

Federated Login to TeraGrid

Jim Basney
jbasney@illinois.edu

Terry Fleury
tfleury@illinois.edu

Von Welch
vwelch@illinois.edu

National Center for Supercomputing Applications
University of Illinois
1205 West Clark Street
Urbana, Illinois 61801

ABSTRACT

We present a new federated login capability for the TeraGrid, currently the world's largest and most comprehensive distributed cyberinfrastructure for open scientific research. Federated login enables TeraGrid users to authenticate using their home organization credentials for secure access to TeraGrid high performance computers, data resources, and high-end experimental facilities. Our novel system design links TeraGrid identities with campus identities and bridges from SAML to PKI credentials to meet the requirements of the TeraGrid environment.

Categories and Subject Descriptors

K.6.5 [Management of Computing and Information Systems]: Security and Protection—*Authentication*

General Terms

Security

Keywords

PKI, SAML, identity federation, grid computing, TeraGrid, MyProxy, GridShib, Shibboleth

1. INTRODUCTION

TeraGrid¹ is an open scientific discovery infrastructure combining leadership class resources at eleven partner sites to create an integrated, persistent computational resource. TeraGrid serves over 4,500 researchers from over 300 colleges, universities, and research institutions in the United States. TeraGrid resources are allocated to researchers by peer review. Researchers must authenticate to TeraGrid resource providers and charge their usage to project accounts. TeraGrid supports authentication via passwords, SSH public keys, and X.509 certificates.

¹<http://www.teragrid.org>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

IDtrust '10, April 13-15, 2010, Gaithersburg, MD

Copyright 2010 ACM ISBN 978-1-60558-895-7/10/04 ...\$10.00.

In this article, we present the design and implementation of a new system that enables researchers to use the authentication method of their home organization for access to TeraGrid. Participating in the InCommon Federation² enables TeraGrid to accept authentication assertions from U.S. institutions of higher education, so researchers can use their existing campus login to authenticate to TeraGrid resources.

This federated login capability brings multiple benefits:

- It mitigates the need for researchers to manage authentication credentials specific to TeraGrid in addition to their existing campus credentials. Simplifying researchers' access to TeraGrid helps them to better focus on doing science.
- Reducing or eliminating the need for a TeraGrid password eases the burden on TeraGrid staff, by reducing the number of helpdesk calls requesting password resets and avoiding the need to distribute passwords to researchers in the first place.
- Using the campus login to access TeraGrid helps to integrate campus computing resources with TeraGrid resources. Researchers should be able to easily combine resources on campus with resources from TeraGrid and other national cyberinfrastructure. Harmonizing security interfaces across the infrastructure is a positive step towards this goal.
- Federated login enables the provisioning of TeraGrid resources according to campus-based identity vetting and authorization. TeraGrid resources could be allocated to a university class or department, and TeraGrid could rely on the university to determine who on their campus is authorized to use the resource allocation (e.g., who is enrolled in the class or who is a department member), thereby eliminating the need for per-user accounting by TeraGrid staff and giving the campus greater flexibility and control in managing the TeraGrid allocation.

Federated login is being applied in many environments to simplify authenticated access to resources and services. In this article, we focus on the unique challenges we faced in implementing federated login for TeraGrid. A primary technical challenge was the need to support multiple usage models, from interactive browser and command-line access

²<http://www.incommonfederation.org>

to multi-stage, unattended batch workflows. Another challenge was the need to establish trust among campuses, TeraGrid members, and peer grids (such as Open Science Grid³ and the Enabling Grids for E-science⁴) in the mechanisms and procedures underlying the federated login capability. In the remainder of the article, we discuss these and other challenges and present our solution in detail.

2. BACKGROUND

Before presenting the federated login capability we developed for TeraGrid, we first provide background information about the previously existing TeraGrid authentication architecture and the InCommon Federation.

2.1 TeraGrid Authentication Architecture

The TeraGrid allocations process provisions TeraGrid user accounts and assigns TeraGrid-wide usernames and passwords, which grant single sign-on access to TeraGrid resources. Our work, which we describe subsequently, leverages this existing architecture without modifying it in order not to disrupt access for existing users.

2.1.1 TeraGrid Allocations

As described in the Introduction, TeraGrid resources are allocated to researchers by peer review. Principal Investigators (PIs) submit proposals for resource allocations to a resource allocations committee, which consists of volunteers selected from the faculty and staff of U.S. universities, laboratories, and other research institutions. All members serve a term of 2–5 years and have expertise in computational science or engineering. Each proposal is assigned to two committee members for review. The committee members can also solicit an external review. After several weeks of review, the entire committee convenes to discuss the relative merits of each proposal and award time based on availability of resources. To apply, the PI must be a researcher or educator at a U.S. academic or non-profit research institution. Proposals are judged on scientific merit, potential for progress, numerical approach, and justification for resources. Allocations are typically awarded for one year, though multi-year allocations may be granted for well-known PIs. PIs can submit renewal or supplemental proposals to the committee to extend their allocation.

PIs are instructed not to share their accounts with others. Instead, they use the Add User Form on the TeraGrid User Portal⁵ to request accounts for their project members. PIs can also use this form to remove project members. PIs submit name, telephone, email, and postal address information for the users on their project. For users on multiple projects, each project PI must complete the required information separately for each user to request the user to have access to the project’s resources. The PI is notified by postal mail whenever a user is added to their project. All users are required to sign the TeraGrid User Responsibility Form, which educates users about secure and appropriate computing practices.

When a PI’s proposal is accepted, or when an active PI requests an account for a project member, TeraGrid allocations staff members enroll the PI or project member in the TeraGrid Central Database, assign a TeraGrid-wide user-

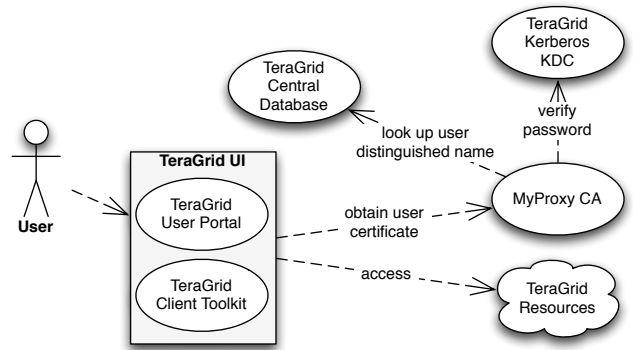


Figure 1: TeraGrid single sign-on provides certificates for secure access to TeraGrid resources.

name and initial password to the researcher, and send the username and password via postal mail to the researcher. The letter distributed with the initial password instructs the researcher to change the password and store the letter in a secure place. If the researcher forgets the password, he or she can call the helpdesk and request that the password be reset to the initial value. If the researcher has lost the letter with the initial password, he or she can call the helpdesk and request that a new letter be sent to their postal address on record. Alternatively, a researcher can reset his or her password via the TeraGrid User Portal, which authenticates the request via the researcher’s registered email address. In the future, TeraGrid researchers will be able to set their username and password when they request an account, eliminating the need for passwords to be sent via postal mail.

