

**That TF-EMC2 Task Force
Deliverable
Version 0.4**

**Building trust in grids by means of PKI
Issues for interoperability with other infrastructures**

Author: Licia Florio (Terena)

Contributions: Diego Lopez (RedIRIS), David Groep (NikeF), Christoph Graf (SWITCH), Reimer Karlsen (DFN-PCA)

Abstract

This document tries to summarise some of the discussions about the differences/analogies of Grid-PKIs and NREN-PKIs. The discussion took place both on the TF-EMC2 (the TERENA Middleware task force) and on the EUGridPMA (European Grid Policy Management Authority, the body that coordinates a Public Key Infrastructures for use with Grid authentication middleware) mailing lists.

Being the topic of general interest it has been common opinion that a summary of the issues tackled was helpful.

This document doesn't aim at being an exhaustive document about how PKI is implemented in Grid applications, but it aims at providing an answer to some of the issues that are more frequently asked.

1. Introduction

All Grid applications make heavy use of X.509 certificates to allow users to access and share resources normally owned by different institutions and located in different parts of the world. On the other side NRENs have set up and have tried to deploy their own PKIs within the academic community, mainly universities. The current situation is that Grid PKIs implementations are usually being run as parallel PKIs to the NREN ones.

This document tries to explain the reasons that led to this situation.

One of the obvious observations is why there is still need to run several concurrent PKIs if an organization or NREN has a working PKI whose authentication and authorization procedures are compliant with the Grid policies.

To start answering the question some background might be helpful. At the time the first Grid infrastructures were built, the NRENs were not able to provide PKI infrastructures which were able to satisfy the Grid requirements. This forced Grid

researchers to set up their own certification authorities (CA), which when Globus came along were demanded to be Globus-compatible.

Also a description of what the Grid requirements and what the policies are should be provided before going into more technical details.

Grid-focussed PKI is driven directly by requirements from end-users and relying parties, and their requirements for usability, stability and security. One of the key issues, for example, is end-users in research organisation, faculties, etc. needing effective access to the PKI even if their home organisation at large is not deploying or supporting any PKI.

Grid policies in this context mean the set of minimum requirements that the EUGridPMA has defined and which are based on operational experience. The full document is available at: <http://www.eugridpma.org/>

This list which has been based on operational knowledge is constantly updated during the EUGridPMA meetings to match the demand of both Grid middleware changes as well as identity providers.

2. EUGridPMA. Minimum requirements relevant to trust links

With the proliferation of CAs, the EUGridPMA felt that some coordination was necessary and that having multiple CAs for the same purpose was just a waste of resources. Furthermore Grid processes work on the bases of mutual trust, which is more difficult to establish if the number of parties involved grows without control. What said above can explain why one the minimum requirements expresses that all Grid applications should be dealt by only one CA per country, which has to go through an accreditation process (which mainly consists to verify that the CA policy satisfies the minimum requirements) within the EUGridPMA.

A wide network of Registration Authorities (RA) for each CA is preferred. The RAs handle the tasks of validating the identity of the end entities and authenticating their requests, which will then be forwarded to the CA. The CA will handle the actual tasks of issuing CRLs, signing Certificates/CRLS and revoking Certificates when necessary.

Only in a few cases the NREN-PKIs were able or interested (most of the NRENS got involved in Grid only in a later stage) to match the minimum requirements, this is the case of CESNET and GRNET for instance. In other cases a new CA was set up in most of the case by institutions dealing with research in Physics, involved in one of the initial Grid projects.

The minimum requirements say also that only a well defined namespace may be served and that this name space must not clash with that of any other CA. The underlying cause of this requirement is due to the X.509 specification which states that two different people cannot have the same DN and Grid authorisation

relied upon this requirement, which in practical terms means that the authorization decisions in the Grid are based on the subject DN only.

The mistake (if any) was to limit the CAs to only issuing names in a particular namespace. It is understandable why this was introduced as it simplifies CA operation (name authentication in particular), but in principle there is nothing in X.509 to stop a Dutch CA issuing a certificate to a non-Dutch entity and putting the non-Dutch global name in it. And conversely, there is nothing to stop a CA issuing a certificate to someone who lives in a country different from the one the CA operates in. But this makes the work of the CA much more difficult, since it has to validate that the person is entitled to claim that name, and usually it is easier to just check locally issued names rather than ones issued in different countries.

The net result is that namespace coordination should now take place in a CA coordination group like the EUGridPMA and IGF (International Grid Federation).

