Reply to "Which Certificate Authority Should LIGO Use?"

Jim Basney and Scott Koranda

LIGO DCC LIGO-T0810038-x0

December 11, 2008

Abstract

During two teleconferences and exchanges of emails the authors, along with Michael Helm and Dhiva Muruganantham from Energy Sciences Network (ESnet) and John Volmer from the DOEGrids Certificate Authority (CA) Policy Management Authority, discussed technical, organizational, and operational issues around the question of which certificate authority the LIGO project should use and what role ESnet staff and resources, including the DOEGrids CA, might play in the LIGO authentication and authorization infrastructure. The discussions were enormously helpful in clarifying the technical ramifications of the LIGO use case requirements and led to a fuller exploration of the solution space than would have otherwise happened. The authors conclude that the LIGO use case requirements, aimed at significantly reducing the burden on LIGO scientists and users, make it difficult to architect a solution that involves ESnet in general and the DOEGrids CA in particular. The most appealing solution identified that meets the LIGO use case requirements is for LIGO to operate its own short lived credential services (SLCS) with deployments located at the major LIGO Data Grid (LDG) sites to meet the requirements for robustness against network outages.

1 Introduction and background

LIGO wishes to remove the burden from users of requesting, retrieving, and managing X.509 digital certificates and the associated private keys. A new authentication and authorization infrastructure design included deploying MyProxy [1] servers at all LIGO computing sites, and storing X.509 certificates and long-lived private keys in the MyProxy repositories. The private keys were to be stored unencrypted so that the MyProxy server could generate and issue a proxy certificate to the authenticated user. The certificates and keys stored in the MyProxy repositories were to be managed by the infrastructure and administrators so that end users no longer needed to request, retrieve, and manage the certificate and key files. In detail the design held that:

- A private key for each user is stored unencrypted in the MyProxy repository. The private key is protected only by the file protection mechanisms of the operating system.
- Users would authenticate to the MyProxy respository using a LIGO credential (Kerberos). This authentication step assures that a LIGO user only has access to his own private key.
- The private key material was to be generated not by the user but by the centralized infrastructure (a privileged account able to generate X.509 certificate requests and associated private keys, sign the request using the private key of a CA, and deploy the signed certificate and private key into the MyProxy repository).

Storing unencrypted in a central repository the associated private keys for X.509 certificates issued by the DOEGrids CA violates the DOEGrids Certificate Policy (see section 2.1.2 of version 2.9 of the policy document [2]), and currently LIGO users are issued X.509 credentials signed by the DOEGrids CA. Further, the proposed design conflicted with requirements from The Americas Grid Policy Management Authority (TAGPMA) [3] and compliance with requirements dictated by the International Grid Trust Federation (IGTF) [4]. Specifically these conflicts arose in the design:

1. The private keys are stored unencrypted.
2. System administrators and privileged users have access to the private keyts.
3. The private key material was not generated by the user but instead generated for the user on a machine to which the user does not have access.

Three options were initially identified for moving forward with the new infrastructure: 1) start a dialogue with the DOEGrids CA to determine if the existing policy could be amended to support the desired LIGO use case or perhaps another subordinate CA might be created and managed by the DOEGrids CA to support the LIGO use case. 2) Continue a dialogue with the TeraGrid about the TeraGrid CA issuing credentials to LIGO users to be stored in LIGO MyProxy repositories. 3) Deploy a LIGO CA. The document *Which Certificate Authority Should LIGO Use?* was written by the Authentication and Authorization Subcommittee of the LIGO Computing Committee and distributed to the Open Science Grid (OSG) [5] and DOEGrids Policy Management Authority (PMA) [6]. The document served to begin the dialogue between LIGO and the DOEGrids CA and in particular the dialogue between the authors of this letter.

After a series of emails, primarily between Mike Helm from ESnet[7] and Scott Koranda from LIGO, to clarify some technical details and the LIGO use case requirements, a teleconference was held on September 3 with Mike Helm, Scott Koranda, Dhiva Muruganantham (ESnet), and John Volmer (PMA). During that conference we realized it would be helpful to have Jim Basney on future call since he is the lead developer and architect for MyProxy, as well as serving in OSG as the Security Policy Officer and on the DOEGrids PMA, and having substantial experience with authentication and authorization issues in the grid space. On September 15 the entire group met during a teleconference and continued the discussion of which solutions best met the LIGO use case requirements and what if any role might the DOEGrids CA and ESnet staff play in the solution.

