

Interoperable PKI Data Distribution in Computational Grids

*Massimiliano Pala**, Dartmouth College, USA
Shreyas Cholia, Lawrence Berkeley National Laboratory, USA
Scott A. Rea, Dartmouth College, USA
Sean W. Smith, Dartmouth College, USA

ABSTRACT

One of the most successful working examples of virtual organizations, computational grids need authentication mechanisms that inter-operate across domain boundaries. Public Key Infrastructures (PKIs) provide sufficient flexibility to allow resource managers to securely grant access to their systems in such distributed environments. However, as PKIs grow and services are added to enhance both security and usability, users and applications must struggle to discover available resources-particularly when the Certification Authority (CA) is alien to the relying party. This paper presents how to overcome these limitations of the current grid authentication model by integrating the PKI Resource Query Protocol (PRQP) into the Grid Security Infrastructure (GSI).

Keywords: Authentication, PKI, Resource Discovery, PRQP

AUTHENTICATION IN VIRTUAL ORGANIZATIONS

Computational grids provide researchers, institutions and organizations with many thousands of nodes that can be used to solve complex computational problems. To leverage collaborations between entities, users of computational grids are often consolidated under very large Virtual Organizations (VOs).

Participants in VOs need to share resources, including data storage, computational power and network bandwidth. Because these resources are valuable, access is usually limited, based on the requested resource and the requesting user's identity. In order to enforce these limits, each grid has to provide secure authentication of users and applications.

Erroneously granting access to unauthorized or even malicious parties can be dangerous even within a single organization---and is unacceptable in such large VOs.

Moreover, the dynamic nature of grid VOs requires the authentication mechanisms to be flexible enough to easily allow administrators to manage trust and quickly re-arrange resource-sharing permissions. Indeed, VOs are usually born from the aggregation of already existing organizations and constitute an umbrella that groups the participating organizations rather than

replacing them. Authentication must allow individual organizations to maintain control over their own resources.

The Problem. When participating in a VO, an organization must solve the problem of securely identifying resource requesters that come from outside its boundaries. PKIs offer a powerful and flexible tool to solve the potential authentication nightmare. Nonetheless, grid and VO administrators are still striving to find an acceptable solution to address interoperability issues that originate from the way VOs differ in policies, infrastructures and resource control.

Consider the situation where access to grid resources is managed via a Web portal. SSL mutual authentication can be enabled at the portal to implement strong authentication based on grid-approved PKI credentials. To do this, the portal administrator needs to set up the SSL Trust List to only allow credentials from approved CAs; the portal also needs to know how to validate the entire trust chain for that credential (that is, the end entity certificate presented, its issuer and the issuer's issuer, and so forth) up to the approved self-signed grid trust anchor.

To do this validation, the portal needs to know how to access services such as the location of the CA certificate and revocation data for each of these intermediate CAs. However, the portal cannot count on having pre-configured details for them. Even if it did--or if the information was packaged in each end entity certificate---this information may change over time, rendering this critical data stale. Having some way to dynamically discover service entry points of interest for grid-approved authorities (or indeed, the very authorities themselves) would solve a number of issues and would also provide for more flexible implementation options for the grid authorities, potentially lowering the costs of future service changes, and facilitating the future offering of additional services.

Our Proposed Solution. In order to help VOs to more efficiently address PKI interoperability issues we propose the adoption of the *PKI Resource Query Protocol* (PRQP) which enables discovery of resources and services in inter and intra PKI environments. We also propose an enhancement to the PRQP and we discuss its integration into the Grid Security Infrastructure (GSI).

AUTHENTICATION IN GRIDS

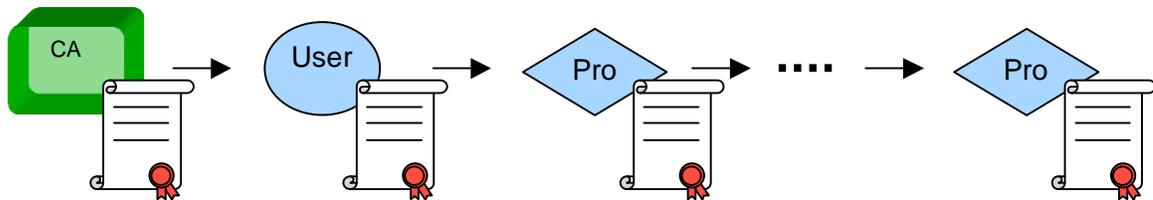
According to Ian Foster, a *grid* is a system that “coordinates resources that are not subject to centralized control, using standard, open, general-purpose protocols and interfaces, to deliver nontrivial qualities of service” (I. Foster, 2002). In order for the grid computing model to be successful, users and VOs must access a wide variety of resources using a uniform set of interfaces. Given that most resource providers have their own security policies and schemes to begin with, grids must overcome the challenge of integrating a wide variety of authentication mechanisms to achieve this kind of resource sharing. Without a common authentication layer, Virtual Organizations and resource providers are forced to adopt ad hoc schemes to achieve integrated resource sharing. However, the adoption of arbitrary schemes discourages information sharing and collaboration among researchers, and essentially makes the grid model unworkable.

The Globus Toolkit and its underlying Grid Security Infrastructure have become the de facto standards for building grids in research and academic communities. They provide applications, VOs and resource providers with a secure and standard means to perform authentication across organizational boundaries. GSI is built on top of a PKI layer and uses standard X509 v3 certificates for authenticating principals and granting access to local resources.

