

Trust Model and Credential Handling for Job Forwarding in the SAM-Grid/LCG Interoperability Project

Tummalapalli Sudhamsh Reddy, David Levine, Gabriele Garzoglio, Andrew Baranovski, Parag Mhashilkar

Abstract—At Fermi National Lab, the Dzero experiment has produced and continues to produce petabytes of experimental results, necessitating the use of large computational clusters and grids to assist in the data analysis. SAM-Grid is an integrated data, job and information management system. SAM-Grid is used to partially meet the distributed computing needs of the Dzero experiment at Fermi. Participants of the Dzero experiment may also access resources available via the European LHC Computing Grid and the Open Science Grid infrastructures.

In order to take advantage of these facilities, a project to allow interoperability between these grids was developed, here we present the credential handling process and the trust model in the critical job forwarding component which has been used to forward jobs from SAM-Grid to LCG and OSG. Lastly, we detail the need for and the role played by MyProxy in this scenario.

Index Terms— Grid Computing, Job Forwarding, X509 User Credentials, MyProxy

I. INTRODUCTION

Most current Grid implementations use middleware based on Globus [1] and Condor [2]. The SAM-Grid [3, 4] system is built on Globus and Condor-G [5]. SAM-Grid is a meta-computing infrastructure used by the Dzero [6] experiment at Fermi National Lab [7]. SAM-Grid is used to manage the globally distributed computing and storage resources. It provides distributed data, job and information management services.

The Dzero accelerator experiment has produced and continues to produce petabytes of experimental results, this necessitates utilizing as many computational clusters and grids as we can build and share to analyze this data.

The Dzero Virtual Organization (VO) [8] is a group of

people from around the globe and the sharing of computing and storage resources available to assist in doing analysis of the data that is recorded by the experiment. The SAM-Grid system typically relies on SAM-Grid specific services to be deployed on remote sites to manage computing and storage resources. The deployment of SAM-Grid specific services requires special agreements with each resource provider, and is a labor intensive process. Some members of the Dzero VO have access to significant computing resources in Europe via the Large Hadron Collider (LHC) Computing Grid (LCG) [9]. Therefore, allowing Dzero users access of these resources in Europe via LCG. The SAM-Grid/LCG interoperability project [10] was undertaken to enable Dzero users to access the LCG resources, while retaining some features of the SAM-Grid system which are critical for the experiment, such as data handling and the user-friendliness of the client interface. This bridging of grids is beneficial to both SAM-Grid and LCG since it provides Dzero users access to the LCG resources and provides LCG with representative, real data intensive applications to test their computing infrastructure. As Dzero is slowing down in the future, LCG, OSG and other grids will begin to process Atlas and CMS data as it becomes available.

The interoperability project is based on job forwarding from SAM-Grid to LCG. The users submit jobs to SAM-Grid which then forwards the jobs to the LCG system via “forwarding nodes”, which have, a LCG client interface available. The job then executes on an LCG worker node and the results are pulled back to the SAM-Grid via the “forwarding nodes”. The data handling for the entire process is done by SAM-Grid.

In this paper we discuss the trust model and credential handling in current Grid systems based on Globus and why an extended model is required for the SAM-Grid/LCG system. The rest of the paper is organized as follows: Section 2 contains a view of the related work. Section 3 presents the trust model and credential handling process in the current Grid systems. Section 4 contains the motivation for an extended model in SAM-Grid/LCG system. Section 5 presents the proposed extended model. Section 6 contains the conclusion and future work.

II. RELATED WORK

The issue of Grid security has been of major interest [11, 12]. Globus provides Grid Security Infrastructure (GSI),

Manuscript submitted on March 21, 2006.

T.S.Reddy* is with the SAMGrid Team at Fermi National Accelerator Labs, on leave from the Computer Science Department, University of Texas at Arlington (e-mail: reddy@fnal.gov).

David Levine* is with the Computer Science Department, University of Texas at Arlington (e-mail: levine@cse.uta.edu).

Gabriele Garzoglio, is with the SAMGrid team at Fermi National Accelerator Labs (email: garzoglio@fnal.gov).

Andrew Baranovski, is with the SAMGrid team at Fermi National Accelerator Labs (email: abaranov@fnal.gov).

