54th EUGridPMA and IGTF meeting

Tuesday, 25 January, 2022 09:30

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

Thanks to all those that joined the on-line sessions of the 54th EUGridPMA+ joint meeting. As we are hopefully moving closer to in-person trust building again in May, many people participated in this hopefully last fully virtual event (I know I said that last time as well…), spread over all time zones and from all our regional PMAs. Plus a large attendance form the AARC, GEANT Enabling Communities, EOSC ISM, SURF, 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/54 or linked therefrom. In this summary:

- Updates from the Americas and the Asia Pacific
- RCauth.eu distributed CA, IP Anycast on the Internet, and HAproxy
- WISE SCI - adding data protection to the PDK policy development kit
- Assurance and MFA - in refeds, fim4r, and for service providers
- Privacy and data protection - in WLCG and for the WISE SCI 'task force on privacy and data protection'
- WLCG server credential working group updates
- Attribute Authority Operations guidelines - for proxies and authorities
- Enabling Communities with GEANT's eScience Engagement
- Making Things Better - and innovate!
- Operational matters
- Attendance

The next 55th EUGridPMA+ meeting is scheduled for **May 23-25th (Monday morning till noon on Wednesday)** and will (likely) be **in-person** in **Garching near Munich, kindly hosted by Jule Ziegler at LRZ**. As a backup, virtual participation will of course still be possible, although must of the trust building and innovation happens more naturally and productively in a face-to-face setting.

Hope to see you soon!

Best, DavidG.

Agenda and Next Meeting

The next meeting will be in-person in Garching! From <start> till <end>.

TAGPMA update

Derek Simmel (PCS, TAGPMA Chair) reviewed the developments in the Americas - the slides are available in the agenda. Key messages from the presentation include:

- The scheduled move of services of CILogon to the AWS East region is experiencing into some technical hurdles and is hence delayed. The move is still scheduled to continue, although the HSMs will remain on-premise at NCSA given the high cost AWS is charging for cloud-HSM slots.
- There will be a *new* DigiCert IGTF accredited set of issuing CAs, based on their general public trust PKI, although a private CA could also be made available under a different CP. For TAGPMA, Jenny, Derek, and Paula do the accreditation review based on the Classic (CEDAR assurance) profile.
- DigiCert will continue to be able to provide service to Open Science Grid (and WLCG) sites based on their new offering
- The XSEDE MyProxy CAs are being migrated to newer servers and hardware, since there are HSM driver issues for the current hardware when moving to newer (EL7+) based distributions. XSEDE is supporting this move as the key customer, using it for back-end authentication.

XSEDE as a project is now officially ending in August 2022 - and is waiting now for a successor, under the "Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS)" programme, for which proposals have been submitted. There is no news from the NSF yet, although when there will be no timely response, the current XSEDE infrastructure will likely be extended again.

- A next IGTF/TAGPMA meeting in-person could be at the Internet2 TechEx, now scheduled for December 5-8, 2022, although no location has yet been announced.
- The proceedings of the October WoTBAN&AZ 2021 and recordings are available (no autotext transcript yet, though). There are two other related workshops on this topic that have even more information: the OSG Token Transition workshop, and the SciAuth workshow at the NSF Cybersecurity Summit.
- The retirement of the OU attribute from the subject name in public trust certificates, following CABF ballot SC47v2, is being discussed in TAGPMA, as it affects the InCommon IGTF server CA that is joint IGTF and webtrust, operated with Sectigo as the issuer. Sectigo will sunset OUs fairly soon, with official end of issuance for all by September 2022. Issued certificates will remain valid (for at most 398 days as usual).

APGridPMA

Eisaku Sakane shared the updates from the Asia Pacific region. The lack of in-person meetings makes it a bit challenging. But progress on token-based AAIs in progressing apace , with an IAM working group at APAN and sharing of IAM and federated IdM experience, such as from the GakuNin trust framework by NII.
There are a few changes in the CA trust fabric (with KEK CA changing manager), and a need to re-connect with CNIC, SDG, IGCA and AusCert.
The 29th APGridPMA virtual meeting will be (virtually) co-located with ISGC on 20-25 March, alongside the security workshop, and the 30th meeting co-located with APAN in fall 2022. Hopefully ISGC 2023 is again in person.