In a distributed environment, it is important to maintain traceability back to the individual entity matching a given certificate. The task of identifying users is distributed across various grid CAs throughout the world. These CAs are accredited and audited by the International Grid Trust Federation and its three regional Policy Management Authorities. A list of accredited CAs is maintained by the IGTF and distributed to relying parties throughout the world.

Grid CAs issue users a PKI certificate, including a public key linked to the private key controlled by the grid subscriber. These certificates may either be long-lived (typically issued by classic grid CAs) or short-lived (typically issued by online CAs such as SWITCH (SWITCH, 2008) or MyProxy-based CAs (NCSA, 2008)) depending on the use case. The IGTF maintains different authentication profiles to manage CAs with different qualities of service, for the benefit of relying parties.

Figure 1 – Chain of Trust in grids environment. The usage of Proxy Certificates allows the user to delegate tasks without exposing her private key—since each Proxy Certificate has its own unique keypair.



A resource provider or virtual organization relies on these CAs to be able to identify a given user. As such, if an end entity is able to present a valid certificate that is signed by a CA trusted by the relying party, the entity can be authenticated (of course, the end entity also needs to prove knowledge of the private key). GSI authentication is mutual (GLOBUS, 2008)—if a user wishes to access a service, both the user and the service must be able to present signed certificates to each other. The respective signing authorities must be trusted by the entity on each side of the transaction. Allowing the user and the service to have certificates signed by different CAs is the key to establishing cross-realm trust in grids. This also eases usability and scalability—the user need maintain only a single individual credential (single point of identity) no matter how many services she wishes to use. In order to improve usability, a user of grid services can sign a *Proxy Certificate (PC)* on his or her own behalf.

In general these proxies contain a slightly modified version of the user's identity (to indicate that it is a proxy certificate), a new public key, and a very short lifetime. These proxy credentials can then be used to access applications, or further delegated to application servers to perform actions on behalf of that user, without having to expose the user's original long-lived credential and private key—thus practicing the security principle of “least privilege.”

Most GSI-based grid applications can recognize PCs and will trust the credential as long as the chain of trust leads back to the original user and a trusted CA. A detailed scheme of the whole chain of certificates involved in identity verification is shown in **Figure 1**.

Additionally, VOs will often deploy a *Virtual Organization Management Service (VOMS)* (V. Ciaschini, 2004) that assigns roles to user certificates. The VOMS service will generate and sign an *Attribute Certificate* that contains one or more *Fully Qualified Attribute Name (FQAN)* strings, linked to the user's subject DN, which the user will embed in a X.509 proxy certificate as an X.509v3 extension. This FQAN defines that user's role within the VO. VOMS proxies can be used to manage roles and levels of access to resources, while using the same identity principal (user certificate) across the grid.

PKI RESOURCE DISCOVERY IN GRIDS

To use these more general PKIs, applications must be capable of finding and using services and data repositories provided by Certification Authorities. Unfortunately, even the retrieval of the list of revoked certificates (CRLs) is still a problem when dealing with CAs from different hierarchies or loosely coupled PKI meshes.

Grid PKIs can become rather complex, and the number of grid CAs accredited by the Policy Bodies (which are relatively young) is expected to grow in the near future. Indeed, as long as policies and common practices are established and well-understood, the number of accredited CAs should increase in the number of hundreds, thus increasing the need for a standardized solution for a PKI resource discovery system.

Figure 2 – Example of distributed info file within grid communities. Notice how some of the distributed information have no equivalent pointers in standard X509 certificates.

```
#
# @ (#) $Id: 1c3f2ca8.info,v 1.5 [...] $
# Information for CA DOEGrids
#   obtained from 1c3f2ca8 in DOEGrids/
alias = DOEGrids
url = http://www.doe grids.org/
crl_url = http://pkil.doe grids.org/CRL/1c3f2ca8.r0
email = trouble@es.net
requires = ESnet
status = accredited:classic
version = 1.16
shalfp.0 = 2D:7C:01: [...] :F8:90
```

Current Data Distribution. Currently, the mechanism for querying the trusted providers is fairly simple: administrators and users download a trusted CA distribution. This can either happen as part of a manual process, or it can be included within the grid software distribution (such as the Open Science grid software stack). This packaged data consists of a set of accredited CAs. (Accreditation is done by peer review in the various policy bodies.)

Because of the need to provide users and administrators with additional data besides the CA certificates, the downloaded package includes extra files. In particular, for a given CA, the package typically includes the following static information: the *CA certificate*, the *.info file*, a *CRL URL file*, a *namespaces file*, and a *signing policy file*.

The *.info file* contains general CA information along with contact information (including a URL). Applications can use information in the *.info file* to contact the CA. An example of a distributed *.info file* is shown in **Figure 2**. Some of the information distributed in this file (e.g. url, email or status) is required by applications and users to find details about the CA. The *CRL URL file* contains a URL pointer from where one would download the CRL. All accredited IGTF classic CAs provide this file. Sites and users build revocation lists by periodically querying the information in the CRL URL file and downloading revocation lists from the CRL url for *each* CA. This means that many grid software installations in the world are downloading these large CRLs from the CA providers at regular intervals. From what we have seen, this has often created denial of service conditions for certain CAs.

The *namespaces file* defines the *Distinguished Names (DN)* namespace that the CA is authorized to use; the *signing policy file* defines the rules for the signing policy of that CA. The *namespaces file* and the *signing policy file* may contain overlapping information from a policy point of view (although only the *signing policy file* has an implementation in software). Although this information could be embedded into a CA's certificate, the need for updating this data periodically led to the creation of the *.info file* and bundling it together with the certificate.

