CILogon: Secure Access to National-Scale CyberInfrastructure

Submitted to the National Science Foundation in response to
"Strategic Technologies for Cyberinfrastructure (STCI)" (NSF 06-7231)

Principal Investigator

James Basney
Senior Research Scientist
National Center for Supercomputing Applications (NCSA)
University of Illinois
1205 W. Clark Street
Urbana, Illinois 61801
Tel: (217) 244-1954
Email: jbasney@ncsa.uiuc.edu

Co-Principal Investigators

Randal Butler, National Center for Supercomputing Applications
Von Welch, National Center for Supercomputing Applications

A Project Summary

The emergence of Cyberinfrastructure (CI) that supports the global research community by linking together researchers, computational resources, instruments, and data, has had a significant impact on the science and engineering research communities. At the heart of CI is trust, between collaborators, organizations, providers, users, applications and services. Security services, specifically authentication and authorization, are foundational services that facilitate trust between those entities. Without the establishment and maintenance of trust, collaborative relationships served by CI are jeopardized. As the role of cyberinfrastructure increases, so do the challenges for the authentication infrastructure. Initially interactions were between a single user and a single resource within a single administrative domain. With grid computing we saw the growth of multi-organization collaborations among national and state centers, supporting entire communities and collections of resources distributed geographically. More recently projects such as NSF TeraGrid and LIGO are actively extending the breadth and depth of these collaborations

The goals of the proposed *CILogon* project are to maintain and provide critical enhancements for CI security technologies developed at the National Center for Supercomputing Applications (NCSA) and to foster science and engineering by helping additional communities build secure CI on these services. The MyProxy, GridShib, and GSI-OpenSSH security technologies were developed through great effort by the NSF community and represent a major success story for NSF by supporting science and engineering research throughout the world. *MyProxy* is the de facto grid credential management service used worldwide. *GridShib* bridges campus to grid identity management systems, facilitating seamless access from campuses to the NSF computational centers, observatories and other major research equipment and facilities construction supported projects. *GSI-OpenSSH* provides a single sign-on remote login and file transfer capability using grid security. *CILogon* leverages NSF's considerable previous investment and will ensure the continued community-driven development and support of MyProxy, GridShib, and GSI-OpenSSH in support of NSF funded collaborative research.

Intellectual merit: The security services of the *CILogon* project address critical security needs of scientific and engineering researchers, as evidenced by their wide use. The project team has a proven track record of providing quality software services to a variety of science communities, along with strategic leadership in the area of security. Team members are active participants in the NSF TeraGrid and Open Science Grid projects and are actively engaged with other NSF grid projects including LTER, OOI, NVO, LIGO, and WATERS.

Broader impacts: The *CILogon* project proposes to enhance the NSF CI for research and education by providing robust, well-supported software services that facilitate secure access to CI. These services are used in communities such as magnetic fusion (NFC), climate research (ESG), high-energy physics (WLCG), and computational chemistry (GridChem), as well as in CI provided by TeraGrid, EGEE, Open Science Grid, NERSC and NCSA. The project has a strong collaboration with Internet2 regarding the integration of Shibboleth with grid security (via *GridShib*) to bridge security of higher education campuses to computational grids, broadening the impact of grid infrastructure to the educational community.

B Project Description

The goals of the proposed CILogon project are to maintain and provide critical enhancements for CI security technologies developed at the National Center for Supercomputing Applications (NCSA) that form the foundation for secure cyberinfrastructure used by NSF science and engineering communities and to foster science and engineering by helping additional communities build secure CI on these services. These services are MyProxy, GSI-OpenSSH, and GridShib. MyProxy is a credential management service that provides key usability and integration capabilities [7]. MyProxy is in use by more than 200 grid sites worldwide including key U.S. communities such as the NSF TeraGrid (www.teragrid.org), GridChem (www.gridchem.org) [14], and Open Science Grid (www.opensciencegrid.org). GSI-OpenSSH has over 50 large-scale deployments worldwide (for example, in the NSF LIGO project). GridShib is playing a key role in TeraGrid and other grid communities' exploration of how to effectively integrate campuses with computational grid security.

NCSA leads development and support activities for these important security services. However, all direct funding for these activities ended September 30, 2007. The loss of this funding without replacement has led to a critical loss of mass for grid security activities at NCSA and an effective end to its security contributions to the wider grid community. The only other funding for grid security work at NCSA is from TeraGrid for TeraGrid-specific integration activities and from Open Science Grid (OSG) for PI Basney's participation in the OSG Security Team. In this proposal, we request funding to revive the MyProxy, GSI-OpenSSH, and GridShib projects, so that we can continue to support the ongoing use of this software by NSF science and engineering communities. We have discussed this proposal with an OCI Program Officer who indicated that it was within scope of the STCI program and that other avenues for support will not be forthcoming in the near future.

Given the support requested in this proposal, we expect these technologies to continue to have broad applicability; however we will ensure the focus of our work by collaborating (as described subsequently and in the letters of support) with a handful of large NSF supported research projects such as the Laser Interferometer Gravitational-Wave Observatory (LIGO), Long-term Ecological Research (LTER), Ocean Observatory Infrastructure (OOI), the National Virtual Observatory (NVO), and the proposed WATERS network. Target communities such as these will be selected based on their readiness level to take advantage of adoption of our technologies and to represent a broad spectrum of the NSF science and engineering directorates. In addition, the PIs are already strongly engaged in the Open Science Grid and NSF TeraGrid programs, which will continue to ensure the applicability of the services for the general CI communities as well.

The remainder of the Project Description is organized as follows. In the next section (B.1), we introduce the three technologies in the proposal (MyProxy, GSI-OpenSSH, and GridShib). Then in section B.2 we describe the impact on science and engineering that these technologies have had to-date. Section B.3 then presents our proposed work in detail, with a discussion of the project plan following in Section B.4. We conclude with a summary in Section B.5 and results from prior NSF support in Section B.6.

B.1 Technology Summaries

The technologies in this proposal are exemplars of open source projects. Their source code is freely available (through the dev.globus.org infrastructure supported by the NSF CDIGS

project), they are distributed with open sources licenses certified by the Open Source Initiative (OSI), and they are based on standardized protocols and interfaces ([6], [13], [21], [26], [29], [34]). While NCSA staff have done the bulk of the development and maintenance, there are significant contributions from the broader community (e.g. one-time password support from NERSC [12], Pubcookie support from University of Virginia [25], and high-performance networking support from the Pittsburgh Supercomputing Center [27]). The technologies are included in a range of open source distributions used to support science and engineering, including the Globus Toolkit, the EGEE gLite Middleware, the Virtual Data Toolkit, the Coordinated TeraGrid Software and Services, and Univa Globus Enterprise.

