EUGridPMA 49: IGTF, EnCo, and EOSCH ISM meeting

Wednesday May 13 - Friday May 15, 2020

Dear all:

The 49th EUGridPMA and 1st fully-virtual PMA meeting has now started. Looking at the screen makes one realize all the more poignantly that trust is also about human relations. I'm glad that the some of the spontaneity and sparkle of the in-person trust building remains also in hyperspace!

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 that are attached to the agenda pages at https://eugridpma.org/agenda/49 and linked therefrom.

The next meeting will be September 7-9, 2020 (and will be followed by an operational security meeting in the EOSCH context). This meeting is scheduled to be in Amsterdam (Science Park), the Netherlands, hosted by Nikhef. Please reserve these days (fully) in your calender!
However, due to the still-uncertain situation:

- ample possibility for remote attendance will definitely be provided.
- the room for the meeting is very large (approx. 20x 15 metres), which should be ample for the number of attendees expected and still keep the prescribed 1.5m distance
- at the moment, gatherings of up to 100 people in the Netherlands are foreseen to be possible in September
- please do not make definite travel plans yet. First of all, it is unclear which modes of transportation will be operating from which departure places, nor if sudden lock-downs reoccur. That will become clearer over the coming period.
- In case the meeting cannot be in-person, it will be fully remote again

In January 2021, we will meet in Garching near Munich, kindly hosted by LRZ and Jule Ziegler!

Operations in a connected, and yet disconnected world

With national or organisational lock-down measures affecting many, challenges emerge for both continuous operation (such as CRL issuance) as well as for identity vetting. The first one is particularly important for off-line and air-gapped operations that need staff present to operate the key material. The challenges of identity vetting affect all who implement higher-assurance operations for, in particular, end-user credentials.
Most authorities are handling it pretty well for the moment. Some initial trouble with issuing the CRL occurred in Latin America, India, and Bulgaria, but there are now all addressed.
The larger CAs with enough staff separated into A and B teams that do not physically meet each other in order to retain resilience in case one group gets infected. When one does get infected, though, e.g. because one person visited a data centre, then the rest of the team may be held in quarantine for days until the suspicion is cleared.
Similarly, it now pays off to have remote heads for your HSMs accessible from trusted places, yet changing processes 'in a whim' must not be done lightly. There has for now been no case where practices actually had to be changes (making such changes would break an audit anyway).
In several cases, the CA either has been designated as organisation-critical, or it is co-located with national critical infrastructure and thus access for staff to the facility is guaranteed that way. Of course a designation only as organisation-critical cannot circumvent government measures, but that situation has not occurred.

There are many factors to consider going forward, and having a solid business continuity/disaster recovery plan in place is non-trivial, and making changes to implement one is hard to do if a lockdown in already in place. But since the lockdown may be periodic, once it is lifted (even if only for a while) is makes sense to consider backup plans. Jens' presentation at
https://indico.nikhef.nl/event/2336/contribution/4/material/slides/1.pptx
shows all the relevant considerations, also for air-gapped CA operations. The cost of proper BC/DR measures, which is clearly necessary as we see now, may also be a factor to consider whether the

cost per subscriber is still working out effectively and efficiently.

Remote identity vetting

Vetting of identity to a sufficient level (typically: Kantara LoA2, REFEDS Cappuccino, IGTF BIRCH/CEDAR) was already an issue before isolation measures came into being. In the ELIXIR AAI context, Mikael Linden has been working with SisuID to perform stronger identity proofing based on photo-ID documents scanned with a smart-phone app in combination with facial recognition/comparison services. This can be made into a completely self-managed solution for the applicant, as shown in the screencast videos "Registering SisuID" https://voutu.be/gAxYgh4GR8g and "Using SisuID in ELIXIR Beacon Network" https://voutu.be/OMawkudVMpo
After enrolment, the app itself *becomes* the 2nd factor for authentication.

Apart from identity vetting and giving the name through an OpenID Provider, it can in addition give other details from the ID document (e.g. nationality, even if use and interpretation thereof may be infrastructure-dependent).