The process of enrolling a new user into the TeraGrid Central Database also assigns a unique certificate subject distinguished name to the user. The distinguished name includes the user’s first and last names, with an optionally appended serial number in case of name conflicts. The database management system ensures that distinguished names are uniquely assigned and are never re-assigned to a different user.

As described later, our federated login solution relies on the fact that the TeraGrid Central Database contains a record for every TeraGrid user, as well as the fact that every TeraGrid user has a TeraGrid-wide username and password.

2.1.2 TeraGrid Single Sign-On

The researcher’s TeraGrid-wide username and password enables single sign-on access to all TeraGrid resources. Researchers can use TeraGrid single sign-on from the TeraGrid User Portal (TGUP) and from the command-line (via the TeraGrid Client Toolkit). Upon entering their username and password, researchers obtain a short-lived certificate from a MyProxy⁶ Certificate Authority (CA) [1, 6] operated by NCSA. Researchers use this certificate to authenticate to remote login, data transfer, batch job submission, and other services. Furthermore, researchers can delegate a proxy certificate [15] to remote login sessions and batch jobs, allowing those sessions/jobs to access resources on their behalf. Figure 1 presents the TeraGrid single sign-on system architecture.

⁶<http://myproxy.ncsa.uiuc.edu>

³<http://www.opensciencegrid.org>

⁴<http://www.eu-egee.org>

⁵<https://portal.teragrid.org>

The TeraGrid PKI consists of CAs (including the NCSA MyProxy CA) operated by TeraGrid member institutions and other partners. TeraGrid resource providers accept a consistent set of CAs to facilitate single sign-on across the TeraGrid resources. The TeraGrid Security Working Group reviews requests to add or remove CAs and operates by consensus across the TeraGrid members. According to the policy of the working group, new CAs must be accredited by the International Grid Trust Federation (IGTF),⁷ the de facto standards body for defining levels of assurance for PKIs in production academic grids around the world. As discussed subsequently, IGTF accreditation was an important step in deploying a new federated CA in TeraGrid in support of single sign-on with federated login.

TeraGrid runs a Kerberos domain to validate usernames and passwords. Kerberos is not typically exposed to end users directly but is instead used by other services (such as the MyProxy CA) as an authentication service.

2.2 InCommon Federation

The InCommon Federation enables users to use their local identity, assigned by their campus, to access services such as academic publications and educational materials, and to collaborate with partners outside the borders of the campus. InCommon facilitates the adoption of standard policies by federation participants on technology issues, legal issues, and acceptable uses of identity information. Several U.S. federal agencies (e.g., NSF, NIH) have joined InCommon, and national-scale infrastructures such as the Ocean Observatories Initiative⁸ are exploring its use. InCommon promises to provide a standard interface to the differing campus identity management systems and allow outside leverage of local identities without the need to understand the nuances at each campus.

Many federation members use the Shibboleth⁹ software for expressing and exchanging identity information between organizations. Shibboleth allows organizations to federate identity information. In practical terms, this means a user from one institution can authenticate at their home institution and have the resulting identity (identifier and/or attributes) made available to a second institution for the purposes of accessing resources at that second institution. Shibboleth is commonly used in privacy-preserving applications, where access to resources is granted based on the user's attributes (e.g., "University of Illinois student") without requiring disclosure of the user's name or other identifying information. For example, many universities partner with online content providers to enable students to access journal articles using Shibboleth attributes. Shibboleth implements the SAML Web Browser Single Sign-On protocols,¹⁰ which work well for browser-based applications but do not translate directly to the command-line, complex-workflow, unattended/batch processes that make up a significant proportion of TeraGrid computing workloads.

As of January 2010, the InCommon Federation includes over 200 universities, representing over 4 million users. Of the 38 institutions that each represent over 50 TeraGrid users, 24 (67%) are currently InCommon members. While

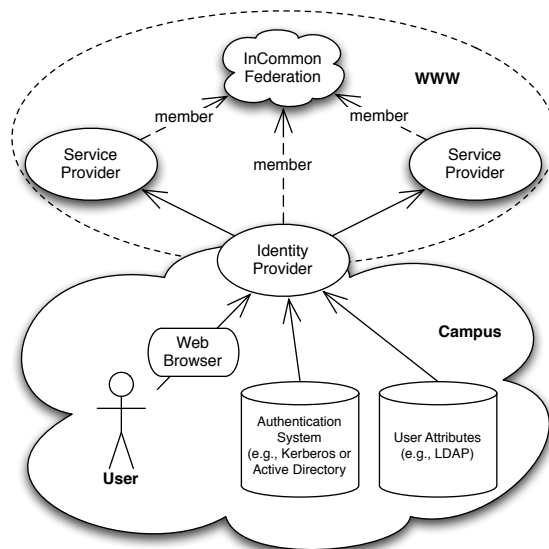


Figure 2: The InCommon Federation defines standard behavior, attributes, and protocols. The campus identity provider converts the user's campus identity into standard SAML format for access to web services.

InCommon membership continues to grow, many TeraGrid users come from campuses that are not (yet) InCommon members. InCommon member ProtectNetwork¹¹ operates an open identity provider that can provide logins for these users.

As depicted in Figure 2, the operational components of the InCommon Federation are the *identity providers*, *service providers*, and the *federation* that brings them together. Identity providers convert the user's campus identity (identifier and/or attributes) into the standard SAML format, providing single sign-on to multiple service providers and supporting anonymity, pseudonymity, and other privacy controls. SAML identity providers rely on campus authentication systems (such as Kerberos) and attribute stores (such as LDAP) to authenticate users and provide identity information. Service providers consume SAML assertions from identity providers to determine a user's identifier and/or attributes for making access control decisions and providing a personalized user experience. SAML metadata, distributed centrally by the federation, identifies the federation members and provides public keys, resource endpoints (URLs), and other information about the members that helps identity providers and service providers establish trust and interoperate.

