

# Bringing Federated Identity to Grid Computing

Jeny Teheran  
jteheran@fnal.gov

Dave Dykstra  
dwd@fnal.gov

Mine Altunay  
maltunay@fnal.gov

Scientific Computing Information Security  
Scientific Computing Division  
Fermi National Accelerator Laboratory  
Fermilab, P.O. Box. 500  
Batavia, IL, 60510-5011

## ABSTRACT

The Fermi National Accelerator Laboratory (FNAL) is facing the challenge of providing scientific data access and grid submission to scientific collaborations that span the globe but are hosted at FNAL. Researchers in these collaborations are currently required to register as FNAL users and obtain FNAL credentials to access grid resources to perform their scientific computations. These requirements burden researchers with managing additional authentication credentials, and put additional load on FNAL for managing user identities. Our design integrates the existing InCommon federated identity infrastructure, CILogon Basic CA, and MyProxy with the FNAL grid submission system to provide secure access for users from diverse experiments and collaborations without requiring each user to have authentication credentials from FNAL. The design automates the handling of certificates, so users do not need to manage them manually. Although the initial implementation is for FNAL's grid submission system, the design and the core of the implementation are general and could be applied to other distributed computing systems.

## CCS Concepts

•Security and privacy→Authentication •Security and privacy→Usability in security and privacy •Computer systems organization→Grid computing

## Keywords

CILogon, InCommon, MyProxy, identity federation, federated authentication, grid computing.

Publication rights licensed to ACM. ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of the United States government. As such, the United States Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.

CISRC '16, April 05 - 07, 2016, Oak Ridge, TN, USA.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-3752-6/16/04...\$15.00.

DOI: <http://dx.doi.org/10.1145/2897795.2897807>

## 1. INTRODUCTION

FNAL<sup>1</sup> (also known as Fermilab) is a laboratory specialized on high-energy particle physics research. Diverse collaborations/experiments produce petabytes of data and use storage and computing elements at distributed computing facilities to achieve peta-scale production for processing of events and simulation data. The distributed computing infrastructure is based on grid computing, where a collection of independent hardware, software, and data resources work together on a common goal<sup>2</sup>. The scientific collaborations making use of the grid resources typically include multiple research institutions with different identity management policies and security infrastructures [1]. A collaboration is translated into the grid world as a Virtual Organization (VO). The concept of VO refers to a defined set of researchers and institutions that work towards a common scientific goal and agree on resource-sharing policies over grid resources [2]. A VO manages its members' access control rights, roles and privileges internally and uses the Virtual Organization Membership Service (VOMS) server to maintain a registry of members' certificates along with any specialized roles and privileges each member is allowed to take on. Fermilab operates and manages the VOMS servers for its hosted VOs.

Access to FNAL grid resources is restricted to authenticated and authorized users. Users must currently follow a registration process to become an authorized user and obtain appropriate credentials to use grid resources. Users must present a valid X.509 certificate to access compute and storage resources. In general, FNAL supports authentication via Kerberos, Single-Sign-On passwords (called "Services" passwords), and X.509 certificates. New sets of authentication credentials are created around the clock to enable scientists to run grid jobs at Fermilab.

In this paper, we present the high level design of a new scheme that integrates federated identity management principles and InCommon<sup>3</sup> infrastructure services with Fermilab grid resources. The new design allows a user to use the identity credentials from his/her home institution to transparently obtain an X.509 certificate to request access to grid resources at FNAL. The design makes use of the CILogon<sup>4</sup> service, which is a member of the InCommon federation. It enables on-demand generation of X.509 certificates

<sup>1</sup> <http://www.fnal.gov/>

<sup>2</sup> <http://www.oracle.com/technetwork/topics/grid/whatsnew/gridtechwhitepaper-0305-130088.pdf>

<sup>3</sup> <http://incommonfederation.org>

<sup>4</sup> <http://www.cilogon.org>

based on authentication from select Identity Providers (IdPs) operated by InCommon-member institutions and organizations. The grid submission at FNAL, like most grid interfaces, is based on the Linux command line and scripting, so the design makes use of the Enhanced Client or Proxy (ECP)<sup>5</sup> capability of IdPs. Proxies of the certificates are also stored in MyProxy<sup>6</sup> for later use by the grid submission system on behalf of the user.

Authentication based on federated identities brings benefits to FNAL and its scientific user communities. Federated identity management improves the user experience by reducing the number of account registrations and access credentials to be obtained and maintained by the end user. It also eliminates the need for the hosting institution to manage registration processes and maintain identities for its scientists. Furthermore, our new design fully automates and streamlines the credential management workflow, removing the need for manual intervention from the user. The design and the core of the implementation could also be used at other institutions that are part of the InCommon federation with other command-line based grid submission systems. For any access control system, consideration of security risks vs benefits to the users and operational environment is crucial. Therefore, a security analysis of the new design and a discussion of risks are also presented in the paper. We discuss the various layers of security controls put in place and how they ensure that the new design does not increase our risk profile.

## 2. BACKGROUND

Before presenting the new design for federated access at FNAL, we provide background information about the existing authentication scheme, as well as grid submission systems in general, and the grid submission system used at Fermilab. We also provide information about InCommon federation, CILogon, Identity Providers with ECP, and MyProxy.