RCauth HA distributed CA

For the RCAuth.eu operational team, Jens, Nicolas, Mischa jointly presented the slides attached to the agenda page.

The RCauth online CA is a credential conversion service based on MyProxy, akin to the ones used for climate research data transfers for instance, and issues credentials at the IOTA/DOGWOOD assurance level. It combines code from CILogon (Jim Basney et al) for the delegation service and MyProxy, with a filtering WAYF and a (range of) back-end issuing systems leveraging various kinds of hardware security modules. Distributed over three sites, it shares common key material, a joint multi-master (Galera) issuance database, and issues a single CRL - from three places. The key reconstruction, discussed at length in previous PMA meetings, worked (with the tooling and the three independent transfer methods with XORed elements), although there were some technical challenges to ensure the secrets never touched a disk and only used in memory.

This is now implemented as a fully distributed multinational solution. Several high-availability models - spanning multiple sites, countries, and administrative domains - have been proposed, and the IP anycast layer-3 solution (BGP announcements from multiple sites) is now operational, with AS5408 and AS1104 announcing the HAproxy address. Some ongoing developments (a move form Quagga to Bird for one of the HAproxy locations) at times still impact the announcement, but given the multi-origin anycast setup, at least one AS has kept announcing the prefix. For the moment, it's IPv4 only (a single /24), and it's done from two distinct autonomous systems: GRNET and Nikhef.
There are other solutions, like the DNS-based short-TTL one that Leif Johansson and Niels van Dijk used for InAcademia, where the HA proxy is feeding a DNS based failover solution to find the closest 'alive' node, provided your TLD DNS already is highly available. But the BGP anycast solution is significantly easier to roll out, since 'the internet' does all the heavy-lifting for you, and anyway you would need a highly-available DNS TLD name server anyway. And anycast was the simplest, hence started there.

The 'HAHA' proxy makes it look like 'a single service' as seen from the user. If you get 'clean' IP space (so no unexpected RPKI locks, and the ability to add some route objects in the RIPE DB), then getting that to work is rather trivial. And it's definitely easier if you use IP space from your own LIR since then you fix all your own issues, although with space from a friendly LIR (in our case SN-LIR) it also works, although you have to keep asking to fix RPKI, reverse DNS, and to get the proper mnt-lower and mnt-routes rights on the inetnum object. Of course, for IPv4 use at least a /24 and for IPv4 at least a /48, otherwise most Tier-1 networks will drop your advertisements. And make sure there are no more specifics out there.

The monitoring logic is probably the most complex, since it has to factor in the availability of all components, but compensate for any local HA that may already be present. It is now based on Icinga - but there cannot be multiple Icingas monitoring the same box, so GRNET and Nikhef use their internal site Icinga to check monitoring data (this is in the site-specific configuration data). Nikhef holds the central logging for all instances, and that will stay inside the site, and the back-end have no outbound connectivity except to their own delegation servers.

All software is available on Github (in the project https://github.com/rcauth-eu/), with only some site-specific elements kept in a (keybase-encrypted) local git tree. By now, the suite of software has become 'rather complex'. How to effectively shared the knowledge gained with the community? It's now presented mainly here in the IGTF and in EOSC forums. TO make it more visible, some conferences can be considered as well.

The same solution now deployed for RCauth.eu can also be used for "Indigo IAM", where the OIDC providers also need high availability. In request of UK-IRIS (and probably later WLCG), Francesco Giacomini at CNAF, responsible now for Indigo IAM, is looking into this. And when used within a single organisation, BGP within a single autonomous system of course also works, and is simpler (just a regular L3 BGP high availability setup within that single AS).