B.1.1 MyProxy

MyProxy (http://myproxy.ncsa.uiuc.edu/) [3] combines a credential repository with an online certificate authority to allow users to securely obtain grid credentials when and where needed for access to secure grid services. MyProxy is mature software that has been used by the grid community since 2000 in projects such as EGEE, EU DataGrid, Earth System Grid [8], FusionGrid [9], LHC Computing Grid, NASA Information Power Grid, NEESgrid, Open Science Grid, and TeraGrid. A substantial effort is required at NCSA to support this community, for testing submitted code patches and integrating them into new software releases, maintaining up-to-date documentation for new features and new uses of the MyProxy software, facilitating discussions on mailing lists, and supporting MyProxy's use with different versions of OpenSSL, the Globus Toolkit, and other MyProxy client implementations, such as the Java CoG MyProxy client, as they are updated. The quality of this effort is demonstrated by the results of a recent "insider, independent, first principles vulnerability assessment of MyProxy" conducted at the University of Wisconsin-Madison, which concluded: "Overall, relatively few vulnerabilities were found in MyProxy and those that were found did not compromise the certificates and their passphrases managed by MyProxy. No design flaws were found in MyProxy." Given MyProxy's critical nature in the security infrastructure, continuing maintaining the software to this high standard is critical.

MyProxy Usage:
TeraGrid: 21,744 requests from 775 users in July 2008.
LCG: 230,000+ requests/day.

B.1.2 GSI-OpenSSH

GSI-OpenSSH (http://grid.ncsa.uiuc.edu/ssh) is an enhanced version of OpenSSH (http://www.openssh.org) that adds support for grid authentication (i.e. X.509 proxy certificates [29] and delegation [30]), providing a single sign-on login and file transfer service for grids. The X.509 support is transparent to users, giving them the look and feel of vanilla OpenSSH to which they are accustomed. Java clients are also available, which can be embedded in web pages, giving users the ability to access grid resources using a standard web browser. GSI-OpenSSH is used on over 50 sites worldwide, including TeraGrid, NEESgrid, LIGO, UK National Grid Service, TIGRE, and the LSC DataGrid.

GSI-OpenSSH is mature technology; however ongoing maintenance is required to keep it up to date with OpenSSH releases and to respond to OpenSSH vulnerabilities. Enhancement requests from communities are also common. Some past examples include: at the request of the NEES CyberInfrastructure Center team at the San Diego Supercomputer Center and the Virtual Data Toolkit (VDT) team at the University of Wisconsin, the GSI-OpenSSH team incorporated modifications for improved network throughput over long and high bandwidth links developed at the Pittsburgh Supercomputing Center (http://www.psc.edu/networking/projects/hpn-ssh/). At

the request of USCMS T2 Nebraska, support for Globus Authorization Callouts [23] to allow integration with PRIMA and GUMS (from the OSG VO Services project) was integrated into GSI-OpenSSH. At the request of LIGO, additional configuration options for controlling delegation policies were added.

B.1.3 GridShib

"GridShib" (http://gridshib.globus.org/) is the common name given the suite of software designed to federate campus and grid identity management. GridShib accomplishes this by providing interoperability between the Shibboleth system and the Globus Toolkit grid computing middleware [16][17].

Shibboleth (http://shibboleth.internet2.edu/) is a security service developed by Internet2 to support identity management on the web by enabling campus users to access web resources (such as digital libraries) outside of their campus using their campus accounts and avoiding their having to obtain additional accounts and passwords. GridShib extends this functionality by allowing those campus users to seamlessly access computational grids. This includes both the conveyance of a user identifier from the campus as well as any relevant attributes from the web world into the grid. Identity federation is relatively new and rapidly developing; effort is required to maintain compatibility with new and evolving identity federation standards such as SAML 2.0 and OpenID that are being adopted by the campus and grid communities.

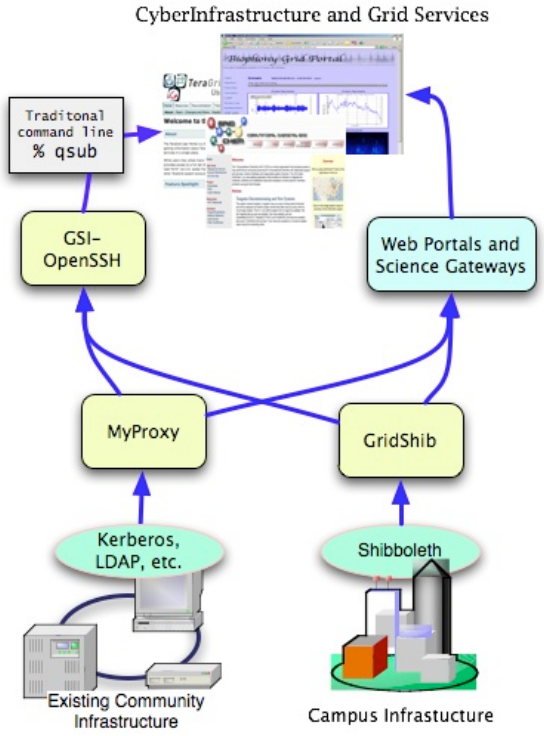

Figure 1: Proposal technologies (MyProxy, GridShib, GSI-OpenSSH) and their role in allowing user access through both web and command line technologies to cyberinfrastructure and the grid.

B.2 Current Science and Engineering Impact of Technologies

The technologies supported by this proposal have a huge positive impact on science and engineering research with their enabling of secure access to cyberinfrastructure across a broad range of projects. For example:

Laser Interferometer Gravitational-Wave Observatory (LIGO): The LIGO Data Grid distributes a terabyte of data from the inferometer sites to computing centers around the world, making it available to the scientific community for analysis. The scientists use GSI-OpenSSH for remote access to this data and LIGO is in the process of deploying MyProxy to support secure, long-running analysis jobs.

LHC Computing Grid: The LHC Computing Grid, which spans over 200 sites and 100,000 CPUs, is using MyProxy in their credential renewal service for long-running jobs [22], serving over 230,000 requests for credentials each day.

Dr. Goasguen at Clemson University: Dr. Goasguen at Clemson University used GridShib software to enable a digital production arts program (http://www.fx.clemson.edu/) to use grid

resources for Maya rendering jobs as part of his cyberinfrastructure class.

Mid-America Earth-quake Center portal (MAEviz): MAEviz provides capabilities to decision makers who allocate resources to respond to earthquakes and other natural disasters. MAEviz federates a number of information sources, including potentially sensitive information from utility companies. This demands strong access control, and MyProxy plays a crucial role in the security architecture, bridging between the existing user database and grid authentication to allow for strong, usable, single sign-on access to the portal, its Java web-start client applications and back end data repository [15].

