AARC TREE PDK may have several variants dependign o the desired collaboration structure. And like UK-IRIS needed to describe 'what it is'. **What Are We**?" as a fundamental question to answer.

Collaboration *governance* next to *policy*. Which is a prerequisite for policy and security. And this should then lead to a document that empowers a security officer to act. But this is again for large communities, not for a 10-person meshup. That also depends on (data) risk, where for high value data people will nw have realised it, but for smaller risks, also local home org/university policies will get ignored jusy to get it to work (before the researchers retires :).
And there is always a weiging of interests, so even if the stick is there, it is mainly to 'encourage' behaviour, not to use in anger.

At some point there will be a 'proxy for the long tail' (in NFDI, as we see in eduTEAMS) where all the policies are standard, but there is no 'choice' per community. The smaller the community, the more needs to be standarised. Like in SRAM as well.

Actions emanating from this discussion:

- capture stories from smaller communities
- separate the PDK into policies and procedures as ancillary documents
- look at the top-level policy to see if it recasts into foundational principles
- split the policies (and procedures) by topic area

DavidG: review and draft something by July.

Policy frameworks for PII 'as a result of Infrastructure use'

The model today in EGI is based on Pretty Binding Not Quite Corporate Rules, and used for sharing accounting data. The fact that is is correlated with (use) information and logging so it more than just the subject name (sub, or DN).
At the time we were assuming that DPCoCoV2 would be a GDPR Code of Conduct, but that did not quite work out due to the monitoring body. Now the current policy on PII in EGI is referencing the Directive, rather than GDPR.
We could just change it to refer to GDPR, but should we improve it? And it may be suffering from formatting issues in the PDK.

This is about access to PII for accounting that is shared between sites, which is transferred around. But the discusion in the WIE SCI WG failed since there were voices that were legal purists who preferred consent, but we know consent does not work. And also contracts do not scale, do the legal approach will in practice not work either. And the risk is limited anyway, so the BCR-like model is the only one that worked.
So if the SCI WG process does not work, we should do it elsewhere. Since WLCG is using this all over the place, no only in the accounting records. And the operations team in WLCG were going to find out where personal data was used, but they have not come back since one year ago. Can we just share aggregated (per VO per country per gruop) data rather than personal data? Might still be enough for showing use of the infrastructure!

For the EGI ticket, with we are periodically reminded of, we will:

- **just change the refs to 95/46/EC Directive to GDPR** (minor change)
(it worked for years, so is there a real issue?)
- and the BCR-like approach was inspired by the GDPR anyway

In AARC TREE we should come up with an updated policy? Alongside an updated privacy notice template!

Would a joint controllership model (like the DPO propsoed for NWO-I, as used in the VW group) work simlers and pretty much like the BCR-like model? Or is that even more paperwork? It is still too legal, with corporate lawyers getting involved and with too mny bilateral interactions. And consent is gaining traction again (even if it is 'forced' consent).

"some commercial IAM services already have a single tickbox for accepting terms of use and the data privacy policy" [Nicolas]

There is no monitoring activity - and what happens if/when something goes wrong. CoCo is self asserted.

In REFEDS & GEANT the attention for GDPR seems to been laid to rest for a whole, probably because we now hae NIS2 taking centre stage, and the Data Act, and AI act :) So we only need to worry about the latest regulation :)

Still the discussion withthe GEANT DPO (Ana Alves or Magdalena Rzaca) might still be useful if the output is useful and scalable.

collecting consent in proxies

If the user revokes consent at the community proxy, how does that propagate to the infrastructure proxies?
What can we use in terms of 'delegting' the collection from the resoruces to the proxy/proxies as a 'service' that the RPs outsource to the proxy (for whih they will then be informed and/or collected and pushed to them).
Will depend on the community model (WLCG and HPC will be very different here).

Recap and evolution of "G040" AUP and Privacy Notice model

AARC-G040 "preliminary recommendations for the LS AAI" presented initial ideas on how to show terms-and-conditions and privacy notices for dynamic proxies. What does the current proxy landscape look like, and what are the current practices, e.g. in SURF SRAM on triggering notice presentation?
What should we keep, and what should we question in G040? Whom to ask for requirements, and how? This item was not further discussed directly during the meeting, but links to the consent collection and privacy notices.

