# EUGridPMA 50, jointly with
# AARC Community, IGTF, GEANT EnCo & EOSChub ISM

*Nikhef Amsterdam, Monday September 7th - Wednesday 9th, 2020*

Dear all:

The 50th EUGridPMA Amsterdam PMA meeting is now over. Combining both an in-person and virtual setup (and one of the first in-person meetings in months) was a positive experience. And it was good to also meet in-person again. I would like to take this opportunity to also thank the Dutch National e-Infrastructure coordinated by SURF for their support, that - together with Nikhef - has helped enable this combined meeting. And I'm glad that the some of the spontaneity and sparkle of the in-person trust building remained also for the video-participants. Your stamina is to be commended!

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/50 and linked therefrom.

The next meeting will be January 2021, and this meeting is tentatively scheduled to be in Garching near Munich, Bavaria, hosted by Jule Ziegler of LRZ. However, due to the still-uncertain situation:

- ample possibility for remote attendance will definitely be provided
- in case the meeting cannot be in-person there, we may have to relocate

## WISE Security for Collaboration amongst Infrastructures SCIv2 updates

Uros reviewed the state of the SCI assessment/maturity model and how help can be provided when filling in the assessment scoring sheet. And at the same time the annotated version also gives hints as to what the 'good practice' implementation could be. Actually, these two aspects of the same coin are difficult to match in the same document, and the 'good practice' recommendations are likely better addressed in an evolution of the AARC Policy Development Kit.
As a matter of order, it was agreed to start up a series of targeted video/phone-meetings (1hr per two weeks) to make progress both on the assessment guidance and on giving good practice advise in the WISE SCI working group (on which DaveK is going to take action)

*The comments column of the sheet was intended to be filled in by its 'users' (i.e. infrastructures), either with a reference to existing documentation (which would imply maturity level 2+), or by a description of the operational implementations (for 'level 1'). This would constitute a self-assessment of the security maturity of the 'infrastructure' - the entity filling in the sheet.*

The group proceeded to review some of the specific elements of the assessment sheet:

- OS3 "Security Plan": in practice almost nobody actually has a complete one, in the sense of having a (top-level) document that actually captures all of the security-relevant elements of an infrastructure as a whole - although partial documentation usually exists. This does point to the fact that maybe the requirement can be satisfied in other ways, relying more heavily on examples than documentation here. "you want answers more than documents", was one of the statements made.
  To some extent this is similar to the 'breadcrumbs' we did not need for our assurance profiles, and the need for (ISO27k-inspired) documentation similarly is mainly needed for infrastructures and environments where no prior shared understanding exists. But in its current form, there is a difference between 'the plan' and the 'shared direction'.
  *In the guidance/FAQ document, suggest to give more examples of how this was done rather than instructions.* The peer-review process of the self-assessment can compensate for the documentation of the plan, since peers will have - to some extent - this shared understanding.
- The difference between PRU* and PRC* should be clarified, in that the scope and the entity responsible for implementing the PRU and PRC items is different. The Community needs to implement the measures to make RPU happen, and the Infrastructure should watch over the PRC items.
  *This needs to be clarified in the guidance, since if the SCI group itself is confused, the readers certainly will be …*
- Guidance about the individual elements can also be taken from elsewhere. For example the construction of the AUP, and how to both write and implement it, is provided for in AARC-I044 (https://aarc-community.org/guidelines/aarc-i044/). The way PRU2 was written meanwhile was overly restrictive, in that it required *the community* to run a registration process. The process necessary to make the user read the AUP can of course also be implemented in different ways, as described in I044, e.g. by having the shared community management platform do that before the user even joins any specific community hosted thereon.
  All PRU* requirements are actually about the AUP, but these are directed at the community and the infrastructures, *not* at the users. So the "User" requirements are *about* the user, not something end-users need to do.
- For the Data Protection (DP*) requirements, it should be clarified that this concerns *only personal data collected as a result of users using the infrastructure*, and *not* research data or content that happens to be personal data (like research data about people). The DP* requirements thus address the same "access management" personal data that the GEANT DP CoCo is concerned with (and come to the AARC/GEANT CoCo &c groups for guidance)

