

56th EUGridPMA Geneva 2022 meeting

Tuesday, 4 October, 2022 09:30

Dear all:

The 56th EUGridPMA+ CERN "2022" meeting is now over. Over a period of three days, we discussed a range of trust, technical and community engagement topics, ranging from joint WebPKI and IGHF trust, Attribute Authority operations security for the AARC Community and EOSC, and Assurance and FIM4R topics and the next challenges for the GEANT Enabling Communities task in 2023 and beyond.

I would like to take this opportunity to deeply thank Hannah Short and CERN for hosting the meeting - and much more besides. For those of you who missed out on the local hospitality by being remote: 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/56> and linked therefrom.

The next EUGridPMA+ meeting will be **February 13-15 2023**, again at CERN in Geneva, Switzerland, and likely followed immediately by a FIM4R meeting on the 15th and 16th. Keep this in mind when booking transport and accommodation. Avoid booking hotels at an exaggerated rate: not all suppliers have yet realised that there will *not be* a Geneva Car Show and still try to get enhanced prices. These will come down. We encourage in-person participation, but of course remote attendance will be facilitated.

Kind regards,
David Groep.

In this summary:

- Attribute Authority Operations Guideline (G071) assessment model
- A list of trusted token issuers
- Self-signed roots and changes in RedHat9 and Firefox 103+
- Name uniqueness and CA/BF Baseline Requirements in Joint Trust scenarios
- End-user client-auth certificates BRs from the S/MIME working group
- Trust store models and token transition
- Anycast for highly available stateful services (RCauth.eu)
- HPCI High Performance Computing Infrastructure authentication developments
- Enabling Communities in GEANT5+ - and for ISGC
- John's Soapbox - on digests, revocation, and paranoia
- Operational matters and self-assessment process
- TAGPMA updates
- Attendance

Attribute Authority Operations Guideline (G071) assessment model

With the increased variety of attribute sources, from community AAs and infrastructure proxies in the AARC Blueprint Architecture model, and the continuing deployment of AA proxies in the research infrastructures and the European Open Science Cloud, protecting these stores of information and trust nodes in the infrastructure is increasingly important. The growing number of proxies also poses challenges for scalability of trust. The "Guidelines for Secure Operation of Attribute Authorities and issuers of statements for entities" (<https://aarc-community.org/guidelines/aarc-g071/>) provide recommendations on how to operate a trustworthy attribute authority or proxy - and this guideline has been endorsed recently (summer 2022) by the AARC AEGIS group. The next step is to test G071 in practice with some of the major proxy operators, and establish an effective trust model around it - where our IGTF peer-reviewed self-assessment

methodology 'comes in handy'. A process that is simple (e.g. akin to the SCI spreadsheets) both allows for rapid assessment (e.g. for a token issuer trust mark), but also helps peer operators to learn from each other.

What should be assessed is the trustworthiness of the AA (and proxy), and *not* any trustworthiness of upstream IdPs. Those IdPs are out of scope for the G071 assessment.

In order to gain experience with the G071 assessment process, both were reviewed collaboratively during the workshop itself, but for each of them a different sub-set of G071 was assessed. The draft review sheets are linked from the agenda and will continue to evolve as the assessment progresses.

In the AAOPS Workshop day, we reviewed two operational proxies: the UK-IRIS and the WLCG proxies, both technically based on Indigo IAM but operating in different environments. The WLCG IAM is run by CERN, and is closely integrated with the CERN SSO system and the CERN HR accounts system since by design all WLCG users have a CERN account and can take assurance and identity vetting from there. The UK-IRIS one is operated by STFC, and leverages a rather more federated identity system that spans the STFC facilities. The IRIS one is thus more 'federated' than the WLCG one.

A template sheet is available at <https://edu.nl/88dwf> (google-hosted). The partial sheets for IRIS and WLCG are linked from the agenda page.

