

CA-based Trust Model for Grid Authentication and Identity Delegation

Status of this Memo

This memo provides information to the grid community regarding trust issues involving reliance on Certificate Authorities and X509 Identity Certificates to authenticate people and services. Distribution is unlimited.

Copyright © Global Grid Forum (2002). All Rights Reserved.

Abstract

This paper examines the basis for the trust a relying party places in an identity certificate signed by a CA. It is intended to be an informational document and is the result of extensive discussion on the security WG mailing list about current and best PKI authentication practice. It does not cover authorization issues or trust relationships between the CA and its RAs, or between the end entities and the resource providers.

- Abstract 1
- 1 Background 1
- 2 Grid Environment 2
 - 2.1 Grid requirements for authentication 2
 - 2.2 Using PKI for authentication 2
 - 2.3 What does "trusting" an identity certificate mean 3
- 3 Role of a Grid CA 3
 - 3.1 Elements of Distinguished Names 3
 - 3.2 Identity Vetting 4
 - 3.3 Significance of the Subject Name 4
 - 3.4 Virtual Organization Authority vs. Grid CA 4
 - 3.5 Cross-signing and Subordinate CAs 4
 - 3.6 Kerberos CA (KCA) 5
- 4 Delegation of Identity 5
- 5 Responsibility of the subscribers 5
- 6 Technology for Protection of Private Keys 6
 - 6.1 FIPS PUB 140-2 - May 2001 6
 - 6.2 Smart Cards, PKCS 11, PKCS 15 7
 - 6.3 On-line Credential server 7
 - 6.4 Off-line CA 8
- 7 Conclusions 8
- Author Information 8
- Intellectual Property Statement 8
- Full Copyright Notice 8
- 8 References 9

1 Background

This paper was suggested at GGF4 after it was announced that the Grid Forum was not going to take any responsibility (and thereby incur liability) to either vet CPs or audit CAs. The probability of unaudited CAs raise the

issue as to what trust can or should be put in certificates signed by unvetted and unaudited CAs, which are hopefully run by competent people with good intentions. This paper is an attempt to clarify the issues faced by a relying party when deciding to trust a grid CA and the certificates that it issues.

The scope of this paper is limited to examining the basis for the trust a relying party places in an identity certificate signed by a CA. It also presents some of the working group's discussion of the authentication needs of a grid and how use of PKI identity certificates can fulfill them. It does not cover issues such as using these certificates to authorize use of resources, trust relationships between the CA and its RAs, or between the end entities and the resource providers.

2 Grid Environment

"A computational grid is a hardware and software infrastructure that provides dependable, consistent, pervasive and inexpensive access to high-end computational capabilities." [2] Typically the grid resources are provided by various organizations and are used by people from diverse sets of organizations. A grid may support (or define) a single virtual organization or it may be used by more than one virtual organization. Individual pieces of hardware may be used in more than one grid and people may be members of more than one virtual organization. The different resources in a grid may have different access policies, including how they authenticate and authorize users. However, if there are no common or overlapping authorizations among the resources, they do not form a usable grid. Thus in this paper we will assume that for any set of resources there are some common authentication/authorization procedures.

2.1 Grid requirements for authentication

Users, hosts and services need to be able to authenticate themselves in the grid environment. Experience in using grids for remote computations has demonstrated the need for unattended user authentication in addition to interactive authentication. Unattended authentication of users is needed when a user is making frequent requests to remote servers and does not want to repeatedly type in a passphrase, and when a long running job may need to authenticate itself after the user has left. Servers specific to a single host may need to be started at system boot time and run with their own or the host's identity. Some services may need to be started periodically on many different hosts and be able to authenticate themselves with a known identity.

Basically, authentication between two entities on remote grid nodes means that each party establishes a level of trust in the identity of the other party. In practical use an authentication protocol sets up a secure communication channel between the authenticated parties, so that subsequent messages can be sent without repeated authentication steps, although it is possible to authenticate every message. The identity of an entity is typically some token or name that uniquely identifies the entity.