The ultimate goal of the guidance/FAQ is to help people fill the SCI assessment sheet, and the "how" to meet the requirements is better answered by the Policy Development Kit (which is also undergoing an update). In the assessment sheet, the infrastructure reps are going to write down what processes they have chosen to address each item (either by documentary reference or by description). This help is likely best provided in a hyperlinked spreadsheet, better than in a separate document - and should be significantly more lightweight than the list of NIST, or ISO, controls. "*if you have a document point to it, otherwise, describe the process in the [how column of the] sheet*". And explicit questions are simpler to answer either in a separate row or rendered as a 'hovering infobox'.

The benefit is ultimately in a shared understanding of what SCI is bringing about amongst the infrastructures, and it's the discussion that probably brings the most value, not the box-ticking exercise. Encouragement to join the WISE community is therefore encouraged, although we realise that the initial EOSC baseline of just having contact information is probably

what we'll have to work with for many of the parties involved for the time to come.

## AARC Policy Development Kit revision

The AARC Policy Development Kit (PDK) aims to structure policy components and help infrastructures, services, and communities on their way to quickly get a complete policy suite. At the same time, because of its comprehensive nature, it has a steep learning curve for communities that have not been previously exposed to the need for policies, or to the 'JSPG' structure. As part of the (UK-IRIS supported) effort to make the PDK more accessible, Ian Neilson made the comparison between the top-level policy in the AARC PDK and the comparable policies in some of the other infrastructures: EOSChub, UK-IRIS, EGI, and the AARC PDK - using a per-sentence appoach

https://indico.nikhef.nl/event/2490/contribution/8/material/0/

While all of these have inherited the JSPG structure, there are differences (e.g. EOSChub addressing collaborating infrastructures more than the others who scope is a single infrastructure), both in terms of detail as well as in the exact wording used. For example, the AARC PDK version appears leaner (with more 'breadcrumbs' missing), but presupposed a lot of other policies and structure.
Based on the structuring that Ian has done for UK-IRIS, we can now build a 'modular' PDK where infrastructures and communities can compose their own policy suite based on individual components - as long as a guide as to *why* a component can be included should be provided alongside the actual text.
The limitation of the current analysis is that all policies considered share the same JSPG parentage. Through the WISE community, it would be instructive to compare policies that do not distribute responsibility to the participants but follow a different course. The PRACE or DIRAC/UK infrastructures could provide useful examples here - and it is good to do that mapping in any case.

A top-level policy for infrastructures need not be exactly the same for all infrastructures for the infrastructures to be interoperable. A top-level policy is likely the place where they will diverge the most - which is fine as long as the components are interoperable and it is clear which components are, and which are not included in any particular infrastructure.

The PDK, starting with the top-level policy, will be restructured along the lines drawn by Ian.

## CA updates I: TCS G4

Getting IGTF compatible certificates from TCS is now possible when your place (or country) name has exclusively IA5String (ASCII) characters *and* all your eligible users can authenticate using eduGAIN. If either of these does not hold true, you are still in trouble. Experiments in .SE are still ongoing, but - based on experience in NL - is is possible to work around the non-ASCII character issues by essentially duplicating the work at the subscriber end and registering *two* organisations for the very same entity: one with the UTF-8 name, one with an ASCIIfied version. This is known to work, and - by now - has also been confirmed by Sectigo.
Generating end-user credentials works now for both RSA and ECC, and in either case the user can choose to upload a CSR or to have the key pair generated on the CA side and delivered as a PKCS-12 blob. It does have a potential downside that the CA has seen the key (but of course the CA could as easily have generated other keypairs), which a client-side solution to the lack of KEYGEN based on JavaScript does not have. UKeScience has a JavaScript model now, for instance. However, PKCS#12 is more compatible across platforms (like mobile, IE, and text-mode browsers). For infrastructures that wish to develop a hosted key management platform (like the portal CESNERT developed, or a MasterPortal does in an RCauth scenario), using the API is likely a better choice. It is a pity that getting IGTF certs through the API does not quite work, as is retains RDN attributes that can be removed via the GUI, but where the necessary functionality lacks in the API (!). Or at least in the API that TCS subscribers can use.