- the UK government "azure" level authentication does it in a similar way, with just scanning the picture of a photoID (and not yet reading the digital content with NFC)
- *Account linking* remains a challenge, yet it is also a separate risk. Governments and banks have an advantage here since they are allowed to process 'social security number' (personal number) like identifiers to link to back-end databases (either databases they themsleves have with that number, or verification in government databases)
- SisuID will be reading the NFC data in a next version
- Although government photoID may look somewhat different in each country, most of the first page and the NFC readable data is ICAO standardized, and that can be leveraged by a single app for all countries
- There are others experimenting with this (SURFnet with SURFsecureID), or using it in production, e.g. the Dutch government for DigID (yielding eIDAS Substantial even without the photo check!), online banks (AEGON), etc.
- Some apps, like the AEGON one, do the facial comparison inside the app (apparently)
- the photo-matching for SisuID is now done with a commercial provider (of which there are several). the rest of the system is fully open source

Use cases abound. CERN has a similar need even for resetting the other multifactor in case staff cannot come on-site. Research and e-Infras have the same, and we see the same being offered as a service by NRENs like SURF. So can a *shared* system be designed and offered? There are some challenges then:

- technologically: do these things work and scale? This could be tested in the GN34 "Trust and Identity Incubator" (Mikael is already talking to Niels van Dijk about this one
- is it better to host our own one, or buy it as a service from a (public or private) provider? E.g. from Signicat (Norwegian) or ReadID/Innovalor (Dutch)?
- can we build a business model around a single shared service? That last one is interesting, since relying parties are distributed across many countries, and the key use case is vetting of people that are - by design - further away. Having national boundaries in such a scheme would be very detrimental. So the model must not allow for people to 'fall between the cracks'. And organisations must be willing to contribute to the scheme (like ELIXIR is doing now as well).

There is a lot of interest to proceed: start with the exploration in the T&I Incubator (already started), and explore the business model with also research infrastructures as stakeholders.

Changes to existing practice for remote vetting

For the existing CAs, those in the UK, NL, and SI have updated their CP/CPS to accommodate remote vetting and more streamlines ways to send documents between applicants, RAs, and CA.

TCS Gen4 status for IGTF certificates

As of May 1st, the 4th generation TCS service using Sectigo as a back-end provider is the only option available by TCS. The IGTF requirements have been (ruthlessly yet effectively) been prioritised, and it is now not impossible to get compliant IGTF OV SSL server and MICS Personal and MICS and classic robot client certificates from the service. The end-user client certificate enrolment portal works nicely for IGTF client certificates (and mostly for the email-only profile as well).
All other elements of the service remain much more challenging, and IGTF SSL server certificates highly prone to inconsistencies. All working elements of the service, as seen from an subscriber organisation standpoint, are mentioned in the slide deck.

ECC certificates and trust chains

As an important result of TCS Gen4, EEC certificates are now first-class citizens, and it will not be long before users will "just to try it' get an ECC client cert and try to use it in the infrastructures. How traditional middleware reacts to this is only partially known. Some softare will ignore ECC trust anchors until an actual ECC client comes along. Others will see ECC immediately (like voms proxy init tools), and need a modern version to deal with it (EMI UI 3.7 or higher seems to work).
There has been limited testing, but keeping ECC out of the distribution much longer will result in client failures for those that use ECC. And some groups, including e.g. OSG and BrianB in the US, are keen to move anyway.

The plan is to introduce the ECC trust anchors in 1.106, unless the designated testers Mischa and Uros find showstoppers. Then we'll see what happens. WLCG is of course concerned for breaking things, but we cannot fully know beforehand (as we saw with the retirement of the expiring AddTrust CA).

RCauth.eu distribution update

The cmobination of STFC, GRNET and with Nikhef in a supporting role, is working with support from EOSChub to deploy the distributed RCauth.eu service that will be highly-resilient and available. There have been some effects from the COVID19 lockdowns on progress, as physical installation of HSM hardware is delayed and initialisation operations deferred. The slide deck presents the detailed plan going forward in order to complete the work in EOSChub.

Enabling Communities with Trust and Identity

The GN4-3 project supports key communities that enable smooth interoperability between research and e-Infrastructure communities through significant enabling investments in WISE, FIM4R, REFEDS, IGTF, and the AARC community. Maarten Kremers reviewed the work of Enco and its activities, and progress in the many areas covered (as described in the minutes attached to the agenda item itself) enabled new work in the policy areas for WISE SCI, the security coordination challenge SCCC Joint WG, Assurance Profiles, Sirtfi, AA Operations Guidelines, Policy Development Kit, and OIDC for research federations.

The guidance from AARC is still hard to be 'just picked up' by new communities and infrastructures: the question "OK, and what do I *do* now?" is often not easy to answer. The AARC-in-Action use cases were conceived to answer that by giving examples that 'closely matched' common community use cases, but although the use cases are there, they do not answer detailed technical questions.
Even just keeping up to date with new AARC developments is non-trivial, since a 'subscribe to updates'-like button is not there. It is also not simple to answer, since in many cases taking AAI 'as a service' from an existing provider may be the better choice.

The proposal by Hannah to give some specific technical flows with real examples ("if you are in this flow, the stuff that you should send over the wire looks like this") was much appreciated. Alongside the 'decision tree support' flow: if you want this, do this.