2.2 Using PKI for authentication

One commonly used identity token is provided by an X.509 public key certificate. An X.509 "public key certificate" is defined in the I-TUT standard [16] as "*the public key of a user, together with some other information, rendered unforgeable by encipherment with the private key of the certification authority which issued it.*" [Section 3.3.44]

In practice an X.509 certificate as defined by the PKIX X.509 Certificate Profile RFC [6] contains a public key, a subject name in the form of a multi-component *Distinguished Name* [DN] and a validity period and is signed by a trusted third party called a *Certification Authority* (CA). While the standards documents refer to these certificates as "public key certificates" or "X.509 certificates" this document also uses the term "identity certificates" to emphasize their use to securely identify an entity in a grid environment. The X.509 certificates are used with the TLS (SSL) security protocol to make a secure authenticated connection between two parties. X.509 certificates are exchanged between entities, usually in a phase where no security has yet been negotiated (no data integrity or confidentiality). The certificates are first tested for validity by checking the expiration dates, possible revocation, acceptable key usage and if they are signed by trusted CA.. If the certificates pass all these checks their public keys are then used to build a challenge handshake to prove that each entity that sent a certificate has the corresponding private key. Passing these tests gives each party a level of confidence that it has established a secure connection to the party represented by its prospective certificate. The identity token could be the public key which is unique by virtue of being a member of a large sparse space or the certificate subject name, which should be unique within the certificates issued by a single CA. Typically the subject name is used rather than the public key because it is human readable and can be expected to name the individual or service provided in some reasonable manner.

There are other public/private key based schemes such as PGP keys [5], SSH keys [1] and the SPKI [4] keys and protocols. SPKI focuses on authorization certificates more than identity certificates. SSH is mostly just a

private/public key mapping with no real attempt to provide global names. The PKIX/X.509 scheme has a smallish set of trusted third parties (CAs) to sign identity certificates that contain a subscriber's public key. This improves the scaling properties of public key distribution since only the CA's public key needs to be distributed in an out-of-band secure manner. In systems without a trusted third party, such as PGP, each key holder must find some secure way of establishing the association of his identity with his public key, to each party with which he wishes to establish authenticated communication. In the X.509 infrastructure, the individual subscribers' public keys can be transmitted in public key certificates as part of a TLS connection handshake and can be accepted as valid if these certificates are signed by a trusted CA.

Another feature of the X.509-TLS infrastructure is that it supports multiple independent CAs. In a grid each site may choose which CAs it will accept for binding DNs and public keys. At this point most of the current grid tools are built on GSI or https, both of which use X.509 certificates for securely establishing a grid identity.

2.3 What does "trusting" an identity certificate mean

Basically the only assurance that is provided by using TLS with mutual authentication is:

Every time this public key /DN is presented through the secure protocol the same private key is known on the other end of the connection.

If knowledge of the private key is assumed to be restricted to the entity the CA originally associated with the DN, then the party on the other end of the connection is the entity named by the DN.

These assurances depend on the correctness of the TLS and certificate validation implementation at all the sites that take part in establishing a secure connection, the diligence of the individual in protecting the private key and the Certificate Policy (CP) and Practices (CPS) of the trusted CAs. In order for a relying party (e.g. a grid resource provider) to trust identity certificates signed by a given CA, it must trust that:

- Only the CA has access to the CA's private key, so that no other entity can sign certificates with that key (see section 6)
- The CA always associates each DN and public key with the same individual or entity. i.e., the DN and the public key are permanently assigned to this entity (even after revocation or expiration, they cannot be reused by another entity)
- The CA did acceptable job of verifying that the holder of the certificate is the individual named by the DN (discussed further in section 3)
- Only the named entity associated with the DN has control of the certificate's private key. What this means for non-human entities must be carefully defined by policy. (Discussed in section 7)

How these objectives are going to be achieved should be defined in the Certification Policy [CP] and Certification Practice Statement [CPS] documents of each CA. If the CA has been independently audited then a relying party can trust that the auditor agrees that the CA is correctly implementing its CP. If the CA is unaudited, the relying party must trust that the operators of the CA are, in fact, following their avowed policy. The CP should spell out the obligations of a certificate holder, to keep the private key secure, but both the CA and the relying party must trust the user to follow the procedures. If each of these assertions is true then the relying party can trust that the person or other entity that has connected to it via a secure protocol is the unique entity that is named in the X.509 certificate. Thus the certificate can be used to make a secure and authenticated connection.

