52nd EUGridPMA and IGTF meeting

Monday, 7 June, 2021 13:30

Dear IGTF, EnCo, EOSC ISM, EUGridPMA and AARC community members!

Thanks to all those that joined the on-line sessions of the 52nd EUGridPMA+ joint meeting. As we are hopefully moving closer to an in-person trust building environment later this year, up to 30 people participated in this hopefully last fully virtual event, 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!

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/52 or linked therefrom. In this summary:

- APGridPMA and TAGPMA updates - and a WoTBAN&AZ-2 workshop in September
- Risk scenarios and the desire for LE in a WLCG host cert context, and working group setup
- WISE SCI developments, finalising and publishing the guidance and assessment sheet
- RCauth.eu distributed operations: towards operations
- Enabling Communities: progress and next steps for e-Science Global Engagement
- EOSC Security baselining process and future plans
- Attribute Authority Operations "G048" bis
- AARC PDK Service Operations policy – evolution in UK-IRIS and towards the EOSC
- Operational updates
- Jens' Soapbox: How To Make Things Better
- Attendance

The next 53rd EUGridPMA+ meeting is scheduled for Tuesday 28 till Thursday 30 September 2021, will likely be **in-person** but with the option for remote participation in case you cannot make it, and hopefully in **Garching near Munich, kindly hosted by Jule Ziegler at LRZ**.
Meanwhile, you can virtually enjoy TNC21 (June 21-25), the FIM4R Assurance workshop (June 17th), REFEDS (june 16th), the TNC21 Security Day (June 14th), or the EOSC Symposium (June 15-18).

Hope to see you at one of these places soon!

Best, DavidG.

APGridPMA: developments in the Asia Pacific region

The APAN meeting is virtual in August, and the APGridPMA will be colocated with that event. That will also have the chair election for the APGridPMA. The catch-all services from the ASGCCA is expending with more countries particiapting and some countries actively devolving their CA to ASGCCA.
MyIFAM has been decommissioned and the any users moved to ASGCCA catch-all.
Also AusCert moved to Sectigo as the CA provider - status of any link to AAF is unknown at the moment, but may be in place (since contact to AusCert is currently minimal)
At APAN52 (Aug 2021) there will be discussion to collaborate with TCS for regional CA services (and with Jisc?)

TAGPMA: an update from the Americas

Same folk, same place, but lost one operational member since NERSC SLCS has been discontinued (and their representation is now via XSEDE).
CILogon CA front-end is moving to AWS US East from NCSA - but the AWS HSM is still too expensive (1k $/mo). The secure link thus go back to on-site NCSA. The TAGPMA has had no objections to the move - to be finalised at the TAGPMA meeting on Tue June 8th.
There is less participation from the Latin-American region, and the presence is a bit spotty on the English-language call. The pandemic is clearly having a big impact, which is continuing.

A new WoTBAN&AZ workshop in fall is in the works by Derek - maybe even in-person at PSC with a TAGPMA F2F around September 2021.

A growing desire (by WLCG and GridPP at least) to use LE for host certs for storage end points

The WLCG as a relying party has become aware of the desire for the use of host certificates from LE, a DCV-encrypt-only CA. In 'DOMA' (storage and data management activities) there is a move to use of web-based protocols that are also used by end-user via browsers in e.g. GridPP. Also since OSG does use LE without additional controls. There are perceived issues because of the ease of use of ACME and wildcards - as described in Dave's slides. In WLCG, there is no consideration yet for any specific risk assessment nor has there been consultation of the WLCG security groups, but this is pushed by technology (and aiming to do this via the WLCG management board).
But the assurance and trust issues do not go away as technology changes. TAGPMA started a working group ~2 years ago, including compensatory controls, but that did not really start with OSG just used LE without compensatory controls. And the use case thus 'went away'.
Guidance is clearly needed, also from the security groups - who were not consulted at all till now.