TACAR (Terena Academic CA Repository) and *IGTF* register and distribute this information to users and sites as follows. The accredited CA sends the trust anchor information directly to the IGTF/TACAR through a TERENA officer or a TERENA TACAR trusted introducer.

The IGTF packages and distributes the official CA package. Relying parties download the IGTF package every time there is a new release (approximately once a month). Relying parties are encouraged to verify this against the TACAR repository. Then, based on the information within the downloaded package, relying parties download the CRL from the CRL URL on a daily basis.

Ultimately, in most cases, this relies on a very static “cron-based” process. There are several improvements to this that can be made by PRQP that would replace this type of static file and crontab based access with something more dynamic, and query driven.

Other Solutions. To publish pointers to data, a CA could use certificate extensions such as the *Authority Information Access (AIA)* and the *Subject Information Access (SIA)* (R. Housley, W. Polk, W. Ford, and D. Solo, 2002). Regrettably the lack of support built into applications and the difficulties in updating extensions in certificates clash with the need for flexibility required by today CAs.

To overcome the problem with updating the pointers, it is possible to use SRV records (A. Gulbrandsen, P. Vixie, and L. Esibov, 2000) in DNS (P. Mockapetris, 1987). Although interesting, the problem with this solution resides in the lack of correspondence between the DNS structure, which is built on a strictly hierarchical namespace, and PKIs where there are no requirements for the used namespace.

Other solutions are either overly complicated to solve our problem---e.g., Web Services (F. Curbera et al, 2002) uses *SOAP* (A. Karmarkar et al, 2003), *WSDL* (E. Christensen et al, 2001; R. Chinnici, et al., 2005) and *UDDI* (L. Clement et al., 2004)) or they are specifically targeted to local area networks---e.g., *Jini* (W. Edwards, 2000; K. Arnold, 2000), *UPnP* (UPnP Forum, 2008; M. Jenronimo and J. Weast, 2003) or *SLP* (M. Jenronimo and J. Weast, 2003; E. Guttman, C. Perkins, and J. Kempf, 1999; E. Guttman, 1999).

TRUST AND CERTIFICATION POLICIES

The use of a standardized and well-established technology such as public key certificates has enabled applications such as browsers to facilitate ease of use within grids. However, especially when integrating credentials from different authorities, an important aspect to consider is the policies under which those credentials have been issued. Although a PKI potentially provides the benefit of strong binding of identities to public keys, the strength of that binding is really dependent on the policies and practices followed by the issuing authority, and the subscribers.

A CA is a trusted third party entity which issues digital certificates for use by relying parties. In a certificate, the CA attests that the public key matches the identity of the owner of the corresponding private key, and also that any other data elements or extensions contained in the certificate match the subject of the certificate. The obligation of a CA (and its registration authorities) is to verify an applicant's credentials, so that relying parties can trust the information contained in the certificates it issues. If a relying party trusts the CA and can verify the CA's signature, then it can also verify that a certain public key does indeed belong to whoever is identified in the certificate (as long as they accept this, the end entity is fulfilling its responsibilities with respect to protecting the private key). If the CA can be subverted, then the security of the entire system is lost; likewise, if an end entity is negligent, then the security and trust associated with their particular credential could be lost.

The degree to which a relying party can trust the binding embodied in a digital certificate depends on several factors. These factors can include the practices followed by the certification authority in authenticating the subject; the CA's operating policy, procedures, and security controls; the scope of the subscriber's responsibilities (for example, in protecting the private key);

and the stated responsibilities and liability terms and conditions of the CA (e.g warranties, disclaimers of warranties, and limitations of liability). The processing of information contained in these multiple complex documents for the purpose of making a trust decision about each PKI involved is too onerous for the average user. Relying parties therefore usually accept recommendations from trusted accreditation bodies about the relative trustworthiness and suitability of credentials being issued by a particular CA. For grids, those accreditation bodies are the three regional PMAs that constitute the IGTF. *TAGPMA* is the accreditation authority for the Americas (covering a geographical region from Canada to Chile).

TAGPMA conducts peer reviews of grid CA operations. A grid CA can be accredited as a grid credential issuer after TAGPMA reviews their *Certificate Policy (CP)* and *Certification Practices Statement (CPS)* to ensure that the practices implement the policies and that the policies are equivalent to standard approved grid profiles. Once approved, the CA and associated information is packaged for official distribution for IGTF relying parties. Re-review of a CA is conducted on a periodic basis to ensure they are still compliant with the standard grid profiles.

Not all grid CA accreditation applicants are able to map their existing policies and practices to an approved IGTF profile. However, a relying party may still wish to accept the credentials of such a CA operator based upon their own assessment of trustworthiness of the CA. In order for the relying party to make a local trust decision, they should consider the statements by the CA published in their CP and CPS and also review any other relevant security or trust-related documentation. Currently this information is generally not readily available to a relying party from the CA's certificate, nor can a relying party or potential subscriber easily find the URI for the application or revocation of credentials from such CAs. A mechanism for publishing and updating this information would greatly enhance the flexibility, and usability of potential grid PKIs. The PRQP is a perfect candidate for providing such functionality.

ALLOWING FOR BETTER INTEROPERABILITY BETWEEN GRID PKIS

Effective authentication frameworks that make use of certificates potentially require many different services provided by accredited CAs such as OCSP servers, CRL repositories, timestamping services, etc. As a consequence, certification authorities need to be able to provide these services and to enable applications to discover them.