### 2.1 Current FNAL authentication scheme

At Fermilab, the computing environment is divided into two sections: general computing environment (GCE) and open science environment (OSE). GCE hosts all interactive resources and other services expected from the laboratory for its employees. OSE, on the other hand, only includes the grid resources that are open to Fermilab scientific collaborators, both computing and storage resources. In OSE, there is no way to interactively access resources (i.e. no interactive shell access); the only access mode allowed is through batch job submissions and data requests. Both environments are segregated from one another through multiple layers of security controls, and it is not allowed to directly access GCE from the OSE. In other words, once a user is allowed to use grid resources, the user is confined to a "sandbox" that he/she cannot escape to access GCE. To request access to GCE resources, a user must perform Kerberos authentication. Login to interactive Linux nodes, only contained in GCE, is allowed through Kerberized SSH or by multi-factor authentication. In order to access OSE grid resources, users must provide a valid X.509 certificate. The grid resources only accept certificates issued by International Grid Trust Federation (IGTF)<sup>7</sup> accredited Certificate Authorities (CA). For convenience to its users, Fermilab operates such a Certificate Authority, Fermilab Kerberos Certificate Authority (KCA) accredited by the TAGPMA<sup>8</sup> branch of IGTF.

<sup>5</sup> <http://www.cilogon.org/ecp>

<sup>6</sup> <http://grid.ncsa.illinois.edu/myproxy>

<sup>7</sup> <https://www.igtf.net>

The users cannot use their Kerberos tickets to directly access grid resources, since OSE requires X.509 certificates. Although a subset of users at FNAL utilize certificates issued by other accredited CAs to access the grid, this is an arduous process because the primary job submission system hosted at FNAL (described below) completely automates the use of KCA certificates but requires several manual steps for other kinds of certificates. Furthermore, some data access services are programmed to only accept KCA certificates. Access to additional services, such as email, document repositories, corporate tools and workstations is done via a separate set of password-only credentials in a domain named "Services".

A Fermilab user obtains Kerberos credentials by running `kinit` once a week and entering his/her Kerberos username and password. The user then can obtain a Fermilab KCA-issued X.509 certificate by utilizing his/her Kerberos ticket through the `kx509`<sup>9</sup> tool. The KCA certificates are valid to access web services and grid resources [3]. The issued certificate is tied to the Kerberos principal and is stored in `/tmp/x509*`. A KCA certificate's lifetime is equal to that of its original Kerberos ticket, which defaults to a week. A user who has submitted jobs and wants to submit again after the Kerberos ticket has expired has to go through the same steps to get a new KCA certificate.

The KCA server software has been around a long time and is losing its software support this year, so that has added urgency to finding a new more modern solution. In addition, it takes a significant amount of resources to maintain the accreditation for any Certificate Authority. Bringing user-friendly access control mechanisms to grid resources that are not solely dependent on KCA certificates is one of the goals of the design presented in this paper.

### 2.2 VOMS: Virtual Organization Membership Service

VOMS is a user management system that maintains a registry of all VO members, their roles in the VO and the access privileges corresponding to those roles. Every user in the VO has his/her certificate Distinguished Name (DN) registered in VOMS, and VOMS maps those to a set of capabilities defined through attributes: groups, roles and generic attributes. If a user has multiple certificates, she/he can register multiple Certificate DNs under his/her account. When a user wants to submit jobs at FNAL, a VOMS X.509 certificate proxy<sup>10</sup> must be provided. The VOMS proxy contains the trusted attributes for the user and these will be used to authorize actions on behalf of the user [4]. A VOMS proxy is obtained by using the command line tool `voms-proxy-init`. The user authenticates with the VOMS server through a valid X.509 proxy of the original certificate, known as a "grid proxy". The VOMS server validates the incoming grid proxy and returns a VOMS proxy that confirms that the user is in the VO and contains the authorized roles of the user.

The grid proxy and VOMS proxy implement the standard Grid Security Infrastructure (GSI) [2] delegation capability. The grid proxy contains the user's identity with a new set of keys and it is signed with the user's original certificate public key. It has a much shorter lifetime (12 hours to 1 day) than a certificate (up to 13 months). Grid proxies enhance user security in the event of a security compromise because they expire quickly. A VOMS proxy is generated much the same way as a grid proxy: it has a different

<sup>8</sup> <http://tagpma.org>

<sup>9</sup> <http://www.umich.edu/~x509>

<sup>10</sup> <http://www.ietf.org/rfc/rfc3820.txt>

set of keys and is signed by the user's original certificate. An important difference is that a VOMS proxy holds an additional set of VO-granted attributes for the user that specify the user's membership in the VO and the user's roles and privileges in the VO.

Fermilab operates and manages VOMS servers for its hosted VOs. The VO members internally determine the access policies for the VO scientists. Fermilab staff makes sure that these policies are implemented and enforced by the VOMS server. As we will discuss later, in the event of an incident or malfunction, Fermilab can temporarily override VO management decisions or ask the VO managers to change their policies.