Authorization and Tokens: updates from the GUT and the WLCG TTT

Matt Doidge - "Introduction to, and update on the Token Trust and Traceability WG"

The Unified Token Profile and the WLCG Tansition To Tokens (TTT) working group are progressing. Matt Doige gives updates on https://twiki.cern.ch/twiki/bin/view/LCG/WLCGTokensGlobusWG and (potentially) Mischa Salle on the Grand Unified Token profile.

The Token Trust and Traceability WG (TTT) was instantiated last summer from members of the WLCG and EGI communities. A somewhat "spiritual" successor to the successful Traceability and Isolation WG, but not (solely) a WLCG Working Group, and aiming to produce outputs usable across communities. It is working in conjunction with the Grand Unified Token (GUT), and WLCG AuthZ WGs.

There is no sudden change of requirements with the advance of tokens:

- Authorisation decisions need to be traceable.
- Authorisation tokens need to be "triggerable on" (and some method of banning)
- Users and Admins need trust in the system.
- Guarantees of security and integrity

and this needs to be documented, since it is hard to trust what you don't understand. And at the moment most of the token work is developer-focussed, rather than being focussed on production (as seen in e.g. WLCG's DC24).
This should have been done already before, but for the moment need to demystify tokens, and find out 'what to look for' in logs. And then follow-up in incident response. Basically translating what we know hw to do in production into the token work, with recommendations and 'best practice'.

Tokens also do tend to end up in places where they should not be shown in plain text, e.g. in logs where they are 'secrets' and more than just user info. With users sharing access tokens as part of debugging, where it is not clear that they are confidential. Sharing in redirect URLs is similarly bad since GET requests are logged in full in log files (and in trouble tickets)

Also, MaartenL distributed tickets to a bunch of sites to review configuration and fin out how the issuers are configured and collect information. There are some issues around distributing issuer trust anchors, and where to securely source them from. By using GGUS at least you have a trusted line, but it does depend on authenticting the originator (in this case Maarten's account).

But at the same time an opportinity to close old loopholes and get more secure, and fix e.g. overly-long life times for tokens. If that works out.

Please get in touch if you want to get involved: m.doidge@lancaster.ac.uk
and https://github.com/TTT-WG

Now there is an opportinity to do it better, but making that work out more securily is also a challenge (no 'temporary solutions' then?). Token life times are not only informed by many constraints, and the resilience of token isues and the current 'dev oriented' deployment in their current state. But without work on making resilient token issuers and investing effort in making those work, users will pressure opertions to bypass security sine the tokens issuers don't work reliably.

But we know that client certificates do not work easily in browsers, but they are fine in the back-end. But WLCG has committed to tokens regardless of their deployment current state and reliability. The immature token capabilities were shown in the WLCG Workshop (as in DC24) was useful in clarifying some issues, but there is an appetite there from the developers there to have a more mature architecture. But the architecture at the moment for tokens is just not as mature for the next couple of years (with some moving targets).

Are there operational solutions that do not involve developers (like horizontal scaling, multi-site distribution, anycast, stability of platforms and HA). Those are being looked at, partially at the sysadmin level (see the slides from WLCG IAM workshop), but scaling it still evolving (e.g. to K8S). The hackathon on Indigo IAM is ongoing today in paralle (29-30 May '24).

See https://indico.cern.ch/event/1369601/contributions/5923591/attachments/2855860/4994809/IAM@DataChallenge24.pdf

But it needs wrk from all sides, also to look at more efficient token workflows. So not using that many tokens to begin with. With a balance betwene the CMS superpowerful token, and the LHCb extreme of having lots of very narrowly-scoped tokens. Somewhere, there is a balance. Without revocation you don't want supertokens, except when they can be limited like in MyToken (MarcusH). Or address that one by scoping limitations (with 'bounded' tokens). And there are soo many different tokens with different capabilitities (refresh, access, mytokens, &c&c) , that at the operational side and in the community there is not enough understanding to engage at an appropriate and safe level. This one has to be fixed.

New opportunities with trust marks in OpenID Federation as well, when used correcly. How would be e.g. use trust marks in SKA to deal with differences between countries (even if now there is just a singel thing). As seen with e.g. FTS still not working for SKA with tokens (even when working with pre-releases).