Because the need to distribute PKI related data and pointers to services is of primary concern in grids, each grid environment defines its own specific format and solution. Although this might temporarily solve specific issues within a specific grid community, it does not encourage the exchange of information and interoperability with other organizations.

It is to be noted that because of the customized nature of current solutions, specific extensions to applications must be developed in order to be able to operate in such environments.

The PKI Resource Discovery Protocol. The notion of a discovery protocol for PKIs first appeared in in our earlier paper (M. Pala and S.W. Smith, 2007), which proposed the *PKI Resource Query Protocol (PRQP)*¹ to provide pointers to any available PKI resource from a particular CA.

The PRQP (M. Pala, 2008a) has been already discussed in the IETF PKIX working group. The updated version of the PRQP specification, which will include modifications proposed in this paper, is scheduled to be published as an Internet Draft on the experimental track. In PRQP, the client and a *Resource Query Authority (RQA)* exchange a single round of messages where the client requests a resource token by sending a request to the server and the server replies back by sending a response to the requesting entity.

¹ The subsequent description here of the PRQP protocol is derived from our earlier paper (M. Pala and S.W. Smith 2007).

The client may ask for the location of all the services provided by a CA by not specifying any identifier in the request. Alternatively, the client can request the address of one or more specific services by embedding one or more Object Identifiers (OIDs) into the request. The resources might be items that are (occasionally) embedded in certificates today---such as URLs for CRLs, OCSP, SCVP or CP/CPS locations---as well as other items, such as addresses for the CA website, the subscription service, or the revocation request.

The Resource Query Authority. In PRQP, the server is called the *Resource Query Authority* (RQA). An RQA can play two roles. First, a CA can directly delegate an RQA as the party that can answer queries about its certificates, by issuing a certificate to the RQA with a unique value set in the *extendedKeyUsage* (i.e. *prqpSigning*). The RQA will provide authoritative responses for requests regarding the CA that issued the RQA certificate. Alternatively, an RQA can act as *Trusted Authority (TA)* (“trusted” in the sense that a client simply chooses to trust the RQA’s recommendations and assertions). In this case, the RQA may provide responses about multiple CAs without the need to have been directly certified by them.

In this case, provided responses are referred to as *non-authoritative*, meaning that no explicit trust relationship exists between the RQA and the CA. To operate as a TA, a specific extension (*prqpTrustedAuthority*) should be present in the RQA’s certificate and its value should be set to TRUE. In this configuration the RQA may be configured to respond for different CAs which may or may not belong to the same PKI as that of the RQA.

Security Considerations. The PRQP provides URLs to PKI resources, therefore it only provides locators to data and services, and not the real data. It still remains the client’s job to access the provided URLs to gather the needed data, and to validate the data (e.g., via signatures or SSL).

Because of this consideration, both the NONCE and the signature are optional in order to provide flexibility in how requests and responses are generated.

Also, it is then possible to provide pre-computed responses in case the NONCE is not provided by the client. If an authenticated secure channel is used at the transport level between the client and the RQA (e.g. HTTPS or SFTP) signatures in requests and responses can be safely omitted.

Distribution of RQA address. The distribution of the RQA’s address to clients is still an open issue. There are three possible approaches. A first option would be to use the AIA and SIA extensions to provide pointers to RQAs. We believe that by using these extensions in certificates to locate the RQA, one could provide an easy way to distribute the RQA’s URL. The size of issued certificates would be smaller than embedding all the pointers for CA’s resources, thus providing a more space efficient solution.

The second option is applicable mostly for LANs, and consists of providing the RQA’s address by means of DHCP. This method would be mostly used when a trusted RQA is available on a local network. These two techniques can then be combined together.

Ultimately, the RQA’s address can also be embedded directly into application software distributions. This approach could be adopted in grids and VOs where a centralized software distribution system is in place. At each software update, the RQA network address can be updated as well. If the distributed software is not signed by a trusted authority, this approach could be subject to serious security threats, e.g. distribution of an altered package by a malicious attacker where the configured RQAs are not the “official” ones. Besides the security considerations already discussed above, the trust level in the application’s RQA configuration should be not less than that put into browser or operating system certificate stores.

INTEGRATING PRQP INTO GSI

In our work, we analyzed the security requirements of grids and the current challenges in distributing pointers to data for authentication. To ease the administrators' burden and to provide a more efficient way to distribute resource locators, we extended the PRQP specification with

Figure 3 – ASN.1 description of the CA identifier added to the *ResponseMessage* in our modified version of PRQP

```
BasicCertIdentifier ::= SEQUENCE {
    subjectNameHash      OCTET STRING,
    issuerNameHash       [0] OCTET STRING  OPTIONAL,
    serialNumber         [1] SerialNumber  OPTIONAL }
```

grid-specific support. In particular, this work aims to provide an interoperable method to distribute information about services provided by CAs. Although some solutions already exist in the computing grid environment (e.g. the monthly IGTF/TACAR update), our work addresses the problem by providing a more standardized solution that would allow for better interoperability between organizations (as discussed earlier).

The GSI is part of a larger bundle of tools provided by the Globus toolkit. The security layer is built on top of the OpenSSL library, a widely used open-source library.

We developed a PRQP library, server and client application that can be integrated into existing PKI software. We also simplified and enhanced the PRQP messages in order to better support grid needs. We also integrated PRQP into OpenCA's LibPKI (M. Pala, 2008b) which will be the core library for OpenCA Next Generation Certification Authority software (M. Pala, 2008c). At present, the developed software is available as a stand-alone application (M. Pala, 2008d). In the future, we plan to integrate a PRQP client and a PRQP server into the Globus Toolkit.