### 2.3 Pilot systems and grid submission systems

Making effective use of a wide diversity of grid resources is complicated and generally requires both a pilot-based workflow management system (WMS) such as GlideinWMS<sup>11</sup> and a grid submission system that submits grid jobs to the pilot system. Large VOs generally have their own pilot and grid submission systems, but smaller VOs often share such a system. The Scientific Computing Division at Fermilab runs a GlideinWMS pilot system and a grid submission system called Jobsub that is shared by a number of VOs hosted at FNAL including NOvA, Mu2e, and MicroBooNE.

A pilot-based WMS is a pull-based WMS that creates an overlay pool of compute resources on top of the grid. GlideinWMS is a pilot-based WMS that creates on-demand a dynamically sized overlay HTCondor batch system [5] [6] on grid resources to address the complex needs of VOs running scientific workflows.

Grid submission systems provide users with interfaces for submitting and managing jobs, including being responsible for renewing short-lived VOMS proxies that are used to authorize the jobs. Grid jobs can be queued for days or weeks and can run for many hours, but VOMS proxies are kept to short lifetimes (12 to 24 hours) for two reasons: 1) keeping them short-lived limits their value to attackers and 2) if a user is removed from a VO, the renewal will fail and user's jobs will be stopped. For these reasons, grid submission systems have to periodically renew the proxies.

### 2.4 Jobsub: FNAL grid submission system

Jobsub is an integrated grid submission system developed at Fermilab for simplifying job submission for users, providing authentication, job scheduling and output retrieval [7]. It uses GlideinWMS as its pilot system. Jobsub is a client-server system. Jobsub Client is the interface for users to submit jobs. Through a command called `jobsub_submit`, the user specifies details about the executable file, input files and grid protocols to perform file transfers as well as the VO to submit to. `jobsub_submit` uses X.509 credentials to authenticate the user to the Jobsub Server over a secure https connection. The Jobsub Client can be run anywhere, not just at Fermilab. If the user does not possess a valid X.509 certificate or proxy but does possess a Fermilab Kerberos ticket, as an extra convenience `jobsub_submit` invokes `kx509` to retrieve a KCA certificate based on the Kerberos ticket. If the authentication is successful, Jobsub Server will submit the user's jobs to the grid.

The user certificate or proxy itself is not transferred to the Jobsub Server, and even if it were, scientific event processing and simulation usually include jobs with lifetimes longer than the validity period of the certificate or proxy, making proxy renewal a

necessity. Currently, Jobsub Server takes care of this by maintaining an extra "Robot" Kerberos keytab for every FNAL user. A Robot Kerberos keytab is a credentials file that is stored on the Jobsub Server and is used to automatically generate KCA certificates and proxies on behalf of the user. Whenever one of those proxies nears the end of its lifetime, the Jobsub Server generates a new Robot certificate proxy to prevent jobs from expiring. The keytab files are only owned by the Jobsub server and a user cannot access them. The generated Robot certificate has a different key pair and a different certificate DN than a user's personal certificate. As a result, both certificates need to be registered in the VOMS server under the user's account so that both have access to the user's roles and privileges.

**Error! Reference source not found.** shows how the user interacts with Jobsub Client and the steps that the Jobsub Server takes:

1. User submits a job to the Jobsub Client using the KCA certificate previously obtained.
2. After the user submits a job, the Jobsub Server uses the Robot Kerberos keytab to get a KCA certificate with `kx509`.
3. The Jobsub Server then executes the command `voms-proxy-init` which contacts VOMS to obtain a VOMS proxy based on the Robot KCA certificate, and it uses this proxy to submit jobs.
4. The Jobsub Server uses the same process to periodically renew the VOMS proxy to prevent it from expiring before the jobs are finished.

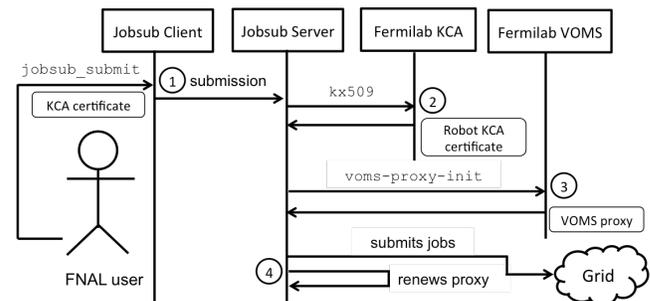


Figure 1. Job submission using Jobsub

Jobsub also supports users who do not have an FNAL Kerberos account and so cannot obtain KCA certificates, but it is a relatively difficult process. First, the user must manually register their certificate DN with the VOMS server, which is automatically done for KCA certificate holders. Second, the user must obtain a secondary service certificate, also register its DN in VOMS, and ask the Jobsub server operations team to manually install it; all that is also automatically done for KCA certificate holders. The service certificates then are used by Jobsub server similarly to the Robot KCA certificates to obtain and renew VOMS proxies. Third, some of Fermilab's data access services are not currently programmed to recognize non-KCA certificates, so users cannot utilize them without their service certificate DN also being registered separately in those systems.