2 On the question of a change in DOEGrids CA policy

LIGO asked if the DOEGrids CA certificate policy [2] might be amended to allow private keys to be kept unencrypted in a MyProxy repository. An examination of the details of the policy revealed that the language used in the policy overreaches since it requires private key material to be kept encrypted without any reference to how the key material is stored. Apart from the question of storage in a MyProxy repository, the language

was identified as being too strong because private keys may be safely kept on devices like key or smart cards, and it is unclear if the keys are (strictly) kept encrypted while stored on the cards. The DOEGrids PMA staff that are part of the discussion intend to address this language with the entire PMA in the future.

More relevant here is that the discussion led to a further clarification of the lifetime of key material and how private key material is generated and then stored in different scenarios for the LIGO use case. It is clear that the fundamental policy and restrictions that the DOEGrids CA must enforce in order to meet its own requirements and maintain its own accreditation require that private key material for end entity (i.e. user) certificates must only be generated by the end entity and not in a central location by a centralized administrator. Further the private key material must not be stored in a repository like MyProxy unencrypted so that administrators have access to the unencrypted private key material.

These restrictions do not permit LIGO administrators to generate private key material on behalf of users (i.e. generate certificate requests), orchestrate the DOEGrids CA signing the requests, and storing the private keys and certificates unencrypted in a MyProxy repository to be used later to generate proxy certificates. We cannot see any path forward where the DOEGrids CA adjusts its policy to meet the LIGO use case requirements without compromising the DOEGrids CA current accreditation with The Americas Grid Policy Management Authority (TAGPMA) [3] and compliance with requirements dictated by the International Grid Trust Federation (IGTF) [4].

3 On the question of a subordinate CA

LIGO asked if ESnet might manage a CA subordinate to the ESnet Root CA and use it to sign certificate requests and have the certificates and associated long-lived private keys stored in a LIGO controlled MyProxy repository.

There is a related precedent. The FusionGrid CA [8], which is subordinate to the ESnet Root CA just as is the DOEGrids CA, does issue certificates with the long-lived private keys kept unencrypted in a MyProxy repository. The FusionGrid CA and the MyProxy repository are managed by ESnet, but it is important to note that this service is hosted only on ESnet computers connected to ESnet networks and wholly controlled by ESnet. Further the FusionGrid CA is *not* accredited by TAGPMA nor does it conform to any authentication profile offered by IGTF.

The LIGO use case requires that the MyProxy repositories be located at each of the LIGO Data Grid computing centers so that in the case of a network outage users could still access credentials from the local server. We cannot see any path forward where ESnet staff manage a subordinate CA service(s) at many LDG sites on hardware and networks not under the control of ESnet, nor would we even recommend such a solution, due to the challenges of remotely managing services for another collaboration.

4 LIGO short lived credential services

The discussions identified short lived credential service(s) (SLCS) [9] as an attractive component of a solution architecture that meets the LIGO use case requirements. With a SLCS the key material is generated by the end entity at the time he wishes to obtain a credential. The public key material (certificate request) is sent over an encrypted channel to the SLCS (LIGO would mostly likely use MyProxy to provide the SLCS, although the GridShib [10] CA might also be considered at some point) and provided the user has authenticated to the SLCS and is properly authorized the service signs the request and returns a short-lived X.509 certificate (for example with a lifetime of one week). The private key material is not sent over the network at any time.

In the LIGO use case the short-lived certificate and associated private key could immediately be used to generate a RFC 3820 proxy credential and both the short-lived X.509 certificate and associated private key would be destroyed. The user would use the proxy credential to authenticate to LDG resources. When the proxy expires the user would again create new key material and obtain a new short-lived X.509 certificate from the MyProxy server and then again immediately generate the proxy credential. It is straightforward to "wrap" the necessary command-line tools so that a LIGO user only "sees" the command `ligo-login` and enters only a password necessary to obtain the credential used to authenticate to the MyProxy server (a Kerberos ticket). The user would not "see" the generation of the key material, retrieval of the signed short-lived X.509 certificate, or the subsequent generation of the proxy credential.