Parag Mhashilkar is with the SAMGrid team at Fermi National Accelerator Labs (email: parag@fnal.gov).

which is widely used by most of the Grid implementations (such as SAM-Grid, LCG, and Open Science Grid [13]). All communication between systems in Globus is done via GSI enabled interfaces. GSI provides tools and libraries for authentication and authorization using standard X.509 certificates [14, 15], X.509 proxy certificates [16], SSL/TSL protocols [17], and public key infrastructure (PKI). The X.509 certificates are used because they provide flexibility to a resource or service provider so that they can trust another organization's Certificate Authority (CA) without depending on the policies of their own organization.

The X.509 proxy credentials allow the users or services which possess a valid X.509 public key certificate to delegate some or all of its privileges to another entity or service. The entity to which the credentials are delegated can then assume the identity of the certificate holder to accomplish some tasks, such as authentication or to establish a secure communication channel with other entities. The delegation of credentials to another process is done via a secure channel.

All of the technologies discussed thus far are well known technologies which have well-tested open source implementations. The main goals that GSI was designed to meet are dynamic delegation of privileges to other entities or dynamic services, and for repeated authentication across the grid.

Some other alternatives [18] to GSI are Kerberos and infrastructures based on secure shell. While these alternatives are reasonable alternatives, they are not widely used.

MyProxy [19] is open source software which provides an online credential management repository and an online certificate authority. MyProxy allows users to securely store and retrieve credentials from its online server. A user stores his X.509 credentials on the server using the tools provided by the MyProxy client interface and retrieve the credentials when they are needed from anywhere.

MyProxy has been used to delegate credentials to services such as a web portal, and also for trusted servers to renew the proxy of the user, so that long running jobs do not fail because of expired credentials. MyProxy allows users to obtain their credentials over the network without transferring their private keys.

III. TRUST MODEL AND CREDENTIAL MANAGEMENT IN CURRENT GRID SYSTEMS.

In many of the current Grid systems such as SAM-Grid, security is handled via the GSI infrastructure. A user requests for a certificate from a CA, which is trusted by all the resource providers. The CA grants the users with X.509 credentials, which is in a public key encoded in certificate format containing the following details:

- A subject name identifying the user or entity that the certificate represents
- The public key of the subject
- The identity of the CA which has signed the certificate, which certifies that the identity of the subject maps to the public key in the certificate.

- The signature of the CA.

The private key associated with the public key certificate is typically stored on the local system in an encrypted format based on a pass phrase. An alternate scenario of handling the private key is to have the private key stored on a smart card.

Once the user has obtained the certificate he can use it to generate X.509 proxy credentials, which are short term credentials in contrast to the long term credentials obtained from the CA. The proxy credentials that are generated by the GSI libraries have one major difference from the credentials obtained from the CA, that the issuer or signer of the proxy credentials is not a CA but rather credentials obtained from the CA or another proxy certificate.

The short term proxy credentials are generated from the long term credentials by generating a new public-private key pair. The public key is then encoded in an X.509 certificate request format and the user's long term credentials are accessed to sign this certificate request. The proxy certificate and the private key associated with the proxy certificate are then stored in a file on the local system.

When a user wants to access the resources of the Grid, the proxy credentials are created to delegate some of the privileges of the user to an entity or service over the network without the exchange of private keys. This process requires that the network connection be secure to prevent third parties from tampering with the messages. Generally, in the systems based on GSI, the policy of least privilege delegation is followed. Globus does not allow the delegation of fully privileged proxies

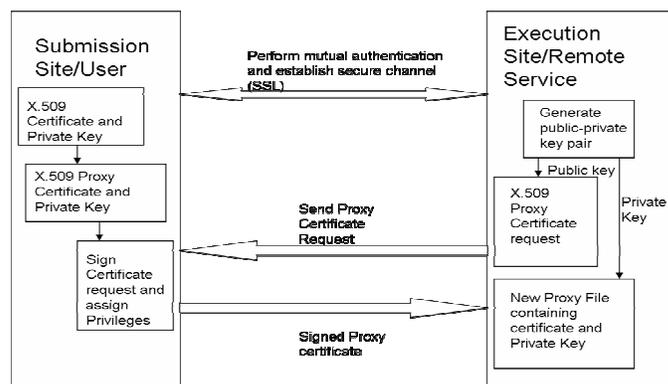


Figure 1 Delegation of proxies over the network.