<sup>11</sup><http://www.uscms.org/SoftwareComputing/Grid/WMS/glideinWMS>

## 2.5 GUMS: Grid User Management Service

GUMS is a grid identity mapping service, which translates VO-granted attributes (VO-specific roles and privileges) into UNIX-based user identities (UNIX usernames) and file access privileges. GUMS has a mapping logic that associates certificates and VO attributes with UNIX user accounts. It provides better control and security for access to the site's grid resources<sup>12</sup>. At FNAL, when a job submission is accepted, Jobsub Server sends a mapping request to the Fermilab GUMS Authorization server. The request includes identity information of the user: the DN (Distinguished Name) and the FQANs (Fully Qualified Attribute Name). Using the DN and the FQAN, GUMS server checks if the user is allowed to utilize the grid resources and returns a valid UNIX username, under which the user's job will be run. If the user or the VO which the user belongs to has been banned, then the GUMS server will return null and the job will not be executed. Fermilab GUMS communicates periodically (every 6 hours by default) with the VOMS server to keep updated information about the VOs and users [8].

## 2.6 InCommon Federation

The InCommon Federation, operated by Internet2, is the identity federation for education and research institutions in the U.S. InCommon enables Identity Providers to manage and share user identities with Service Providers in a secure and trusted framework. InCommon's identity management federation serves 8 million end-users (IPEDS data; October 2014)<sup>13</sup>. Through InCommon, users can access resources at various service providers by using a single set of authentication credentials from their home organizations. InCommon uses OASIS SAML-based authentication and authorization systems (such as Shibboleth) to enable scalable, trusted collaborations among its community of participants<sup>14</sup>.

## 2.7 CILogon

CILogon is an X.509 Certificate Authority relying on the InCommon federation to authenticate users to provide certificates. Based on SAML, CILogon issues X.509 certificates for users who are authenticated against their Identity Providers [9]. In other words, CILogon links user identities released by their home institutions' IdPs with X.509 certificates. The certificate issuance process is fully automated and a user can obtain a certificate within a few seconds. CILogon is also accredited by the TAGPMA branch of the IGTF, the same body that accredited the Fermilab KCA. CILogon is a unique CA in that its identity vetting process is based on the SAML assertions given by InCommon IdPs. Other accredited CAs either manually vet a user's identity (2-3 days process) or can only automate the process within a single administrative domain such as Fermi KCA.

CILogon operates multiple Certification Authorities with different levels of assurances<sup>15</sup>. We only use CILogon Basic CA in our design. In this paper, whenever we use the term "CILogon", we refer to CILogon Basic CA.

## 2.8 Identity Providers with ECP

Federated identity infrastructures were first built to be convenient for web browser-based applications, but SAML v.2.0 has an extended capability called "Enhanced Client or Proxy" (ECP) which allows for the exchange of identity information outside of

the context of a web browser. It is very useful for command line interfaces such as grid submission systems. Many Identity Providers in the InCommon federation do not yet have ECP enabled, but its use is growing, and Fermilab has configured a Shibboleth IdP with ECP. The FNAL IdP supports both Services domain authentication with username/password and Kerberos authentication. The FNAL IdP is currently connected to the CILogon service so that users registered at FNAL can authenticate against the FNAL IdP and get an X.509 certificate issued by CILogon instantly.

## 2.9 MyProxy

MyProxy is an online credential repository service, where users can store their credentials and from where the credentials can be securely retrieved for later use. MyProxy offers access to stored credentials via various authentication methods in a web browser interface or through a command-line client. Users can define access control policies for later access. It can safely store user credentials for grid submission systems which need to renew user credentials for long-running jobs [10].

## 3. FEDERATED IDENTITY FOR THE GRID

The implementation of federated identity authentication at FNAL will enable users from diverse collaborations in the U.S. to submit jobs to the Fermilab grid resources without registering and obtaining a FNAL identity. Using CILogon and the InCommon federation, users will be able to authenticate against their home institutions' Identity Providers and obtain a valid X.509 certificate to request access to grid resources. Currently, there are very few IdPs with ECP support and our system only allows FNAL IdP to provide the CILogon certificates. However, our technology allows adding more IdPs in the future as they provide ECP support and also meet Fermilab's security requirements that are discussed in Section 5.1 below.

In this section, we present the details of our solution which consists of using CILogon as a Certificate Authority, leveraging the FNAL IdP ECP interface to automate the certificate request and utilizing MyProxy capabilities for proxy renewal.

### 3.1 Authenticating the user

Below are the first steps, for authenticating the user, as illustrated in Figure 2:

1. When a user wants to submit jobs, he or she runs `jobsub_submit`. If there is no valid X.509 certificate or proxy in `/tmp/x509*`, `jobsub_submit` gives an error and tells the user to run `cigetcert`<sup>16</sup> and try again.
2. The user runs `cigetcert` which is a new command that establishes communication with CILogon to obtain a certificate and requests authentication through an IdP with ECP on the user's behalf. `cigetcert` has no specific knowledge about the FNAL job submission system, it is general for all InCommon users, so `jobsub_submit` tells the user to run `cigetcert` with an option that specifies the name of the Jobsub

<sup>12</sup> <https://www.racf.bnl.gov/Facility/GUMS/1.3/index.html>

<sup>13</sup> <https://www.incommon.org>

<sup>14</sup> <http://www.internet2.edu/products-services/trust-identity-middleware/incommon-federation/>

<sup>15</sup> <http://ca.cilogon.org/loa>

<sup>16</sup> <http://redmine.fnal.gov/projects/fermitools/wiki/cigetcert>

server. The Jobsub server hosts a file over https with a well-known name (cigetcerto.txt) that contains FNAL-specific parameters, including the institution name for the FNAL IdP, the name of the MyProxy server and the lifetime of the proxy to store there. The proxy lifetime is multiple weeks, as explained further below. `cigetcert` maps the institution name to a URL for the IdP by reading a mapping file from CILogon. (Reading the `cigetcert` options file and institution mapping file are not shown in the figure since they are implementation details.)