Note that it is not strictly necessary for the end-entity short-lived X.509 certificate from the SLCS and the associated private key to be destroyed if and when a RFC 3820 proxy credential is created from them. The end-entity short-lived credential itself can be used directly with the Globus Grid Security Infrastructure

(GSI) client tools deployed across the LIGO Data Grid. The specific details for how and when the short-lived X.509 credential is used to create a RFC 3820 proxy credential will depend on the use cases and specific user experience that LIGO decides to support.

The SLCS needs to be backed by a CA that can automatically sign certificate requests, provided that the user has successfully authenticated to the SLCS and is authorized to obtain a short-lived certificate. The MyProxy service provides the necessary functionality:

> Starting with v3.0, the MyProxy server includes the ability to act as a Certificate Authority (CA), signing certificates with a configured CA key on request for authenticated users that don't already have certificates stored in the MyProxy repository. Users can run myproxy-logon to authenticate and obtain a certificate from the MyProxy CA when and where needed, without needing to store long-lived keys and certificates in the MyProxy repository or elsewhere.
>
> The MyProxy CA has been developed to meet the requirements of the Short Lived Credential Services X.509 Public Key Certification Authorities Profile of The Americas Grid Policy Management Authority, a member of the International Grid Trust Federation. The NCSA MyProxy CA has been accredited under the Profile.

Because of the LIGO use case requirements the authors see no path forward using a SLCS backed by a CA that is managed or otherwise associates with ESnet and the DOEGrids. We (Jim Basney and Scott Koranda) also explored in email the possibility of leveraging the NCSA short-lived credential service (`C=US, O=National Center for Supercomputing Applications, OU=Certificate Authorities, CN=MyProxy`) [11], but we could see no path forward that would not require NCSA to renegotiate the accreditation of the CA with the TAGPMA and the IGTF.

Our recommendation, given the LIGO use case requirements, is for LIGO to operate SLCS(s) backed with CAs owned and managed by the LIGO project.

5 Technical challenges of the LIGO SLCS approach

The primary technical challenge of the SLCS solution architecture discussed above, aside from the well known challenges of deploying certificate authorities and the related infrastructure, is to develop the client tool `ligo-login` and meet the LIGO use case scenario requirements. The `ligo-login` tool must itself be robust against network failures that prevent it from contacting any particular MyProxy SLCS. If it cannot contact the server configured as its primary server it must seamlessly, without any input from the user, continue to try and contact the other LIGO SLCS servers in its configuration. It must properly handle communication protocol "time outs" and "hangs", with appropriate visual clues for the user, so that the user understands that the tool is continuing to try and contact the next available server in the configured list of servers.

While meeting these requirements does call for careful design and coding, it does not call for special techniques or new research. The `ligo-login` tool can be developed using well understood techniques for developing these types of user tools.

Note that because each SLCS server in the proposed solution architecture issues X.509 certificates from a unique range of serial numbers, there is no need for any type of replication between the servers. MyProxy does have the capability for replication, but when used as a SLCS it is not usually configured for replication since users can simply obtain a new short-lived credential from a secondary server if the primary server is unavailable for some reason.

6 On accreditation of a LIGO SLCS

The IGTF does provide an authentication profile for a Short-Lived Credential Service, and as noted above the NCSA SLCS is accredited by the TAGPMA and meets the minimum requirements dictated by the IGTF for a SLCS. Note that the accredited NCSA SLCS is a service backed by a single CA run from a single geographic location.

To meet the LIGO use case requirement of robustness against network outages LIGO would need to deploy multiple SLCS services across the LIGO Data Grid. Most likely each instance would be backed by a CA with a unique DN, for example

- Caltech: `DN=orgDN=ligoOU=Certificate Authorities, CN=LIGO Caltech 1`
- MIT: `/DN=org/DN=ligo/OU=Certificate Authorities, CN=LIGO MIT 1`

- LIGO Hanford: `/DN=org/DN=ligo/OU=Certificate Authorities, CN=LIGO Hanford 1`
- LIGO Livingston: `/DN=org/DN=ligo/OU=Certificate Authorities, CN=LIGO Livingston 1`
- PSU: `/DN=org/DN=ligo/OU=Certificate Authorities, CN=LIGO Penn State University 1`
- UWM: `/DN=org/DN=ligo/OU=Certificate Authorities, CN=LIGO UW-Milwaukee 1`
- Syracuse: `/DN=org/DN=ligo/OU=Certificate Authorities, CN=LIGO Syracuse 1`
- AEI Hannover: `/DN=org/DN=ligo/OU=Certificate Authorities, CN=LIGO AEI Hannover 1`
- AEI Golm: `/DN=org/DN=ligo/OU=Certificate Authorities, CN=LIGO AEI Golm 1`
- Cardiff: `/DN=org/DN=ligo/OU=Certificate Authorities, CN=LIGO Cardiff 1`

Each of these CAs could sign requests for short-lived certificates in the same namespace, ie.