3 Role of a Grid CA

A grid CA is defined as a CA that is independent of any single organization and whose purpose is to sign certificates for individuals who may be allowed access to the grid resources and hosts. A grid CA is substantially different from a traditional organizational CA which only signs certificates for members of its organization and is closely linked with the authority that defines who those members are. Those certificates are then used to access resources within the organization. There are two implications of this difference: one in the format of the DNs and the other in the methods of vetting user identification.

3.1 Elements of Distinguished Names

In identity certificates issued by an organizational CA, the DN often contains a number of attributes taken from the organization's X.500 or LDAP directory, e.g. OrganizationalUnit, Location and Email. There is often an underlying assumption that the X.509 certificate is stored in the directory entry for the user. An organizational CA is in the

position to find existing LDAP entries, verify the correctness of the name elements, issue certificates for such a user, and store the certificate back in the LDAP entry. As a result of this paradigm a DN could have several vetted components. A grid CA breaks this paradigm by being independent of its subscribers. Even in the organizational environment there was a problem with putting too much information in a DN, since whenever any part of the information changed, e.g. an employee changed departments or got a new email address, the certificate would need to be reissued.

Since a grid CA is independent of the organizations to which its subscribers belong, it does not have a way to verify much information about a subscriber or to know when such information changes. The prudent approach for a grid CA to take is to put as little information in the certificate as possible. A minimal set that has been chosen by the several grids [9,10,11,12] is an Organization element that identifies the grid to which the CA belongs, a class of certificates such as *person*, *host*, *service* which is intended to be used when storing and retrieving certificates in the grid CA's publishing directory and a Common Name that reasonably identifies the entity for which the certificate is issued. An E-mail address can be added as an alternative name for the sake of convenience, but not for identification.

3.2 Identity Vetting

Since the operator of a grid CA does not personally know the persons who are requesting certificates or have access to a trusted directory of such users, it must rely on Registration Agents (RA). These are individuals who are likely to know first or second hand a subset of subscribers. If the users of a grid can be grouped by actual or virtual organizations an RA may be chosen for each such organization and given the responsibility to approve requests from members of that organization only. The rules for establishing member identities should be published by each RA, and the procedures for verifying the identities and certificate requests should be consistent between all the RAs and approved by the CA.

3.3 Significance of the Subject Name

There has been much discussion at the working group meetings and on the mailing list as to what "meaning" should be ascribed to the Subject Name. On the one hand, most CA's specify that the Common name component of the subject name should be an official and recognized name for the person who requested the certificate and the identity vetting process should assure this. On the other hand, the name should be treated by a relying party as just an identity token that can be used to assure that the entity making the current connection is the same entity that has used this token before. In either case before the name is added to any lists that authorize access to resources, the name must be checked by the authorizing party against some other database or Virtual Organization authority to see what rights should be allowed for the holder of this identity. Using just the subject name for authorization is not safe, since subject names are only guaranteed to be unique within the domain of a single CA. So either both the subject name and the issuer (CA) name must be used, or some other means must be used to limit what namespaces may be signed for by which CA's.

3.4 Virtual Organization Authority vs. Grid CA

Another subject that was discussed was the role of a Virtual Organization (VO) as a trusted third party. If an entity is authorized to use resources because it is a member of a VO, the relying party needs to verify that a token belongs to a member of the VO. In this case, it might be more efficient to have the VO issue the certificates in the first place. Then any entity that holds a certificate from the VO could be assumed to be a member of that VO. This approach has two drawbacks. First, not every user of grid resources is a member of an accepted VO, so having VO CAs does not eliminate the need for a broader grid CA. Second some persons are members of more than one VO and they would end up with a certificate from each VO. This could be seen as a good thing, since it would allow a user to act in different capacities (or roles) by using different identity certificates. However, managing different certificates is also a burden on the user, especially in view of the primitive tools available for certificate handling. A more serious objection was the merging of authorization into the concept of identity and authentication. The general consensus was that we should remain with the principle of using X.509 certificates purely for authentication of identity and handle authorization as a separate problem.