WISE SCI Security for Collaboration among Infrastructures - evolution

Dave Kelsey and Ian Neilson presented the current state and planned evolution of the WISE SCI working group and the Policy Development Kit (PDK) it curates (slides will be attached to the agenda).
A continuing concern is the lack of engagement by non-Europeans, and particularly the lack of US participation where work is happening, but mostly on a local level. There is a definite need for new infrastructures *and communities* to be engaged (and hopefully the US in the successor of XSEDE and its "ACCESS" successor)

Updating the structure of the templates to accommodate a broader range of infrastructures will be a significant challenge - service operations and data protection, for one. Adding more FAQs may help, but does make it more complicated. The AUP baseline was simple compared to those. Meanwhile, the feed-back of Trusted CI on the AUP Baseline has been captured, but it's not in a structured format. It could make it appropriate to issue a 'version 2' of the Baseline AUP that includes that feedback.

There is already a whole body of knowledge on GDPR from Andrew Cormack, but it's not always considered by the service providers. Or the service even puts out obviously illegal privacy notices because they don't bother to read the templates and guidance. In the meantime, the user may not even know the service they will be accessing, as it can be hidden behind a layered set of other services.

WISE may have to focus on guidance to minimize risk for service provides in such a layered stack -- and the research communities should take responsibility for personal data that's inside their own research data and not put that to the infrastructure implicitly. The REFEDS Data Protection Best Practice document (formerly DPCoCo v2) should be taken in as well, and a specific REFEDS branding (like branding by WISE) will help its adoption. Some service providers inappropriately commented "why should I follow a *GEANT* Data Protection Code of Conduct v1, I'm not GEANT" - something that one EOSC core service provider commented when asked to consider joining the long-standing REFEDS DPCoCov1 scheme.

Anyway, also for discussion on privacy and data protection, WISE is explicitly open to non-European participation - there might be elements of e.g. FERPA or HIPAA to consider … There is a one-hour WISE SCI working meeting every 2 weeks on Monday's for WISE SCI progress. Please join.

Assurance, REFEDS, and FIM4R

Jule presented the progress at the REFEDS assurance. Last year the MFDA subgroup was created to cater for the NIH (US National Institutes of Health) health case use cases. The REFEDS MFA group has produced the recommendation document (now on the web pages), but the work is also continuing in clarification of what is the current good practice.

For this, the group has been re-chartered, to extend the guidance and make it easier to understand and implement - although for now the updates are still in a 'bulleted list' form. But in weekly calls good progress is being made, and next week's meeting will probably show the plan for finalisation the new guidance. This would include: requirements on the factors to be used, with recommendations on the most appropriate ones to be added. And some of the edge cases: what if the validator is down (should authentication pass as valid, or not?), and what to do for factor re-set.

The results will be presented to the full REFEDS Assurance working group, at which point also the charter of that WG can be reconsolidated.

The main REFEDS group continues to progress on "2.0" of the REFEDS Assurance Framework RAF:

- part on identity proofing is being revisited, since the current ones only references external documents and standards (Kantara, NIST, IGTF, eIDAS) and this makes it rather hard for the reader to decide what is needed,
- thus, a document that puts such guidance in-line could be helpful for adoption, taking from the existing external references, making it more understandable and implementable.
- also, the RAF mainly carries the IdP perspective now - this should be complemented with recommendations for the SP side: "which assurance framework profiles should I deem acceptable based on my risk profile?"

  Potential input to the SP risk-inspired assurance side can come from e.g. EGI/WLCG "acceptable assurance" documents (https://documents.egi.eu/document/2930) and WISE risk assessment models (https://wise-community.org/risk-assessment/, and that Urpo is also working on in an EOSC context, and the WISE RAW WG? UK-IRIS is also doing risk assessment, for now trying it with specific products and services (as Ian Neilson is doing for the Indigo IAM risk assessment). And there's of course the ISGC Assurance paper of Jule et al. A use case, based on e.g. a FIM4R use case, can then be documented in a follow-up paper (which is slightly less hand-waving).

