51st EUGridPMA meeting (virtual)

Monday, 8 February, 2021 09:20

Dear all:

Thanks to all those that joined the on-line sessions of the 51st EUGridPMA PMA meeting!
I'm glad that the some of the spontaneity and sparkle of the in-person trust building remained also for the video-participants. In total 30 people participated, spread over all time zones and from all our regional PMAs. Plus a large attendance form the AARC, GEANT Enabling Communities, EOSC ISM, and IRIS communities. Thanks for joining!

The next meeting will be the 52nd EUGridPMA meeting, again joint with our sister communities, and - unless something changes for the better - again a virtual meeting. It will be three half-day sessions again:

Monday June 7 from 13.30 - 17.00 CEST (11.30-15.00 UTC)
Tuesday June 8 from 09.30 - 12.30 CEST (07.30-10.30 UTC)
Wednesday June 9th 09.30 - 12.30 CEST (*ibid.*)

In this summary, I'll try to give an impression of the main discussions and results. As usual, much is also contained in the slides and ancillary materials and documents that are attached to the agenda at https://eugridpma.org/agenda/51 or linked therefrom. In this summary:

- APGridPMA meeting and ISGC
- Certificates for container-based and self-service entities - CERN CA update
- WISE SCI guidance and the Baseline AUP
- Enabling Communities - eScience Global Engagement
- Attribute (and other) authorities Operations Guidelines (AAOPS, or "G048 rev2")
- OpenID Connect Federation development
- Assurance Profile evolution and implementation
- TAGPMA and WotBAN&AZ Token-based authorization
- Jens' Soapbox - complexity has to go somewhere
- Operational matters & self-assessments

APGridPMA meeting

Everyone is welcome to join the upcoming virtual 27th APGridPMA Meeting:

- 11am - 3pm CET on 23 March - co-located with ISGC 2021
- the ISGC Security Workshop is scheduled on 22 March (2 sessions)

Certificates for container-based and self-service entities

Hannah Short presented the updated CERN CP/CPS that aims to better accommodate container-based services that are not really 'machines' (real or virtual), but rather containers that get spun up in an OpenShift container management service. Such containers, like an Indigo IAM instance as 'wlcg-iam.web.cern.ch' have a proper domain name, but would appear to fall outside the host definition since they do not appear in the local (CERN) network database.
This is actually common, and the DCV process for example does not distinguish it, as long as the domain names used of course are properly controlled - so for a web-service based container like wlcg-iam.web.cern.ch, the owner of web.cern.ch (and/or cern.ch) would have to allow that.
If so, the term used for these can be many things - industry still revers to Server SSL certificates even if things are not servers, nor using SSL (but TLS). And the discussion on what a subordinate or an intermediate CA is, is equally muddled (as seen on the LAMPS list).
All CERN container names will be domain-name based (so not use the legacy Kerberism *service/hostname*), and this is perfectly fine.
The changes proposed throughout the CERN CP/CPS (as attached to the agenda) are approved as-is,

and the new CERN CP/CPS can take effect immediately.

WISE SCI and the AUP

Various groups jointly work on the aspects of SCI, and the alignment between AARC, IGTF, IRIS, GN43, EGI, EOSC, WLCG, and others is very close - addressing the whole range of SCI issues:

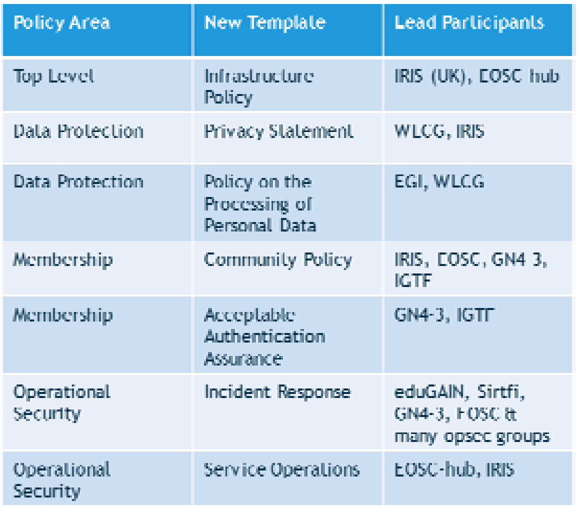

Policy Area	New Template	Lead Participants
Top Level	Infrastructure Policy	IRIS (UK), EOSC hub
Data Protection	Privacy Statement	WLCG, IRIS
Data Protection	Policy on the Processing of Personal Data	EGI, WLCG
Membership	Community Policy	IRIS, EOSC, GN4 3, IGTF
Membership	Acceptable Authentication Assurance	GN4-3, IGTF
Operational Security	Incident Response	eduGAIN, Sirtfi, GN4-3, EOSC & many opsec groups
Operational Security	Service Operations	EOSC-hub, IRIS

The AUP version 1 is no longer a draft, and is now available from the WISE web site - the version from Feb 2019 is materially unchanged. It is not directly uploaded (for technical issues with the WordPress site) but the WISE web points to the proper version on the GEANT wiki.
There is a slight discrepancy with the TrustedCI template, which would need reconciliation in order to be fully compatible (specifically the clause 10, using "responsible" instead of "liable" which maybe has a specific legal US meaning to some readers there, although Fermilab was fine with the WISE version?) - work is ongoing in this area. The US is still discussing, updated in the coming time.

WISE SCI implementation guide

The WISE SCI v2 was endorsed by many, which is one of its strong points - although not everyone may remember. It defined five areas with specific requirements, and each of them can have a maturity level associated with it (0..4) that can be used for self-assessment. The self-assessment can also be used an inspiration for peer infrastructures if these are shared within the (trusted) group, extending on the Policy Development Kit, in a Google document (attached to the agenda).
While terms are re-used if relevant from similar frameworks (NIST, ISO), but SCI does not specifically redefine terms (using more their dictionary meeting), and does aim to be a self-contained document. The ISO/NIST frameworks does target single administrative domains, where SCI also specifically tries to address federated environment. The collaborative environments does not make a good match and, say, accrediting the EGI federation to ISO27k will never be a good match. SCI does aim to address to those issues.
So for example, information sharing for incident response does make an important element for SCI in the 'collaborating' ("C") sense.
But of course SCI should not be confusing terms with respect to their NIST/ISO definitions.

The document was revised and improved during the meeting, starting with section IR1 (Incident Response contact information). Anyone with the link can comment!
The document should not turn into the ultimate reference guide, so for some more detailed guidance (like on how to do incident response, IR2) it need not contain a comprehensive procedure nor a complete set of references. If needed, readers can either look themselves for courses like TRANSITS, talk to SCI, etc. "Talk to your local security team" (if it exists) does not quite help since the audience for the SCI document is (also) the local team. The same applies also to some of the OS requirements like the security plan - so an introductory comment ("what does this document help you with") at the top may clarify this scope.
For IR3, since people are not likely to converge on a single tool for collaboration, the text is necessarily a bit generic. Some of the mechanisms are actually in place and tested (like for the Sirtfi exercises), but most of that is either related to eduGAIN or (on a smaller scale) the IGTF RATCCs. But adding those may well be too detailed, and the Sirtfi one is specifically scoped to eduGAIN as well.

We could add endless links here, but is that really helpful? In the "how", the main aim is "To be consistent with the SCI framework should be able and willing to collaborate with others", with "able" means having both the tools/resources and the mandate to act, but also specifically including 'willingness'. So "authority, resources, and sufficient priority".
And you cannot have IR3 without IR4 on information sharing, and you need to explain how to then share information - which is what all the wording there provides (including "TLP").

Enabling Communities - eScience Global Engagement

Enabling Communities is the outward-looking activity of the GEANT project when engaging with research users, including well-known activities like FIM4R, WISE, AEGIS, the IGTF, and REFEDS. Trust and Identity Outreach captures both the (internal) 'business development', but also AEGIS, and broad AARC community support.