Yet the use of ACME (or API access) and wildcards for the use case here are not incompatible with IGTF, as is shown by TCS (or InCommon IGTF). TCS does give at least a REST API (and soon ACME), as well the ability to use wildcards - and does more than just the http-01 and dns-01 that LE offers.

The IGTF provides additional assurance measures, just like many R&E federations that are getting there as well (with Baselining, Sirtfi, and REFEDS assurance), and whether that holds for communities will be varying per community. Most of the assurance there is however directed at users, not specifically at hosts - or agents, for which host credentials are often use in their role as 'semi-robots' and automated agents.

There is a clear need for an IGTF WG again to do the risk assessment and communicate that expertise back (like a link back to a responsible party). The use cases also need to be assessed and the solution may already be out there, and significantly distinguish technological problems from the trust issues.

- Where can be improve usability without sacrificing trust authentication and assurance?
- How can we communicate how the trust has helped and build back that trust - and then jointly agree what we need to do (including some DOMA people, so that there is no duplication of discussion)

How long would it take to see host certs show up as users - since in LE specifically the eKU does allow their use as clients(!). How long would it take to see DCV hostnames show up as community members, e.g. as 'cheap anonymous robots' or taking on client agent roles?!

And even if the WLCG MB would 'trust' LE (or DCV in general), that does not mean that all the WLCG sites would accept it, esp. for sites that host multiple communities and will and should not trust it since trust *interoperability* is a prerequisite for a generic infrastructure -- which could result is fragmentation of WLCG.

Many of the trust arguments can also be found in
https://www.eugridpma.org/meetings/2018-05/summary-eugridpma-2018-05-karlsruhe.txt

Whitebox/greybox adversarial pen-testing leveraging a DCV CA, (social) engineering, and pen-testing may help clarify what the risk level is and which elements are needed. Demonstration helps rather than talking in abstract concepts.
There have been EGI SSC tests as 'users' based on some communities, but of course for now these all use OV certs. Maybe historic real cases can be identified where host or robot identity has been used for traceability.

Some of this is also coming in with the use of more non-grid specific storage solutions, where the public web+IGTF trust is needed. On the Mozilla list there was a blog about not letting the Mozilla root store certs be used for anything else than web browser (as is re-using 'the store' for non-web purposes). The root store is only applicable to browser trust and not for other trust - and Ryan Sleevi wants an agile system more than a secure system… and that would be a big argument as to not re-using LE for other purposes, as the WLCG/OSG use case may divert from plain browser use cases (https://blog.mozilla.org/security/2021/05/10/beware-of-applications-misusing-root-stores/)

Enrolling typo-squatted (and route-table-squatted) DCV/LE certs can easily be enrolled in a community - and the current BIRCH/CEDAR/OV certs at least provide assurance and traceability. Acting as a client is then the more dangerous case (or you need to do DNS - Amazon Route 53 - and route table - Bitcoin wallet diversion - attacks - both are commonplace).
But the 'abuse cases' are not sufficiently known in the WLCG community, and even sysadmins do not realise all the risk or distinguish between AuthN and AuthZ - and also they can get carried away by fancy new technology and 'ease of use'.
Most of the risk is with the use of host certs as clients, and that LE asserts eKU wClientAuth!

We might have promoted protocols like TAMP to make updates easier in a browser. Would browsers implement dynamic trust lists (like the EU trust list) just like Adobe did recently? GEANT already sent a letter to the EC to plead for regulation of CABF (https://www.geant.org/Resources/Documents/CAB_GEANT_Statement-GPiesiewicz.pdf)

WG terms and scope:

- what is the actual problem statement
- position paper listing the considerations
- list of pros and cons, of what is currently provided and how that relates to incidents
- assurance discussion of host certificates (since till now it always 'just worked' in case of incidents), also when viewed as a client credentials if that is a technically allowed purpose
- risk assessment and 'beyond technology' assurance
- list of current capabilities and possibilities (ease of use and availability of trustworthy solutions)
- how much client work is still CLI based, and what percentage is web-based 'joint' trust cases?
- additional compensatory controls - that is linked to the (token) assurance debate
- are the use cases really up to a system like LE where agility is more important than security or operational stability
- (software or interface) sharing for easing processes (e.g. using ACME key-id) and a consistent API

Working group members:

- Scott, Uros, Jens, Dave, Hannah, DavidC, DavidG, Eric Yen
  + 'DOMA people' + people from the AARC community list: ELIXIR
- start with a small team to get a first set going, then announce to AEGIS
- Sven, for a 'creative use' of LE/DCV and ongoing input

WISE SCI and the PDK

The WISE meeting a couple of weeks ago (https://events.geant.org/event/733/) highlighted the SCI work again, as shown in the slides linked to the agenda.
The PDK top-level policy has been re-written and improved by UK-IRIS, but the community policies are less mature, and anyway would need to be re-cast as good-practice documents for use for the EOSC communities - in the infrastructure that likes to say 'yes'.

The WISE SCI mature assessment document has a couple of open issues, and some specific contentious issues like the 'security plan': https://docs.google.com/document/d/1O2UTrKD70erpmO5DVIgn_1xpFX3NfVae_BGKPHoFuWo

- on OS3 "Security Plan": what to some extent was meant is that the infrastructure should at least have *thought about* the issues listed there - not intending to be extremely burdensome and unattainable. Not a many, many page document, and e.g. neither EGI nor WLCG would have such a document.
  Maybe distinguish betwene what the ideal end-result would be, but a description of the process and the way to get there (a plan or roadmap of how to get there) is probably equally valuable. And some words are in themselves ambiguous, since e.g. 'architecture' is - in other contexts - is more like a service portfolio and a set of boxes that *could* be used and put together.

This is more like a list of things that a security has to think about, and ask here to describe - maybe in one line or so - how things are actually done. "Do you do risk assessments?" "Yeah."
Currently the text is so vague, that the result might be 'useless'.
If we agree on the 'why', than this is about putting knowledge on (virtual) paper, like a TWiki. This is 'part of the whole' overview.
Even following the SCI framework or an ISO27k document structure as a 'table of contents' is already a good start (like doing that for an org and finding out what is actually done).

Or: do we drop this in SCI v3, since we rely on 'knowing' that the security structures are in place. Would that translate to new infrastructures (not only the conventional ones like EGI, OSG, &c)? Maybe the elements are implicitly covered in some of the others (like OS7 covering AAI elements in OS3).

But OS3 is about *document what you do*, but even without documentation there is some level of trust.

In that, it is attempting to establish a baseline - and it might go away or be less complex once all the other elements are put in place (and documented). It is in that respect more a statement like 'we take security seriously'. That's why it was deliberately 'vague'. Some examples can be added now and improve in v3.

- IR1 - include pointer to a list of tools to send communications challenges

Publish fairly soon (also as input to EOSC maturity assessment), including the spreadsheet. That needs a relatively quick WISE SC approval process. Then also OSG can use it, and distribute it via the wise community web (and Zenodo?)

CA Updates I: RCauth.eu

The RCauth on-line IOTA DOGWOOD CA is a distributed operation between three sites sharing a common key, which has been distributed and cloned securely across the three sites. The various parts have been transported using three different means (postal, keybase, in-person, signal, SMIME/PGP, &c).
The actual work of building a distributed set-up has been more complex than initially thought, which is to some extent surprising since all technological components were all well known. The result is a long series of Jira tasks.
For the network there is an interim solution using an OpenVPN over the public (GEANT) internet instead of a dedicated low-latency lambda, but the network remains local and within GEANT. But even today the latency is already good enough for the current purpose and load.
Extensive testing with fail-over and databases has pushed a move to Galera over secure VPN - also showing that the database sync is faster than the client HTTPS negotiation (so a client can negotiate to multiple sites and not notice the fail-over). The resilience can thus be made to work - using an additional HA Proxy layer which is fully meshed. The OpenVPN routing can be via any of the participating sites, regardless of which single link goes down.
The final move to production will also be a name change for most of the public facing names, removing "Pilot" (but not from the CN RDN of the CA name itself).