Community policies and the PDK

The current top-level PDK document is very 'thin' and appears to the uninitiated as almost content-free. It was designed on purpose for Snctfi based communities where much of the 'meat' had

already been taken care of by an established infrastructure. So it defines context and index, and was never aimed at the top-level e-infrastructures (UKIRIS, EGI, SURFDNI, &c).
A new guidance document is useful to actually define more of the details concretely, give guidance as to what *an infrastructure* should do and have, and make that into a WISE document.

The WISE Community is an essential element in this puzzle to ensure it appropriately reflects the diversity of infrastructures in Europe, and through WISE gains adoption and endorsement from the whole ecosystem. Dave can initiate the WISE work under the SCI WG, contributors come from the UK, SURF (Maarten), DavidG, and Uros.

Ian reviewed the Community policy suite and proposed the "Combined" community security policy at <https://docs.google.com/document/d/1LxfkL9mIhtKNpkGRCuzxM7mQXJA14eysZmEIIcSABYs>. The current DPK version is split in two elements, but with many communities either searching for an integrated 'solution' to what they have to do, or actually have no direct way of doing their own policy (or wanting to!), the split makes it unnecessarily complicated.

Scoping the new policy will be important: it is for mid-sized and large communities that have a significant number of members and a non-trivial structure. For small communities, this will have to be taken care of in the 'business relationship' between the provider of the membership service and the principal investigator for the community. This same push is seen in the feed-back on the G048 AAOPS guideline.

In addition, terminology will need to be explained. For exmaple, the 'management contact' for a community apparently often difficult to define, yet in many cases the PI receiving the resource grant *is* the manager of the community and is - because of granting processes - already responsible. Such guidance should be made more explicit before dropping into the document itself.

Details were noted in the document during the meeting.

WISE SCI, the future for Security for Collaborating among Infrastructures

The WISE (WISE Information Security for E-infrastructures) SCI working group has a new slate of work, including the community policies, an update to the SCI assessment framework itself ("SCIv3"), alignment with Sirtfi (also version 2) and developments such as "Service Layer at the Edge" (SLATE) that changes the way communities and community service management overlays interact with sits and resource providers. DaveK reviewed the future of all the SCI activities (as seen in the slides attached to the agenda).
One of the more obvious elements that is not mentioned in SCIv2 is the existence of identity federations as a source of authentication and identity information - since the concepts in SCI are focused around the resource-providing infrastructures (mostly in their role relying parties). This is something that of course needs to be addressed. Similarly feedback from the SLATE folk (e.g. Chris and Tom) was highly valuable.
Could Sirtfi version 2 be actually incorporated in to SCIv3, so that it aligns again? The timeline for both pieces of work is different, and incorporation might result in delays, so a good alternative is to refer to Sirtfiv2 from SCIv3, thus decoupling the timelines.

To keep the work manageable, the focus of the SCI WG in WISE will be on the community membership policies first.

Attribute Authority Operations (G048): review of Infrastructure feedback

The major infrastructures, through the AEGIS group, have given valuable feedback to the AAOPS Guidelines (AARC-G048) that was published last year and presented to AEGIS in January. Feedback was predominantly from GEANT, EUDAT, and XSEDE, and can be reviewed on the google doc linked from the G048bvis evolution wiki
https://wiki.geant.org/display/AARC/Attribute+Authority+and+Proxy+operational+security
The comments fall into roughly three categories, and are summarised in the slides presented (https://indico.nikhef.nl/event/2336/contribution/15/material/slides/0.pdf):

- questions around logging, auditability, integrity, and the need to participate meaningfully in incident response
- questions about relationships between the involved parties: communities, AA operator, and relying parties/service providers, and which party has recourse to which other party (or is left holding risk without commensurate recourse)
- to what extent communities are always autonomous, or whether the hoster takes on and controls more of the information (e.g. if there even a community info web page?)
- lack of clarity about the operational environment requirements (the text appears to drive towards on-prem hosting, which was not intended)
- specific 'parameters' have values in the document that are either too specific or are not substantiated with a rationale
- legacy in wording that makes it unclear (e.g. are multiple AAs providing redundancies and thus must have the same content, or are they 'stacked' and augmenting eachother, and thus should have just non-contradictory release? The document was never clear here!)
  Similarly, there are named examples that should be generalised.

The comments were reviewed and commented on, resulting in a rough consensus and a lot of ideas to improve the text.

The detailed deliberations are - in as far as possible - captures in the hand-written notes (also available on-line), and a smaller editorial group consisting of Uros, Maarten, DaveK, and DavidG will draft responses to the comments, preferably after having had a meeting with Christos for GEANT, Sander for EUDAT, and Jim for XSEDE, to clarify the comments and put them into context.