In addition, some useful clarifications to G071 were identified:

- in AN-1, it was unclear which identifiers should be chosen in accordance with the AARC Guidelines. This should be clarified as "... Identifiers *for subjects and attributes* should be chosen in accordance with ..."
- some requirements actually contain multiple statements (e.g. AN-1, AMR-1, AAS-2). In the assessment sheet, these should be separated into multiple items.

A list of trusted token issuers

Currently many relying parties explicitly add trusted token issuers to a list or configuration (e.g. in a HTCondorCE config file, in a configuration database for a proxy, &c). With a growing number of token issuers, and a more diverse range of relying parties, e.g. in the European Open Science Cloud, there is value in having a curated source of trusted token issuers.

The situation is of course slightly different from the one we had with per-country/region authentication sources, since in principle there is a single token issuer per community (as per the AARC BPA 2019) and thus a smaller number of issuers to be configured at the RP side (e.g. for WLCG there are only 4 issuers). Yet, with more RPs and a significantly larger number of issuers, the scaling challenges are comparable (e.g. in the EOSC) and slightly more process and transparency is welcome.

Considerations for such a trust list:

- the tokens (JWTs) issued have a proper place to include assurance information (acm, acr), hence there is no need to separate the list of trusted issuers by assurance level
- the OpenID federation work is partially solving a similar issue. However, it is more complex to deploy and has over time seen significant delays in coming to a final specification. Currently, there is not yet an adopted OIDCfed spec, nor are there full implementations yet. The Token Trust List is a 'policy' based, simplified alternative to OIDCfed, which is relevant for at least as long as there is no OIDCfed widely deployed. In that sense, it is comparable to the IGTF PKIX policy bridge vs. a technical PKIX bridge (like the 4BF bridges)
- It is useful to have some meta-data alongside the list of token issuer end-points. The 'native' format for this would be JSON
- The meta-data should align where possible with the OIDCfed metadata
- there should be a single list (IGTF wide), which can be taken and filtered by RPs based on accepted communities
- OIDC provider trust is necessary for cross-provider access use cases in, e.g. the EOSC - since for community trust there is in addition an explicit customer relationship with the (community) issuer

- JWKS may or may not be useful (or can be retrieved later based on the trusted URLs)

A trust list should have a defined set of criteria that drive inclusion. The requirements for inclusion were discussed, and there is rough consensus that

- the enrolment process should be clear and relatively quick (order weeks, not months)
- inclusion can be based on G071 peer-reviewed self-assessment, but mandating that could work against the 'rapidity' of the enrolment process. Hence, G071 compliance should be seen as a 'trust mark' (encoded in e.g. a `policy_uri` in metadata or 'entity category' in the meta-data)
- the IGTF trusted web site could host the list (aggregating where useful or needed)
- having a security contact (email) in the metadata is useful (akin to Sirtfi)

The procedures for adding token issuers to the list were also discussed, but no conclusion has yet been reached. It could be managed on (e.g.) github with pull requests (requires technical expertise), through an enrolment form (used together with registration in other systems, e.g. EOSC) in a way that is like the REEP service, or other means.

The list can also be populated by aggregating other information sources, such as the EOSC list of issuers, national hubs (like SRAM in the Netherlands), and direct submissions. This is effectively the 'out-of-band' equivalent of OIDCfed cross-signing.

Some software support on the RP side (like a filtering tool to generate a list of token issuers based on certain JSON filtering criteria) may be necessary or useful at a later stage.

This is also discussed on the AppInt ("Technical AAI Forum for Architectures & Application Integration") list, and interested people SHOULD subscribe at

<https://lists.geant.org/sympa/info/appint>

Self-signed roots and changes in RedHat9 and Firefox 103+

While all intermediate and end-entity certificates in this PKI are using SHA-2 based digest algorithms, a significant number of self-signed root certificates are still using SHA-1. This is not dissimilar from the public WebPKI, where there are also a sizeable number of self-signed roots that are SHA-1 based (and some of these roots are common to both the public web trust and the IGTF). For security and integrity purposes, the digest algorithm used for self-signed roots is immaterial, since these certificates (and hence their public key) are explicitly trusted.