More information about the site-specific setups can be found in the presentation. As of April 1st, RCauth.eu is also co-supported by the EOSC-Future project (and no longer EOSChub).

Enabling Communities: progress and next steps for e-Science Global Engagement

Maarten shared the ongoing work in the GEANT 4-3 Enabling Communties activities, highlighting the work on Assurance (there will be a REFEDS Assurance Framework 2.0 update!), and a FIM4R paper on the use of assurance.
A preprint of the ISGC paper on practical assurance is available on Zenodo:
https://dx.doi.org/10.5281/zenodo.4916049

In preparation for Sirtfi v2 a community survey was concluded and lots of responses received. These are now being evaluated. Multiple people per IdP responded at times (it was sent to all IdP contacts, either via the national federation or directly, depending on national preference of the fed operator.

And join the FIM4R Assurance workshop on June 17th (1600 CEST)

EOSC Security baselining process and future plans

The EOSC security activities as part of the new EOSC-Future project started on April 1st. The presentation has most of the security details, and most of the other discussion was about EOSC in general.
For the security stance of the EOSC, MISC IoC sharing should be a base for strengthening the distributed security in the EOSC (and a SOC-like entity later) as a focus. The classification of research and edu orgs as 'import entities' in the context of the NIS2.0 directive will likely be a boost for the research services that will emerge in the EOSC, even through the implementation of the NIS2 directive in national legislation will take a few years, this could be a strong push. Regulatory compliance is usually a good incentive for compliance even in R&E.
And WISE templates could help new service providers!

Not all problems can be resolved by EOSC security coordination - we need to manage expectations
There will be a baseline for all of the EOSC (as also seen in the current AAI report), but more 'security quality flags' could be considered for showing in the marketplace portal and equivalent entities. Authentication strength requirements inbound for services can also be embedded in the service operations policy, and then both sending/brokering parties and the receiving service should ensure, say, BIRCH/Cappuccino assurance inbound (and therefore not DCV/LE as a client). This risk assessment should be done for services and shared. Think also of DCAU where then upstream user attributes and auth need to be used, not host creds from 3rd party client.

Attribute Authority Operations v2

The document was reviewed and comments resolved *up to and including* to section 3.5. The pending comments there were resolved and clarified, also on further clarifying external hosting. Basic security for the signing key also was clarified, specifically that plain-text storage of keys in databases is not quite appropriate (even keycloak recognised that as an issue in https://issues.redhat.com/browse/KEYCLOAK-3445)
The extensive discussion on OpenID Connect clarified that there are by now very few public clients, although not all confidential clients are that trustworthy given the simple registration process in many cases. It still applies to React-style sites and JavaScript …
We'll schedule a dedicated meeting to finalise the document.

AARC PDK Service Operations policy – evolution in UK-IRIS and towards the EOSC

The most recent UK-IRIS (https://www.iris.ac.uk/) policy is the *service operations policy* - the one that also got the most feed-back from the UK community since it was explicitly presented to the service owners and site managers. UK-IRIS also has a broad range of types of services, and some authenticate via a central IAM, others do not, and the diversity of services is rather broad. As such, it could even resemble the EOSC to some extent.
The supplementary aim of the IRIS work included making this policy stand-alone and more independent of the rest of the AARC PDK suite of document templates, and differentiate between conceptual requirements and the specific implementation measures (as explained in the slides). This also makes it more palatable to a loosely structures infra like IRIS (and EOSC).
The UK-IRIS draft (with its comparison with the PDK v1 version) is linked to the agenda (and the 'grey text' is what has been removed, and the crossed-text is what may have bene lost entirely in the new wording).