Making Identity Assurance and Authentication Strength Work

Determining the appropriate assurance level based on a risk assesment of the service being protected is an approach complementary to that taken by the providers of authentication services that can provide assurance at just one, or a few, levels, or provide step-up services. But many services in the research world are not easily assessed, and the operators of those services not versed in (or have time for) doing a full risk assessment. Also, many services are 'similar', and traditionally the result has been that most chose the 'default', medium, assurance level. In e-Infrastructures like WLCG, EGI, OSG &c, that all protected similar compute and data services until now, this resulted in use of IGTF BIRCH/CEDAR, REFEDS Cappuccino, and a 'F2F or equivalent' assurance being used everywhere.
Now that there are additional profiles, and (AAI proxy) software implementations capable of dealing more easily with multiple assurance levels, it can be more specialised per service.
The Assurance Paper for ISGC, aiming at giving use-case based guidance for assurance taking into account existing risk-based approaches as well as a 'grouping' approach to common assurance levels per service cvategory, was reviewed. Use cases were discussed, including both 'lower levels' like in the VO portal policy suite from BiG Grid and EGI, as well as 'higher levels' that employ step-up services for biomedical cases.
But then, we also see a push for lower assurance levels based on purely individual (per-site/per-centre) considerations based on cost, that do not include cross-domain and federated risk assessments into account when pushing for 'cheap' assurance (just "because that is what we can do"). This at times makes it not even comparable to the '6-point assurance baseline' from AARC MNA3.1.
Assurance at any decent level should at least make it possible to effectively suspend/block credentials that have been compromised (or can be created at-will) in order to stop persistent attacks against services.
But it's hard to 'measure' required assurance unless one can 'valuate' the service used.
And even if ISGC is deferred again till 2021, Jule will continue as the primary editor for this very important paper and get it ready!

Jens' Soapbox: 'organizing people, groups, & work in a composite ecosystem'

Our work area is a large and complex space, spanning Policy, Technilogy, and Operations. But will all of them ever meet? And do they meet in a structured way, or because of organised, or haphazard overlap of people between all these groups? And why are people attracted to some

work areas and not others, or why does sharing of software almost never seem to work?
The slides from Jens' soapbos give a unique insight into the complexity of our ecosystem. And as a side effect show why the security group appears to be so overly large, whereas in fact it's just the same small set of people populating all those different groups and committees (and whether that's something good or bad, one can debate endlessly over drinks).

But in the end: is topical coverage complete, and is it consistent? Is it still possible to het a comprehensive overview and 'know what you need to know' when it gets to dealing with incidents and operational evolution? Are we speading ideas merely (or even because) of attendance in meetings, even if participation is not always very active? Or are we lacking elements because all those groups and committees are not in the 'right place' in the ecosystem diagram on slide #5?
We must look for the gaps!

For general communications, lists like the igtf-general, or dg-eur-ca, or refeds, are fine, yet sending operationally (or security) sensitive information may not work out well. There were options, like NCSA's SLES encrypted mail service, but through disuse they typically because unmaintained and defunct rapidly. And it is true, that in the last ~20 years so specific channel has actually been needed. A trusted list of contacts in maintained (in the PMA internal databases), and can be used by key people, but it's not automated. But when it *was* needed, it worked.
And as an alternative, one can announce on the public list that confidential communication is about to take place, and then invite people ad-hoc into a platform. And then it's at least (for a time) in active use and working.
Unused stuff will deteriorate and be unusable when it's actually needed. So there will not be a new list …

Operational matters

- The self-assessment status was reviewed for RDIG, MD-Grid, and UKeScience. The UKeScience peer review is now complete. There are still open issues left for RDIG and MD-Grid.
- A largish number of authorities will need to complete a self-assessment in the very near future. Please review your status at https://www.eugridpma.org/members/internal/display

We thank the large group of people who participated persistently in the sessions by video: Adeel-ur-Rehman, Anders Waananen, Baptiste Grenier, Bill Yau, Cosmin Nistor, Daniel Kouril, Dave Kelsey, David Groep, Dennis van Dok, Eisaku Sakane, Eric Yen, Hannah Short, Ian Collier, Ian Neilson, Jan Chvojka, Jens Jensen, Jule Ziegler, Lidija Milosavljevic, Maarten Kremers, Miroslav Dobrucky, Mikael Linden, Mirvat Aljogami, Mischa Salle, Nuno Dias, Scott Rea, Uros Stevanovic, Jan Jona Javorsek, and John Kewley.