3. CILogon then returns a SAML Authentication request, which `cigetcert` passes on to the IdP.

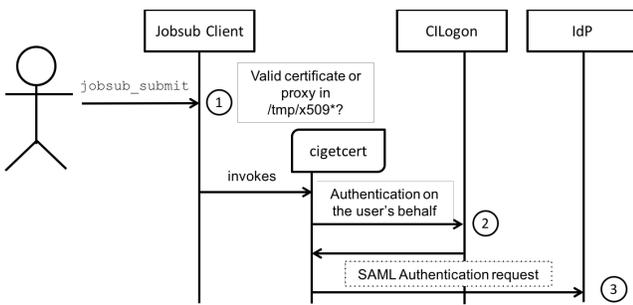


Figure 2. Authenticating the user

### 3.2 Obtaining an X.509 certificate

The steps to obtain an X.509 certificate are shown in Figure 3.

4. The IdP requests Basic authentication.
5. `cigetcert` prompts the user for username and password.
6. If the authentication is successful, the IdP returns a SAML Assertion to `cigetcert`.
7. `cigetcert` forwards the SAML assertion to CILogon to complete the certificate request.
8. CILogon issues a certificate for the user and returns it to `cigetcert`.

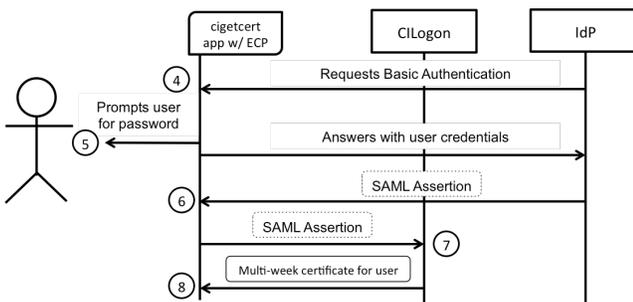


Figure 3. Obtaining the certificate

An IdP can authenticate its users in a number of ways including username/passwords, Kerberos tokens, one-time-challenges and so on. Some IdPs can offer multiple authentication methods simultaneously as an optimization. The FNAL IdP and

`cigetcert` support using Kerberos credentials in addition to the Services username/password authentication. `jobsub_submit` looks for Kerberos credentials when there is no valid certificate, and if they exist, it automatically invokes `cigetcert` rather than asking the user to do it. Since a Fermilab user with a Kerberos account will most likely have his/her Kerberos ticket ready on his/her computer, such a user does not need to provide any authentication credential, and it is as convenient to them as the old Jobsub system. Other users who are using a different IdP or don't have a Kerberos account must use an alternative authentication method such as username/password authentication.

### 3.3 Storing proxies

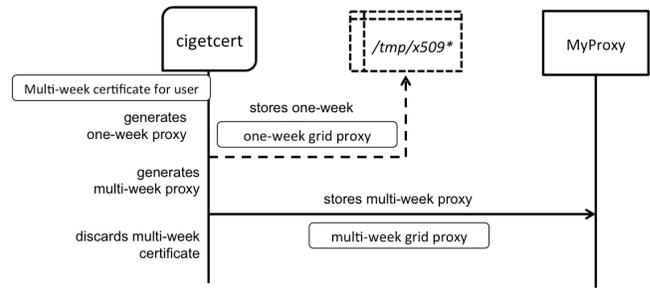


Figure 4. Storing proxies

At this point, `cigetcert` holds a multi-week certificate for the user. It generates a proxy and stores it on the MyProxy server. It then generates a 1-week long proxy that it stores on the local disks and discards the long-lived certificate. Unlike the former authentication scheme, where KCA only issued short-lived X.509 certificates and had to have “Robot” certificates, the MyProxy server takes care of the long-running jobs by constantly allowing proxy renewals. The IGTF rules allow long-lived user regular certificates only if they are stored encrypted on a disk, so our system stores a long-lived X.509 proxy encrypted in MyProxy and only stores a short-lived proxy on the local disk. This enhances the security of our system since user client machines may be more likely to be compromised than the secured MyProxy repository. As a result, a user has to renew his/her certificate each week as the certificate expires on local disk since the user cannot submit new jobs without a valid certificate. However, once a job is submitted, even if it runs longer than the lifetime of the user's proxy on the local disk, Jobsub Server will continue to renew the proxy from the MyProxy server and continue to manage the job. The steps are shown in Figure 4.

9. Using the multi-week certificate, `cigetcert` generates a one-week grid proxy and stores it in `/tmp/x509*`.
10. `cigetcert` generates a multi-week grid proxy and stores it in the MyProxy server at Fermilab.
11. In compliance with IGTF requirements, `cigetcert` discards the long-lived certificate.