- the REFEDS Assurance Framework (RAF) has been published, and its promotion is taken care of by Jule, who is also leading the work on journal papers to make it better known and provide implementation guidance, examples and use cases, and best practice. There are challenging use cases including the US NIH that needs assurance signaling for their applications - and they could be the 'launching example' for the RAF. You need SPs requiring it in order for IdPs to consider adoption.
- WISE SCI is working on the guidance doc (as we already reviewed during this meeting)
- Sirtfi has inspired the eduGAIN Security Incident Response Handbook (being formalised in the eduGAIN Steering Group), which is co-supported also by EnCo
- the Policy Development Kit top-level policy has been progressed also based on UK IRIS experience, and that is being fed back into the PDK document suite
- AAOPS - has its own session in this meeting as well
- SCCC WG inspired the eduGAIN federator-level comms challenge executed by the eduGAIN Security Team
- The Blueprint Overview graphic designed by Hannah bring together all guidelines in a single leaflet centered around the BPA architecture graphic as part of T&I Outreach. This has already been very well received and gives a great overveiw to the AARC document suite.

Meanwhile, GN4-3 is halfway its project life time, and the workplan will be updated in a dedicated working session later.

Attribute (and other) authorities Operations Guidelines (AAOPS, or "G048 rev2")

[https://docs.google.com/document/d/1-hbqSpQegm7UaC_wupFzFMm19Q024UPkG-8Jwokkmzc]

Splitting the document into a separate one for push and one pull. By providing different profile docs it becomes clearer to pick specific guidance docs, but then some implementers may miss the higher-level picture.
What can also be concerning about separate documents you may end up with a lot of overlap between the two documents. Splitting may loose the coherency and adds lot of technical protocol details. Will formatting the document 'side by side' make that easier to read (two columns instead of two rows, one for each model). That would likely help!
Splitting into protocol-specific docs may end up in much duplication.
That would allow even services and infras that do not yet give in to a proxy-only model to still be served. And not all the world will move to a proxy (and for many cases provisioning into a directory is very relevant.

Shortening the document may help readability - refer to other best practice docs for operational implementations.
"Use any one of these best practice guidance, and be able to answer these questions… from WISE SCI and Sirtfi"

Some implementers also indicated that concrete ("clear") requirements are easier to implement than abstract guidance, even if that means that you become more prescriptive (and may not be as burdensome if it re-uses some existing standards and best practices).

Requirements for on-prem and public cloud providers as hosters shold be as equivalent as possible, and a list of questions (inspired by e.g. the WISE SCI and Sirtfi items) are likely easier to use when assessing cloud proviudes than formal standards (and asking everyone for ISO27k audits is certainly too heavyweight).

Reviewing the (most contentious) section 3.4 on operational requirements on the AA, the "Can we comply with this?" remains a question for implementers. Can e.g. CERN OpenShift meet this? Does OS-level virtualisation like a docker count as virtualisation? Does a Pod count?
The functional requirement is "The hosting environment should prevent any compromise from spreading between virtualised guest environments", so we can add that as the high-level functional requirement. Larger commercial cloud providers provide and released documented descriptions and audit reports on how they achieve that - self-managed hosting can do the same is the management of the environment is properly done.

This can then be aligned for both physical and virtual environments, as done in the new text in the document:
"*Implementers of AAs SHOULD use placement policies to ensure physical and/or virtual separation of sensitive and non-sensitive services, containers, or VMs to reduce the risk of cross-compromise. In all cases, the environment itself must be protected according to current best practice, and a risk assessment of the environment should be performed[e.g. based on the WISE SCI and Sirtfi requirements], taking into account both the integrity of the AA as well as the requirements of the communities hosted on the AA and the relying parties receiving attributes.*"
This then replaces former items 1 and 2.
The statements about transparency don't quite belong here and move to the later (publication and audit) sections.

The actual changes are committed in the document, as linked above. The in-depth review focused on section 3.4, the security environment operational requirements and key management (up to and including 3.4.1, item 5).