Other (non-IGTF) issues are also still open, as listed in the presentation
https://indico.nikhef.nl/event/2490/contribution/3/material/slides/0.pptx

## Security and the European Open Science Cloud in the future

The European Open Science Cloud (EOSC) will be a main structuring element of the European Research Area in the 2020s, and many complementary activities, project, and entities are working on creating elements of the EOSC ecosystem, both at the 'core' as well as in the ecosystem as a whole. And, while the exact nature and future of the EOSC remains still a subject for debate, it will have an impact on trust and security - and on the way security and integrity can, or cannot, be maintained. Following on from the whitepaper "Trust Coordination for Research Collaboration in the EOSC era" (https://doi.org/10.5281/zenodo.3674677), the EOSC-Future has incorporated some of these ideas in its foreseen operational activities.
In the discussion following the presentation and the white paper, the risk-inventory approach method and the use of baselining were reconfirmed as useful key elements - as long as recourse to an actionable operational core team of CSIRT/forensic experts is available for hard cases. But with the low-barrier-to-entry encouraged by the (Rules of Participation) working groups, real and urgent risks will remain commonplace. At least a baseline including contact information, and hopefully an agreement that confidential data shared during incident response will be treated as confidential by those receiving it, a base level of integrity may be maintained. Through the task forces of the Architecture WG, additional security maturity elements can be proposed and incorporated.
Baselining activities in ecosystems like eduGAIN and InCommon, as well as the *Sirtfi* work, has shown that - at least over time - adherence to a baseline can be achieved and mature organisations are willing to join in common agreements. The use of 'trust marks', combined hopefully with (user-review or AI supported?) scoring, may encourage participants in the ecosystem to adopt good risk management and security practices. Hopefully even pushing with the rope that links the

providers to the EOSC will have some effect, through setting 'good practice' examples.
Matt confirmed meanwhile that the onboarding requirements for service provider sin the eosc-portal.eu today (and in the foreseeable future) will be very lightweight: at least TRL7, having a contact address, and having a place to send issues/bugs. Hopefully, we can add to it at least the requirement to treat confidential data confidentially …

## Security Communication Challenge Coordination SCCC JWG

Whilst some communities are actively proving and validating contact and response information, this is still far from commonplace. The SCCC Joint Working Group maintains a (for the moment small) list of infrastructures and groups that conduct such challenges (at https://wiki.geant.org/display/WISE/SCCC-JWG), and some federations that are also part of eduGAIN run them themselves (like SURFcert against the SURFconext participants), but it could do with a reinvigoration.
There are tools, including a self-monitoring mailer for automated testing, available, and offering that either as a toolkit, as a service, or both, may help getting more readiness assessments done, as well as making them easier to do.
When deployed in a federation context, these should likely be run by the federations themsleves, but making the fact that these are done more public can help 'pour encourages les autres'. This *need not* be cast as a 'security' thing so as to make it more palatable for federations - to be discussed with Davide Vaghetti and Marina. Having Hannah's experience with Sirtfi there as well, will almost certainly help participation and acceptance. Also the yearly challenges by SURFcert can act as a good example to follow.
Meanwhile, the IGTF will run its next RATCC5 soon - it has been a year since.

## OpenID Connect for (Snctfi) infrastructures

One of the bigger challenges facing OpenID Connect Federation (OIDCfed) today, is the large amount of fluisity in the specification itself. While lots of changes are being made to the OP and RP integration with OIDCfed, also the spec itself frequently changes - not making it any better or worse, but just changing - which is beginning to bite adoption.
There have been several successful interoperating implementations of the spec draft, between the path validation codes from Henri and Jouke, but recent changes have invalidated that again (e.g. the change to have roots-of-trust *not* having an authorityHint, compared to a previous version where such a trust root was identified by having *itself* as an authorityHint). Other software implementing a different version of the draft spec then did not include loop protection (something you will need anyway to prevent also non-trivial loops), and exploded.
In addition, the OIDCfed draft specs add many new elements, which - while worthwhile in themselves - tend to address increasingly smaller market segments. Although it is nice to see that almost all elements of X.500 & PKIX are now being grafted onto OIDCfed.
It would be good to freeze the current spec, take the minimum viable subset, and call it a "version-1" standard that people can actually implement and deploy. New work can then go into an evolution of the standard.
This idea will be taken to the implementers' list for OIDCfed, the de-facto place where standards discussions happen since on the original OIDF list there is no more traction for - what is in the end - a very important spec and development for the research community and our multi-lateral federation concept. Jouke will initiate this, with Maarten and through Mike & Roland.
Basically, having an extension mechanism, and a way - immediately in v1 - to indicate 'critical' extensions that implementation that do not understand the extension make the chain distrusted, would be a nice way to progress. A phase-space-complete spec can then be done later.