We observed that self-signed roots that use the SHA-1 digest are no longer accepted in RedHat 9, unless the 'LEGACY' policy is set, or `update-crypto-policies --set DEFAULT:SHA1` is used. However, at the same time SHA-1 based self-signed root certificates as shipped in, e.g., `ca-certificates-2022.2.54-90.2.el9.noarch.rpm`, and (via that package's post-install scripts) installed using `p11-kit`, are accepted without issues and fully trusted.

We have not been able to identify in which way the system trust store (accepting SHA-1 self-signed roots) and non-system trust stored ("`X509_CERT_DIR`", *not* accepting such SHA-1 self-signed roots) differ. The expected behaviour is that a custom trust store of CA certs can also include self-signed SHA-1 roots, using the system-provided OpenSSL, and where the trust store is configured either in code (i.e. an `X509_STORE` created with `X509_STORE_new()` and then populated with `X509_STORE_load_locations()`) or by setting `X509_CERT_DIR`. To maintain overall security, this should otherwise retain the secure settings of RedHat 9, not allowing SHA-1 for intermediate and end-entity certificates.

This behaviour shows up, with tickets in ATLAS and OSG, where file transfers fail.

Also Firefox 103 changed its acceptance for self-signed SHA-1 roots, when they are in the Software Security Device rather than a built-in object token. The error message there is even more confusing (`bad_cert_domain`), and moreover when used in combination with HTST cannot be overridden to act normally again.

Hannah has kindly offered to open a ticket from CERN with RedHat, and we are awaiting an answer.

However, any resolution (even if RedHat were to do the right thing and make the digest immaterial for self-signed roots) will take a lot of time to percolate through the update system.

Meanwhile

- all IGTF CAs that can be recommended to re-issue their roots with a SHA-2 (e.g. SHA-256) digest
- Maybe also DigiCert can provide a private variant of the DigiCert Assured ID Root that is SHA-256 signed (given that this hierarchy was recently used for the new SHA256 ICAs and the Swiss)?
- the two unused intermediates (DigiCertGrid) ICAs should be withdrawn (there is a new hierarchy in place)
- all CAs listed in the presentation should consider re-issuance (<https://indico.cern.ch/event/1181342/contributions/5076241/attachments/2520809/>)
- DavidG will in the end contact them as well - some can already be changed in the 1.118 upcoming distribution (this has meanwhile been done for GridCanada)

It affects the following roots:

ASGCCA-2007, ArmeSFo, BYGCA, CESNET-CA-Root, CNIC, DFN-GridGermany-Root, DZeScience, **DigiCertAssuredIDRootCA-Root** (joint trust), DigiCertGridCA-1-Classic, DigiCertGridRootCA-Root, DigiCertGridTrustCA-Classic, GermanGrid, GridCanada (done), IHEP-2013, KEK, LIPCA, MARGI, **QuoVadis-Root-CA2** (joint trust), RDIG, RomanianGRID, SRCE, SiGNET-CA, TRGrid, UKeScienceRoot-2007, cilogon-basic, cilogon-silver, seegrid-ca-2013

Name uniqueness and CA/BF Baseline Requirements in Joint Trust scenarios

Global uniqueness of the subject DN, as provided by the IGTF model, is essential for persistently identifying end-points in the IGTF trust fabric. At the same time, having these certificates trusted for WebPKI is also critical for securing connection to services that are accessed by both browsers as well as automated agents. Such as storage services: end-users can retrieve files via their browser, but automated large-scale research transfer services connect to the endpoints to do data placement (e.g. for CERN's worldwide Large Hadron Collider Computing Grid) and must authenticate the endpoints based on their certificate subject name. To identify the endpoints, unique naming is necessary, which is what – for, e.g., TCS and InCommon - the subject DN domainComponents supplies.

There are in fact no other attribute types in the RDN that can take this role, because other attributes are either (rightfully) completely regulated based on subscriber information, *or* are not consistently stringified across platforms (in particular across OpenSSL and BouncyCastle. Experiments with either 'info', or 'unstructuredName' failed.