3. APPROACH

Recall that our goal is to enable TeraGrid researchers to use the authentication method of their home organization for access to TeraGrid. We achieve this goal by implementing a federated login capability that leverages the InCommon Federation to provide a bridge from campus authentication to the existing TeraGrid authentication architecture. In this section, we present the details of our developed solution,

⁷<http://www.igtf.net>

⁸<http://ooi.oceanleadership.org>

⁹<http://shibboleth.internet2.edu>

¹⁰<http://saml.xml.org/saml-specifications>

¹¹<http://www.protectnetwork.org>

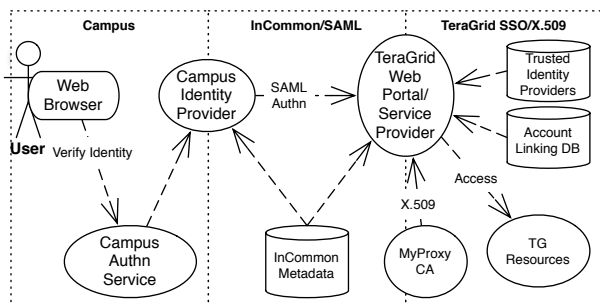


Figure 3: Federated login to TeraGrid relies on translation of credentials between the campus domain, InCommon, and the TeraGrid single sign-on system.

which at its core combines *account linking* and *credential translation*. Our solution builds on the InCommon Federation and existing TeraGrid authentication architecture described in the previous section.

Figure 3 shows a conceptual overview of the credential translation processes. The translation at left between the campus domain and InCommon is handled by Shibboleth or a similar SAML identity provider. The translation at right between InCommon and the existing TeraGrid single sign-on system constitutes our contribution and the focus of this paper. This translation uses the account linking process to bind SAML identities to existing TeraGrid identities.

3.1 Account Linking

The *account linking* process binds the researcher’s campus identity, conveyed via InCommon/SAML, to his or her TeraGrid identity, as stored in the TeraGrid Central Database (TGCDB). When the researcher visits the TeraGrid federated login web site, which implements a standard InCommon SAML service provider using the Shibboleth software, he or she sees a prompt to select an InCommon identity provider (i.e., the researcher’s home campus) in order to initiate authentication. The Shibboleth software redirects the researcher to the selected identity provider, where the researcher logs in. The identity provider then redirects the researcher back to the TeraGrid site with a SAML authentication assertion, according to the SAML protocols. At this point the account linking component is activated. It first searches the account-link database (actually a table in the existing user database) for an entry matching the researcher’s authenticated campus (SAML) identifier. If found, the entry identifies the TeraGrid username linked to that campus identity, allowing the researcher’s TeraGrid login to proceed. If no entry is found, the federated login site prompts the researcher for his or her TeraGrid-wide username and password. If the username and password verify (via the TeraGrid Kerberos service), the federated login site creates a new entry in the account-link database linking the TeraGrid account with the campus identity. Then the researcher’s TeraGrid login can proceed with that TeraGrid-wide username. When the researcher returns to the site at a later time, the account-link entry will be in place, so the researcher will be able to log in using his or her campus identity without being prompted again for a TeraGrid-wide username and password.

It is important to note that the account linking process does not replace the TeraGrid allocations process. Rather, the account linking process relies on the allocations process for identity vetting and authorization of TeraGrid users. The federated login capability provides only a new authentication method for vetted TeraGrid researchers.

TeraGrid users may link identities from multiple identity providers to their TeraGrid account, allowing researchers associated with multiple research institutions to log in to TeraGrid using whichever identity provider is convenient at the time. However, to avoid account sharing (which is a violation of TeraGrid policy), researchers may link at most one identity from each identity provider with their TeraGrid account. For example, a professor may not link his or her graduate students’ campus identities with his or her TeraGrid account. Instead, the TeraGrid policy requires each professor, graduate student, etc., to obtain their own individual TeraGrid account. After login, TeraGrid users may view and delete their account links.

Account links expire one year after creation, at which point the user is required to perform the account linking process again, to re-verify the binding between the user’s federated identity and his or her TeraGrid account. This periodic verification of the binding protects against stale or re-assigned campus identities (e.g., when a student graduates). When federating with each campus, TeraGrid staff members confirm with the campus operators that campus procedures ensure that identities are never re-assigned within a one year interval.

3.2 Credential Translation

The account linking process facilitates a browser-based, federated login to TeraGrid systems. However, as discussed previously, a significant proportion of TeraGrid use cases and workloads are command-line, complex-workflow, and/or unattended/batch processes, which are not well supported by browser-based authentication (i.e., SAML Web Browser Single Sign-On). So, the TeraGrid federated login employs *credential translation* to convert the browser-based credential to a credential that supports these use cases.

Specifically, the TeraGrid federated login converts the authentication assertion, provided by an InCommon-member identity provider, to an X.509 certificate, provided by a certificate authority (CA) trusted by TeraGrid. TeraGrid has a significant investment in a certificate-based single sign-on infrastructure. Support for certificate-based authentication in remote login (GSISSH), job submission (GRAM), and file transfer (GridFTP) protocols enables today’s interactive TeraGrid use cases. Furthermore, proxy certificate delegation [15] enables complex, multi-tier workflows and batch processing in TeraGrid.

Through TeraGrid’s federated login capability, TeraGrid researchers can use their campus login to obtain certificates for web and desktop applications. After federated login, the TeraGrid web site presents a menu of options. Researchers can launch remote login and file transfer applets in their browser, authenticating with a certificate loaded into their browser session. Additionally, researchers can launch an application that delivers a certificate to the local filesystem, ready to be used with desktop applications such as those provided by the TeraGrid Client Toolkit. Implementation details are provided in later sections.

In summary, the researcher’s federated login to TeraGrid

requires multiple credential translation steps. First, the local campus identity provider translates a local campus credential (such as a Kerberos username and password) to a SAML authentication assertion as specified by InCommon. Then, TeraGrid's federated login system translates the SAML assertion to an X.509 certificate. Finally, TeraGrid resource providers translate the certificate to a local resource login (i.e., a Unix account).

3.3 Trust Establishment

Establishing trust is critical to successfully bridging from campus identity providers to TeraGrid resource providers. Deploying the TeraGrid federated login required negotiation with InCommon members (to release identities to TeraGrid) and accreditation of our CA by IGTF (so the certificates will be accepted by TeraGrid members).