OpenID Connect Federation

The last interop verification session was in December, using both Henri's and Jouke's implementation of the standard.
Meanwhile, the standard itself is becoming more stable, although details are still changing. This does not direct affect path validation, and these are largely focused on easing implementation by Connect2ID and Roland Hedberg.
There will be a new draft coming out based on the results of the interop fest around February 2021. These will not be the last changes, but no major overhaul is expected any more.
http://lists.openid.net/pipermail/openid-specs-ab/2021-February/008058.html

There was an official vote at the OpenID Foundation, but that was a long time ago, and the biggest changes since have been in the section of communication between OPs and RPs - which is more prominent in Henri's service implementation. By the time the spec is finished, there should be at least two full implementations ready.

Assurance update

The REFEDS RAF framework is getting some additional love and care to make it more populaer, both by promoting eduPersonAssurance as an attribute, but also at least try to get ePAssurance also an optional element to the REFEDDS R&S attribute set. The discussion there circles around whether it could be optional, or must be made mandatory to ensure consistent update.
The best would be to have a couple of strong use cases to drive adoption, like Sirtfi worked with CERN pushing it. For this year, we will have a two-pronged approach: a review of the profile documents to see how they are used today (and see how they have been established) - and work on a mapping between the standards that are in AF and what is ongoing 'outside'. The other direction is look at a test capability for ePAssurance and SFA/MFA: there is a sole SWITCH-provided tool but nothing central on eduGAIN technical pages. There is nothing about assurance there now.

These are both actions of the REFEDS RAF WG today.

In addition, Jule et al. are working on a paper to explain which use cases and which requirements could drive assurance, and how to select the appropriate assurance level that fits their need and risk profile for the service(s). Also identity providers guidance will be included, as to how they could implement it. This was already presented in an early form at the EGI conference. In general, if you are a university or so, you already have a lot of assurance elements, and the challenge is not to re-do the vetting, but to use the existing vetting and assurance and 'just' assert the proper RAF values. This will be presented at the ISGC conference next month, and this is currently in progress. The paper will be published in ISGC as well.
Meanwhile, even an employment process (like I-9 in the US) may seem good, but there is no way of auditing it from the credential issuer side - and an attacker will always target the weakest process step here.

The profiles are very high-level, and there is not much practical guidance for implementers … What about 'fail-open' or 'fail-closed' (e.g. if DUO is down, esp. for US institutions), which is getting more common with cloud-based services that can fail in uncoordinated ways? At the same time, institutions tend to be risk-averse and don't want to take on a (presumed) liability for what they signal in ePAssurance.
The GEANT T&I Incubator also evaluated a possible central solution for ID vetting using (typically NFC-using) apps.

Also WLCG is going through a risk assessment and assurance discussion, led by DaveK: assurance is probably the most prominent element of the changing AuthN landscape - with higher-assurance services requiring assurance. So the NIH is for example to ask for both MFA as well as higher assurance.
In some circles there is confusion that with new (token) technologies, it would be only technology that is needed, and assurance is (un)happily ignored - that sentiment of course misses the point. You cannot just do low assurance since the cost (on the service providers) due to e.g. incident response and handling and 'whack-a-mole' games is just too costly (unless you charge more for it, of course …). So also an accreditation scheme for things like token issuers will be at least as important as the technological move.

As an example, the LZ dark matter experiment (UK+US) has two major data centres (US+UK), and the US physicists want to access the UK data centre, but these US researchers cannot get any Silver/Cappuccino level assurance, but only basic identifiers. The current solution is to re-do the validation remotely from the UK, but it remains that the US just dropped out of the joint assurance schemes - and the universities are not ready to provide assurance and hence CILogon Silver remains out of reach of them.

It's good to see that e.g. the mail from Tom Barton that reconfirms that service providers are actively asking for assurance. 'NIH could be to Assurance what CERN has been to Sirtfi'.