In June, we became aware of proposed changes in the CA/BF baseline requirements that would disallow the use of the domainComponent attribute, and a lot of discussion ensued between the IGTF community and CA/BF. This has resulted in several really important positive improvements. First of all, a communications channel has been established to discuss these matters. Secondly, we worked with the server validation working group to come up with improved wording (and discussed the validation requirements) for the domainComponents when used in the joint-trust scenarios. With PR#392 accepted, the new BR explicitly allows multi-valued domainComponents in the subject DN (<https://github.com/cabforum/servercert/pull/392>).

All conversations with Clint, Dimitris, and Tim (in alphabetical order :) were extremely useful. Thanks a lot for all your help, also on behalf of our relying parties & the research communities! The PR was merged on Thursday October 8th.

End-user client-auth certificates BRs from the S/MIME working group

There are also changes coming up from the CA/BF S/MIME working group, and the proposed changes - also again affecting domainComponent - would make joint S/MIME (mail) trust and client-auth incompatible. This is rather annoying, but in this case there is no true 'joint trust' required, since the mail and authentication use cases are more disjoint than in the server SSL case.

The changes in S/MIME BR could be accommodated, e.g. for TCS, by

- re-grafting the MICS eScience Personal, MICS eScience Robot Personal, and IGTF eScience

Robot Email onto a new, private-trust, root

- split the use cases for the Robot Email in (1) authenticating robots (on the new private root), and (2) email/list robots (which do not need uniqueness beyond what the S/MIME BRs would provide)

This would of course require an update to the TCS (and InCommon) CPS documents, but that can be done. The PMA endorses this general direction and will look willingly to a 'quick fix' once it becomes clear that that is needed. It would introduce a few more trust anchors (and a new separate root) in the IGTF distribution.

Trust store models and token transition

In the discussions in this EUGridPMA meeting, we also identified - for a (long-term) future - that there are actually three different trust contexts, but at the moment we use a single trust anchor store for all purposes:

- transport-layer protection, where the connection is initiated by an entity that has independently-verifiable foreknowledge over the expected domain name of the networked end-point to which it will connect. For this transport-layer protection, in principle DCV is sufficient (augmented with DNSsec and/or DANE)
- agent-mediated trust, where a third party (broker, agent) mediated the connection, and thus (at the only common party in the transaction) knows the identity of the parties that will actually establish a trusted connection (third-party copy scenarios for storage, where the client/FTS knows or could know the identity of both end-points). However, the protocols today (including FTP, GSIFTP) do not allow for this information to be shared consistently
- eye-ball trust, where a (human) end-user with a browser connects to a (foreknown) end-point and then authenticates with a personal identity (client-auth)

Today, in some protocols (like https) the transport-layer protection and client-auth can be separated. In Apache httpd, this is e.g. the different between the SSLCertificatePath (client-trust) and the SSLCertificateFile and SSLCertificateChainFile (which send the network-endpoint credentials to the client). But otherwise, today these three different stored cannot be disentangled in software. The problem will likely go away by itself on the move to token-based authentication, since that implicitly separates transport-level trust from agent- and client-trust. Backporting this notion to the current PKIX-based model is however very complex, and given the evolution of the infrastructure not the proper way to go. So until at least 2026 (CERN WLCG transition date), and likely later, the trust stores cannot be disentangled. And other communities, like DUNE, may want to move earlier to tokens!

DavidC, AlistairD, and MaartenL will follow up in the WLCG AuthZ context. Others welcome!

In the end, of course also tokens need to be trusted. For the 'community' tokens (e.g. from IAM, checkIn, eduTEAMS, &c) this can be done based on interoperable standards and a (to be discussed) trust list. However, (storage) systems themselves often also have their own proprietary authZ system, e.g. by way of API tokens, bespoke grants, or capabilities (thing of e.g. AWS S3), and then the tokens will need to be translated.

This should be done preferably through a plug-in approach in the data management systems, i.e. in GFAL, FTS, and Rucio, to translate 'our' JWTs into something the target system can understand. or that can be used on e.g. the HPC systems in the US that contribute resources in a bespoke way.