3.3.1 Campus Federation

When TeraGrid became a member of the InCommon Federation, it was not automatically entitled to obtain authentication assertions from InCommon-member identity providers. First, TeraGrid needed to register its federated login service provider with the federation, so its information would be included in the federation metadata, enabling it to be recognized by identity providers. This registration is a lightweight task, requiring only a few minutes of effort.

Following that registration, and of significant effort to arrange, the identity providers need to configure their local policies to release identity information to the TeraGrid's federated login service. Specifically, the federated login service depends on receiving a persistent user identifier from the identity provider via the `eduPersonPrincipalName` (ePPN) or `eduPersonTargetedID` (ePTID) attribute defined by the `eduPerson` specification.¹²

In our effort to have identity providers release ePPNs or ePTIDs to TeraGrid, we encountered three categories of identity providers:

- The first type of identity provider was willing to release ePPNs or ePTIDs to any InCommon-member service provider by default. In this case, after reviewing the published policies of the identity provider, we asked a TeraGrid user associated with that identity provider to help us with testing. After a successful test (i.e., a valid assertion with ePPN or ePTID was received), we added that identity provider to the supported list.
- The second type of identity provider was willing to release ePPNs or ePTIDs on request. In this case, we sent email to the contact address found in InCommon Federation metadata, explaining our application and requesting the needed attribute. Once we received a reply that our request was approved, we proceeded with testing as in the first case.
- The third type of identity provider required local sponsorship and review of our request. In this case, we sent a list of TeraGrid PIs affiliated with the institution to the identity provider contact and worked with them to identify sponsors and follow the local approval process. For some of these campuses, the review is still in progress or stalled.

¹²<http://middleware.internet2.edu/eduperson>

Since federating with campuses was a manual, campus-by-campus process, and there is no method to discern what behavior a campus would present until they were engaged, we focused our efforts on campuses with over 50 TeraGrid users. Of the 38 target institutions, 24 (67%) were InCommon members. To date, we have successfully federated with 16 of those. We have also federated by request with 11 additional campuses outside our initial target list, bringing our current total number of supported campuses to 27.

3.3.2 PKI Federation

Translating SAML authentication assertions from InCommon members to certificates accepted by TeraGrid resource providers and peer grids required us to deploy a certificate authority (CA) and obtain accreditation of the CA from the International Grid Trust Federation (IGTF), to satisfy TeraGrid Security Working Group policies. The IGTF consists of three regional Policy Management Authorities (PMAs). The Americas Grid PMA (TAGPMA)¹³ covers the U.S. region.

Worldwide participation in the IGTF ensures that certificates issued by accredited CAs can be accepted by TeraGrid and peer grids around the world. While today's academic SAML federations are national in scope, with limited international inter-federation, translating SAML assertions to internationally accepted certificates supports international science projects such as the Worldwide Large Hadron Collider Computing Grid (WLCG).¹⁴

The IGTF currently supports accreditation under three CA profiles: Classic, Member Integrated Credential Services (MICS), and Short-Lived Credential Services (SLCS).¹⁵ For Classic CAs, subscriber identity vetting is performed by registration authority (RA) staff persons. In contrast, MICS and SLCS CAs leverage an existing identity management system for vetting certificate requests. We pursued accreditation for our federated CA under the SLCS profile, since our CA leverages the TeraGrid Central Database and identity providers in the InCommon Federation.

SLCS CAs issue short-lived certificates. The short certificate lifetime acts as a countermeasure against credential theft and misuse. The maximum lifetime of one million seconds (or about twelve days) was determined through a requirements-gathering process in the Global Grid Forum [12] and was later incorporated into the SLCS profile.

IGTF profiles require that CAs operate according to community standards. Each CA must publish a Certificate Policy and Certification Practices Statement (CP/CPS) according to RFC 3647 [7]. NCSA's CP/CPS documents are published on the NCSA CA web site.¹⁶ Certificates and Certificate Revocation Lists (CRLs) must conform to RFC 5280 [8] and the Open Grid Forum Grid Certificate Profile [10]. Additionally, since SLCS CAs are online and automated, and therefore subject to network-based attacks, the SLCS profile requires that the CA private key be protected in a FIPS 140 level 2 rated hardware security module [13].

The TAGPMA review process includes a presentation to the TAGPMA membership at a regularly scheduled meeting and a checklist-based review of the CA's policies and operations, followed by a vote for acceptance by the TAGPMA

¹³<http://www.tagpma.org>

¹⁴<http://lcg.web.cern.ch>

¹⁵http://www.tagpma.org/authn_profiles

¹⁶<http://ca.ncsa.uiuc.edu>

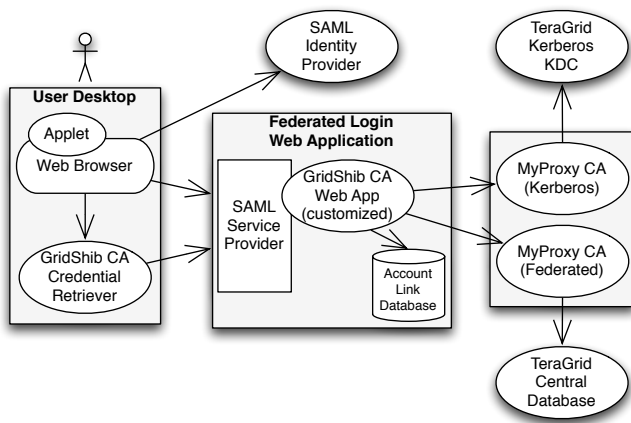


Figure 4: The TeraGrid federated login system provides certificates, issued by a MyProxy CA, for web and desktop applications. The web application binds campus identities to TeraGrid identities via an account-link database.

membership. NCSA began the TAGPMA review process for the federated CA in March 2009 and obtained certification in May 2009. NCSA has been a TAGPMA member since 2005, and this was our third CA to be accredited via the TAGPMA process. Approved CAs are included in the IGTF CA distribution, as well as the TERENA Academic CA Repository (TACAR).¹⁷

3.4 System Architecture

Figure 4 presents the components of the TeraGrid federated login system. The federated login web application is a SAML service provider, which consumes SAML authentication assertions from InCommon-member identity providers, via the Shibboleth software implementation. The web application has a local PostgreSQL database that stores the account linking information. We decided to (initially) maintain this information in a local database separate from the TeraGrid Central Database (TGCDDB), to obtain local database performance and simplify the initial implementation. However, we plan to migrate it to the TGCDDB (also PostgreSQL) when we integrate the federation functionality with the TeraGrid User Portal (see Section 6.1).