The requirement for a well-defined namespace per CA prevents at the global level that a particular subject name is - by accident - linked to two different entities. Of course, the same entity is free to collect any number of subject DNs from whatever source he/she wishes. This does not pose a problem for the infrastructure of the authorization policies.

A CA is free to issue certificates in any other namespace, but *for use in the grid/* (and for those uses only) those specific namespaces must be enumerated. Client-side grid software is capable of enforcing these namespace restrictions on a per-CA bases (using a "signing-policy" file), and will implicitly disallow certificates with subjects issued outside of that defined namespace. This check prevents the subjectDN collisions referred to earlier.

Therefore there is nothing in the Grid certificates that makes them in any way "grid specific". The (minor) differences are in the lack of support for hierarchies, namespace, the requirements to issue timely CRLs, and specific interpretation of their nextUpdate field (the last being an OpenSSL implementation decision as well). Such a type of certificate cannot contain additional attributes for authorisation and this matches the Grid approach to keep the authorisation and authentication as two separate things.

Making a good existing PKI service, like the one ran by one of the NRENs "grid-compliant" is a quite viable and probably also the cheapest solution.

3. Certificate path validation: the SSL and the Grid ways

The trust path verification of intermediate certificates as done by the OpenSSL suite (used both by the Grid middleware and most of the PKIs and applications deployed by NRENs) was discussed within the EMC2 group.

In many protocols (including SSL/TLS and S/MIME) a set of intermediate certificates is supplied by the peer along with the end entity certificate. If this chain contains sufficient intermediate certificates then only the root needs to be included in the trusted store, otherwise any certificates missing from the path to the root need to be included. In SSL/TLS the standards require the including of the whole chain (with the root CA being optional) so any peer that doesn't do this is broken or mis-configured.

During the SSL handshake the whole chain should be presented, if not, the SSL implementation is faulty, therefore only the root is needed in the local store to verify the chain. In the case that the subordinate CA certificates are taken from the handshake, there would be still the problem of getting the CRL and/or the OCSP responder location.

OpenSSL accepts everything the peer provides, not including a final self-signed certificate, and then it searches the local certificate store to complete a chain up to a self-signed certificate. If the peer provided the self-signed cert, then in addition it must match what was found in the store.

Putting an intermediate CA certificate in the store is not sufficient to let the handshake succeed, but having a trusted root CA certificate in the store and a complete chain (with all certificates but root cert) provided by the peer during the handshake will work. In this case the difficult part will be to negotiate/set the trust for root certificates and maybe some specific intermediate CA certificates as well as the on-the-fly retrieval of revocation information of certificates used in the constructed certificate path

The Grid Security Infrastructure (GSI) does implement to some level its own certificate chain validation. This is needed to support the use of "proxy" certificates (the short-lived certificates based on a temporary, user-generated, keypair that in turn is signed by the entities long-term credential). Since these certificates normally violate the "CA:False" basic constraint in the certificate extension, default OpenSSL validation fails.

Code inspection shows that the modifications required to adequately process proxy certificates only imply changes to the certificate validation code and not to the certificate chain retrieval process. The validation callback does check whether a CRL is available (there must be a CRL for every CA cert) and this is checked only against the local cert store.

Also a "signing_policy" file per CA (with namespace constraints to counter the lack of global namespace coordination) must be there and must be in the local cert store, thus limiting the trust domain of an intermediate or root certificate in the local store based on the subject namespace. It is expected though that some

future version of OpenSSL will implement the validation of proxy certs, now that these have been standardized in RFC3820.

This support should also incorporate the use of OCSP to evaluate revocations and the extension of the trust path calculation, so intermediate (not self-signed) certificates could also be incorporated as roots of trust.

4. Future work

It seems clear that there is enough interest on the NRENs side to understand how they could support Grids, for instance providing certificates.

Also a channel of communication between the two groups (NRENs and Grids) has been activated and it is expected to improve in the future.

NRENs are invited to accredit their CAs to the EUGridPMA and to attend the CAOPs (Certificate Authority Policy) Working Group which meets during GGF. This would be a way for NRENs to better understand Grid requirements and to make sure that some decisions taken by the Grid community do not diverge from what the NRENs believe is a the good way to proceed.

5. References

<http://www.openca.org/ocsp/>

<http://ejbca.sf.net>

<http://www.openvalidation.org/>

<http://www.openssl.org>

<http://www.terena.nl/tech/task-forces/tf-emc2>

<http://www.eugridpma.org>