This work is well underway, with ongoing lively discussions.

Meanwhile, we see adoption of the current RAF increasing, with also DFN moving from their own reliance document to the RAF (https://doku.tid.dfn.de/en:aai:assurance).

Maarten proceeded to present the FIM4R updates, though still needs to reach out to the community. A wider consultation meeting would be timely. A (virtual) WISE spring meeting could be one target, but an in-person joint WISE+FIM4R meeting early autumn would equally make a lot of sense. "Start work in the Spring, then a whitepaper by that time on "Assurance requirements for the FIM4R".

Maarten will set up a date selector to find a coordination meeting (since Hannah is not currently available) and will not wait too long with this. Combining with IGTF Garching meeting would also be an option! The informal steering committee will be polled to float some ideas.

But then, for all this work, there is a definite need to prioritize!

The biggest challenge will remain the deployment of all the new schemes. One of the aims of FIM4R writing this down is to convince the IdPs that there is a concrete and urgent need. For InCommon with the baseline requirements that did work out. Guidance document should be clear enough for it to be understood, but not too specific that is gets written off as 'undoable' by the providers. The REFEDS assurance group is suffering from the same - they are at times 'too expert'. But an organisation may suddenly find itself in a regulatory scheme where e.g. MFA is pushed in on the short term. But then where it gets pushed may not be in the place of highest risk. Or it looks MFA, but then in practice uses a long-lived token.

WISE SCI - privacy and data protection

The current WLCG (and AARC PDK) privacy document (2019), at https://wlcg-docs.web.cern.ch/?

dir=policy/security, ties in with the WISE SCI working group discussions earlier in the meeting. The next WISE SCI policy development kit policy was there decided to be privacy notes and data protection. There is a policy framework already in the PDK that proposes a framework to tie service providers together. This was used in WLCG under the 'pretty binding not quite corporate rules" argument, and the binding of the SPs in WLCG was used as the clincher argument in WLCG - and it worked both under the Directive and GDPR. This was to be replaced by the CoCo v2, but that one did not materialize.
In discussions with Andrew Cormack the set of policies was identified as a pretty good reasonably set. It does not necessarily meet all the formal requirements, but it certainly implement the model and spirit in a way that is actually scalable - which lots of bilateral contracts do not.

Now if we document what everyone does in different infrastructures (who all face the same issues, and all the processing agreements would never work out), we could come up with a best practice that captures our current good practice. We minimize the risks everywhere, and that there are controls in place to meet minimization requirements, and we don't ever process more than is needed. And it's not any form of sensitive data - usually just name, institutional email, and affiliation. That information is anyway willingly published in author list etc. For the majority of the researchers, people do want to publish their names.

There are thus two aspects to consider: the privacy notice and the policy framework

Federated Privacy notice for WLCG

WLCG management approved the 2019 notice, and considered applicable to all WLCG service anywhere in the world - and it could be applied as-is by each service. But then CERN brought out its own rules ("OC11") and in that all services at CERN that hold personal data should then service (via their service manager) a record of data processed, and then CERN generates the notice for you for your CERN service. And that includes all the information therein, but does not quite follow the WLCG one. And the data processing office at CERN is snowed under, so that there are hardly any services with a privacy notice, and meanwhile the WLCG boards decided to deferred the WLCG notice as well until CERN gets its act together. WLCG did generate a list of all services that should have a privacy notice, and they now push for the WLCG privacy notice, and just pressure CERN DPO office to complete the CERN-hosted services. There is, luckily, a remarkable alignment between the WLCG and standard CERN policy. Only in retention period differs a bit (18 months for WLCG, based on the need for annual accounting; CERN default is 13 months). And the WLCG privacy notice has a section on user's rights, and the CERN notice is lacking one.
The conflict of precedence also occurs at other institutions - including the designation of the controller. And on the WLCG level, the lawyer-recommended approach was to have a list of all privacy notices for all services in WLCG listed on a page. That would be an infinite list that would anyway be always out of date.