The references are probably more specific to the constituency, and in a template structure these should be open to re-writing by the user of the template (e.g. remove the UK specific references when used in e-Brains), rather become an overly long list of best practice references.

We may now have 'lost' elements, so it should be reviewed also with a fresh pair of eyes, and under the auspices of the WISE SCI look how this would apply to the US side (Trusted-CI) for instance (which have a different view of liability, for instance).

For this service-centric document, slight modifications are not harmful (unlike in the AUP) - in the end it's a scaling issue, and between services and infrastructures a (limited!) amount of mapping can be dealt with. Especially when one version is simple more strict than another, or just has a specific header text. But we can surely highlight the benefits of making them compatible!
For WISE, we should focus on the 10 core points (without the preamble and refs) for sharing and requesting comments.

This is all also highly relevant for the EOSC Exchange service on-boarding (and as a baseline it would - when adopted and endorsed - be applicable as part of the RoP). To a large extent, it's about posture and 'do no harm' - the UK IRIS version is about the minimal things you should be doing as a service!

The evolve this work, we need to have a follow-on meeting, under the **WISE SCI working group** and supported by EOSCF and others, based on the new (google) doc with just the 10 points - and make the other versions available alongside in the same folder.

For the EOSC baseline (after the first 'Sirtfi-only' initial baseline by August 2021 or so), we should move fairly quickly, but that process takes much longer since it has the to go through the EOSC processes, maybe leveraging the joint AAI and Security 'TCB' process and taskforces.

Operational matters & self-assessments

- Self-audit RDIG CA: started in 2016 and we're *almost* there but not quite. The last outstanding issue has probably been there for a long time and misses the description of the binding - which is operationally correct as far as we know, but it would be good to have this properly described in the CP/CPS itself.
- Self-audit TR-GRID: nearing completion, and it only needs a double check. A latest CP/CPS will be sent shortly to the whole list.

Jens' Soapbox "Q series": How To Make Things Better

Let's not be surprised … but there are unintended consequences lurking around every corner, and with the increasing compleity, fewer experts around, and the misguided organisations forcing job rotation scheme make all of that worse by making expertise leak away. "PEBKAC" is a major risk … and biases are misleading and dangerous.
To get a more complete idea of the complexities, look at Jens' slides!

A technical forum is lacking and we (as a global community) are now dwindling in expertise since people tend to drift. As seen in TAGPMA, people move away, and then unintentional things happen.
In the commercial world, there are lots of training requirements that cannot be waved away (these are audited), but that needs a lot of resources - and that makes it hard for the smaller IGTF CAs. That does lead to a bit more risk.
We then need the forums in place so that the larger participants can share information and also inform the smaller-scale operations.
This needs to be global, and maybe in a more colloquium-style setting rather than the standing meetings. But that needs people to make room in their agenda.
We see the same issue for trust itself - and we may now be spending time on correcting misguided ideas and correcting those, rather than spending time on proper training/education.

We are a trust community, and at least half of that is 'community'!

Attendance

We thank the following people for the extended attendance and stamina for sitting through the virtual meeting: Bill Yau, Lidija Milosavljevic, Mischa Salle, Marcus Hardt, David Groep, Burcu Ortakaya, Jule Ziegler, Maarten Kremers, Hannah Short, Baptiste Grenier, Miroslav Dobrucky, David Crooks, Eric Yen, Sven Gabriel, Jens Jensen, Feyza Eryol, Mirvat AlJogami, Uros Stevanovic, Cosmin Nistor, Jana Zrakova, Ian Neilson, Scott Rea, Ralph Niederberger, Reimer Karlsen-Masur, Ian Collier, Nicolas Liampotis, Derek Simmel, Jan Chvojka, Adeel Ur-Rehman, Nuno Dias.