The web application interacts with two MyProxy CA instances (via the simple MyProxy protocol [2]) for verifying TeraGrid passwords and obtaining short-lived certificates. The first MyProxy CA instance was already in existence (certified by TAGPMA in March 2007) serving TeraGrid single sign-on. It verifies the user's TeraGrid-wide username and password and issues short-lived certificates. In the federated login application, we use this MyProxy instance to verify TeraGrid (Kerberos) passwords at account linking time. Since the web application already contained MyProxy client libraries, using the MyProxy interface to Kerberos rather than interacting with Kerberos directly simplified the web application. The second MyProxy CA instance is the new federated CA, certified by TAGPMA in May 2009. It issues certificates based on federated login. It trusts the federated login web application to properly validate SAML

¹⁷<http://www.tacar.org>

authentication assertions (using Shibboleth) and map campus identities to TeraGrid usernames. The web application sends the authenticated TeraGrid username to MyProxy, which issues a short-lived certificate corresponding to that username. The web application authenticates to MyProxy using its own trusted certificate. The federated MyProxy instance will only accept requests properly authenticated using that certificate. Both MyProxy instances map TeraGrid usernames to certificate subject distinguished names via the TGCDDB.

When the TeraGrid user launches one of the browser applets that require a certificate for authentication to TeraGrid resources, the federated login web application, via the MyProxy API, generates a new RSA keypair associated with the user's web session (via state in the web server referenced by a session cookie) and issues a certificate request containing the RSA public key to MyProxy, which returns a short-lived, signed certificate for the user to the web application. The applets can then access the private key and certificate for authentication on the user's behalf. Similarly, when the TeraGrid user selects the credential retrieval desktop application, the browser downloads and launches the application via Java Web Start [11]. The desktop application then generates a new RSA keypair and issues a certificate request to the web application, which passes it to MyProxy and returns the signed certificate to the desktop application, which writes the certificate and private key to the filesystem for access by TeraGrid client applications. The credential retrieval application and components of the web application are re-used from the GridShib CA software as developed by the GridShib project [18].¹⁸

3.5 Current Status

The TeraGrid federated login service¹⁹ is in production, supporting logins from 27 institutions. After accreditation by TAGPMA in May 2009, the site entered a friendly-user beta testing period, where we solicited test users from each supported campus to try the service and give their feedback. We announced the service to all TeraGrid researchers via TeraGrid News on September 1, 2009.

As of February 2010, we have 72 entries in the identity-mapping table from 21 (of the 27 available) institutions, and we have issued over 800 certificates. The most popular application is the remote login GSI-SSHTerm applet,²⁰ followed closely by the credential retrieval desktop application.

4. SECURITY CONSIDERATIONS

Security was a primary consideration throughout the design and deployment of the federated login service. We highlight security considerations of particular interest in this section.

4.1 Trust Architecture

Adding federated identity to the TeraGrid single sign-on model gives rise to two meaningful changes to the trust relationships in the TeraGrid security architecture.

First, the InCommon identity providers add a new set of trusted entities. Identity providers are trusted to correctly

¹⁸<http://gridshib.globus.org>

¹⁹<https://go.teragrid.org>

²⁰<https://sourceforge.net/projects/gsi-sshterm>

authenticate users, disallow the reuse of identifiers, and adhere to other basic policies, as discussed in the following section. Identity providers also play a role in incident response as discussed in Section 4.6.

Second, the federated MyProxy CA outsources authentication to the web front-end. In the current TeraGrid User Portal, a user presents a username and password, which are passed to the MyProxy CA for validation before issuance of a credential. In the federated identity model, the web application presents just a username to the MyProxy CA and authenticates using a trusted certificate specific to the web application instead of the user. The MyProxy CA trusts that the web application has done appropriate authentication of the user. This increases the ramifications of a compromised web application.

The MyProxy CA could be modified to require and validate some proof that the web application actually authenticated the user. One way to provide this validation could be to implement SAML delegation.²¹ The ShibGrid project [14] modified MyProxy to validate SAML authentication assertions obtained by the web application. While that implementation does not use SAML delegation, it provides some additional protection. This capability could be added to the TeraGrid service, but it would increase the complexity of the solution.

4.2 Peering with Identity Providers

As discussed in Section 3.3.1, federating with campus identity providers is a manual process. Identity providers decide whether they are willing to release user identifiers to the TeraGrid service. Likewise, TeraGrid staff members, in their role as administrators of the federation service, decide whether to peer with a given campus identity provider. The federated login service is explicitly configured with a list of trusted identity providers (i.e., not all InCommon-member identity providers are automatically accepted). Our review process confirms that the identity provider: (1) serves TeraGrid users; (2) is operated by a known and respected organization; and (3) operates a trustworthy authentication service and provides globally-unique and non-reassigned identifiers, so that subscribers are uniquely identified.

So far, the issue of identifier re-assignment has blocked us from peering with a few campus identity providers. Our annual verification process allows us to support campuses that re-assign identifiers only after a one year or greater hiatus period. We have found in some cases, campuses will re-assign identifiers more quickly for a subset of their population (e.g., undergraduate students and/or visitors), and we are working with those campuses to identify a method to distinguish between those identities that meet our requirements (i.e., those not re-assigned more quickly than our threshold) and those that don't. InCommon's new Identity Assurance program²² may help with this issue.

4.3 Disallowing Account Sharing

As discussed in Section 2.1.1, TeraGrid policy forbids account sharing. This policy is primarily for clarity during incident response, since multiple users sharing an account complicates the process of determining if suspect account activity was performed by the authorized account holder or

by an unauthorized party using the stolen password of the account holder. To enforce this policy, we allow only one identifier per identity provider to be linked with a particular TeraGrid identity.

4.4 Web Application Security

We use multiple methods in the web front end to protect against web-based attacks. The web front end accepts connections only via HTTPS, which provides certificate-based authentication of the service to the web browser and privacy of network data (including SAML assertions, cookies, and certificate requests). To protect against cross-site request forgery (CSRF) attacks, the GridShib CA software uses standard anti-CSRF mechanisms (cookies and hidden form fields) to ensure that web sessions follow an approved workflow, i.e., requiring the user to always visit the login page before requesting a certificate, so a malicious site can not redirect the user's browser directly to the certificate-request form to force a malicious certificate issuance.

The account-link database is configured to allow only local access, and anonymous read access to the database is disabled. The username and password for accessing the database is stored outside publicly accessible web space, and is readable only by the web server process. This configuration gives the server-side web application read and write access to the database while preventing all client-side web access.