Based on the WLCG document (https://wlcg-docs.web.cern.ch/?dir=policy/security) the meeting then reviewed the most pertinent points.

Doing it all in a fully formal way is infeasible, but there are good inspirational examples, like the LSAAI use of 'consent' for transfer - here for the user-office to WLCG transfer (since there is no imbalance of power, officially), and anyway CERN (with OC11), like the European Organisations themselves (!), are not subject to GDPR proper.

The policy does not address unintentional release of information - mainly implicit knowledge gained because of indirect factors (like collating access requirements to infer more sensitive information, or even the mere fact of 'access to a facility or location is permitted for this person')? Should the controller protect, and to which extent, against such effective 'de-anonymization'? It's akin to identification (or mis-identification!) in medical population studies. A few pieces (as few as three) are sufficient for de-anonymization. These are not considered for WLCG for now.

For now, WLCG, like most of the GEANT REFEDS best practice, is based around controller-to-controller transfers (and we keep in mind the 'traditional travel agent' model and performance of contract).

The one thing to consider is to add a scoping statement, and words conveying the 'augmentation' of any institutional notices? Or conversely, in the institutional one: "For the worldwide LHC Computing Grid, the WLCG Privacy Notice augments this statement." And whoever is behind the currently-stated mail address "wlcg-privacy@cern.ch" is facing the challenge to find the proper controller to forward any requests - since there are many controllers in the WLCG scheme. The notice can only be shown once on enrolment and renewal, not for each individual service.
And the local institution should at least check whether it can meet the WLCG policy requirements before joining WLCG. That is for the institutions to make sure it is valid! Policies should not contradict each other, akin to *Snctfi*. The GEANT endorsed CoCo v2 would have been so much simpler!

But WLCG actually does work - even without all those contracts :)

SCI task force on privacy and data protection

WISE SCI will be looking at data protection guidance consolidation. There are (some pre-GDPR) documents from the DPK and other sources:

- AARC PDK: https://docs.google.com/document/d/1QseGQVzUQqvosqhjkF2qlHUI4Swlhgb8oDe8N6NWcqE/edit
- EGI SPG: https://docs.google.com/document/d/11S5UrCytHdeh4mNQc3btvZPW_ox_QgSBx0lII-XhKoI/edit (evolved using the GDPR terminology, and beware of cases where GDPR has now defined a formerly un-encumbered term like "DPO")
- WLCG and EGI 'version 2': https://documents.egi.eu/document/2732

and there is a need to consolidate and progress also the Policy Development Kit to current GDPR era. These are all documents that binds all the participants in EGI (c.q. WLCG) together, yet all based on the 'pretty binding not quite corporate rules' (we can throw entities out, but there are no statements regarding liabilities and paying fines).

A comparison of these three (or more) document in a table to identify overlap and missing elements is useful, and then update the guidance (for the PDK). Whether to propose a new template or a best way to fill existing good templates also needs to be decided - so join the WISE SCI WG to join in this discussion.
Following the discussions on the AAOPS - do not explicitly put in terms like "18 months", but *do* put in the way in which this number is to be determined. Just "timely" is also not enough. The EGI SPG (google) doc has most pertinent comments in it, collected over the past year.

Also the REFEDS DP Best Practice v2 (formerly CoCo v2) can be used to our profit here. Just like using Article 85 where we can (more or less creatively) use it, depending on the specific Member State :)

Everyone is warmly invited to the WISE SCI biweekly meetings (currently on Monday afternoon, for one hour), and Dave Kelsey will send the invitation to the IGTF mailing list.