Reducing complexity, like we did for VOMS 20 years ago - which has many capabilities that were never, ever, used!

And an automatic sun-set clause for exceptions (like the 72hr proxies that never were reverted).
This would be a novum for WLCG ...
And have this tested by the next SSC Security Service Challenge.

What we can do now:

- engage some more key players to join the TTT from IGTF and AARC
- part of EnCo and AARC

For PKIX proxies the policies and best practices were kind-of piecemeal, but there is no comprehensive documentation. The TTT group is now doing the holistic approach. And get everybody ivolved in the discussion, both security and operation. Robust conversation is good, as long as the outcome is high quality and constructive!

AARC Policy: token life time and revocation guidance

(Nicolas and Marcus)
Slides: https://docs.google.com/presentation/d/1N9w-7-a8Ulhh0_wT2BZSYUMn9Ojo2VLE

Review of tokens properties and their verification models, including the various token properies, and the blance of refresh vs revocation vs validity periods.
MyTokens have live time depending on the subject and the scope in combination,a nd you can encode all of this inside. This is like an extension to Refresh tokens in OIDC - would be nice if it was picked up there as an RFC extension.

On the lifetimes (slide #7) there is not yet an AARC recommendation, but that would be a good policy outcome from the policy team. We could at least capture current practice and then share community good practice.

Lifetim based on risk analysis (like GFD.030 with Dane Skow) where it was inspired by the security requiremetns. Also depends on the use case, where e.g. WLCG use off-line varified long lived tokens, where revocation is impossible, and hence the isk a lot larger. The phasespace is huge. This presentation gives the comprehensive overview, and the mitigation measures nedd to match the risk assessment. The deployment model is therefore critical for the validity periods (scope, audience, verifiability, token-type).
No one-size-fits-all.

Refresh tokens needs rotation, and e.g. public clients need stricter rotation policy. And other other balanced in slide #8! Or auto-extend as it is being used (which is equivalent to inactivity timeout).
And Risk-Based-Access model as an addition to inactivity timeouts (which is implementation decision). So take into account Geo-IP, client user-agents, inactivity, context, and access patters (AI/ML assessments)?
And have validity depedent on the model used?

And can you make revocation actully trigger notification and action on the RP side (like stopping and inspecting workflows and running jobs).

- Can we encourage implementayion by making this best practice a distinguishing factor between implementations? So customer pressure on the solutions?
- Document good practices what communities are currently doing. Via the questionnaire in AARC TREE WP3 or ask that later. What do the communicaties have **configured** with the AAI providers? Doing the one-on-one interview with teh policy/security experts!

Eucalyptus or Elm, that's the Question (GTS) and TAGPMA update

(Derek Simmel)

(see slides from Derek and the updated Assurance Document).

Draft ELM profile integrated with IGTF assurance document.

- Similar in many ways to DOGWOOD
- Need to ensure unique issuer namespace

And we need unique naming, as per the IGTF Federation document:
"3.1 Management and communication of identifiers
On accreditation, a specific subject name space or set of subject name spaces is allocated to each PMA member for its accredited authorities. This name space must not overlap with any existing name space already assigned to an existing PMA member for any AP, assigned by any of the regional PMAs within the International Grid Trust Federation"

The only allviating circumstance would be resirtction in key usage and hard-forbid TLSClientAuth. But we need the uniqueness. Like "/DC=com/DC=googletrustservices/DC=gcc/CN=*fqdn*". It could even come from their existing CA, since it is an RPDNC constrain on the RP side.

- This would be the host-equivalent of RCauth for users.
- it should not be called ACME, since it is actuallt just "constrained DCV", and would not be ACME
- similarly, things like 90-days are not ACME, but implementation (LE/CABF) specific requirements.

Naming: the PKI rendering of ELM could be "CDCV" or "DCV-only".

Can the assurance profile not just refer (also) a specific version of DCV as per CABF guidelines to make it easier?

Anyway, once the DC/prefix naming is addressed, the rest can be fgured out. The keyUsage limitation is HighlyDesirable(tm).

But in the naming, make sure AMCE, ELM, DCV, and the ACME http-01/dns-01/EAB profiles are clearly distinguished.