3.5 Cross-signing and Subordinate CAs

A relying party may want to evaluate independently each CA that it is going to trust before installing its certificate. However, an alternative trust model exists where a CA will sign another CA certificate once it has verified that the other CA enforces policies that are as secure as its own. A subordinate CA is similar except that it only has a CA certificate signed by the parent CA, where a cross-signed CA has either a self-signed CA certificate or a CA certificate signed by a parent CA. The advantage of this model is that a relying party can accept subscribers from

more CAs without having to individually evaluate and install all the CA certificates. However, it is important for a relying party to know if a trusted CA signs for other CA's certificates. Otherwise the relying party may be unexpectedly widening its trust domain. Note that not all path validation software recognizes CA chains, but may insist on having the certificates for every trusted CA.

3.6 Kerberos CA (KCA)

Another solution to facilitate user access to the grid is to run a CA that will create a short-term X.509 credential based on a Kerberos identity. Thus Kerberos users can transparently use grid resources without having to apply for a long-term grid identity. A relying party should be aware if a trusted grid CA either cross-signs or is a parent of a KCA, in order to decide if this model of user identification is acceptable. The relying party is heavily dependent on the quality of the Kerberos KDCs' security. The automatic issuance of certificates requires that this CA be on-line at all times.

4 Delegation of Identity

As was previously mentioned in section 2.1, grid experience has shown the need for unattended authorization. It has also demonstrated the need for both local and remote processes to run on behalf of a user and with his authorization rights.

These issues have been addressed by the Globus Grid Security Infrastructure (GSI) by allowing for short-term proxy certificates, stored with unencrypted private keys, to which a user has delegated his identity. These certificates are correctly formatted X.509 certificates, except that they are marked as proxy certificates and are signed by an end entity rather than a CA. The choice of the life-time of proxy certificates requires a compromise between allowing long term jobs to continue to run as an authenticated entity and the need to limit the damage that can be done by a stolen proxy. Proxy certificates with restricted rights are another way of limiting the damage done by a stolen proxy.

Authorization software run by relying parties must be able to recognize proxy certificates and back up the certificate chain until the end-entity certificate is reached in order to do the authorization based on that identity token. Such software may also want to enforce policy decisions based on the lifetime of the proxy, or how many levels of delegation have been done. While restriction of proxy rights may make a site more secure, it will likely break some grid software attempting to run at that site.

The delegation of credentials may take place on the machine on which the original credential resides, or may take place between two machines in different administrative domains. In the latter case, the delegation expands the trust relationships to include an additional domain (and the delegation software that runs there). The relying party should be aware this, but in the absence of secure DNS, it is difficult to include trusted domain name information in a certificate chain. The Grid Forum document on proxy certificates [8] discusses this matter in greater depth. Although several schemes for including trace delegation information in the proxy delegation chain were discussed, no standard was agreed upon. At the current time, only the number of times a proxy has been delegated can be deduced from the chain of delegated proxies.

5 Responsibility of the subscribers

In the trust model listed above, it was the obligation of the subscriber to keep complete control over the private key. If there is any evidence that the private key has been compromised, the certificate should be revoked. In the case of certificates issued to people, the holder of the certificate who believes it has been compromised has the responsibility to ensure that the certificate is revoked. The subscriber should be aware that any actions taken with a stolen key will be attributed to her, thus it is in her interest to quickly revoke a compromised certificate. In addition, if an RA is presented with sufficient evidence that a key has been compromised, he should revoke the certificate. However, for certificates issued to hosts or services it is not so obvious who is responsible for the key. A CA that issues such certificates should make a clear statement of what the host or service certificate is supposed to represent and how the private key is to be protected. For example, a CP may contain the following requirements for subscribers:

- For a personal certificate:
 - ❑ Only the named user has control of the certificate's private key
 - ❑ When the key is stored on a computer, it is always encrypted with a passphrase at least as strong as commonly accepted guidelines for account passwords. Preferably, the passphrase is significantly longer than an eight character password and is not in any dictionary or common phrase book. A strong passphrase

is highly desirable because of the inconvenience of revoking a compromised key and certifying a new one.