The steps involved in this process are (shown in figure 1):

1. The user connects to the target site and, both the user and target site perform mutual authentication, the user using his proxy credentials and the target using its public key certificate. After the mutual authentication process, a secure channel is established using the SSL protocol.
2. The user then expresses the objective of delegating some of his privileges to an entity or service at the remote site.
3. The remote site generates a new public-private key

pair, and encodes the public key in a certificate request format. The certificate request is then sent over the network to the user.

4. The private key of the proxy credentials of the user is accessed and is used to sign the certificate request. The user also assigns the rights he wishes to delegate.
5. The new proxy certificate is then sent back over the network to the target site, where it is stored in a file along with the private key associated with the public key encode in the certificate. This new proxy credentials are then available to the entity or service to accomplish some task.

In this model the remote sites can authenticate the user based on the certificate chain present in his proxy certificate. Since the proxy chain ends with the name of the CA that has signed the user's long term credentials, the sites can authenticate a user based on the fact that they trust the CA or not.

IV. MOTIVATION FOR THE EXTENDED MODEL

The model described in the previous section works well for most current Grid systems. SAM-Grid also uses this model. This model fails when one tries to do job forwarding as in the case of the SAM-Grid/LCG system previously described.

SAM-Grid views all execution sites as batch systems. The jobs are given to the execution site head nodes, which in turn submit the jobs to the worker nodes via the local batch systems. Therefore, from the SAM-Grid point of view the "forwarding node" is the execution site head node and LCG is the batch system. The job lands on the "forwarding node", which submits the job to a worker node via the LCG grid system. Since, in the general SAM-Grid model, we do not need a fully privileged proxy to submit a job from the head node of the cluster to the worker nodes, we only delegate a limited privileged proxy certificate over the network.

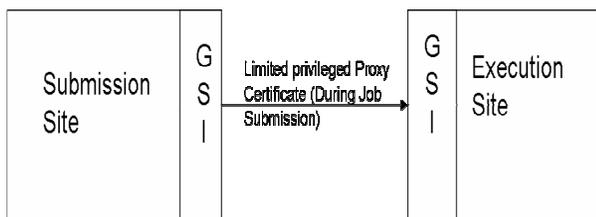


Figure 2 Delegation of Proxy credentials in SAM-Grid

In the general SAM-Grid model (as shown in figure 2), a user delegates some of her privileges over the network to another entity or service. Therefore, in SAM-Grid when a user submits a job, she is first authenticated on the remote site and then limited privileged proxy credentials are delegated to the remote site. The user can then use these proxy credentials to run her job. But, in the SAM-Grid/LCG system the user submits a job to SAM-Grid and his job is sent to the forwarding node. The forwarding node then submits the job to the LCG system using a LCG client interface.

The proxy credentials that are passed to the forwarding node are limited delegated proxies, i.e. the proxy credentials have limited rights. Since the proxy credentials available at the forwarding node have limited rights, they can't be used to submit jobs to LCG.

Hence, an extended model must be used, which utilizes the existing model based on GSI technologies, which are deployed at various sites to solve this above problem. One important constraint of the extended model is that it should not require additional software to be installed on the resources provide by the LCG system.

V. THE MODEL USED FOR SAM-GRID/LCG INTEROPERABILITY PROJECT

GSI does not provide a mechanism to delegate full privileged proxies over the network. Therefore, to overcome this problem, while still maintaining interoperability with the existing GSI systems deployed at the various execution sites, we have developed two models based on the existing GSI systems. One of the models was prototypical, clearly insecure and not appropriate for a production system. We describe both the models and explain the reasoning behind the process.

The first model is based on transferring the fully privileged proxy credentials via a secure channel. The second model is based on using MyProxy software.

In the first model, the user's pass their generated proxy credential file location as an input sandbox parameter in their job description file. During job submission the contents of this location are copied via a secure channel to the execution site. On the execution site, we use the proxy file to submit jobs to the LCG system. This solution was used in the initial test bed system, and was later rejected. The solution was rejected because the proxy file contains both the public and private keys, and it is generally a poor design where the user's private key is exposed to the network.

The second solution, currently implemented in the SAM-Grid/LCG system, is based on using GSI and MyProxy. MyProxy can be used to store the user's credentials, and then retrieve them from anywhere on the Grid. The model is centered on the idea, that the users store their credentials at a MyProxy server, and during job submission fully privileged proxy credentials are obtained from the MyProxy server and are used to submit the jobs to the LCG system.