In the following sections we describe how we modified the PRQP messages for grid integration and how we envisage the integration of PRQP into existing grids by proposing a deployment scenario within the TAGPMA community.

Our Modified PRQP. Originally intended for general purpose PKIs, the original PRQP had responses containing the fields *version*, *NONCE*, *PKIStatusInfo* and *ResourceResponseToken*.

The *Version* specifies the version of the message protocol. The *NONCE*, a random large integer, binds the response to a specific request in signed responses in order to defend against reply attacks. The *PKIStatusInfo* carries the response status and, in case of error, a description of the cause. The *ResourceResponseToken* is a more complex data structure that carries the URLs of the requested services.

Because of special need of grids' PKIs, we modified the protocol in two different ways: we modified PRQP responses, and we defined new OIDs to support grid specific pointers. We now discuss these changes.

Efficient PRQP Response Caching. The PRQP is efficient and simple in design. By using a software implementation on commonly available hardware, a client requires around 12ms to retrieve a signed response from a Resource Query Authority.

Moreover, in the Grid environment, the protocol can be executed as little as once per day by the authentication framework, thus making the already small overhead introduced by PRQP negligible. In order to efficiently cache PRQP responses, we propose a change in the protocol.

In the original protocol, only PRQP requests carry an identifier for the CA. This identifier is used by the RQA to identify the CA whose pointers are requested by the client. Although efficient, the client would not be able to identify the CA that the response refers to by simply looking at the response.

Table 1 - Newly Identified OIDs for Grid Operations. Of particular interest are the Grid specific pointers that enable an RQA to provide Grid specific information to applications. It is also to be noted that some of the proposed PKIX Identifiers refer to services that are not yet standardized.

	<i>OID</i>	<i>Text</i>	<i>Description</i>
PKIX	{id-ad 1}	ocsp	OCSP Service
	{id-ad 2}	caIssuers	CA Information
	{id-ad 3}	timeStamping	TimeStamping Service
	{id-ad 10}	dvcs	DVCS Service
	{id-ad 11}	scvp	SCVP Service
General PKI operations	{id-ad 50}	certPolicy	Certificate Policy (CP) URL
	{id-ad 51}	certPracticesStatement	Certification Practices Statement (CPS) URL
	{id-ad 60}	httpRevokeCertificate	HTTP Based (Browsers) Certificate Revocation Service
	{id-ad 61}	httpRequestCertificate	HTTP Based (Browsers) Certificate Request Service
	{id-ad 62}	httpRenewCertificate	HTTP Based (Browsers) Certificate Renewal Service
	{id-ad 63}	httpSuspendCertificate	Certificate Suspension Service
	{id-ad 40}	cmsGateway	CMS Gateway
	{id-ad 41}	scepGateway	SCEP Gateway
	{id-ad 42}	xkmsGateway	XKMS Gateway
	{eng-ltd 3344810 10 2}	webdavCert	Webdav Certificate Validation Service
	{eng-ltd 3344810 10 3}	webdavRev	Webdav Certificate Revocation Service
Grid Specific	{id-ad 90}	accreditationBody	Accreditation Body URL
	{id-ad 91}	accreditationPolicy	Accreditation Policy
	{id-ad 92}	accreditationStatus	Accreditation Status Document
	{id-ad 95}	commonDistributionUpdate	Grid Distribution Package
	{id-ad 96}	accreditedCACertificates	Certificates of Currently Accredited CAs

We added a CA identifier in the PRQP response message. This identifier allows the client to tie the information received from an RQA to a CA without the need to cache the sent request as well. The new CA identifier structure definition is shown in **Figure 3**. By adding the new data structure, we introduced a small overhead in terms of response size, however this modification simplifies response caching on PRQP enabled clients. Moreover, because the CA identifier does not change, its contents can be pre-computed, thus it does not add any significant computational burden on the server.

Defining Grid-Specific Pointers. In order to better leverage PRQP in the Grid environment, we defined a set of object identifiers (OIDs) that enhance PRQP with the ability to provide grid-specific data distribution. Because grid communities organize themselves in VOs that accept

common authentication profiles (such as those of the IGTF), it has been easy to analyze the requirements and identify the needed enhancements to PRQP.

Besides identifying the OIDs for general PKI operations (e.g., HTTP based or browser-specific services, CA “communication gateways”, etc.)², we also defined some Grid-specific pointers (see **Table 1**).

The *accreditationBody* and the *accreditationPolicy* pointers can be used to specify the bodies and the policies (or profiles) under which a CA has been accredited. In addition to these, we also defined the *commonDistributionUpdate* and the *accreditedCACertificates* OIDs. These identifiers can carry information about pointers to the most recent Grid distribution data (the former) and to the set of accredited CA certificates (the latter).

One interesting feature of PRQP is its flexibility. It can provide CA management with a dynamic model to add services or, if needed, to switch to newer and more efficient ones. This feature becomes of primary concern in grids where currently grid-specific services have not been standardized yet.

CAs can leverage this feature of PRQP in order to provide dynamically updated information about its accreditation status to applications by using the *accreditationStatus* pointer. This set of grid-specific pointers can also facilitate more flexible trust options from the VO's perspective, in the set of CAs it chooses to trust. For instance, besides the generally accepted IGTF distribution, these pointers also allow a VO to specify a set of additional CAs that the VO wishes to trust locally (that the VO has vetted itself for use within the community), by simply specifying an additional local distribution maintained by the VO or any entity it delegates this responsibility to (e.g. refer to the additional non-IGTF accredited CAs that are accepted by TeraGrid).