There are plenty of *use cases* where OIDCfed would be very helpful to have - and G052 is just one of them. The EOSC federation will likely be another. And we have a few OPs to go around to try it out with. On the RP side, Uros' OIDCAgent might be a good RP to start as well.

## Attribute Authority Operations Security - AARC G048 rev 2

The revision of "Guidelines for Secure Operation of Attribute Authorities and other issuers of access-granting statements" was extensively discussed, based on the comments provided through the AEGIS group by Jim Basney, Christos Kanellopoulos, and Sander Appweiler. The 'rev2' version (https://docs.google.com/document/d/1-hbqSpQegm7UaC_wupFzFMm19Q024UPkG-8Jwokkmzc) aims to generalise some of the guidance, and clarify those points where the examples actually obfuscated the intent of the document - e.g. in the use of virtualised or (cloud) hosted environments.
The direction of the changes now proposed in rev2, in general, has been:

- towards explaining the *why* more than the *how*, in particular for naming of communities, attributes, and semantics
- taking out explicit (time) indications and instead referring to existing standards and recommendations. For example, the "no more than 24 hour" validity is more appropriately expressed in terms of "as short as reasonably possible" and reference to RFC 6750 for bearer tokens (suggesting 1 hour or less), to typical OIDC scenarios (15 minutes), defaults in the VOMS system (12 hours), and some edge cases (like communities setting it to even 72 hours). And the (RFC6750 stated) reasoning behind it "reducing the impact of them being leaked".
- clarifying that operational controls can also be implemented by contracts and agreements.

But the key element is that AA in the BPA are highly valuable resources and should be properly protected. It is not quite the place to take shortcuts and expose communities hosted on the platform.
*During the meeting, items up to section 3.4.1 directive 2 were reviewed.*

## RCauth.eu distributed operation

RCauth.eu is an accredited IOTA CA that takes the (Sirtfi+R&S) entities in eduGAIN as well as qualified BPA proxies as a source of identity and issues, short-lived, PKIX certificates to end-users. These certificates are then typically managed in a infrastructure credential management system (a MasterPortal, WaTTS instance, or standard MyProxy). For resilience, it is engaged in moving to a fully distributed redundant setup (with a common governance structure of SURF, EGI, STFC, and GEANT and the transition work supported by the EOSChub project). Jens, Mischa, and Nicolas described the distribution of key material and the move towards a highly-available setup spread between Didcot, Amsterdam, and Athens:
https://indico.nikhef.nl/event/2490/contribution/13/material/slides/0.pdf
Much of the key distribution strategy has already been discussed previously - and the system is now in a state that the remainder (using the actual original key material) can be done in a safe way.

For the HA setup in the end the load balancing shall ensure a persistent connection of clients to one delegation service (DS) since it retains state therein, and a HA-Linux layer + round-robin DNS RRs will work if the latency is small enough.

Anycasting the DS network would be better, since BGP is a much more robust protocol, but it is also more involved to set up - so a simpler HA setup will be tried first.