### 3.4 Submitting jobs and renewing proxies

The final steps of submitting jobs and renewing proxies are shown in Figure 5:

12. For the next week, the user can submit jobs with `jobsub_submit` using the grid proxy stored in `/tmp/x509*`.

13. The Jobsub Server retrieves a new short-lived proxy from MyProxy, runs `voms-proxy-init` to have the VOMS server turn it into a VOMS proxy, and submits that with the job. Periodically the Jobsub server repeats this step during the lifetime of the job. It is important to note that the user never has access to the VOMS proxies because they are only stored at the Jobsub Server.
14. After the one-week proxy has expired on local disk, the user must start from step 1 to obtain a new certificate to submit new jobs.

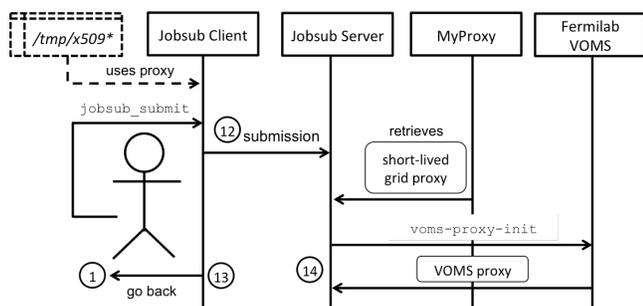


Figure 5. Submitting jobs and renewing proxies

### 3.5 Automating the user registration workflow in the FNAL system

For users who utilize the FNAL IdP, their registry in the VOMS server is automatically handled. Since Fermilab operates the IdP which controls the attributes released to CILogon, a script easily computes the certificate DNs that CILogon will assign to the users' certificates. This allows registering these certificate DNs in the VOMS server without requiring any manual intervention by the user. This is very much like the old method: with the Fermilab Kerberos CA, a script computed each user's certificate DN and automatically registered them in the VOMS server, in addition to the Robot certificate DNs.

For users that authenticate with non-FNAL IdP in the InCommon federation, registering in VOMS still has to be a manual step because their certificate DNs cannot be computed. The user must apply to the VOMS server with his/her certificate so that the VOMS server can register their certificate in the database. However, we think that this manual step can further be improved: a user can delegate authority to a separate web service that can contact CILogon to get a certificate for the user and register the DN of the certificate in VOMS on behalf of the user. OAuth protocol allows implementing such a delegation service<sup>17</sup>.

## 4. RELATED WORK

Multiple projects around the globe are supporting federated access through CILogon and enabling secure access to computing resources and scientific data with existing user credentials. The projects presented have similar implementations with CILogon and

ECP to issue certificates for enabling collaboration and resource sharing around the globe.

### 4.1 LIGO

LIGO<sup>18</sup> (Laser Interferometer Gravitational Wave Observatory) is a physics experiment aiming to detect gravitational waves. Like FNAL, a user had to register at LIGO before accessing resources such as general computing and data analysis facilities. A LIGO identity consists of a username and password and is linked with a Kerberos principal. However, with the increasing need of collaboration with researchers throughout the world, provisioning and managing identities for each user became a challenging task that meant more funding and additional administrative load [11]. The identity federation has enabled LIGO users to obtain certificates from CILogon. In this way, users stick with their home institution's authentication credentials to access LIGO resources. The LIGO Identity Provider which supports ECP allows LIGO users to utilize other research projects resources, like the European Virgo project. Instead of getting an additional identity at Virgo project, LIGO users access Virgo resources with their LIGO identity and vice versa.

The tool that LIGO uses to manage certificates<sup>19</sup> is LIGO-specific and is not intended to be general purpose like `cigetcert` is, and it does not automatically store proxy certificates in MyProxy.

### 4.2 LTERN AND DataONE

LTERN<sup>20</sup> (Long Term Ecological Research Network) and DataONE<sup>21</sup> (Data Observation Network for Earth) are an example of mutual collaboration in enabling secure access to multiple resources for users in both projects. Due to the importance of DataONE, many LTERN users are required to access the distributed resources offered through the DataONE science network. In this case, the LTERN IdP allows LTERN users to access resources with their existing credentials. The LTERN IdP also supports ECP for easy command line access.

## 5. SECURITY CONSIDERATIONS

Throughout the design of this federated grid access system, security was always a key factor. In the next section, we discuss the important aspects we took into consideration during the design.