WLCG server credential working group evolution

The Worldwide LHC Computing Grid (WLCG), supporting the CERN LHC experiments, has set up a working group to study the evolution of credentialing of servers (hosts), the "WLCG Resource Trust Evolution TF", late 2021. WLCG was seeing "an increasing number of resources it would like to use that are available with host certificates issued by a non-IGTF CA, and where WLCG cannot influence the service provider’s choice of CA". The notes from the initial discussions of the task force are at https://docs.google.com/document/d/1Sl0C_q-lGMCifChmFArHjsGzdnd-RM7O7jbpsGa8XRw

David Crooks presented the state and progress. Items that were discussed include:

- the TF has recently started, and it yet to arrange for a next meeting and a regular meeting slot
- first item of business is to discuss and review the problem statement - which is scoped to WLCG as a whole, including also, but not exclusively, its sites in the US
- there are other communities and their associates resources, acting as relying parties, that decide on their own independent sources of trust, but sites in WLCG are usually shared, and per-community trust cannot be easily done.
  WLCG (as well as the general infrastructures) appreciates some guidance as well.

- "Whilst maintaining adequate assurance" is also an important element to be considered in the problem statement, but it is not yet explicitly stated therein.
- 'host' certs from free (as in 'gratis') CAs usually *also* have clientAuth asserted, and not only the requisite serverAuth, in the keyUsage field, so can be abused outside a host server and server-transport security context.
- a number of US sites appear to have had trouble replacing existing OSG certs, and some do not have a good InCommon deal yet. But it should be noted that DigiCert will now put up a new IGTF joint trust CA independently of InCommon that will be available in the US again.
- That solution (a joint-trust IGTF + webpki issuer) does not help for 'purely cloud' hosted services, like S3 from commercial providers, that are then used also by the WLCG experiments (as these cloud providers often have their own CAs embedded in the provisioning and certificates are linked to their provisioning systems in e.g. AWS, Azure, or Google Cloud). Of course it is rather straightforward to fix by adding your own (wildcard) certificates, but this would need a bit of work to make it work with a dedicated CAs and users don't want to invest in getting that to work.
- Technically, with SNI (support by e.g. HAproxy already at the RAL T1, but it works anywhere and is a standard feature also of Apache and Nginx) this could simply be solved. The server can then present a specific certificate depending on the hostname that's requested in by the client from the server, even on the same end-point. And clients can just use a different (CNAME) name pointing to the same resource, possibly with an HAproxy inbetween, given that HAproxy is probably the simplest to configure. And SNI is a standard and could do that (i.e. it is not technology specific).
- There are some WLCG providers that don't have access to a joint CABF-IGTF trusted CA like GEANT TCS or InCommon IGTF Server certs.
- There is no way to distinguish uncontrolled web certs (like LE) from the controlled ones, since there is no unique namespace (RPDNC) that can be defined for those. Hence adding them to the IGTF distribution in the same way as the others does not make sense, since name uniqueness is the binding element for assurance.

It seems worthwhile to separate out those elements of the problem statement more clearly. Jens could contribute to the WG also for this.

Both technical solutions can be provided (like: SNI/HAproxy can solve most technical issues, of course on a per-host basis for performance/bandwidth reasons), as well as a WLCG position on what could interwork and send the recommendations: how to provide infrastructure resources to communities that have different security requirements, and how does that impact other communities.
Also from a WLCG perspective, the 'intent' is not to weaken things, although there are some more cavalier approaches by some in that community. The question to IGTF is also to make some in WLCG ask for help make 'good quality decisions' by WLCG.
But at the moment, some WLCG groups are actively working around things and pushing without such considerations, which is a risk in terms of trust and coherency (also to other communities and service providers).

With a move to tokens, client trust and channel/transport trust is better separated, and so some of the complexity is going away once everyone uses tokens. But then of course the tokens need to be trustworthy as well. And this technology is not quite there yet.

But also today (with or without SNI in front of a service), some elements of client and server trust can be separated. The certificate used for the host cert (which is the use case here) is validated by the client (and not another server or service!), and the trust store on the service remains using uniquely name-spaced IGTF controlled trust anchors, since these are only used to validate clients and are independent of the chain presented by the server to the 'outside world'. Only on the (community-specific) client side, the additional public webpki trust anchors would then be needed, NOT on the server side.