NSF TeraGrid: MyProxy and GSI-OpenSSH form the core of the TeraGrid's single sign-on solution, offering users a means to access any TeraGrid resource with a single authentication. The GridShib software is a key component of the security architecture supporting TeraGrid science gateways, allowing TeraGrid to scale to orders of magnitude more users than would otherwise be possible with traditional HPC user management schemes while providing comparable security [33].

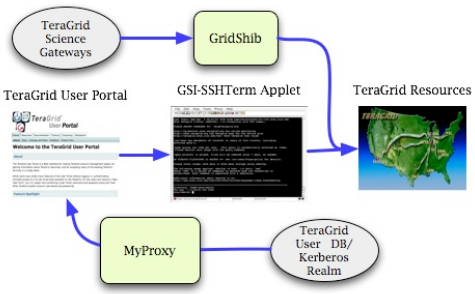

Figure 2: MyProxy and GridShib's role in supporting TeraGrid.

Australian Access Federation (AAF): AAF is developing and deployiong a security infrastructure for collaboration within and between higher education and research institutions in Australia. AAF uses MyProxy to bridge their users from the Australian Shibboleth Federation and their X.509-based grid services.

Earth Systems Grid (ESG): ESG is a distributed data storage and management system that has been running in production mode since 2004. The system is accessed primarily through a web portal interface (http://www.earthsystemgrid.org) by thousands of users across the globe. MyProxy is at the core of its security capabilities.

Long-Term Ecological Research (LTER) Network Office (LNO): The LNO is responsible for infrastructure that integrates the 26 research sites comprising the NSF-funded LTER program. In conjunction with staff on this proposal team, they developed a prototype infrastructure for allowing LTER scientists to easily find and analyze LTER acoustic sensor data, which was based on MyProxy [11].

NERSC: NERSC is a long-time user of MyProxy as well as a collaborator on its development. NERSC uses MyProxy to bridge between their hardware-based one-time password system and their grid services, giving their users ease of use while meeting Department of Energy security standards [12].

The **UK National Grid Service and National Research Council of Canada** developed a Java client that works with the GSI-OpenSSH service and integrates MyProxy authentication (http://sf.net/projects/gsi-sshterm/). This gives scientists a means of accessing the grid that requires no special software on their local computer. TeraGrid subsequently adopted this work thanks to the effort of the proposal team.

B.3 Proposed Work: Facilitating Secure Access to National-Scale CI

We propose the continued development and support of MyProxy, GSI-OpenSSH, and GridShib, driven by community requirements. Building on our experience providing security services to grid projects to-date, we have identified target communities with new requirements

for these services, who we will engage with directly in the proposed program of work, and we have identified three areas in which improvements to the current security services would make a significant impact on these and other science and engineering communities.

In the next Section (B.3.1), we propose a program of targeted outreach to five identified communities during the project period that will serve as drivers for the development of security services and tools for the broader science and engineering research community. We introduce each community and describe unique requirements of that community that drives our work. Then in Section B.3.2 we propose improvements to *campus interoperability*, enabling researchers to access advanced cyberinfrastructure using their existing campus logon, without requiring them to manage additional credentials or overcome other barriers to entry. In Section B.3.3 we propose improvements in the area of *web integration*, making browser-based access to different types of cyberinfrastructure services more seamless. Then in Section B.3.4 we propose *operational security* improvements that will make the cyber security infrastructure more reliable and trustworthy for scientists and address the needs of CI deployers and operators. These improvements, taken together and directed by the community requirements, will continue to move the broader community closer to a seamless authentication infrastructure for CI.

B.3.1 Outreach and Deployment

B.3.1.1 Outreach Strategy

Capturing and understanding the needs of the general scientific and engineering communities represents an important challenge for our proposed project. We discuss in some detail throughout this proposal the kinds of communities that we are targeting to support along with their highest priority needs. We do this as a way to guide our advancements in the sincere hope that the needs of a few targeted communities will be representative of the broader community of researchers, and our work therefore beneficial to the wider set of researchers. As witnessed by our letters of support, and supported by our detailed discussions with several other communities, it is clear that these communities understand their needs and priorities. Over the course of the project lifetime however there is a critical need to remain engaged with the "end user" to once again ensure that our work is acceptable and useful. Such hard learned lessons are derived from previous cyberinfrastructure building efforts including NEESgrid, NSF Alliance, and the LTER Pilot Study.

Specifically our approach to outreach will be to: 1) hold a series of focused workshops targeted at a select set of projects, yet open to all NSF funded "grid projects", 2) implement and support a CILogon collaboration wiki for community support, information dissemination, and discussions, 3) public capability demonstrations, and 4) participation in CI coordination efforts such as the recently emerging Federation of Environmental Observatory Networks (FEON). This section discusses how we propose to effectively engage with a small list of targeted communities as well as mechanisms for disseminating our work to, and gathering requirements from, the broader community of scientific and engineering researchers.

Targeted Workshops: Following a model successfully used by the NMI Grids Center and suggested in [28], we will seek out emerging grid communities through a series of workshops that will have the goals of disseminating our work, soliciting feedback from the broader community, and identifying target communities for in-depth engagement. The first workshop, which will occur in the first four months, will include an initial list of communities that we have already contacted, including participants from the Laser Interferometer Gravitational-Wave Observatory (LIGO), Long-term Ecological Research (LTER), Ocean Observatory Infrastructure

(OOI), the National Virtual Observatory (NVO), and the proposed WATERS Network. In addition to presenting our first year plans, we will invite science and engineering community representatives to present their work and the challenges they face. Through discussions at these workshops we will improve and document our plans, and we will ensure they meet the needs of the research communities.

We will hold a second workshop by the 16th month to re-engage directly with the science and engineering teams to present our work, learn about their experiences, and plan for the remaining year of activity. We will hold both of these meetings at pre-existing events such as the American Geophysical Union (AGU) so that we might draw in additional research communities.

Collaboration Wiki: Throughout the project lifetime we will work to maintain a close working relationship with our target communities. Workshops provide time for focused discussions; however much effective collaboration can also happen between workshops. We propose to tackle this communication challenge by establishing a collaboration mechanism where experiences can be exchanged and discussion can occur. We have found wikis to be a highly effective mechanism to support community collaborations by providing an easy means for content dissemination, modification and discussion. While each of the existing security technologies we propose to advance and support has pre-existing support wikis, we propose a higher-level CILogon wiki that will provide these communities with a place to document and discuss their general security experience. It is our hope that, in addition to providing a venue for the dissemination of our work, these discussion sites will spark the broader dissemination of experience and knowledge between groups.

Public Capability Demonstrations: In our experience, public demonstration provides the research community with real tangible examples of our work and is essential to increasing the likelihood that it will be successfully accepted and adopted. Demonstrations serve to solicit feedback and to build the community's confidence in the implementation. In addition to supporting any demonstrations made by our five targeted communities involving our work, we will present demonstrations of our security tools and services at various technical, science and engineering community gatherings, such as the American Geophysical Union (AGU).