For the next generation (token) issuers that will have to do assurance, but many orgs decide on the 'necessary' assurance is not based on what is needed or based on the risks, but just on how much they thing the can afford - i.e. not spend money there and thereby move the cost onto others. What RPs need to do is to make it very clear what the requirements are (but that is also hard unless you do both a full risk assessment and a TCO costing exercise). The tokens are not changing that, and in order to make it work with appropriate controls, in the end you will need most of the controls implemented that were there in, e.g. PKIX technology in the first place. That is very obvious in e.g. OIDCFederation, and will become clear to all solutions over time.
Human nature tells us that the right number of choices for a user is between 2-3. Four is too much, and just two results induces friction. When there are more choices, people choose the path of least resistance and drop to the lowest bottom. Maybe we should give them an even number of choices, i.e. we now have "medium" and "high" :) And the service provider should specify the level - and that take into account the cost of a breach (monetary and reputational), vs. potential loss of user base. Maybe there is a role for FIM4R here! To be discussed in 2021 …

TAGPMA and WotBAN&AZ Token-based authorization

Derek's slides at https://indico.nikhef.nl/event/2990/contribution/14/material/slides/ provide a great summary of the workshop that discussed token-based technology transition, while taking into account also assurance, migration, and accreditation processes. For the slides and video-recordings of the whole WotBAN&AZ workshop, see the agenda at https://indico.rnp.br/event/33/
Highly recommended!
The JWTs are used in the context of OAuth2 flows (like the SciTokens of Jim Basney et al.). This also holds in most other flows - but the exact form in which they are issued and employed may not be based in one 'format'. But they are bearer tokens, so can be used in other contexts as well (Jens is putting them into REST web services as bearer tokens). But that would likely still be an OAuth flow, since almost anything is an OAuth flow - and you need to call out to validate the tokens against the AS (the introspection endpoint), but the bearer token RFCs do not make such assumptions. You can validate signed JWTs that contain information, which is still within the OAuth space. This discussion itself makes clear where the understanding of token-based systems is …

There will likely be some follow-up meetings later - and Derek is putting together a general call for presentations.

Jens' Soapbox

Complexity has to go somewhere, and what can appear simple is often not what is needed, and simple solutions over time (as requirements get rediscovered) grow complexity. Even debugging and the use of the right tools in itself is a complexity issue. But, as usual, summarizing the soapbox is hard, so please review the slides.
In the discussion we addressed what is the cost of an actual piece of software, and how much should we invest in either code or its documentation? And adding documentation and testing is a cost - so who is to blame for not doing all of that? Maybe not the student that writes a proof-of-concept …
There is also a cost decision for using open source, in that you have to factor in risk and additional uptake effort. Thus for open source software using the commercially supported version in the end may be cheaper - just depends on how to factor in the cost (and for some: the 'colour' of the money and which budgets allow what kinds of spending). In order to get the right functionality, you need to maintain the integration between the components as well - and that integration is oft expensive.

Operational matters

- the MD-Grid self-audit was completed successfully and can be closed
- for the RDIG CA, Eygene really should send the new CP/CPS (and make that update persistently available), and ensure that practice (good) and document (outdated) align
- the TR-Grid self-audit review was presented - peers are DavidC, Hannah, and Anders for a *quick* round of review

We thank the following people for the extended attendance and stamina for sitting through the virtual meeting: Eric Yen, Eisaku Sakane, David Kelsey, Anders Wäänänen, Ian Collier, Cosmin Nistor, David Crooks, Daniel Kouřil, Hannah Short, Jana Zraková, Jule Ziegler, Lidija Milosavljevic, Maarten Kremers, Marcus Hardt, Miroslav Dobrucky, Mirvat Aljogami, Mischa Sallé, Ian Neilson, Paul Mantilla, Reimer Karlsen-Masur, Scott Rea, Uros Stevanovic, Jens Jensen, Sven Gabriel, Adeel-ur-Rehman, Nuno Dias, John Kewley, David Groep, Bill Yau, . And especially Derek Simmel for joining so early in the morning!