Anycast for highly available stateful services (RCauth.eu)

The RCauth.eu signing service has recently moved to a fully redundant, multinational setup based on IP anycast technologies and a (gallera-based) synchronised database backend over a mesh of OpenVPN protected back-end networks. David reviewed the internet routing protocols and models that make this possible, highlighting that the use of IP anycast is fairly trivial, provided that service managers and internetworking engineers are able to communicate effectively. For now, the RCauth.eu HAproxy front-end services are anycasted from Amsterdam (AS1104, Nikhef) and Athens (AS5408, GRNET), for both IPv6 (2a07:8504:1a0::1/48) and legacy IP (145.116.216.1/24).

This setup does not have any single points of failure any more, and only the BGP convergence time influences failover speed. It does not as-such bring load balancing (since no path-length traffic engineering was done), but that was also not a requirement for this service.

It also demonstrated that IP anycast for high availability is almost trivial to do. The complex bit is in monitoring and the decision process to determine, on the per-instance level, whether a service end-point is actually 'up'.

See <https://indico.cern.ch/event/1181342/contributions/5064592/attachments/2520808/> for details.

HPCI High Performance Computing Infrastructure authentication developments

The Japanese HPCI infrastructure in advancing rapidly towards new authentication models, and new uses cases are driving integration with other communities, such as the Cherenkov Telescope Array (CTA), where data are transferred automatically (using FTS3) between PIC in Barcelona and the HPCI. Access to the GFarm compute and storage facilities can now be done over OAuth-SSHD, which combines both KeyCloak and the OIDC-Agent. In this context, it is worth noting that KIT has added usability improvements to OIDC-Agent, and that the next version of the Putty SSH client for Windows will include a plug-in mechanism to then easily support also OIDC-Agent access. For file transfers between PIC and HPCI: the certificate used to identify the end-points in Barcelona is TCSG4 eScience SSL server, and hence meets the IGTF uniqueness requirements. This works fine with FTS3

With the GakuNin federation, HPCI is introducing a proxy (hub-n-spoke) model to better support assurance use cases for research, and a working group on proxies was set up to guide the development of the 'Orthros' system. Here, identities can be linked, including government and enterprise ID systems, ORCID, etc. It also includes assurance step-up.

This work is very close to the activities of the AARC AEGIS group, and Eisaku, Licia, and Christos will explore ways for HPCI to participate to the AEGIS group and participate in the development of the AARC guidelines for interoperability.

See <https://indico.cern.ch/event/1181342/contributions/5077946/attachments/2522312/> for a great overview!

Enabling Communities in GEANT5+ - and for ISGC

The GEANT's project Enabling Communities (EnCo) task serves as a linking pin, bringing support to a range of eScience and federation engagement activities in the federation and eduGAIN ecosystem. As shown on MaartenK's slides, this has brought not only support for FIM4R, WISE, IGTF, REFEDS and Sirtfi, but GN4-3 also resulted in e.g. the eduGAIN Security Handbook.

- FIM4R is being resuscitated with a focus on assurance, with a meeting coming up before I2 TechEx (Sunday December 4th) in Denver, CO, USA.
- SCCC JWG: there appears to be limited interest in contributing to this - or just a lack of security challenges. Not much effort will be devoted to this one in the future unless things change.
- OIDCfed: if a stable standard emerges, and there are implementations, that remains a worthwhile effort. But for the moment, we will focus more on the policy-bridged OIDC trusted token issuer lists.