CyberInfrastructure Coordination Efforts: At the time of this proposal we know of one emerging U.S. based effort to coordinate the CI between major CI development and deployment activities. The Federation of Environmental Observatories (FEON) has been proposed to help facilitate the coordination between large NSF-funded environmental observatories. We propose to participate in FEON and similar efforts providing expertise on identity management. Co-PI Butler is a co-author of the whitepaper titled “Shared CyberInfrastructure for Environmental Observatories,” which was a significant contributor to FEON’s formation.

B.3.1.2 Targeted Communities

As previously described, we will select a handful of target communities with whom we will engage deeply to ensure strong impact of our efforts. Early target community candidates, with whom we have already been in discussion, include the Long Term Ecological Research Network Office (LNO), which supports the LTER community, the Ocean Observatory Initiative (OOI), the National Virtual Observatory (NVO), the Laser Gravitational-Wave Observatory (LIGO), and the proposed WATERS Network. Several of these communities have already provided us letters of commitment included with this proposal. In the remainder of this section, we discuss each of these communities and their needs.

WATERS Network

The WATERS Network is a relatively new effort that is still in the early planning phases. It is a joint collaboration between a number of environmental and social science research communities and is funded by the NSF Engineering and Geosciences Directorates. The WATERS Network aims to construct a national-scale network of observational and experimental facilities focused on water related science and engineering including systematic water measurements, data storage and curation, modeling and visualization

The WATERS Network represents one of the new emerging CI development projects offering a great opportunity to collaborate with that team early in the design and development of the security infrastructure rather than the more costly and risky approach of trying to later retrofit security solutions into place. We expect the WATERS Network to drive requirements for identity management and federating technologies as a result of their plans to integrate a vast collection of legacy instruments and new WATERS Network instruments that are widely distributed across the U.S. and administered by a broad range of authorities from local, state, and federal agencies.

LTER

The Long Term Ecological Research (LTER) community consists of more than 1,500 scientists and students who are working at 26 different sites distributed from the Arctic tundra to Antarctica and from French Polynesia to Puerto Rico, including ecoregions throughout the United States. The LTER program has been responsible for the discovery of long-term ecological patterns and processes. While LTER represents 26 individual sites, recent plans and work are drawing towards an interconnected national grid. The past autonomy and often very different technical approaches to the CI found at each site has created significant challenges with regard to the design and implementation of network and service level security.

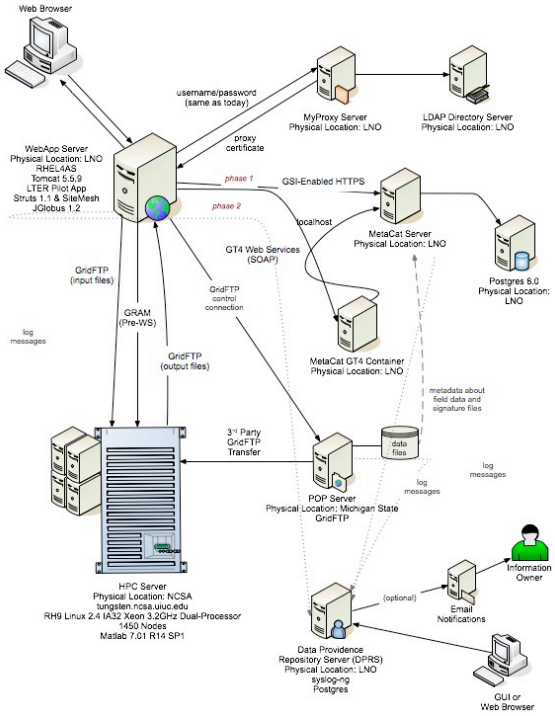

Figure 3: LTER Pilot Study, created by LNO and NCSA, built on MyProxy.

The LTER Network Office (LNO) supports the LTER CI and is responsible for overall LTER CI security. As documented in their letter of support, the foremost need for the LTER community is the continued support of the MyProxy credential management service and the Shibboleth federated identity service. LNO provides storage of and access to data that may be sensitive or proprietary and hence needs an authentication service that federates the distributed LTER sites and even broader LTER user community. This data service will require a flexible and robust credential service, such as MyProxy, that supports short-lived X.509 credentials for both authentication and authorization. The LTER Pilot Study [11] demonstrated the benefits of MyProxy integration across the LTER sites. Ultimately, access to

LTER data services will be best addressed through additional integration of federated identity management services, like Shibboleth and/or OpenID protocols, which will allow the broader community easy, authenticated access.

NVO

NVO (http://www.us-vo.org/) is the National Virtual Observatory, a US-based Virtual Observatory project that is part of the larger International Virtual Observatory Alliance (IVOA). Their goal is to significantly improve the ability of astronomical researchers to find, retrieve, and analyze astronomical data from ground- and space-based telescopes worldwide.

NVO will open up access from what has historically been small, carefully selected samples of objects in one or a few wavelength bands, to the use of multi-wavelength data for millions to billions of objects. NVO will provide simultaneous access to multi-wavelength archives and advanced visualization and statistical analysis tools to researchers across the nation from a variety of institutions.

NVO has already been working on a prototype authentication infrastructure based on MyProxy and Pubcookie (http://www.pubcookie.org/). Their identified high priority security needs include the need to support two-factor authentication, the replacement of Pubcookie with a standardized mechanism such as OpenID or Shibboleth, achieving credential policy compliance with the International Grid Trust Federation (IGTF) for uses beyond the NVO project, and the addition of an authentication strength attribute in credentials.

OOI

The Ocean Observatory Initiative is set to transform ocean science research by building an international scale CI that supports researchers' access to multiple ocean observatories, simulations, analysis services and other researchers, for long-term and adaptive measurements in the oceans. As these efforts mature, the research-focused observatories enabled by the OOI will be networked, becoming an integral part of the proposed Integrated and Sustained Ocean Observing System (http://www.ocean.us). IOOS is an operationally focused national system and in turn will be a key and enabling U.S. contribution to the international Global Ocean Observing System and the Global Earth Observing System of Systems (http://www.earthobservations.org).

The proposers are actively participating in the design of the security architecture for the OOI CyberInfrastructure. The OOI project has identified a strong requirement for federated access to OOI CI, using Shibboleth or similar technologies. OOI will support access via web browsers and rich clients to instrument data. The project team will explore whether the GridShib and MyProxy technologies can meet this need for OOI.