Operational trust and Baseline - Snctfi v2

(DaveK, DavidG, *)

There is AARC-G071 AAOPS, there was the work from UK-IRIS and EOSC Future on the Security Baseline that should be ported back to the "SCI" community, and the inforporation of Srtfi for proxies. This could be the effective evoluion of Snctfi.
Now Snctfi does not quite incorporate the federation concept completely, and it does not inclode the upgrades the Sitfi-v2 saw since.

The transitive trust issue for proxies in a federation, so what does it mean when a proxy makes specific EC claims. That is a complementary aspect of Snctfi (as discussed in Denver). But in Denver also there was no consensus since (CK) make it clear NOT to make proxies transparent - they are authoritative for what is behind them and they should not be poked through. But then could/should Snctfi be asserted as a trust mark on the proxy? Or better not?

As an IdP, there are mixed feelings since some do want to know the whole chain is trusted. On teh SP side: "that's none of your business!". Both of which are true. And decisions are the based both on actal risk and on the benefits/administrative push (as seen with some mixed SP proxies).

- "does Sirtfi apply throughout the back-end of the proxy?"
- "will the services communicate with me in case of an incident?"
  So there is use for a trust mark to be applied to a 'nice' proxy.

But if all SPs behind a proxy adhere to a Baseline (like the IRIS/EOSC one) which is common senseand widely accepted by many service providers, would that then not constitute a new "Snctfi" for the proxy? And if you do ISO27k/27k10, would this be 'implicitly OK'? Or is that too much of a shortcut?

- EOSC Security Op Baseline attached to agenda ...
  Because too many different standards will not work. The *multi-management domain* aspect is the most relevant one here.

Snctfi should say something about proxies, and their different models:

- the community as responsible for linking services o their community in a proxy-as-a-service model
  - What if those services require acceptance of different policies beyond the baseline?
  - Should the proxy track this or the individual service?
- operator of the proxy as to the guardian of connections
- Does that enable the research community to do their work?

Would 'Snctfi' be the service operator/proxy side of the Policy Development Kit?
The Baseline policies will then be the ones needed to satisfy Snctfi.

And the value of services behind the proxy makes many complaintsto away. Be that the Commute reimbursement service, or the NIH granting systems. Then suddenly the IdPs become a whole lot more lenient! :)

For AARC WP2 Policy (service side):

- How to do Sirtfi v2 across collective services, with a single proxy interface
- Baseline Operational Policy for services
- Working across multiple (community + infrastructure) proxies: what do they need to agree on beyond the baseline?
  This would be "Snctfi v2".

On the community sie (task 2): the community should also participate in incident response, part of thei "community-oriented" PDK.

There is an AARC architecture informational guide on how to technical trust proxies and how thy talk to each other (AARC-I058). Of course, there are both technical and policy aspects of how to establish trust.

On the trust model side, thinking of OIDFederation and separating technical bridging and trust bridging: what are the model options, even if we later decide not to separate these roles. If we want a proxy to be opaque (like CK recommends), the you certainly don't want to separate trust and translation. Only once if we have the models clearer, can we recommend one or the other model.

Try and align with Christos before going to TechEx ... :) That would be between now and the M15 deadline for the deliverable in AARC TREE. And present to AEGIS as well, of course.

Jens' Soapbox

"Making use of the security exeprtise we have for the work that we do"

Can we use the token transition to make things for secure than today when we upgrade technology? First we make it work, then later we make it better, as a project progress moethodlogy to use incremental improvement. For example, analysing ip access behaviour to a service and after a whole look for who was legitimate and block the rest. So all the 'others' are firewalled away, and the scheme becomes more managable. But improving security takes effort, and doing that incrementally *also* takes effort and work later, not only at the start.
But "security does not look sexy", and you don't see anything if there are no issues. And as a result, you get feature-rich products, with a security-last approach - as it was not "nice to demonstrate" at the end of the project.
Now this has changed, and security has become more important in the perception of people.

Now, for CA roll-over and extension, people look at their own retirement date and set the validUtil date as appopriate to be 'just past that date'. So do you extend. Or rekey?
What about waiting for a working quantum-resistent signing/key algorithm whic may beome more prevalent in 2028+ or so, and rekey an expiring UKeScience 2* which now expired in 2027 for just a few more years and keep the same ke for now? And then more to QC resistent soon after?

How to focus attention of people to work in shared challenges in absence of hard deadlines and project deliverables?