Thing to consider in the WG is also a set of definitions. What does a "trust store" mean specifically in this context? And is the IGTF becoming a bit like CABforum in settings standards, so advising the trust stores of others?

Enabling Communities and GEANT eScience Global Engagement

Following a quick review of GEANT T&I activities, such as eduroam and eduGAIN, Maarten Kremers

highlighted the connecting role that the e-science global engagement task holds for many community activities such as FIM4R, AEGIS, WISE, the IGTF, REFEDS, and the AARC Community.
Based on the REFEDS Assurance Work, the full paper by Jule Ziegler et al. in the ISGC proceedings has been published, and also the (for now) final version of the SCI guidance document has been released - its evolution has been (and will be) discussed in this IGTF week as well. For REFEDS Sirtfi, a new version of the Sirtfi framework (v2) is upcoming, alongside an eduGAIN Security Incident Response handbook. The SCCC "Security Challenge Coordination" working group in WISE now also has the US folk on board, with a challenge coming up there. And the attribute authority operations guidelines are on the agenda for today.
And for those who want to know more about 'EnCo', there's of course also the Wednesday morning of TNC in Trieste in June - the submission has been accepted! As for WISE workshops, the spring 2022 is being pushed back a bit to allow for a potential in-person meeting.

The planning for GEANT5-1 is to include again a Trust & Identity activity with an 'EnCo' task within it (the planning also has a comparable effort foreseen for now - which again goes through the national NREN). It will be current to discuss effort in May, with submission of the proposal in summer 2022. There are continuing political issues with Switzerland and the UK that could impact Horizon Europe participation.

For more information, see https://edu.nl/ctxxg!

Attribute Authority operations

The document outstanding comments (some from 2020 and 2021) were resolved and guidance improved throughput. The 'commentable' (google) document can be found at https://edu.nl/qf8uu. A consolidated version with adequate formatting for readability will be published shortly, and linked from all of the google doc, the AARC Community pages (G048bis), as well as from the EUGridPMA Wiki (where an old version is now present). In the end, this will be the updated version of G048.

Making things better

Communities thrive on innovation, but making innovation work is not as easy as it may appear. Innovation - making things better, harmonising techonology, documentatin ans sharing best practices - could work as a volunteer activity, but then is often swamped with the realities of day-to-day commitments and the need to work on things that are 'officially funded'. There is certainly pressure to innovate: the need for security and resiliance drives change, external improvements (or changes) trigger adoption and interoperation with new RFCs, and some of the Big Problems stay with us for a very long time. But lack of time, lack of funding, and - at times - lack of lab space hamper all what we do. And then the current challenges of remote collaboration have not even yet been mentioned.
But sharing, collaboration, and joint innovation is what also gives us that Warm and Fuzzy Feeling of Trust. Which is what Jens aptly highlighted in his Soapbox. And there are all the strengths and opportunities for us to exploit!