LIGO

The Laser Interferometer Gravitational-Wave Observatory (LIGO) was designed for the detection of cosmic gravitational waves by searching for ripples in space-time. LIGO supports more than 550 scientists from more than 40 institutions worldwide, who conduct scientific research to analyze and interpret the observational data produced by LIGO. The LIGO Data Grid (LDG) distributes a terabyte of data collected each day at the interferometer sites to computing centers around the world. LDG is already a heavy user of GSI-OpenSSH, which it uses to provide LIGO scientists with secure single sign-on access to all LDG computational resources. In their distributed environment, GSI-OpenSSH greatly reduces the administrative burden for providing this access.

LIGO is the most advanced of our target communities in terms of identifying their security needs, having completed an internal planning project. One future need that they have identified is authentication support for long-running analysis workflows, and they intend to integrate and deploy MyProxy servers at each LDR analysis site to work in conjunction with Condor-G [18] credential renewal support for these workflows. Another future need is better integration between web and grid security; while both use the same underlying X.509 technology, in practice client applications manage credentials very differently, requiring unreasonable expertise by users to shift between web browser and grid clients. LIGO will also be deploying Shibboleth for managing access to web resources and will be evaluating GridShib as a mechanism to give their Shibboleth users seamless access to the grid services in LDG. More details on all these needs may be found in the LIGO's letter of support written for this proposal.

B.3.1.3 General Community Support

As described in the introduction, the goal of our project is to maintain and continue to grow in response to community requirements, critical cyberservices (MyProxy, GSI-OpenSSH and GridShib) developed at NCSA under NSF funding. The reimplementation of existing services, using a different technology or a new standard or research result, is often counterproductive to the goal of providing stable CyberInfrastructure upon which scientific communities can rely. Hence our efforts are focused on existing services by providing maintenance, integration and interoperability, increased usability and operational readiness, and outreach and deployment support. We propose the following activities as key to maintenance and operational readiness:

- *Keeping up to date with new versions of underlying software.* The key to success of all the technologies in this proposal is their interoperability with software of importance to science and engineering communities, namely the Globus Toolkit, OpenSSH and Shibboleth. These underlying software packages commonly produce updates to correct issues and add features. These updates require testing and documentation updates by NCSA to its technologies to ensure continued compatibility. OpenSSH has historically had 2-3 releases per year, each requiring updates to the GSI-OpenSSH patches. Globus Toolkit minor releases are typically slightly less frequent but require testing and documentation updates for MyProxy, GridShib and GSI-OpenSSH. There have been 11 released patches to the Shibboleth 1.3 components (Idp: 1.3[b,c], 1.3.1, 1.3.2, 1.3.3 SP: 1.3[b-f], 1.3.1) since its initial release in mid-2005, requiring testing and documentation updates for GridShib. There have been new major releases of both Shibboleth (version 2.0) and the Globus Toolkit (version 4.2) in 2008, both of which require more substantial efforts to support. We note that without funding, despite best efforts, our ability to maintain these critical technologies in synch with these new major releases is already lagging.
- *Responding to security vulnerabilities both within the cybersecurity services themselves and the underlying software.* While they are thankfully rare occurrences, these vulnerabilities require rapid response from team members. Two OpenSSH security advisories have impacted GSI-OpenSSH in recent years (CVE-2006-5051 and CVE-2008-1483). Two Globus Toolkit advisories impacted MyProxy in 2006. Both advisories (2006-1 and 2006-2) related to insecure temporary file handling and their report provoked a code review of MyProxy and GSI-OpenSSH, with Advisory 2006-2 requiring a new MyProxy release. There were two OpenSSL security advisories in 2006 and while neither required changes to any of the NCSA services, this determination required careful consideration and scrutiny of the code. A Shibboleth vulnerability in 2006 required an update to GridShib deployments. In

addition, many vulnerabilities were discovered in underlying software that do not affect the NCSA security technologies, but still required substantial effort to analyze, determine benign and disseminate that information to the communities to assure them.

- *Responding to user queries and bug reports.* Providing timely responses to user questions and bug reports is essential for fostering and sustaining adoption. The following statistics illustrate these ongoing support activities. There have been 575 posts to the MyProxy user email lists since May 2002, and 123 reports submitted to the project's bug tracking systems since April 2003. There have been 389 posts to the GSI-OpenSSH user email lists since May 2001, and 37 bugs submitted to the project's bug tracking systems since March 2003. As a component in other software distributions, MyProxy and GSI-OpenSSH are frequently discussed on other mailing lists (for example, gt-user@globus.org), and the NCSA team assists with problem reports submitted to the Globus Toolkit, EGEE, and VDT bug tracking systems. Additionally, there have been 621 posts to the GridShib user mailing list since July 2006 and 261 entries in the GridShib bug tracking system. NCSA staff take responsibility for monitoring these mailing lists and providing assistance as required, in addition to investigating and resolving reported bugs.
- *Core Globus security maintenance and vulnerability handling.* NCSA's development of key cybersecurity services makes them significant contributors to Globus Toolkit security through their contributions of code, expertise in design and architectural issues and leadership in handling vulnerabilities.

B.3.1.4 Metrics for Outreach and Deployment

To effectively measure increased usage of these cybersecurity services as a result of our outreach efforts, we will enhance the ability to gather usage metrics for the services in the project through the integration of dev.Globus Metrics project libraries [19], which provide a standard metrics gathering capability used across Globus Toolkit components. We will use this mechanism to collect detailed usage statistics for our services to be included in project reports. This mechanism will also allow deploying grids to collect statistics about usage inside their grids for their own reporting purposes. Usage statistics for these logon services can be particularly valuable for reporting on the number of unique users of a grid's services. We note that a deployer can, if they wish to preserve privacy, disable this functionality, hence preventing the inclusion of their usage in our statistics collection.

B.3.2 Campus Interoperability

The GridShib project has developed and demonstrated mechanisms for bridging campus authentication with grid services to provide a seamless logon experience for the scientist. In discussions with our target community representatives, this capability has foremost interest. In addition to ongoing support for existing capabilities, there are the following areas of improvement needed to meet the requirements that these grids have for campus interoperability.

B.3.2.1 Shibboleth Version 2 Support

With the release of Shibboleth version 2.0 in early 2008, campuses are beginning to deploy federated security services based on the new SAML 2.0 standard. GridShib currently supports SAML 1.1 and Shibboleth version 1.3. Adding support for Shibboleth 2.0 and SAML 2.0 will allow GridShib to continue to provide a bridge between campus and grid authentication as campuses adopt the new software and standard. GridShib will be updated to use the new OpenSAML libraries and generate and consume the new SAML assertion formats. In addition to

campus interoperability, this will allow GridShib to interoperate with emerging standards in the Open Grid Forum based on SAML 2.0 being adopted by the widely used Virtual Organization Management Service (VOMS) [1]. This enhancement will also include support for SAML 2.0 metadata to facilitate interoperability between larger grid deployments and federations such as InCommon and an interface to XACML V2.0 leveraging the Globus implementation of "SAML V2.0 Profile for XACML V2.0". As discussed subsequently in the section on Web Integration, this effort will go hand-in-hand with adding support for OpenID as an emerging complementary protocol to SAML.