THE TRUST MODEL

An interesting aspect of the grid trust model is the presence of a central authority, often embodied by the grid policy management authority. Usually this authority is represented by a federation of authentication providers and relying parties responsible for accreditation of CAs willing to participate in the organization.

The presence of such an authority eases the deployment of PRQP in that it provides a central point where the RQA can be deployed. In this section, we discuss the issues and the benefits arising from adopting RQAs in two different ways: a centralized approach where a centralized RQA service would serve the entire grid community (e.g., IGTF/TACAR) or by adopting a more decentralized approach where participating grids or VOs run their own RQA infrastructures.

Trusting a Central RQA. One possible trust model envisages the use of a centralized Resource Query Authority which would serve all the organizations participating in the grid community.

This model is easily applicable when the VOs and grids share the same set of accredited Certification Authorities. In other words, it best fits organizations where grids and VOs only recognize the same set of accredited CAs (e.g., the ones accredited by the IGTF). In this case, the client application queries the central RQA and finds out the information needed about a particular CA. For this model to work, the central RQA must know the pointers for each and every CA that is recognized by the participating grids. In this case, the RQA should be trusted by all the participating parties. The RQA can be configured to act as a trusted responder or, if every participating CA is willing to certify the RQA's key pair, as an authoritative responder.

It may be unrealistic to expect a policy authority (like IGTF) to operate a central RQA which would require 24x7 support; however, the operation could be delegated by IGTF to one of the more prominent accredited CA sites that are already geared for 24x7 services, or to a community

² A more complete explanation of the non grid-specific pointers is currently submitted for publication.

service point like TACAR. The IGTF would then simply need to require periodic assertions (or audits) to confirm that the central service was operated precisely and integrally.

A Per-Grid RQA Model. A per-grid RQA Model can be adopted in order to provide grids with the possibility to configure multiple CAs which might not be accredited by the larger VOs (or policy bodies). In this model, we envisage the deployment of RQAs on a per-grid basis. From several points of view, this model might seem to be better than having a single RQA operated by the accreditation body (e.g., IGTF/TACAR).

Individual grid infrastructures in practice often support additional CAs over and above those in the standard CA distribution. For example there are a number of TeraGrid CAs that have not been accredited by the IGTF, but have been operational and in use on that grid for some time. Also, VOs or grids may not want to recognize certain accredited CAs for policy reasons. Having an RQA at the level of a computational grid allows that grid to determine its own boundaries of trust, as opposed to using one predefined by the IGTF/TACAR.

This model has another interesting property. It may simplify the problem of distributing pointers to the RQA by simply integrating this with the software stack for that particular grid.

In contrast, having separated authorities for each grid could impact interoperability between grids. Our belief is that the distributed model best fits many environments, especially grids.

As we discuss at the end of this paper, we also believe that future work in the field shall be directed into building a world-wide RQA infrastructure which would act as a DNS for PKIs.

SOLVING THE PROBLEM: TWO EXAMPLES

In this section we provide two different use examples that clarify how the use of PRQP can improve PKIs usability and reliability.

Discovering OCSP. In many cases, the authentication layer of the grid software needs to discover services provided by accredited CAs in order to get the latest information about the validity of a certificate.

Normally the authentication server would receive the certificate to validate from the application and, by using the local configuration, or by looking at the certificate's contents (e.g., extensions), it would try to contact the identified OCSP server.

By integrating PRQP into the authentication server, the reliability of the location of the OCSP URL can be improved. The authentication server could issue a query to the RQA before proceeding to validate the certificate. This would enable the server to discover (a) if the URL of the OCSP embedded in the certificate (if any) is still valid, or (b) if more OCSP servers are currently available, or (c) if another validation service has become available (e.g., a new SCVP server has been setup).