## Authentication and remoteness in the post-pandemic world

Over the past months, a lot of trust has been bestowed based on remote interactions. But many of those remote interactions have followed previously-in-person contacts, trust that has been established in a period where 'walking over' was, also literally, common place. Slowly but steadily, we are moving to a world where also initial trust is established over remote channels, and we are learning how to increase such trust by having multiple independent interactions:
https://indico.nikhef.nl/event/2490/contribution/15/material/slides/
And meanwhile we did learn some significant lessons:

- using multiple ways of identification - like we do now for remote vetting - actually improves quality over the traditional checks of a single document by a single person (possibly unacquainted with the document presented). Trust is augmented when you can correlate.
  Slide #12 clarifies also that different means of communication result in different assurance levels - and this needs to be taken into account in the final assertions
- Documentation (on authentication authority and RA processes) typically falls short of adequate, and remote operations make this much more visible and tiresome
- having 'sequential' links in a social network graph is weakening trust, whereas having many parallel paths in a graph strengthens trust - so the social graph structure is important in assigning trust to an entity (and for identifying cliques like the IGTF has in the Thawte days)

After the current pandemic, it is unlikely that everything reverts back to the old ways. Having compensatory controls and multiple trust levels that are actually tested and strong enough will be important, also when replacing wet-ink signatures with remove video meetings.
Things like checking for coercion during signing are difficult to do when a participant is remote!

*We might borrow good practices from the methods states are now implementing for remote transactions and vetting in their legislation. For the next PMA meeting, consider getting an overview of such changes from a few countries where we know what happens, as well as a view from an eIDAS perspective.*

## Jens' Soapbox: 11111 laws of computing

Complexity has to go somewhere … and any attempt to summarize Jens' Soapbox is both complex and exciting at the same time. And although my patent inability to do justice to the intricacies of what Jens presented in https://indico.nikhef.nl/event/2490/contribution/16/material/slides/ may qualify me as a computer, or maybe as a mobile device, at least writing about it may make it half-human and half-computer. As such, Jens' law of computer 011b may apply, since computers are excited by communicating with humans! And thinking about forking, or threads, also applies to me.
In the end, it comes down to making sure that the appropriate tasks are done by the right entity: and that many humans are bad at doing repetitive tasks, whereas this is something in which most computers excel. And it is on the interface that misunderstanding occurs, and such complexity and misunderstanding is almost impossible to massage away - yet asking the right question in the right way in already a good improvement.
And for as long as computers fail to appreciate a cup to tea, treat computers and humans equally, but not the same!

## Operational and other matters

- The self-assessment status was reviewed for RDIG and MD-Grid. The RDIG CP/CPS update is now ready and the peers will retrieve it for review.
  For MD-Grid, Valentin confirmed and the reviewers verified that all material changes are now complete, and the CP/CPS will be published on the MD-Grid web site shortly. Once done, the MD-Grid self-audit is then completed.
- The NERSC CA, a bit unexpectedly, keeps operating and issuing CRLs
- DigitalTrust - by the way, now owned by Digital14 but a self-managed entity whose structure remains exactly the same - has a new set of trust anchors for distribution that will move from QuoVadis to Sectigo. It will also share the same public-browser trust anchors UserTRUST RSA and ECC with the TCS G4 service.
  The RPDNC namespaces will be updated to reflect that, and the new ICAs introduced, in the next release once Scott sends them over.
- For the next round of self-assessments, authorities should review their status on the internal pages. At least TR-Grid will present in January. For all others: please review your status at
  https://www.eugridpma.org/members/internal/display

We thank our local participants for joining us in Amsterdam and sharing convivial trust building lunches and dinners with us: Jana Zraková, Maarten Kremers, David Groep, Mischa Sallé, Jouke Roorda, Gerben Venekamp, and Sven Gabriel.
And the large group of people who participated persistently in the sessions by video (in random order): Cosmin Nistor, Eric Yen, Hannah Short, Dave Kelsey, David Crooks, Matt Viljoen, Jan Cvojka, Baptiste Grenier, Feyza Eryol, Nuno Dias, Ian Collier, Ian Neilson, Jan Jona Javorsek, Lidija Milosavljevic, Nicolas Liampotis, Miroslav Dobrucky, Mirvat Aljogami, Scott Rea, Reimer Karlsen-Masur, Jens Jensen, Daniel Kouřil, Dennis van Dok, Adeel-Ur-Rehman, Bill Yau, Anders Wäänänen, Uros Stevanovic, and John Kewley.