B.3.2.2 GridShib Support for Globus Toolkit C Services

GridShib provides authorization plug-ins for Globus Toolkit Java services that grant access based on SAML assertions containing campus and community credentials. However, grids today also deploy Globus Toolkit C services, including GSI-OpenSSH and GridFTP, which would also benefit from the SAML authorization plug-ins provided by GridShib. We propose to develop these plug-ins for C services, so that a researcher can use his or her campus login to open a remote SSH login session via GSI-OpenSSH and transfer large datasets efficiently via GridFTP. LIGO is one example of a heavy user of these C services, and LIGO has requested the development of this functionality that will also benefit other grids.

B.3.3 Web Integration

Our target community representatives have also identified a need for improved integration of web browser authentication with grid mechanisms, which we plan to address by developing a MyProxy browser interface and by adding OpenID support to GridShib.

B.3.3.1 MyProxy Browser Interface

As discussed previously when we described LIGO's needs, a significant usability hurdle for providing both web browser and thick client interfaces to CI services is management of credentials. When credentials are created inside the web browser, they are difficult to export for use by non-browser-based applications. Likewise, credentials created outside the browser are difficult to import for use inside the browser. Our prior work with MyProxy and Java Web Start for MaeVIZ [15] provides one solution for migrating credentials from a browser session to a thick client session. We propose to perform complementary work, to make it easier to migrate credentials into a browser session, by providing a MyProxy interface for web browsers. We will develop a MyProxy *web portal* that bridges between the browser and the MyProxy server. This will allow researchers to logon to MyProxy via the portal and obtain X.509 credentials from MyProxy, which will be stored in the browser, so the researcher can then authenticate via HTTP/S to other secure web sites using those credentials. We will leverage our GridShib work to provide for campus-based authentication using Shibboleth and/or OpenID. We will work with grid communities to deploy these mechanisms to integrate browser and grid security.

B.3.3.2 GridShib OpenID Support

OpenID (http://openid.net/) is a lightweight browser single sign-on protocol that is gaining wide adoption across the web, particularly among commercial service providers (e.g. Google, Yahoo, AOL). In the grid arena, the DOE Earth System Grid has decided to adopt OpenID and NVO is also considering use of OpenID. A comparison between OpenID and Shibboleth/SAML is outside the scope of this proposal – we expect both technologies to play an important role in future NSF CI. GridShib currently leverages the Shibboleth Apache module for front-end

authentication. We propose to add support for the OpenID Apache module to GridShib, providing an OpenID to X.509 bridge for communities adopting OpenID.

B.3.4 Operational Security

The following operational security improvements are needed to improve the quality of service provided by cybersecurity services to researchers as well as provide controls requested by security staff responsible for the operation of CI deployments. Avoiding credential expiration and server downtime, as well as providing improved facilities for responding to credential compromise, will improve the reliability of the cyberinfrastructure, allowing scientists and engineers to work more effectively and avoiding the frustration caused by service failures.

B.3.4.1 Credential Renewal

In grids, users obtain short-lived session credentials, with a typical lifetime of 12 hours. Short-lived credentials provide valuable security properties, but in practice it is often difficult to anticipate the length of time a credential will be needed, due to the unpredictable behavior of compute processes and the surrounding computing environment. Short-lived credentials mitigate the risk of their theft and misuse by a malicious party, but grids often must support compute and data movement jobs lasting days or weeks that require credentials throughout their runs. An example of this is the previously discussed analysis workflows in LIGO.

Basney collaborated with Daniel Kouril (CESNET) to develop a credential renewal mechanism based on MyProxy that has been widely used for European grids (EDG, EGEE, and LCG) [22]. We propose to continue their collaboration to extend this work to create a general-purpose proxy credential renewal mechanism for the Globus Toolkit (GRAM, WS-GRAM, RFT) and EGEE's gLite middleware. (Kouril's work will be separately funded by CESNET.)

B.3.4.2 Delegation Auditability

An issue that has arisen in our discussions with TeraGrid and OSG operators, is that in the event of a compromised resource, these CI operators want the ability to reject any credentials that passed through the compromised resource and may have been compromised in the process (i.e., the intruder may have gained a copy of the credential and be using it inappropriately). We will enhance the credential issuing services (MyProxy, GridShib) to annotate credentials they create with both the host where they were created and the host to which they were delivered (i.e. the host from which the user ran the client to request the credential), as well as an indication of the strength of authentication used to obtain the credential (as requested by NVO). We will also contribute enhancements to core GSI security code to implement this same logging in other grid applications (e.g. GRAM), and produce a specification to allow the larger open source grid community to leverage the solution.

B.3.4.3 MyProxy Server Reliability

MyProxy servers are a critical component of a grid's security infrastructure. If the MyProxy service is unavailable, users are unable to obtain credentials for their sessions and long-running jobs are unable to renew their credentials before they expire. This is a concern expressed by both LIGO and TeraGrid as they see themselves becoming increasingly reliant on MyProxy.

MyProxy currently supports a primary-backup passive replication scheme (developed as part of the Dependable Grids ITR project) where the backup can provide limited service if the primary is unavailable. However, this scheme does not provide fully automated fail-over service and does not support replication for load balancing. We will develop a new peer-to-peer active

replication capability for MyProxy servers to provide the enhanced reliability and performance required by next generation grids.

B.3.4.4 MyProxy Updates for IGTF Compliance

One of MyProxy's capabilities is to issue credentials based on authentication against an existing authentication system to bridge that system to grid services (such as in TeraGrid and the LTER Pilot system we described previously). This functionality is commonly referred to as the "MyProxy CA." The MyProxy CA has been developed to meet the requirements of the Short Lived Credential Services X.509 Public Key Certification Authorities Profile of The Americas Grid Policy Management Authority (http://www.tagpma.org), a member of the International Grid Trust Federation (http://www.igtf.net). The IGTF facilitates acceptance of certificates by grids worldwide. NCSA operates a MyProxy CA for TeraGrid that has been accredited under this Profile, and both NERSC and PSC are planning to submit their own MyProxy CAs for IGTF accreditation.

One of the requirements of the Profile is use of FIPS 140-2 level 3 Hardware Security Modules (HSMs) [24] for protection of the CA's private key. The UK National Grid Service contributed modifications to MyProxy software to support HSMs. While the HSM support is now operational for both the UK NGS and NCSA, it has not been documented, tested and made robust to the level of other MyProxy functionality. We propose to develop regression tests for MyProxy's HSM support, using multiple HSMs already available to the MyProxy project at NCSA, and to resolve reported problems with MyProxy's HSM support that currently require work-arounds.