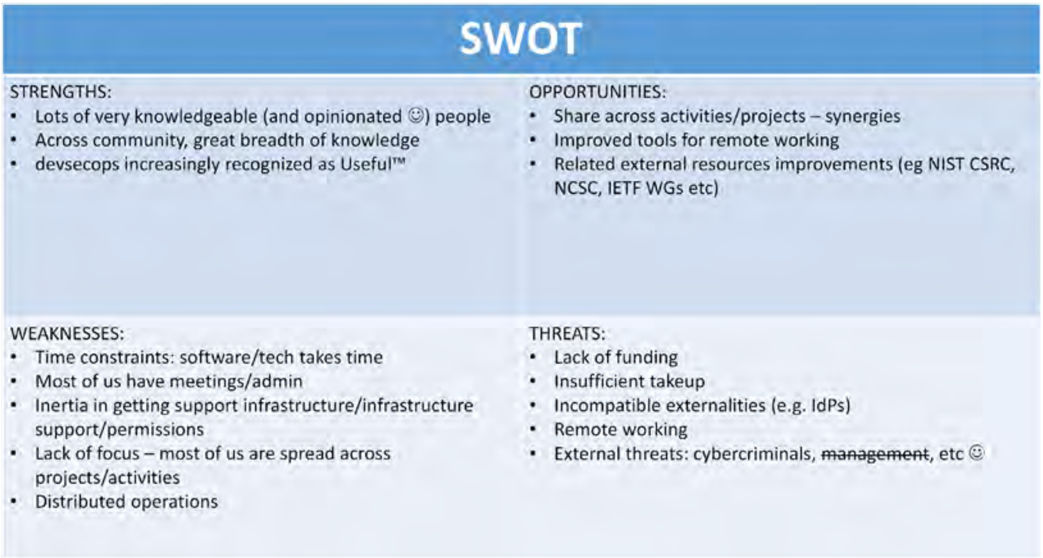

SWOT

STRENGTHS:
- Lots of very knowledgeable (and opinionated ☺) people
- Across community, great breadth of knowledge
- devsecops increasingly recognized as Useful™

OPPORTUNITIES:
- Share across activities/projects – synergies
- Improved tools for remote working
- Related external resources improvements (eg NIST CSRC, NCSC, IETF WGs etc)

WEAKNESSES:
- Time constraints: software/tech takes time
- Most of us have meetings/admin
- Inertia in getting support infrastructure/infrastructure support/permissions
- Lack of focus – most of us are spread across projects/activities
- Distributed operations

THREATS:
- Lack of funding
- Insufficient takeup
- Incompatible externalities (e.g. IdPs)
- Remote working
- External threats: cybercriminals, ~~management~~, etc ☺

The most concrete thing you can to now is to join the IGTF keybase channel (https://keybase.io/igtf) for a friendly and informal way of getting collaboration to work!

Operational matters

- We observe more CAs throwing in the towel - and at the same time a lower response to (autonomously) performing self-assessment and their presentation. Also attendance from low-volume authorities that joined the PMA in the early 2010's is down. Asking them if they do want to continue might be useful, and not implicitly extend the timeframe for performing self-assessments, since pushing them may elicit either a response or closure.
  The timeframe (every 2-3 years) is fine, and those that are still issuing we should be pressuring. "It's not difficult, and one of the few things the PMA actually *has* to do". And it's an effective way to measure response.
  Since there are now quite a few authorities to prod, we prioritize those that also have not joined recently: AEGIS, ArmeSFo, BYGCA, DZ-eScience, HellasGrid (mostly personal certs left), IRAN-GRID, NIIF (ask if for them still relevant or the move to TCS is complete), QuoVadis (better moves to TAGPMA since DigiCert is also there), TSA-GR from Georgia, and the UkrainianGrid CA.
  The Chair, together with the Self-Audit Officer, will take action here.
  In case of non-response, these CAs will be suspended in accordance with the suspension review committee process and the PMA bylaws.
- The suspension review committee henceforth consists of: Jens Jensen (replacing Ingrid on retirement of the GermanGrid CA), Dave Kelsey, Jan Jona Javorsek, and David Groep (who all three were already on there).
- The KIT (FZK) GermanGrid CA will be retiring soon, with subscribers moving to the DFN-offerings (and to GEANT TCS).

Attendance

We thank the following people for the extended attendance and stamina for sitting through the virtual meeting (in random order): Eisaku Sakane, David Kelsey, Anders Wäänänen, Ian Collier, David Crooks, Jana Zraková, Jule Ziegler, Feyza Eryol, Lidija Milosavljevic, Maarten Kremers, Marcus Hardt, Miroslav Dobrucky, Mischa Sallé, Ian Neilson, Reimer Karlsen-Masur, Scott Rea, Jens Jensen, Adeel-ur-Rehman, Nuno Dias, John Kewley, David Groep, Nicolas Liampotis, Jan Chvojka, and Bill Yau. But especially Derek Simmel for joining so early in the morning!