- For a host or service certificate that is only used on a single host
 - ❑ The key may be stored unencrypted only on local storage of the host named in the certificate
 - ❑ It should only be used to represent services run on that host
 - ❑ The system administrator of the host is responsible for revoking the certificate if there is any evidence that the host has been compromised.
- For a service certificate that is used on multiple hosts.
 - ❑ The key may be kept unencrypted on more than one host.
 - ❑ The requester of this certificate is responsible for the security of the key on each host.
 - ❑ Relying parties must be able to find out who the responsible party is.

6 Technology for Protection of Private Keys

Trust in the user certificates depends on how well the CA's, the RAs', and the users' private keys are protected. There exist several technologies and standards relating to key management as well as hardware and software implementations of these standards. Since these standards or products could be mentioned in CA Policy or Practices statements, a brief overview is included here. Since CA and RA keys are more limited in number and more critical than user keys, a more secure and expensive means of protection is likely to be required of them. Convenience of use may be a bigger factor in the options recommended for storing subscriber keys. A relying party should be aware of what technologies are required or allowed for key storage.

6.1 FIPS PUB 140-2 - May 2001.

This is a standard from NIST specifying "Security Requirements for Cryptographic Modules" [15]. It specifies four different security levels and defines standards for key management (including key generation and storage), physical security of the module, cryptographic algorithms and security functions that the module should perform. NIST sponsors a testing and certification system; see the NIST Validation List web page for a list of validated products [17]. Customary usage is to refer to rated modules by the shorthand formula "rated at FIPS-140 level 3".

Several manufacturers produce Hardware Security Modules, hardware-based cryptographic modules, several of which have been issued FIPS-140 validation certificates. They represent only a few of the products in the NIST validation list. Some vendors include Chrysalis (Luna), nCipher (nFast), Rainbow, Baltimore, and IBM. The HSM controls access to the CA signing private key (and other private keys) through a well-known API such as PKCS #11 [13]. The HSM is usually responsible for creating the key pair; depending on FIPS rating and configuration, the HSM may forbid, or severely restrict, export of private keys from the HSM device by limiting or disabling key export functions. The HSM's cryptographic module will use the CA's signing private key to perform operations, such as signing certificate requests or other objects. The HSM will perform these operations through commands issued through the module API (PKCS #11). The CA will never have the private key in plaintext on the computer host. These devices minimize or eliminate the danger to the CA or other server's private keys in the event of a service compromise, operating system compromise, or breach of CA's physical security. The best of these devices have been evaluated under certain conditions and system configurations by NIST as meeting security level 2 and security level 3 specifications. The FIPS-140 certification process is also applied to other devices and software, and may demonstrate a product's level of maturity and quality of implementation of cryptographic algorithms and random number generators.

Security level 1 certification shows that the module executes cryptographic functions in an objectively reasonable way. This is the basic certification for a hardware cryptographic module.

Security level 2 certification requires that the module provide tamper-evidence, meaning the physical device will be sealed with a coating or otherwise provide evidence of physical access, should an intruder attempt to read the contents of the module. The device must also implement an operator role, requiring a basic authorization process before executing certain functions.

Security level 3 certification requires the module to be tamper-resistant; according to the language of the FIPS-140 standards document, it should have a high probability of detecting and responding to attempts to physically access the module contents. Identity based authentication is required before functions are executed, and there are classes of functions assigned to different roles. Level 3 modules do not allow the unencrypted export of keys to the host operating system.

Security level 4 certification requires a high degree of assurance that physical access to the module cannot result in a breach of security. Modules should demonstrate a high likelihood of detecting module penetration, and high

resistance to environmental compromise (such as operation outside normal temperature and voltage levels). Security level 4 devices might be very useful in a physically unprotected or physically demanding location. There are relatively few level 4 certifications.

There are numerous other conditions and capabilities associated with the different security levels. Commercial HSM's typically are used in conjunction with a smart card system (see next section), which implements the roles and authentication requirements of the standard. In addition, these smart card services further restrict the capabilities of the service (eliminating or encrypting plain text key up- and down-loads), and enhancing the manageability of the CA or other service for the CA administrators and other personnel.