The TAGPMA, a federation of authentication providers and relying parties, recently updated the Short Lived Credential Service (SLCS) Profile to require CAs to issue Certificate Revocation Lists (CRLs). Previously, the Profile did not require CRLs from issuers of short-lived certificates. Currently, the MyProxy CA does not support issuance and management of CRLs. We propose adding CRL support, leveraging the CRL functionality provided by OpenSSL, upon which MyProxy builds.

B.4 Project Plan

B.4.1 Project Team

Our project team is composed of expert developers of grid security services who have existing relationships with a number of major grid projects that will benefit from the proposed efforts. Basney, Butler and Welch have a long history of involvement in CI deployment and grid security. Basney has led the MyProxy and GSI-OpenSSH projects since 2001 and currently manages GridShib activities in TeraGrid. As project PI, Basney will be responsible for day-to-day project management and technical directions. Welch led the initial development of MyProxy, GSI-OpenSSH, and GridShib and is a security architect for the Globus Toolkit. He will provide technical advice and assist with outreach for the proposed project. Butler is on the project team specifically to lead outreach activities to new communities as described in detail in Section B.3.1.1. His experience as a PI for the NSF NMI Grids Center and member of the NSF NEESgrid project [28] and his current connections to a number of communities (WATERS, ORION, NEON, LTER) demonstrate his expertise in such activities.

The project members already have strong interactions with TeraGrid and Open Science Grid that ensure a deep understanding of community requirements and effective vehicles for deployments of results. Internet2, Open Science Grid, and TeraGrid are collaborating on a

Campus Partnership to bring a usable and deployable cyberinfrastructure to a large number of University campuses. We will work with the support groups of the client and user libraries from the different software infrastructures (VDT, CTSS) to provide a uniform interface for university users and administrators. Basney and Welch are currently spearheading TeraGrid work on integrating with campus identity management, which we expect will generate additional requirements for our project.

B.4.2 Success Metrics

The software services that form the foundation of our proposed work currently exist and are used by a number of production grids on a regular basis. Hence, we will judge our success based on improvements to these services and on increased adoption due to our enhancements. The metrics to capture these factors are:

- Utilization. As we described in Section B.3.1.4, we will integrate the dev.Globus metrics code into our software services in order to allow for collection of data on actual usage. This will allow us to measure increased usage over time.
- Adoption. Adoption of our results by new communities represents not only an increase in the amount of total usage, but also a fundamental increased benefit to science, because an entire new community is being served. Reported metrics will show if the community was engaged through the workshops or other activities described in Section B.3.1.1.

B.4.3 Milestones and Deliverables

We organize our work plan into biannual milestones and deliverables as follows:

Month 6	Month 12
First workshop held and report produced.	GridShib C support completed.
12 month project plan produced.	GridShib CA OpenID support completed.
Collaboration Wiki established.	MyProxy peer-to-peer replication completed.
Integration of dev.globus metrics completed.	Delegation auditability completed.
MyProxy CRL issuance completed.	
Month 18	**Month 24**
Second workshop held and report produced.	Shibboleth version 2 support complete.
Project plan for final 6 months produced.	Credential renewal service complete.
MyProxy browser interface completed.	MyProxy HSM support improvements.

B.5 Closing Project Summary

CyberInfrastructure is quickly moving away from a subject of research and becoming a critical supporting element of all scientific and engineering research. In the last several years we have seen the emergence of large collaborative teams of researchers complete with sensors and instruments of every kind, as well as a large collection of computational and data resources, all widely distributed both geographically and administratively. CI is the glue that holds these collaborations together and at the very foundation of that CI are the security services [10]. These same security services that provide the trust necessary for collaborations of discovery are also crucial to the resource service providers to ensure the protection of our nation's computational resources. As we continue to move forward with the development of cyberinfrastructure, we increasingly take these underlying services for granted; however, that does not lessen the

dependencies placed upon them by scientists and engineers to protect their intellectual property, to safeguard their collaborations, and to ensure the credibility of their work.

MyProxy, GridShib and GSI-OpenSSH are key blocks in the foundation of our national research cyberinfrastructure. Each of these technologies lives in the shadows, as they should, providing the research community with the ability to collaborate securely. MyProxy is in use by over 200 sites today around the world, and all three technologies are widely distributed via their inclusion in the Globus Toolkit, VDT and other community CI packages and have had a huge impact on science and engineering research.

We have proposed to maintain and provide enhancements for these critical CI security technologies developed at the National Center for Supercomputing Applications and aggressively reach out to NSF-funded science and engineering projects such as LIGO, LTER, OOI, NVO, WATERS Network, TeraGrid, and others by helping these communities build secure CI. The MyProxy, GridShib, and GSI-OpenSSH security technologies were developed through great effort by the NSF community and represent a major success story for NSF by supporting science and engineering research throughout the world.

The project team members each have extensive experience developing and successfully deploying grid security technologies and are regularly called upon to assist in the design of security systems for large NSF collaborative projects. The proposal lays out a balanced approach that includes the fundamental support for these technologies and the advancement of new services and capabilities through focused engagements with scientific and engineering communities who are leading the way towards the development, deployment and utilization of a pervasive national cyberinfrastructure. Finally, the proposal lays out a dissemination and adoption plan for the broader research community. Funds to support continued maintenance and advancements of MyProxy, GridShib and GSI-OpenSSH are depleted and without additional funding they will no longer be supported.

B.6 Results from Prior NSF Support

James Basney: "Integration of the MyProxy Online Credential Repository into the NSF Middleware Initiative Software Infrastructure", award number 0222571, $598,343, 7/1/02–6/30/05, co-PI: Marty Humphrey, UVa. This grant supported the inclusion of MyProxy in the NSF Middleware Initiative software distribution. This grant also supported the creation of an OGSI-compliant MyProxy and a research prototype hardware-secured MyProxy. Publications supported include [3][4][5][20][22][24].

Randal Butler: "NMI: Disseminating and Supporting Middleware Infrastructure: Engaging and Expanding Scientific Grid Communities", award number 0330670, $1,800,000, 10/1/03-8/30/07. This grant funded the GRIDS Center, which designed, developed, deployed and supported a set of reusable and expandable middleware functions and services with emphasis on outreach to scientific communities. Publications supported include [11][28].

Von Welch: "SCI: Collaborative Research: NMI DEVELOPMENT: Policy Controlled Attribute Framework", award number 0438424, $396,060, 12/1/04–11/30/07, collaborative with award number 0438385: Katarzyna Keahey, U Chicago, $595,514. This grant funded initial development of the GridShib software at NCSA and the University of Chicago. Publications supported include [2][31][32].

C References Cited

[1] Alfieri, R., R. Cecchini, V. Ciaschini, L. dell'Agnello, A. Frohner, A. Gianoli, K. Lorentey, and F. Spataro, "VOMS, an Authorization System for Virtual Organizations," in Grid Computing: First European Across Grids Conference, 2004.