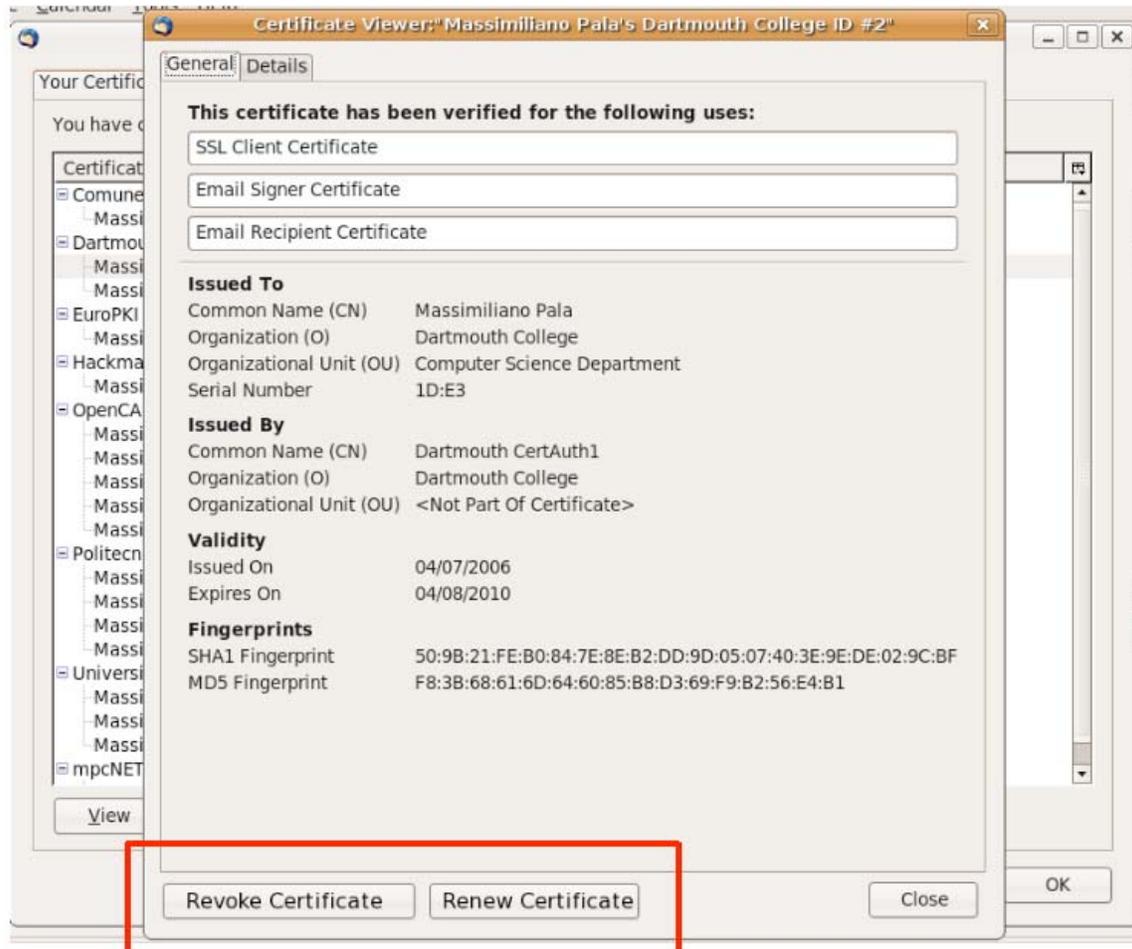
Ultimately the RQA provides the authentication server with the list of URLs of the requested services. The additional step provided by PRQP would enable authentication services to identify classes of equivalent services and use the ones that are provided by the issuing CAs in a dynamic fashion.

Revocation of Short-Lived Certificates. Users accessing grid services often have two different types of credentials: a long-lived credential (i.e., their identity certificate) and a short-lived credential (i.e., a proxy or end-entity certificate with a short lifetime, usually 24 hours). Because the second set of credentials is valid for a very short period, grids often do not require revocation support for these kinds of credentials. Regrettably, requesting certificate revocation is often not an easy task even for long-lived certificates.

In many cases, the URL to request the revocation of a certificate is only distributed to the user by means of an email sent at the time the credential was initially issued, and that email may not be easily accessible when the subscriber has the greatest need for it. When considering short-lived

certificates the situation is even worse. They do not usually present any revocation access point, and it is not usual for grid services relying upon them to check their validity. However, in practice it has been shown that many benefits could arise if such validation data were available, and could be checked by applications prior to reliance on these credentials, in certain contexts where higher assurance is required.

Figure 4 – Example of usable UI where the user is presented with two simple buttons to revoke or renew her own certificate.



PRQP can be utilized to add validation checking, e.g. through the standing up of an OCSP service, and providing pointers to the service via PRQP (for those applications that need it), without the need to reissue any of the short-lived credentials or even changing the profile of the credentials for future issuance.

Applications could even present easier User interfaces (UIs). For example an “Ask for Revocation” button could help the user with her interaction with the CA: the application could automatically access the revocation service by requesting the URL from the RQA. This approach potentially eases key management issues for the subscriber, thus enhancing the certificate usage experience for the user ().

In this way, the same short-live credential can be used with those applications that have a need for higher assurance with respect to checking validation on this class of credential, as well as with those applications where such assurance is not required. Moreover, this can be done without having to change any of the underlying infrastructure built to cater to the latter in

response to the requirements of the former. In this way PRQP facilitates greater flexibility for grid applications.

CONCLUSIONS AND FUTURE WORK

In our work we provide a description of the grid authentication layer. We also provide an overview of the issues that grids and virtual organizations face every day in distributing crucial information that enables the usage of digital certificates.

Our work also analyzes the current status of the PKI Resource Query Protocol and proposes an enhanced version that specifically targets the needs of grids. Additionally, two different PRQP adoption models are discussed in detail.

We believe that PRQP would provide an effective solution to the PKI services pointer distribution issue, especially in grids where a common authentication layer exists. The PRQP introduces a new layer of indirection that allows mapping of PKI resource discovery to network addresses. Today, no existing software provides such a flexible service. In fact, no deployed infrastructure exists that provides an efficient and interoperable PKI resource-discovery service. Although it is possible to provide similar services by using existing protocols, the lack of a standardized (and specific) approach to the problem has led to today's situation where no PKI actually provides a discovery service (or even addresses the issue).

Our future work will be focused on allowing for improved interoperability among RQAs. In particular we are studying two different possibilities. The first one is to enhance the protocol by providing the notion of a *referral server*. In this case if the answer for the queried resource/CA is not found at a particular RQA, the server can redirect the client to a different one. Although easy to deploy, this solution could be somewhat impractical because it requires hierarchies among RQAs.

To overcome this problem we are also investigating the adoption of a Peer-2-Peer (P2P) based approach. A P2P network would map network addresses to services mostly like the DNS maps logical names to IP addresses. The main difference between the DNS and the RQA network would be the absence of a hierarchy and the possibility to dynamically add or dismiss RQAs participating in this P2P network.