The trusted certificate used to request user certificates from the federated MyProxy CA is stored on the web server outside publicly accessible web space and is readable only by the web server process.

Remote login to the web server is restricted to a small set of remote hosts through the use of an iptables-based firewall. Additionally, SSH access is limited to a small number of administrators, who must log in with a one time password (OTP), e.g., by using a CRYPTOCARD token generator.

4.5 MyProxy CA Security

The back-end MyProxy CA is secured according to IGTF standards. The CA private key is protected in FIPS 140 level 2 rated hardware security modules. The servers are located on a dedicated network, behind a hardware firewall with a restrictive policy, with network-based and host-based intrusion detection. The firewall allows network connections to the MyProxy CA instance used by the web application only from the host on which that application resides. System logs are streamed to a dedicated syslog collector host, where they are monitored by the NCSA security team. The CA issues a certificate revocation list (CRL) daily or immediately after any revocation.

4.6 Incident Response

The federated login system architecture provides multiple methods for responding to account compromises and other security incidents. In case a federated identity is deemed suspect, the account link for that identity can be disabled in the account-link database by administrators so it can no longer be used to obtain certificates. In case an identity provider is deemed suspect, it can be removed by an administrator from the list of trusted identity providers so assertions from that provider can no longer be used to log in. Extensive CA logging enables administrators to quickly identify certificates associated with a compromise so they can be revoked.

TeraGrid incident response is coordinated through the se-

²¹<http://docs.oasis-open.org/security/saml/Post2.0/sstc-saml-delegation.html>

²²<http://www.incommonfederation.org/assurance>

curity working group. In response to compromise, TeraGrid resource providers can locally disable accounts, and TeraGrid staff can centrally disable or reset TeraGrid-wide passwords.

InCommon metadata contains operational contact information for each identity provider that TeraGrid security staff can utilize during incident response. Additionally, work is underway in the Committee on Institutional Cooperation²³ Identity Management Taskforce to propose a set of policies and additional available information for incident response in federated identity environments such as InCommon.

Like all IGTF CAs, the federated NCSA CA publishes operational contact information on its home page and in metadata files included in the IGTF CA distribution. The IGTF Risk Assessment Team²⁴ is available for coordinating response to incidents and vulnerabilities impacting IGTF CAs.

5. LESSONS LEARNED

In this section we discuss some of the lessons learned during the deployment of our solution and establishment of trust with identity providers in InCommon.

5.1 Effort for Trust Establishment

As we described previously in Section 3.3.1, while InCommon defines standard (SAML) profiles for identity and attribute transmission and an automated means of metadata distribution, simply being a member of InCommon as a service provider does not guarantee that any particular identity provider will release user attributes to that service provider. Nor does it provide guarantees about identifier persistence in that ePPN identifiers can be potentially re-issued (e.g., after a student leaves the student's identifier could be re-assigned to a new incoming student).

The process of contacting identity providers to arrange attribute release and establish their policies on identifier re-issuance is very time consuming. This manual, campus-by-campus effort will be very difficult to scale to the hundreds of campuses associated with TeraGrid researchers, not to mention the thousands of research institutions in the U.S. from where future TeraGrid users might come.

We look forward to deployment of user-driven attribute release in the InCommon Federation, which would avoid the need for manual policy changes by campus operators. User-driven attribute release, via tools such as uApprove,²⁵ allows users to review and consent to the release of requested attributes when they access the service.

5.2 Testing

Another complexity encountered during attribute release testing was that the identity provider administrators at campuses were rarely TeraGrid users. This meant that only our end users, who are not generally Shibboleth experts, could test the system from end-to-end, as they were the only ones with accounts at both the identity provider and the TeraGrid. Adding a simple test application that could be used by identity provider operators to more fully test the attribute release process, without needing to have a TeraGrid account,

²³<http://www.cic.net>

²⁴<http://tagpma.es.net/wiki/bin/view/IGTF-RAT>

²⁵<http://www.switch.ch/aai/support/tools>

would be a useful addition to this trust establishment procedure.

5.3 Software Issues

A major source of issues during our beta testing period was the lack of constraint as to the contents of eduPerson-TargetedID (ePTID) values. We found significant variety in the formatting and character sets of ePTID values across campuses, which clashed with several assumptions in our software:

- The various ePTID values triggered exceptions in the GridShib CA identifier sanitizing routines, which attempted to sanitize data from the identity provider to protect against accidental or malicious string encoding that could cause problems. These routines were too aggressive in removing “invalid characters”, thereby corrupting the identifiers, and we were forced to abandon such sanitization.
- There was also an assumption in the original software of the identifiers being usable as filenames to maintain an audit record of issued credentials (a requirement of IGTF accreditation). However, some of the characters were meaningful to the file manipulation routines (e.g., forward slashes which represent a path separator under Unix). Hence the approach of using the ePTID was abandoned and instead we used a hash of the distinguished name with a constrained character set.
- Finally, our web site originally displayed the ePTID value to the user after login. While this approach worked with eduPersonPrincipalName values, which are reasonably similar to users' campus usernames and email addresses, we found that the lengthy ePTID string with its broad range of characters distracted and confused users, who expect to see their friendly campus username.

In summary, we have learned to treat ePTIDs as opaque blobs unsuitable for use as a string representation of an identifier and have strengthened the underlying GridShib CA identifier-handling code to support the full range of ePTID values.

6. FUTURE WORK

We consider this work to be just a first step toward enabling federated login to TeraGrid and other U.S. cyberinfrastructure. We envision the following future work.

6.1 Integration with TeraGrid User Portal

The next step for the TeraGrid effort is to integrate federated login with the TeraGrid User Portal (TGUP). Currently, the federated login site is separate from the TGUP, and the TGUP itself requires login with TeraGrid-wide username and password. Integration with the TGUP will provide a more coherent experience to TeraGrid researchers, as well as make TGUP functionality (such as management of TeraGrid allocations) accessible via federated login.

The TeraGrid project is in the process of integrating the Partnership Online Proposal System (POPS)²⁶ with the user portal, which opens up the possibility of federated logins

²⁶<https://pops-submit.teragrid.org>

for TeraGrid proposal submission, potentially eliminating the need for TeraGrid-specific passwords as described in the following section.