GN4-3 will end in December 2022 - long live GN5-1! In January, also Marina will be back as WP lead alongside Licia for T&I, and Maarten can dedicate some more time to other interesting activities again ... such as the proposed activities programme for EnCo in GN5:

- trust policies for token issuers and the process for listing (see above)
- OIDCfed, if that progresses to implementable standards
- Sntctfi and trust in proxies: there are continuing lingering concerns, mainly in InCommon, regarding 'Middle Things' and why to trust them. The judicious use for Sntctfi and G071 may alleviate those concerns, and we can make a push there for transparency - so that lost 'warm and fuzzy feeling of trust' can be restored?
- Trust marks for proxies
- Applnt integration (join the list at <https://lists.geant.org/sympa/info/appint!>)
- non-web use cases (and OIDC-Agent), specifically also for linking with the EuroHPC activities and access to HPC resources

It is important that these activities generate *actionable and operational output*, and thus are driven by tangible use cases. Linking closer to the AEGIS infrastructures and activities (and the follow-up of adoption there) will be helpful!

The output from FIM4R, adding in new communities and using this as input to EnCo also helps to draft more concrete and adopted guidelines.

ISGC Security Day and EnCo Contributions

The International Symposium on Grids and Clouds (ISGC) is an effective forum to reach out to a global T&I and security community, and foster links with the activities in the AP region (and HPCI). Traditionally, EnCo is represented both with a talk during the main conference, and during the Security Day special workshop. In 2023, it will be in-person again (most likely) in Taipei, March 19-24 (Security day on Sunday 19th).

For EnCo Maarten will submit the GN43 achievements talk, tentatively titled "A holistic view of enabling eScience collaboration through trust and identity networking and federation". Or something better. Also a Sirtfiv2 talk would be very welcome.

For the security day:

- linking operational security and federation is a pressing need - these communities are still too separate in terms of both culture and operational readiness
- OpSec for Federation is more than 'just' info sharing - it also includes joint exercises and use of realistic federation scenarios in the exercise scenarios taking into account the actual federation models and the diversity in terms of 'control' federation operators have over their participating entities. This is very country-dependent, and effective opsec should account for these differences in staffing and culture.
- the ISGC Security Day exercise in ~2018 of a federated incident, with Alessandra as the federation operator, was a good example!

The Security Workshop can thus help bridge the gap that is still present between federation structure and some forms of opsec structure.

This can be further discussed also at the GN43 WP5 T&I All Hands meeting at SURF in November 10-11, 2022.

John's Soapbox - on digests, revocation, and paranoia

Quite some discussion during the meeting revolved around the need for SHA-2 self-signed roots. But when you upgrade a root, what should happen to the old one (and the old serial number).

Revocation is immaterial (since you cannot revoke yourself reliably, as has been discussed some years ago already), so this comes down to the relying parties not accepting the old one any more - e.g. because it is no longer in the trust store.

But for old SHA-1 intermediates, would revocation there make sense? Like for the UK eScience 2B ICA? But then, if a digest like SHA-1 is broken and collisions can be forced, then revoking any specific intermediate is also not useful, since revocation is based on issuer+serial, and a new serial could be created and - through collision - appear to be validly signed by the issuer regardless of the revocation state of the original intermediate.

So for the practical case of the UK, "option 2", just re-issue the root with SHA-256, and later move to a new hierarchy if there is a need (e.g. for longer key lengths, or a new crypto algorithm) is good enough.

Meanwhile, it should be noted that almost all VOMS servers use the "--skipissuer" flag nowadays, since so many CAs have gone through a new-issuer-same-subject-DN cycle, such as TCS (3 times now). So that is quite common.

Pushing now for really new algorithms, like post-QC quantum-resistant encryption, is very premature. Not only lacks software support, but even algorithms that may be quantum-resistant may not be very secure in any conventional sense (see e.g.

<https://thehackernews.com/2022/08/single-core-cpu-cracked-post-quantum.html>).

Practical actions for the UK include:

- re-issue the Root with SHA-256 (longer makes no sense given the key length)
- all 2A-series certificates can be removed from the namespaces/signing_policy files

- have an earnest chat with JISC regarding their profit model for TCS, emphasizing the need for large quantities of, specifically, eScience server certs (for storage clusters) and clients.
- Promote the --skipissuer flag for VOMS servers local to the UK

Operational matters and self-assessment process

The new 'assisted assessment and review' model has been implemented for the self-assessment process for MREN and the PolishGrid CA. Under this scheme, both the CA manager and peer-reviewers are collectively and synchronously reviewing the self-assessment results (prepared by the CA beforehand as far as possible), and give immediate feed-back on the results. The results of an 'assisted assessment and review' process should either be discussed in the next plenary meeting, or the summary shared on the mailing list following the customary two-week review period, before the self-audit is actually closed.