ACKNOWLEDGMENTS

The authors would like to thank the IGTF members for their contribution and inspiring suggestions. This work was supported in part by the NSF (under Grant CNS-0448499); the U.S. Department of Homeland Security (under Grant Award Number 2006-CS-001-000001); and the Director, Office of Science, Office of Advanced Scientific Computing Research of the U.S. Department of Energy (under Contract No. DE-AC02-05CH11231). The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of any of the sponsors.

A preliminary version of this work appeared as Pala et al, "Extending PKI Interoperability in Computational Grids," *8th IEEE International Symposium on Cluster Computing and the Grid*.

REFERENCES

- I. Foster (2002), "What is the Grid? - a Three Point Checklist", *GRIDtoday*, vol. 1, no. 6.
- SWITCH (2008), "*SWITCHpki, an X.509 Public Key Infrastructure for the Swiss Higher Education System*." [Online]. Available: <http://www.switch.ch/pki/>
- NCSA (2008), "*MyProxy Credential Management Service*." [Online]. Available: <http://grid.ncsa.uiuc.edu/myproxy/ca/>

- GLOBUS (2008), “*Overview of the Grid Security Infrastructure.*” [Online]. Available: <http://www.globus.org/security/overview.html>
- V. Ciaschini (2004), “*A VOMS Attribute Certificate Profile for Authorization.*” [Online]. Available: <http://grid-auth.infn.it/docs/AC-RFC.pdf>
- R. Housley, W. Polk, W. Ford, and D. Solo (2002) “*Certificate and Certificate Revocation List (CRL) Profile.*” Internet Engineering Task Force: RFC 3280.
- A. Gulbrandsen, P. Vixie, and L. Esibov (2000), “*A DNS RR for specifying the location of services (DNS SRV).*” Internet Engineering Task Force: RFC 2782.
- P. Mockapetris (1987), “*Domain Names - Implementation and Specification.*” Internet Engineering Task Force: RFC 1035.
- F. Curbera, M. Duftler, R. Khalaf, W. Nagy, N. Mukhi, and S. Weerawarana (2002), “Unraveling the Web Services Web: An Introduction to SOAP, WSDL, and UDDI,” *IEEE Internet Computing*, vol. 6, no. 2, pp. 86–93.
- A. Karmarkar, M. Hadley, N. Mendolsohn, Y. Lafon, M. Gudgin, J.-J. Moreau, H. Nielsen. (2007) “*SOAP Version 1.2 Part 1: Messaging Framework (Second Edition).*” W3C Recommendation. [Online]. Available: <http://www.w3.org/TR/2007/REC-soap12-part1-20070427/>
- E. Christensen, F. Curbera, G. Meredith, and S. Weerawarana (2001), “*PWeb Services Description Language (WSDL) 1.1.*” W3C Note. [Online]. Available: <http://www.w3.org/TR/2001/NOTE-wsdl-20010315>
- R. Chinnici, M. Gudgin, J.-J. Moreau, and S. Weerawarana (2005), “*Web Services Description Language (WSDL) Version 2.0 Part 1: Core Language.*” W3C Working. [Online]. Available: <http://www.w3.org/TR/wsdl20>
- L. Clement, A. Hatley, C. von Riegen, and T. Rogers. (2004), “*UDDI Version 3.0.2.*” [Online]. Available: http://uddi.org/pubs/uddi_v3.htm
- W. Edwards (2000), *Core Jini (2nd edition)*, Prentice-Hall.
- K. Arnold (2000), *The Jini Specification (2nd edition)*. Addison-Wesley.
- UPnP Forum (2008), “*Universal Plug and Play Specifications.*” [Online]. Available: <http://www.upnp.org/resources/>
- M. Jenronimo and J. Weast (2003), *UPnP Design by Example: A Software Developer’s Guide to Universal Plug and Play*, Intel Press, ISBN-10: 0971786119, ISBN-13: 978-0971786110
- E. Guttman, C. Perkins, J. Veizades, and M. Day (1999), “*Service Location Protocol, version 2.*” Internet Engineering Task Force: RFC 2608.
- E. Guttman, C. Perkins, and J. Kempf (1999), “*Service Templates and Schemes.*” Internet Engineering Task Force: RFC 2609.
- E. Guttman (1999), “Service Location Protocol: Automatic Discovery of IP Network Services,” *IEEE Internet Computing*, vol. 3, no. 4, pp. 71–80.
- M. Pala and S. W. Smith (2007), “AutoPKI: A PKI Resources Discovery System,” *Public Key Infrastructure: EuroPKI 2007*, Springer-Verlag LNCS 4582, pp. 154–169.
- M. Pala (2008a), “*PKI Resource Discovery Protocol (PRQP).*” PKIX Internet Engineering Task Force, Internet Draft, Experimental.
- M. Pala (2008b), “*LibPKI: the OpenCA’s Easy PKI Library.*” [Online]. Available: <http://www.openca.org/projects/libpki/>
- M. Pala (2008c), “*OpenCA-NG: the Next Generation CA.*” [Online]. Available: <http://www.openca.org/projects/ng/>
- M. Pala (2008d), “*OpenCA’s PKI Resource Discovery Package.*” [Online]. Available: <http://www.openca.org/projects/prqpd/>
- J. Schaad and M. Myers (2007), “*Certificate Management over CMS (CMC) Transport Protocols.*” IETF Draft.
- D. W. Chadwick (2007), “*Use of WebDAV for Certificate Publishing and Revocation.*” IETF, Internet Draft.