6.2 Eliminating TeraGrid Passwords

The account linking process as described so far requires TeraGrid researchers to log in with their TeraGrid username and password at least once per year to maintain the link with their campus identity. This method provides a transition for existing TeraGrid users from daily use of a TeraGrid-specific password to daily use of campus credentials for TeraGrid access, but it does not entirely obviate the need for TeraGrid-specific passwords.

In the future, we plan to integrate account linking with the TeraGrid allocations process, giving TeraGrid researchers the option of never using a TeraGrid-specific password. In this scenario, TeraGrid researchers would authenticate with their campus identity when submitting a proposal for TeraGrid access. A researcher's campus identity will be linked with the proposal at that point, so if the proposal is accepted and TeraGrid access is granted, the researcher's TeraGrid account will be linked with the campus identity when the TeraGrid account is created.

Likewise, project members to be added to a TeraGrid allocation will first authenticate with their campus identity and register a TeraGrid account linked with that campus identity. Then, the project PI will lookup the prospective member's account and add the member to the TeraGrid project. Thus, PIs and other project members will have their campus identities linked with their TeraGrid accounts when the TeraGrid accounts are created, so researchers will be able to access TeraGrid resources using their campus logins without ever having a TeraGrid-specific password. These linked identities could be re-verified each year as part of the allocations renewal process.

It is an open question whether TeraGrid could ever truly eliminate TeraGrid-specific passwords for all users. While we expect many users would prefer to use a federated login, some users may still desire TeraGrid-specific passwords by preference or special requirements.

6.3 Access Based on Attributes

There is a small amount of access to TeraGrid today that is not based on the peer-review process previously described, but is instead granted to a class or workshop for educational purposes. In theory, this access could be granted based on a user's attribute, namely their membership in the class, if it were asserted by their identity provider. Working with campuses to grant access to TeraGrid resources based on such attributes is another area of future investigation.

6.4 Alternative Authentication Technologies

While InCommon and SAML appear to be the most popular technology for federated identity at the home institutions of most TeraGrid users, other web-based authentication methods such as OpenID²⁷ are popular in the commercial space. We plan on investigating the support of these technologies in our federation model.

6.5 CILogon

Expanding federated login to other U.S. cyberinfrastructure is another area of future work. Relying on the TeraGrid

²⁷<http://openid.net>

allocations process for identity vetting restricts the availability of the TeraGrid federated login service to registered TeraGrid users. The CILogon project²⁸ is deploying a modified version of the TeraGrid federated login service that removes the TeraGrid dependencies. The CILogon Service will directly leverage campus identity vetting for certificate issuance. The InCommon Silver Identity Assurance Profile, which maps to NIST Level of Assurance (LOA) 2 [5], provides identity assertions which meet IGTF SLCS profile requirements [3].

Scaling the CILogon Service to serve the national cyberinfrastructure will be a significant challenge. Federating with thousands of U.S. research institutions will require moving beyond the manual campus-by-campus trust establishment process. Providing a usable method for choosing among thousands of available identity providers for a given login is an unsolved challenge. Certainly today's interfaces, where users select their identity provider from a list, will not scale.

7. RELATED WORK

The two areas of related work we find most relevant to the TeraGrid federated login service are (1) similar efforts to bridge SAML and PKI for grids in Europe and (2) TeraGrid's Science Gateways program.

7.1 European SAML-PKI Bridging Efforts

Many European countries have established national SAML federations, with multiple national-scale efforts to link with PKIs in support of cyberinfrastructure.

In Switzerland, SWITCH operates the SWITCHaai federation²⁹ deployed by most Swiss universities supporting e-learning, e-conferencing, and document exchange services. The IGTF-accredited SWITCH Short Lived Credential Service (SLCS) issues certificates based on successful authentication at a SWITCHaai identity provider.

In Germany, the IGTF-accredited DFN-SLCS CA³⁰ issues certificates to users of the DFN-AAI federation³¹ of universities, technical colleges, and research organizations in Germany.

In the UK, JANET, the national education and research network, operates the UK Access Management Federation for Education and Research,³² with over 700 members. The SARoNGS Credential Translation Service [16] issues certificates to users of the UK National Grid Service³³ based on successful authentication in the UK Access Management Federation.

Additionally, the Trans-European Research and Education Networking Association (TERENA) has recently developed the TERENA Certificate Service (TCS),³⁴ which leverages the national SAML-based federations across Europe to deliver certificates to tens of thousands of grid users. Initial TCS partners include the national grid projects and SAML federations of Denmark, Finland, Netherlands, Norway, and Sweden.

²⁸<http://www.cilogon.org>

²⁹<http://www.switch.ch/aa>

³⁰<http://www.pki.dfn.de>

³¹<https://www.aai.dfn.de>

³²<http://www.ukfederation.org.uk>

³³<http://www.ngs.ac.uk>

³⁴<https://www.terena.org/activities/tcs>

Our work to implement federated login for TeraGrid benefited from the examples provided by these related efforts and discussions in IGTF on lessons learned and best practices for bridging SAML and PKI for grids.

7.2 TeraGrid Science Gateways Program

Considering that our work to deploy federated login for TeraGrid is motivated by the desires to make secure access to TeraGrid more convenient for researchers as well as reduce TeraGrid's identity management burdens (e.g., password resets), we find similar motivations for the security design of the TeraGrid Science Gateway program [4, 17]. TeraGrid science gateways³⁵ provide community-based access to TeraGrid resources, typically via web portals with custom interfaces and applications for specific science communities. The gateway program is part of TeraGrid's effort to serve the larger science community, while continuing to provide high-end computing services to a smaller number of leading-edge researchers. TeraGrid's gateways are designed to serve orders of magnitude more users than can be supported by TeraGrid's existing accounting procedures.

To achieve this goal, TeraGrid provides community allocations to gateways. Gateway PIs and staff are registered in the TeraGrid Central Database (TGCDDB), but the gateways manage their own user registration. Gateways access community accounts on TeraGrid resources, with the gateway taking responsibility for isolating its users from one another, so the TeraGrid resource providers are not burdened with managing orders of magnitude more local accounts. Since TeraGrid's federated login capability is based on TGCDDB registration, science gateway users do not benefit directly. However, we hope science gateways will provide their own federated login capability. For one proposal, see [9].

8. ALTERNATIVE APPROACHES

A question often posed is what is needed in order to implement a user authentication solution based entirely on SAML or PKI instead of a SAML to PKI bridge. There are significant components missing for each approach, as we describe in the following subsections, that led us to the bridge approach.