- The MREN review was completed within one meeting with both Lidija and the reviewers, Dave Kelsey and Jan Chvojka. The two-hour videocall was a good experience for all. Only minor issues were identified, and these are to be implemented shortly on the production system. For the MREN CA, this means that when the minor issues in the CP/CPS are addressed, the result is to be endorsed after a two-week period 'silence implies consent' call on the mailing list.
- For the PolishGrid CA, two separate meetings were held due to scheduling possibilities. This also worked, and the result is good after confirmation by Pawel and Feyza.

The next series of self-assessments due is from KENET, TSU-GRENA, BYGCA, GridPK, and KIFU/NIIF. Cosmin will trigger all these in parallel. Meanwhile, since Adeel is on the call, we quickly reviewed the GridPK status. The changes planned by Adeel including resolving the lack of KEYGEN support in browsers (for which Jan @CESNET has a Javascript-based solution available), and preparation for the self-assessment with a target completion date of January 2023. Incidentally, the SHA-1 root was already replaced by a SHA-256 incarnation when the validity period was extended recently. So no change is needed there.

Under the 'old' process, the LIP CA is now fine (endorsed during this meeting), and KENET still needs to respond to the reviewers - it could move to the new model instead to accelerate the process.

For performing the self-assessments, the grading scheme from GFD.169 remains appropriate, but they should be done

- against the proper IGTF Assurance Level (<https://www.igtf.net/ap/>, TAGPMA has sheets available), **and**
- checked against the PKI Technology Guidelines (<https://www.igtf.net/guidelines/pkitech/>).

A discussion on secure destruction of old media (hard drives) ensued. While keeping the hard drives in the safe for future destruction is certainly fine, at one point they need to be got rid of. This should typically include physical destruction (see e.g. the NSA/CSS Storage Device Sanitization Manual Policy Manual 9 – 12 at <https://www.nsa.gov/portals/75/documents/resources/everyone/media-destruction/storage-device-declassification-manual.pdf>).

For the DigiCertGrid intermediates, some are still signed with SHA-1. Since a new set of ICAs is available, these should likely be withdrawn. But Tomofomi should confirm first, of course.

TAGPMA updates

- The XSEDE programme has ended, long live the ACCESS programme!
- Concerns were raised in TAGPMA by Tomofumi on the required re-mapping to ASCII for organisation names. This is indeed incompatible with the EV guidelines, but since there is no need for joint EV+IGTF trust, the OV BR guidelines apply. There, there is room for this, "provided that the CA documents the difference and any abbreviations used are locally accepted abbreviations". (BR v1.8.4, 7.1.4.2.2(b)). So there should not be an issue.
- The WoTBAn&Az workshop is held in a few weeks in conjunction with the NFS Cybersecurity Summit in Bloomington, IN, USA

- The December I2 TechEx series of meetings has a FIM4R+TAGPMA meeting on Sunday, and REDEFS on Monday
- TAGPMA will follow up the transition to SHA-256 for self-signed roots

Attendance

We thank Hannah Short, Eisaku Sakane, Jule Ziegler, David Kelsey, Maarten Kremers, Ian Neilson, Alistair Dewhurst, and David Groep for their in-person attendance at CERN.

And, in random order, also John Kewley, Adeel-ur-Rehman, Mirvat Al-Joghami, Miroslav Dobrucky, Cosmin Nistor, Eric Yen, Jan Chvojka, Lidija Milosavljevic, Christos Kanellopoulos, David Crooks, Licia Florio, Mario Lassnig, Ian Collier, Tom Dack, Maarten Litmaath, and Nuno Dias all managed to survive (part of) the three-day-long videoconference call, for which they are to be highly commended!