MyProxy uses a secure channel to transfer the proxy credentials. MyProxy also provides users with the ability to encrypt their delegated proxy credentials while they are stored in the MyProxy server. The user can specify the constraints on the proxy credentials while storing the credentials on the server. If the user wishes to provide a pass phrase, which is to be used while obtaining the credentials from the server, the pass phrase must be passed to the execution site. In SAM-Grid we do not allow users to specify pass phrases as it would be difficult to automate obtaining of credentials for a non-interactive process. Instead, we provide tools built on top of MyProxy client tools which allow users to store their proxies on the server, and they can only be retrieved by showing proxy credentials that belongs to the user.

The model developed for SAM-Grid/LCG system is described below (Figure 3 gives an overview of the process):

1. The user obtains a X.509 certificate from the CA.
2. The user then delegates his X.509 proxy credentials on the MyProxy server.
3. The user submits his job to the SAM-Grid system.
4. The user's job is sent to the forwarding node, along with a limited delegated proxy.
5. Using the user's limited delegated proxy, the service running on the forwarding node obtains a fully privileged proxy from the MyProxy server.
6. The X.509 proxy credentials obtained from the MyProxy server are then used to submit jobs to LCG.
7. On LCG the job lands on an execution site.
8. A limited privileged proxy is delegated to the execution site, using GSI.

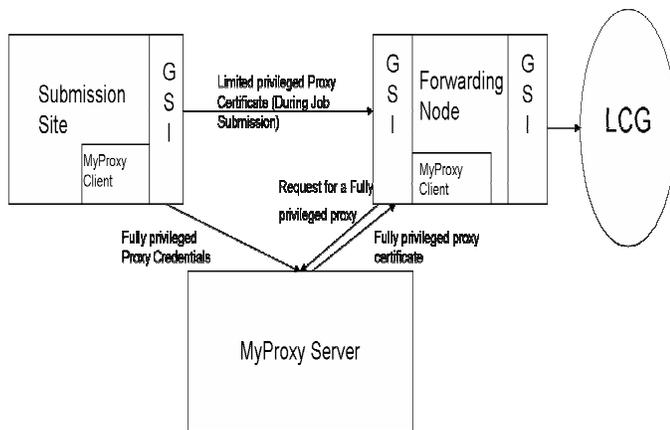


Figure 3 The proposed model for credential handling in job forwarding systems such as the SAM-Grid/LCG system.

Using the system described above, we are able to execute the SAM-Grid job on LCG, using the job forwarding mechanism. The advantages of this model are: user's private keys are never shipped over the network, and this model enables us to take advantage of a feature in LCG, whereby, long running jobs can obtain fresh proxies from the MyProxy server before the proxies associated with the job can expire, which will cause the jobs to fail.

In this model the resource providers must trust the CA that has signed the user's original certificate and also the MyProxy server. The LCG resource providers, who provide the resources to the Dzero VO trust Fermilab to keep the MyProxy server secure. Access to the node running MyProxy server is highly restricted.

Using MyProxy as an intermediate node for storing and retrieving credentials in the SAM-Grid/LCG system imposes some cost on the system. First, the node on which the MyProxy server is running has a minimal set of other applications running on it to prevent malicious persons from entering the node via any other application. Therefore, there is the cost of a dedicated machine.

Average Time to Generate a 512 bit public-private key on the forwarding node	0.052 seconds
Average Time to obtain a proxy on a client within the same network as the server	0.302 seconds
Average Time to obtain a proxy on the Forwarding node (in Europe) from a server in Fermilab	1.351 seconds
Average total time for a job in SAM-grid	10 Hour. (36000 seconds)
Cost of using MyProxy (With respect to time per job)	0.0038%

Figure 4 Parameters regarding the use of MyProxy in SAM-Grid/LCG system.

Secondly (as shown in figure 4), we have observed that the average time to obtain a proxy from the server (which is in Fermilab, USA) to the forwarding node (which is in Wuppertal, Germany) is approximately 1.351 seconds. Most of the cost associated with this process is due to network delay. The average time to obtain a proxy from the server to a client (in Fermilab, USA) was approximately 0.34 seconds. The cost of generating a 512 bit public-private key pair (which is used in proxy generation in Globus) based on the "rsa" algorithm using the SSL libraries was approximately 0.052 seconds on a 1GHz machine.

Since the average time for each job in SAM-Grid is 10 hours, the penalty of 1.351 seconds is negligible.