### 5.1 Authorization Control

CILogon Basic CA accepts identities from multiple Identity Providers registered in the InCommon Trust Federation. The identities released by InCommon members are trustworthy in the sense that they are safeguarded and vouched by respectable research and higher education institutions. We explain below in Section 5.2 the InCommon membership criteria and the operating principles that each member contractually agrees to abide by. However, the second layer of Fermilab security is determined and fully managed by Fermilab's own authorization system. A certificate, or any other authentication credential, is simply not valuable unless it is authorized to access resources. Fermilab controls the authorization decision through the VOMS and GUMS servers. Unless a CILogon certificate is registered in an authorized and valid VOMS server, it cannot gain access to Fermilab grid resources. In addition to the VOMS server, without GUMS server's

<sup>17</sup> <http://www.cilogon.org/portal-delegation>

<sup>18</sup> <https://ligo.caltech.edu/>

<sup>19</sup> <https://github.com/skymoo/macldg/blob/master/source/scripts/li-go-proxy-init>

<sup>20</sup> <http://www.tern.org.au/long-term-ecological-research-network-ltern-pg17872.html>

<sup>21</sup> <https://www.dataone.org/what-dataone>

mapping services, no users can access grid resources because VO specific roles and privileges are meaningless in a production environment. The GUMS servers are solely controlled by the Fermilab operators and in the event of an incident or malfunction, they are fully capable of turning off access for a specific user or an entire VO. Fermilab operators are trained for banning users through periodic training drills and have demonstrated that they are capable of taking this action. This protection allows Fermilab to safely open itself to accept authentication credentials from trusted institutions while fully controlling which one of these credentials can access its resources and under which conditions.

If an authorized certificate is compromised, the certificate will be revoked by CILogon Basic CA and can instantly be removed from the Fermilab VOMS and GUMS servers. Since Fermilab solely controls the authorization servers, it can decide when and which credentials to remove without having to wait for an external third party's decision. As a result, the security of FNAL grid resources are not degraded: we are relying on an external service to issue valid credentials; however, authorization to grid resources is determined and controlled by Fermilab.

## 5.2 Trust relationships in identity federations and risk management

The implementation of federated access includes trusting the identity providers registered with CILogon to implement appropriate registration, vetting and identity management processes. Membership in InCommon is exclusive and a participant must prove its eligibility in addition to accepting InCommon participation agreement and disclosing its own operational practices. The membership criteria defines the following membership categories: two- and four-year, degree-granting academic institutions that are accredited by a U.S. Department of Education Regional Institutional Accrediting Agency, or some national or state accrediting agencies; research organizations related to a particular federal research agency and listed on an official publicly available government listing; sponsored partners that are business, education, and research organizations sponsored by the designated Executive of a current InCommon Higher Education Institution or Research Organization; and international organizations of higher education and any international organizations legally established outside the U.S. Once an institution satisfies the membership criteria, it agrees to the InCommon Participation Agreement and also publishes its operational practices among various other documents. As a result, InCommon members trust identity information released by their peers knowing that they are coming from respectable and vetted research and higher education institutions.

Although InCommon aims to build a trust framework among its members, it certainly does not forbid its members from requiring more stringent requirements from one another should their specific security models need it. A member can choose to form closer ties with a select few InCommon members and accept credentials from this small group only. Currently, Fermilab grid systems are configured to only trust credentials released by Fermilab's own Identity Provider. This is partially because a big chunk of our user base have ties with Fermilab and can use the Fermilab IdP. In the future as our scientific collaborations grow and new scientists without Fermilab accounts join, we will consider adding on new Identity Providers in our ecosystem. As we accept these new Identity Providers into our ecosystem, we will closely examine their operational security practices such as how they operate their Identity Provider systems, how they monitor for incidents and how quickly they can detect credential compromise and notify us. A

thorough discussion of risks in trust federations is presented in [12] [13] and will be included in our evaluation of new IdPs. Our new model shifts some of the security responsibility beyond Fermilab's control, but the user's home institution has a closer relationship with the user and should be able to detect and respond in the event of a compromise of a user's authentication credentials.

Furthermore, credential compromise is a risk that can very well happen to Fermilab's own staff members. Our staff laptops can get compromised, key-loggers may be installed and their authentication credentials can be compromised. The security controls against this risk are: to give the least amount of privileges necessary for each user; segregate and protect sensitive resources by giving access to a minimum number of trusted users through strong access credentials; and monitor for suspicious behavior. We implement all of these controls in our grid environment. First, each grid user only has a very restricted level of access and cannot escape their sandbox in the grid. Second, there are no sensitive or protected resources in the grid environment. Users are only given idle computing cycles and storage capabilities. Moreover, the scientific data produced by experiments are not sensitive and open to public access. The biggest concerns with the grid are (1) accidental or malicious overwrite of the scientific results, which will force experiments to redo their computations, and (2) wasting of compute cycles by a malicious user. For these reasons, we monitor our users' behavior on the grid. When a user harvests compute power noticeably more than his/her peers, this gets recorded in our accounting systems and raises a red flag.

## 5.3 MyProxy security

MyProxy server is a key component in this design: it acts as a credential repository and holds multi-week proxies for every user submitting jobs at FNAL. Accordingly, there are security requirements for MyProxy server installation and also for communication with other components in the system.

The MyProxy application has been audited and shown to be secure. Regarding installation, MyProxy needs to be run on a secured host with limited services and with well-defined rules for user access. The host where MyProxy is deployed needs to be carefully monitored and maintained (e.g. kept up to date with security patches).

`cigetcert` and Jobsub Server act as MyProxy clients; they are the only two components establishing communication with the MyProxy server. The confidentiality of the communication between these components is also essential. MyProxy client and server establish an encrypted channel before passing data. Additionally, mutual authentication is required by MyProxy client; that is, the server is required to present credentials to prove its identity. `cigetcert` uses the user's own certificate to authenticate with the MyProxy server. The Jobsub Server and MyProxy have their own X.509 certificates which they use for authentication. As a result of the use of access control and mutual authentication in the communication between the MyProxy client and server, the risk of an attacker stealing credentials by acting as a fake client or server is minimized.