6.2 Smart Cards, PKCS 11, PKCS 15

Smart cards (or Integrated Circuit Cards) are another technology that is used to store credentials off-line. The private key is kept only on the smart card that is inserted into a reader device when it is needed to sign a document or make a secure connection. A number of manufacturers make the cards and readers. Currently there is a widely used standard, RSA PKCS11 "Cryptographic Token Interface Standard" [13] for signing and cryptographic functions. In the current state of the art, each smart card vendor must provide his own PKCS 11 driver. There is a new standard, PKCS 15, "Cryptographic Token Information Format Standard" [14], which is an attempt to standardize the format in which cryptographic information is stored on smart cards and thus allow a standard reader to handle different manufacturers' cards. A smart card with a PKCS 11 interface is a possible solution for protecting CA keys, RA keys or end user keys. The key would only be vulnerable when the smart card was connected to an on-line computer and unlocked by the PIN. This scheme provides credential portability if every machine on which a user wants to use his credential is equipped with a compatible smart card reader, but makes it impossible to use the credential on any other machine. Therefore, it is unlikely to be required of general grid users, but could be feasible for CA or RA keys.

Smart card implementations are usually limited in on-board CPU power, memory, and bandwidth between host computer and smart card device. The most serious of these today is the limited amount of on board memory. 16K – 32 K of RAM is adequate to store a key pair or two, provide the certificates are reasonably small, but the need to support multiple vendor standards on a single card, store procedures, and possibly use the cards for security services based on other algorithms makes the system designer's job difficult and the security issues complex.

The chips have limited capability and thus the user's cryptographic work must be offloaded to the host computer; these devices are no substitute for cryptographic accelerators and HSM's. It is also difficult to ensure the integrity of the user's private key, either electronically or physically, although most cards have some tamper-evidence features.

Incorporating biometric tests, at some cost in internal memory and system complexity, can enhance the security provided by PIN numbers.

Smart cards may improve the security of certain roles in the PKI, such as CA administrator, RA, or system administrator. When relying on a smart card for additional security, an administrator should be aware that a secure key storage device, even if it does its own PK operations, may still have unintended data presented to it by compromised host software, and may be authorizing transactions the user does not intend. The best that can be hoped for with a compromised host, is that keys stored on the smart card will not also be compromised.

6.3 On-line Credential server

Another suggested technology for storing credentials is an on-line server. This is a secure server on which a user can store her credential protected by a password. When the credential is needed the user connects to the server via a server-side protected TLS connection and types a password that allows the credential server to use the stored credential on her behalf. The rationales for such a server are: it could provide a more secure environment than the user's home machine; it would allow people without home machines such as students to use credentials; and it would allow people away from their home machines to use their credentials. The objections to such a server include the fact that it becomes a single point of attack and potential failure and the security of the user's key is only as good as the security of the server and the password that is used to protect the key. However, the server can vet the user password and refuse one that is too easy to break. In contrast, when the user stores her key on her own workstation there is no central enforcement of the existence or quality of the passphrase.

The trust issues raised by an on-line credential server are similar to those of the KCA mentioned in section 3.6. In each case, there is a trusted, on-line server which issues X.509 certificates on behalf of a user when authorized by that user presenting a password. If a relying party chooses to trust these certificates, it is trusting the credential server software, the machine that is it running on and the strength of the passwords used. However, it only has to trust that the end user can manage to keep a password secure rather than to securely manage an identity credential.

The CP section on subscriber obligations should specify how the user's private key is to be stored. If storage on on-line credential servers is allowed, or specifically excluded, that information should be stated here.

6.4 Off-line CA

A certificate authority can be configured such that the signing engine is available only to the CA administrator, on a host that is never available on any network or by any other means except personal use by the administrator. The host should be kept in a locked, secured facility, and the administrators' accesses and uses carefully logged and controlled. Signing requests are conveyed to the offline CA through removable media, and signed certificates and revocation lists published by writing on removable media and transferring this data back to the "public" part of the CA. While wholly dependent on the administrators' behavior and subject to a variety of theoretical attacks, in practice this is a very good security solution for private key protection for a small PKI.