8.1 End-to-End PKI Solution?

The TeraGrid has a PKI solution in place with its existing single sign-on system as described in Section 2.1.2. However, ideally TeraGrid would not need to issue certificates, but instead would rely on certificates issued by the user's home organization, taking advantage of the in-person vetting that is (or at least could be) accomplished by that organization. However, despite some progress, we are seeing very limited deployment of externally usable PKIs at universities, as compared with the number of universities that have joined the InCommon Federation. It is the broad and increasing adoption of InCommon in the organizations representing TeraGrid users that led us to build on it, rather than any technical aspect of the SAML technology.

Note that users with credentials from trusted certificate authorities at universities that do operate a PKI can bind, through existing mechanisms in the TeraGrid User Portal, the identity asserted by those credentials to their existing TeraGrid account and access the TeraGrid with those cre-

³⁵<http://www.teragrid.org/gateways>

entials. In order for such certificate authorities to be considered trusted by the TeraGrid they must have achieved accreditation by the International Grid Trust Federation as described in Section 2.1.2.

8.2 End-to-End SAML Solution?

To replace the PKI currently in use for single sign-on in the TeraGrid today would not only require that TeraGrid modify a large software deployment base, but would also require addressing functional limitations in SAML, namely:

- Support for clients other than web browsers. Many of the science applications supported by TeraGrid involve desktop applications rather than or in addition to web browsers.
- Delegation support. Our architecture supports authentication on behalf of the user by the web application. It also supports authentication by unattended processes, for example, when the initiating user is offline. (SAML delegation may address this requirement.)
- International federation support. SAML federations have not (yet) reached the global scope of the International Grid Trust Federation as needed to support large grid applications.

Until these issues are addressed, we do not envision a migration away from PKI to be a practical option for TeraGrid.

9. CONCLUSION

In conclusion, we have presented TeraGrid's new federated login capability, which enables TeraGrid users to authenticate using their home organization credentials for secure access to high performance computers, data resources, and high-end experimental facilities. This capability binds campus identities to TeraGrid identities (via *account linking*) and issues certificates based on SAML assertions (via *credential translation*). It is the first effort to leverage federated authentication for access to national-scale research cyberinfrastructure in the United States.

It is our opinion that the world is unlikely to ever settle on a single authentication technology, due to varied technical requirements, as well as significant social and economic issues. Therefore, we believe that the bridging approach described in this article is not simply a short-term hack, but rather an approach that will continue to be required and further refined over time.

10. ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grant No. 0503697.

11. REFERENCES

- [1] T. Barton, J. Basney, T. Freeman, T. Scavo, F. Siebenlist, V. Welch, R. Ananthakrishnan, B. Baker, M. Goode, and K. Keahey. Identity Federation and Attribute-based Authorization through the Globus Toolkit, Shibboleth, Gridshib, and MyProxy. In *Proceedings of the 5th Annual PKI R&D Workshop*, April 2006.

- [2] J. Basney. MyProxy Protocol. Global Grid Forum GFD-E.54, November 2005.
- [3] J. Basney. Mapping InCommon Bronze and Silver Identity Assurance Profiles to TAGPMA SLCS Requirements, March 2009. <http://sl.cilogon.org/incommon-slcs-map.pdf>.
- [4] J. Basney, S. Martin, J. Navarro, M. Pierce, T. Scavo, L. Strand, T. Uram, N. Wilkins-Diehr, W. Wu, and C. Youn. The Problem Solving Environments of TeraGrid, Science Gateways, and the Intersection of the Two. *IEEE International Conference on eScience*, pages 725–734, 2008.
- [5] W. E. Burr, D. F. Dodson, and W. T. Polk. Electronic Authentication Guideline. NIST Special Publication 800-63, April 2006.
- [6] S. Chan and M. Andrews. Simplifying Public Key Credential Management Through Online Certificate Authorities and PAM. In *Proceedings of the 5th Annual PKI R&D Workshop*, April 2006.
- [7] S. Chokhani, W. Ford, R. Sabett, C. Merrill, and S. Wu. Internet X.509 Public Key Infrastructure Certificate Policy and Certification Practices Framework. IETF RFC 3647, November 2003.
- [8] D. Cooper, S. Santesson, S. Farrell, S. Boeyen, R. Housley, and W. Polk. Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile. IETF RFC 5280, May 2008.
- [9] T. Fleury, Y. Liu, T. Scavo, and V. Welch. A Web Browser SSO Model for Science Gateways. In *Proceedings of the 2009 TeraGrid Conference*, June 2009.
- [10] D. Groep, M. Helm, J. Jensen, M. Sova, S. Rea, R. Karlsen-Masur, U. Epting, and M. Jones. Grid Certificate Profile. Open Grid Forum GFD-C.125, March 2008.
- [11] A. Herrick. Java Network Launching Protocol & API Specification. JSR-56, 2005.
- [12] S. Mullen, M. Crawford, M. Lorch, and D. Skow. Site Requirements for Grid Authentication, Authorization and Accounting. Global Grid Forum GFD-I.032, October 2004.
- [13] NIST. Security Requirements for Cryptographic Modules. Federal Information Processing Standards (FIPS) Publication 140-2, May 2001.
- [14] D. Spence, N. Geddes, J. Jensen, A. Richards, M. Viljoen, A. Martin, M. Dovey, M. Norman, K. Tang, A. Trefethen, D. Wallom, R. Allan, and D. Meredith. ShibGrid: Shibboleth Access for the UK National Grid Service. In *Proceedings of the International Conference on e-Science and Grid Computing*, December 2006.
- [15] S. Tuecke, V. Welch, D. Engert, L. Pearlman, and M. Thompson. Internet X.509 Public Key Infrastructure Proxy Certificate Profile. IETF RFC 3820, June 2004.
- [16] X. D. Wang, M. Jones, J. Jensen, A. Richards, D. Wallom, T. Ma, R. Frank, D. Spence, S. Young, C. Devereux, and N. Geddes. Shibboleth Access for Resources on the National Grid Service (SARoNGS). *International Symposium on Information Assurance and Security*, 2:338–341, 2009.
- [17] V. Welch, J. Barlow, J. Basney, D. Marcusiu, and N. Wilkins-Diehr. A AAAA model to support science gateways with community accounts. *Concurrency and Computation: Practice and Experience*, 19(6):893–904, 2007.
- [18] V. Welch, T. Barton, K. Keahey, and F. Siebenlist. Attributes, Anonymity, and Access: Shibboleth and Globus Integration to Facilitate Grid Collaboration. In *Proceedings of the 4th Annual PKI R&D Workshop*, April 2005.