VI. CONCLUSION AND FUTURE WORK

We presented the traditional model of trust and credential handling, based on the principle of limited delegation of credentials, in Globus, which is used in majority of current Grid implementations. We also discuss the motivation for an extended model. Finally, we presented the extended model based on using MyProxy along with existing GSI libraries that can be used in Grid interoperability projects such as SAM-Grid/LCG project and SAM-Grid/OSG projects.

The proposed model has been implemented and tested for production systems. The new model satisfies the requirements of the Grid security system and has enabled Dzero users to utilize the resources in Europe available via the LCG system and additional resources in the US available via the OSG system.

REFERENCES

- [1] I. Foster, "Globus Toolkit Version 4: Software for Service-Oriented Systems", IFIP International Conference on Network and Parallel Computing, Springer-Verlag LNCS 3779, pp 2-13, 2005.
- [2] Douglas Thain, Todd Tannenbaum, and Miron Livny, "Distributed Computing in Practice: The Condor Experience", Concurrency and Computation: Practice and Experience, Vol. 17, No. 2-4, pages 323-356, February-April, 2005
- [3] G. Garzoglio, "A Globally Distributed System for Job, Data and Information Handling for High-Energy Physics"; Ph.D. Dissertation, DePaul University, Chicago; Dec 05
- [4] G. Garzoglio, I. Terekhov, A. Baranovski, S. Veseli, L. Lueking, P. Mhashikar, V. Murthi, "The SAM-Grid Fabric services", talk at the IX International Workshop on Advanced Computing and Analysis Techniques in Physics Research (ACAT-03), Tsukuba, Japan, Dec 2003

- [5] James Frey, Todd Tannenbaum, Miron Livny, Ian Foster and Steven Tuecke, "Condor-G: A Computation Management Agent for Multi-Institutional Grids", Cluster Computing, Springer publications Netherlands, Volume 5, Number 3, July 2002, Pages: 237 – 246
- [6] The D0 Collaboration, "The D0 Upgrade: The Detector and its Physics", Fermilab Pub-96/357-E.
- [7] FNAL: <http://www.fnal.gov/>, 2006
- [8] I. Foster, C. Kesselman and S. Tuecke, "The Anatomy of the Grid: Enabling Scalable Virtual Organizations", International J. Supercomputer Applications, 15(3), 2001.
- [9] P. Buncic, F. Rademakers, R. Jones, R. Gardner, L.A.T. Bauerdick, L. Silvestris, P. Charpentier, A. Tsaregorodtsev, D. Foster, T.Wenaus, F. Carminati, "LHC Computing Grid Project: Architectural Roadmap Towards Distributed Analysis", CERN-LCG-2003-033, Oct-2003
- [10] G. Garzoglio, A. Baranovski, P. Mhashikar, T. Kurca, F. Villeneuve-Séguier, A. Rajendra, S. Reddy, T. Harenberg, "The SAM-Grid / LCG interoperability system: a bridge between two grids", in Proceedings of Computing in High Energy Physics (CHEP06), Mumbai, India, Feb 2006
- [11] I. Foster, C. Kesselman, G. Tsudik, S. Tuecke, "A Security Architecture for Computational Grids" Proc. 5th ACM Conference on Computer and Communications Security Conference, pp. 83-92, 1998.
- [12] I. Foster, C. Kesselman, G. Tsudik, S. Tuecke, "Managing Security in High-Performance Distributed Computing", Cluster Computing, 1(1):95-107, 1998.
- [13] Open Science Grid, www.opensciencegrid.org/, 2006
- [14] CCITT Recommendation, X.509: The Directory – Authentication Framework. 1988.
- [15] R. Housley, W. Polk, W. Ford, D. Solo, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile". RFC 3280, IETF, April 2002.
- [16] V. Welch, I. Foster, C. Kesselman, O. Mulmo, L. Pearlman, S. Tuecke, J. Gawor, S. Meder, and F. Siebenlist, "X.509 proxy certificates for dynamic delegation", 3rd Annual PKI R&D Workshop, Apr. 2004.
- [17] Dierks, T. and Allen, C., "The TLS Protocol Version 1.0", RFC 2246, IETF, 1999.
- [18] R. Butler, V. Welch, D. Engert, I. Foster, S. Tuecke, J. Volmer, C. Kesselman, "A National-Scale Authentication Infrastructure," Computer, vol. 33, no. 12, pp. 60-66, Dec., 2000.
- [19] 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, August 2001.