`/DN=org/DN=ligo/OU=People/CN=Scott Koranda`

The different SLCS and the CAs would be configured to issue certificates into a specified serial number range.

We note that there is *no precedent* for this type of distributed SLCS to obtain accreditation from TAGPMA and no direct authentication profile for this type of distributed network from the IGTF. Our discussions clearly identified the substantial challenge LIGO would face if attempting to obtain accreditation from TAGPMA for such a distributed network of SLCS services.

If LIGO, however, deploys the distributed network of SLCS so that each service and its embedded CA sign requests for short-lived certificates in *different* namespaces then the path to TAGPMA accreditation is easier to see. Each of the distributed LIGO SLCS would be operating individually within the bounds of an authentication profile issued by the IGTF (the same authentication profile that the NCSA MyProxy CA reflects). This configuration would require, however, that authorization services and access control lists such as grid-mapfiles include for each LIGO user the subject name of each certificate that an individual user would obtain from the distributed SCLS since in the LIGO use case users may obtain a credential from any of the distributed services. With each LIGO site shown above deploying a SLCS, every grid-mapfile used within the LDG would need to include 10 mappings for each LIGO user. LIGO will need to investigate and then decide if this model can be accomodated within its infrastructure.

We learned from talking with Michael Helm that there has been discussions within the ESnet staff about attempting a similar but much smaller distributed network in order to provide robustness against network and similar geographically focused failures, but at this time there is no formal plan to proceed and it is clear that obtaining TAGPMA accreditation for such a distributed network of SLCS will require substantial effort.

7 Continued use of DOEGrids CA for LIGO

We note specifically that the deployment of a network of SLCS services within LIGO to enable obtaining credentials for use on the LIGO Data Grid does not prohibit LIGO users from continuing to obtain (in the normal approved way) credentials from the DOEGrids CA for use with the OSG and other regional grids including grids in Europe used by LIGO collaborators from the Virgo project. Until such a time as LIGO is either able to obtain TAGPMA accreditation for its CA deployment or is able to negotiate agreements with OSG and other grids for the credentials issued by LIGO CAs to be used to access resources, it is *critical* that the LIGO VO continue to have access to credentials issued by the DOEGrids CA. We strongly encourage the LIGO, OSG, and ESnet collaboration to continue to evolve so that the DOEGrids CA may continue to provide this essential service to the LIGO user community.

8 Conclusions

We have examined the question of how LIGO can best meet its use case requirements and what if any role the DOEGrids CA, ESnet staff, and the TeraGrid CA (the correct term would be the NCSA SLCS accredited by TAGPMA) can play in a solution architecture. The authors (Jim Basney and Scott Koranda) conclude that given the requirements for the LIGO use case and the current status and relationship of the DOEGrids and NCSA CAs with accrediting bodies, there is no path forward that meets the LIGO requirements and that includes the DOEGrids CA, ESnet staff, or the NCSA SLCS.

We recommend that LIGO research and pursue a path involving multiple SLCS services in some configuration but recognize that obtaining TAGPMA accreditation for such a deployment might be difficult and may not be possible at all, depending on the details of the deployment.

9 Acknowledgements

Scott Koranda gratefully acknowledges the assistance of Jim Basney, Michael Helm, Dhiva Muruganantham, and John Volmer in helping move LIGO forward on this question and their willingness to understand the LIGO use case.

References

[1] grid.ncsa.uiuc.edu/myproxy/

[2] http://www.doegrids.org/Docs/CP-CPS.pdf

[3] http://www.tagpma.org/

[4] http://www.igtf.net/

[5] www.opensciencegrid.org

[6] http://www.doegrids.org/pages/doegridspma.html

[7] http://www.es.net/

[8] https://cert.fusiongrid.org/FusionGrid/FusionGridCAs.html

[9] http://www.tagpma.org/authn_profiles/slcs

[10] http://gridshib.globus.org/

[11] http://security.ncsa.uiuc.edu/CA/