Jobsub Server is the only entity authorized by MyProxy server to retrieve any user's proxy. No other entities have this kind of authorization. Jobsub Server is implemented as a service distributed into multiple machines for scalability and load balancing. These machines run in a secure environment with specific rules for user access. Also, it includes monitoring services and multiple pluggable security modules that interact with banning tools (e.g. GUMS) [14].

## 6. CONCLUSIONS

There are a number of benefits to this design compared to the previous system at FNAL.

1. Improved end-user experience: federated identity management avoids the need for a user to have a login account at the institution hosting the grid submission system. This implementation improves the end-user experience, by eliminating the need for extra account registration and management of an extra set of credentials.
2. Certificate-free environment: Throughout the whole process, the user does not need to manage a certificate; the certificates are all handled behind the scenes. From user's perspective, certificate request processes and public/private key management are burdensome. Within this new scheme, the user only needs to type a password once a week, or if FNAL Kerberos credentials are available, then not even that is necessary.
3. Administrative load diminished: FNAL no longer needs to maintain their own Certificate Authority and manage new identities for users linked to a single experiment or collaboration. Enabling access for new users is easier and faster. FNAL is now ready to host more experiments and build more scientific collaboration.
4. Automatic proxy renewal: The capability of automated proxy renewal through MyProxy for long-running jobs eases the burden on Jobsub Server, by eliminating the need of maintaining a robot Kerberos keytab or service certificate for every user.

In addition, the `cigetcert` tool which is a crucial part of the implementation is simple and can be used by other institutions.

## 7. ACKNOWLEDGEMENTS

Fermilab is operated by Fermi Research Alliance LLC under Contract No. DE-AC02-07CH11359 with the United States Department of Energy.

## 8. REFERENCES

- [1] D. Broeder, B. Jones, D. Kelsey, P. Kershaw, S. Lüders, A. Lyall, T. Nyrönen, R. Wartel and H. J. Weyer, "Federated Identity Management for Research Collaborations," CERN, Geneva.
- [2] I. Foster, C. Kesselman and S. Tuecke, "The Anatomy of the Grid: Enabling Scalable Virtual Organizations," *International Journal High Performance Computing Applications*, vol. 15, no. 3, pp. 200-222, 2001.
- [3] O. Kornievskaja, P. Honeyman, B. Doster and K. Coffman, "Kerberized Credential Translation: A Solution to Web Access Control," in *USENIX Security Symposium*, Washington D.C., 2001.
- [4] R. Alfieri, R. Cecchini, V. Ciaschini, L. dellaAgnelo, A. Frohner, K. Larentey and F. Spantaro, "From gridmap-file to VOMS: managing authorization in a Grid environment," *Future Generation Computer Systems*, vol. 21, no. 4, pp. 549-558, 2005.
- [5] I. Sfiligoi, D. Bradley, B. Holzman, P. Mhashilkar, S. Padhi and F. Wurthwein, "The Pilot Way to Grid Resources Using glideinWMS," *Computer Science and Information Engineering*, vol. 2, pp. 428-432, March 2009.
- [6] P. Mhashilkar, A. Tiradani, B. Holzman, K. Larson, I. Sfiligoi and M. Rynge, "Cloud bursting with GlideinWMS: Means to satisfy ever increasing computing needs for scientific workflows," *Journal of Physics: Conference series*, vol. 513, no. 3, pp. 32-69, 2014.
- [7] D. Box, "FIFE-Jobsub: a grid submission system for intensity frontier experiments at Fermilab," *Journal of Physics: Conference Series*, vol. 513, no. 3, p. 32, 2014.
- [8] K. Chadwick, E. Berman, P. Canal, T. Hesselroth, G. Garzoglio, T. Levshina, V. Sergeev, I. Sfiligoi, N. Sharma, S. Timm and D. Yocum, "FermiGrid experience and future plans," *Journal of Physics: Conference Series*, vol. 119, no. 5, p. 52, 2008.
- [9] J. Basney, T. Fleury and J. Gaynor, "CILogon: A Federated X.509 Certification Authority for Cyberinfrastructure Logon," *Concurrency and Computation : Practice and Experience*.
- [10] J. Basney, M. Humphrey and V. Welch, "The MyProxy online credential repository.," *Software: Practice and Experience*, vol. 35, no. 9, pp. 801-816, 2005.
- [11] J. Basney, K. Scott and W. Von, "An analysis of the benefits and risks to LIGO when participating in identity federations," LIGO.
- [12] R. Horbe and W. Hotzendorfer, "Privacy by design in federated identity management," *Security and Privacy Workshops*, pp. 167-174, 2015.
- [13] U. Kylau, I. Thomas and C. Meinel, "Trust requirements in identity federation topologies," *International Conference of Advanced Information*, pp. 137-145, 2009.
- [14] M. Kirby, "The Fabric for Frontier Experiments at Fermilab," *Journal of Physics: Conference Series*, 2014.