7 Conclusions

Because of the limited trust relationship between the subscribers, relying parties and a grid CA, it is strongly recommended that the X.509 certificate by itself only be used to authenticate a user or a service. In the TLS handshake model, this means that an entity presenting a valid certificate from a trusted CA will be able to connect to the relying party's server. Any authorization to use specific resources should be done explicitly by the server supplying the resources or by a virtual organization (VO) that verifies the identity of the person holding the private key associated with the certificate. The certificate or the contained DN can be used to identify a user who should be granted access as in the case when GSI looks for a DN in a site-specific grid-map-file. However, no authorization should be granted a user simply because he presents a valid X.509 certificate signed by a grid CA. The same caveat applies to users of services. They should check the DN contained in a service certificate to see that it agrees with an expected name in addition to checking that a trusted CA issued the certificate.

Before accepting a CA as a trusted third party, a site should carefully review the Certification Policy and Practices document(s). Items that should be carefully evaluated include: how well the CA and RA keys are protected, how thoroughly the subscriber identities are checked by the RAs; what the obligations of the subscribers are to keep their private keys secure; and whether or not the CA signs subordinate CA certificate or cross-signs for other independent CAs.

Author Information

Mary Thompson, LBNL - mrthompson@lbl.gov
Doug Olsen, LBNL - dloslen@lbl.gov
Robert Cowles, SLAC - rdc@slac.stanford.edu
Shawn Mullen, IBM - shawnm@austin.ibm.com
Mike Helm, ESnet - helm@fionn.es.net

Intellectual Property Statement

The GGF takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the GGF Secretariat.

The GGF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this recommendation. Please address the information to the GGF Executive Director.

Full Copyright Notice

Copyright (C) Global Grid Forum (date). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included

on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the GGF or other organizations, except as needed for the purpose of developing Grid Recommendations in which case the procedures for copyrights defined in the GGF Document process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the GGF or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE GLOBAL GRID FORUM DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE."

8 References

1. D. Barret, R. Silverman, SSH: the Secure Shell, O'Reilly & Associates, 2001
2. I. Foster, C. Kesselman, Ed., The Grid, Blueprint for a New Computing Infrastructure, Morgan Kaufman, 1999
3. T. Dierks, C. Allen, *The TLS Protocol, Version 1* IETF RFC 2246; <http://www.ietf.org/rfc/rfc2246.txt>
4. C. Ellison SPKI Requirements, IETF RFC 2692 1999, <http://www.ietf.org/rfc/rfc2692.txt>
5. S. Garfinkel: PGP:Pretty Good Privacy. O'Reilly & Associates, 1994
6. R. Housley, W. Polk, W. Ford, D. Solo, *Internet X.509 Public Key Infrastructure Certificate and CRL Profile* <RFC3280>, <http://www.ietf.org/rfc/rfc3280.txt>
7. J. Novotny, S. Tuecke, V. Welch, *An Online Credential Repository for the Grid: MyProxy*, Proceedings of the Tenth International Symposium on High Performance Distributed Computing (HPDC-10) IEEE Press Aug 2001
8. S. Tuecke, et al. "Internet X.509 Public Key Infrastructure Proxy Certificate Profile", GGF Information document, draft-ggf-gsi-proxy.04, Aug 2002, <http://www.gridforum.org/security/gsi/draftpggf-gsi-proxy-04.pdf>
9. DOE Science Grids CA <http://www.doegrids.org>
10. CERN CA <http://globus.home.cert.ch/globus/ca/>
11. INFN <http://security.fi.infn.it/CA/>
12. UK HEP Testbed CA <http://www.gridpp.ac.uk/ca/>
13. "Cryptographic Token Interface Standard." PKCS11, <http://www.rsasecurity.com/rsalabs/pkcs/index.html>
14. "Cryptographic Token Information Format Standard", PKCS 15, <http://www.rsasecurity.com/rsalabs/pkcs/index.html>
15. "Security Requirements for Cryptographic Modules". <http://csrc.nist.gov/publications/fips/fips140-2/fips1402.pdf>
16. ITU-T Recommendation X.509 | ISO/IEC 9594-8, "The Directory: Public-Key and Attribute Certificate Frameworks", Draft V4 Feb 23, 2001
17. "Validation Lists for Cryptographic Standards", <http://csrc.nist.gov/cryptval/>