[2] Barton, T., 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 5th Annual PKI R&D Workshop, Apr. 2006.

[3] Basney, J., M. Humphrey, and V. Welch, "The MyProxy Online Credential Repository," Software: Practice and Experience, 35(9):801–816, July 2005.

[4] Basney, J., S. S. Chetan, F. Qin, S. Song, X. Tu, and M. Humphrey, "An OGSI Credential Manager Service," In Proceedings of the Workshop on Grid Security Practice and Experience (UK e-Science Security Task Force), July 2004.

[5] Basney, J., W. Yurcik, R. Bonilla, and A. Slagell, "The Credential Wallet: A Classification of Credential Repositories Highlighting MyProxy," In Proceedings of the 31st Research Conference on Communication, Information and Internet Policy (TPRC 2003), September 2003.

[6] Basney, J., "MyProxy Protocol," Global Grid Forum Experimental Document GFD-E.54, November 26, 2005.

[7] Beckles, Bruce, Von Welch, and Jim Basney, "Mechanisms for Increasing the Usability of Grid Security," International Journal of Human Computer Studies, Special Issue on HCI Research in Privacy and Security, Volume 63, Issues 1-2, July 2005, pages 74-101.

[8] Bernholdt, D., S. Bharathi, D. Brown, K. Chancio, M. Chen, A. Chervenak, L. Cinquini, B. Drach, I. Foster, P. Fox, J. Garcia, C. Kesselman, R. Markel, D. Middleton, V. Nefedova, L. Pouchard, A. Shoshani, A. Sim, G. Strand, D. Williams, "The Earth System Grid: Supporting the Next Generation of Climate Modeling Research," Proceedings of the IEEE, 93:3, March 2005, 485-495.

[9] Burruss, J. R., T. W. Fredian, M. R. Thompson, "Simplifying FusionGrid Security," Proc. Challenges of Large Applications in Distributed Environments (CLADE), 95-105, (2005)

[10] Butler, R., D. Engert, I. Foster, C. Kesselman, S. Tuecke, J. Volmer, and V. Welch, "A national-scale authentication infrastructure," IEEE Computer, 33(12):60–66, 2000.

[11] Butler, R., M. Servilla, S. Gage, J. Basney, V. Welch, B. Baker, T. Fleury, P. Duda, D. Gehrig, M. Bletzinger, J. Tao, D. M. Freemon, "CyberInfrastructure for the Analysis of Ecological Acoustic Sensor Data: A Use Case Study in Grid Deployment," Challenges of Large Applications in Distributed Environments (CLADE 2006) Workshop (associated with the 15th International Symposium on High Performance Distributed Computing), Paris, France, June 19, 2006.

[12] Chan, S. and M. Andrews, "Simplifying Public Key Credential Management Through Online Certificate Authorities and PAM," 5th Annual PKI R&D Workshop, April 2006.

[13] Cooper, D., 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 (Standards Track), May 2008.

[14] Dooley, R., G. Allen, and S. Pamidighantam. "Computational Chemistry Grid: Production CyberInfrastructure for Computational Chemistry". Proceedings of the 13th Annual Mardi Gras Conference, Baton Rouge, LA, Feb. 2005, pg 83.

[15] Fleury, T., J. Basney, and V. Welch, "Single Sign-On for Java Web Start Applications Using MyProxy," Proceedings of the ACM Workshop on Secure Web Services (associated with the 13th ACM Conference on Computer and Communications Security), November 3, 2006.

[16] Foster, I. and C. Kesselman, "Globus: a metacomputing infrastructure toolkit," International Journal of Supercomputer Application, 11(2):115–128, 1997.

[17] Foster, I., "Globus Toolkit Version 4: Software for Service-Oriented Systems," In IFIP International Conference on Network and Parallel Computing, pages 2–13. Springer-Verlag LNCS 3779, 2006.

[18] Frey, James, Todd Tannenbaum, Ian Foster, Miron Livny, and Steven Tuecke, "Condor-G: A Computation Management Agent for Multi-Institutional Grids", Journal of Cluster Computing volume 5, pages 237-246, 2002.

[19] Globus Metrics Incubator Project. http://dev.globus.org/wiki/Incubator/Metrics. December 2006.

[20] Humphrey, M., J. Basney, and J. Jokl, "The Case for using Bridge Certificate Authorities for Grid Computing," Software: Practice and Experience, 35(9):817–826, July 2005.

[21] Hutzelman, J., J. Salowey, J. Galbraith, and V. Welch, "Generic Security Service Application Program Interface (GSS-API) Authentication and Key Exchange for the Secure Shell (SSH) Protocol," IETF RFC 4462 (Standards Track), May 2006.

[22] Kouril, D. and J. Basney, "A Credential Renewal Service for Long-Running Jobs," In Proceedings of the 6th IEEE/ACM International Workshop on Grid Computing (Grid 2005), November 2005.

[23] Lang, Bo, Ian Foster, Frank Siebenlist, Rachana Ananthakrishnan, Tim Freeman, "A Multipolicy Authorization Framework for Grid Security," Proc. Fifth IEEE Symposium on Network Computing and Application, Cambridge, USA, July 24-26, 2006.

[24] Lorch, M., J. Basney, and D. Kafura, "A Hardware-secured Credential Repository for Grid PKIs," In Proceedings of the 4th IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGrid2004), April 2004.

[25] Martin, J., J. Basney, and M. Humphrey, "Extending Existing Campus Trust Relationships to the Grid through the Integration of Pubcookie and MyProxy," 2005 International Conference on Computational Science (ICCS 2005), Emory University, Atlanta, GA, May 22-25, 2005.

[26] Maler, E. et al., "Assertions and Protocols for the OASIS Security Assertion Markup Language (SAML) V1.1," OASIS, September 2003.

[27] Rapier, C. and B. Bennett, "High Speed Bulk Data Transfer Using the SSH Protocol," In Mardi Gras Conference, February 2008.

[28] Spencer, B., Butler, R., Ricker, K., Marcusiu, D., Finholt, T., Foster, I., Kesselman, C. "Cyberenvironment Project Management: Lessons Learned." (September 2006). http://neesgrid.ncsa.uiuc.edu/documents/CPMLL.pdf

[29] Tuecke, S., V. Welch, D. Engert, L. Pearlman, and M. Thompson, "Internet X.509 Public Key Infrastructure (PKI) proxy certificate profile," IETF RFC 3820 (Standards Track), June 2004.

[30] Welch, V., I. Foster, C. Kesselman, O. Mulmo, L. Pearlman, S. Tuecke, J. Gawor, S. Meder, and F. Siebenlist, "X.509 proxy certificates for dynamic delegation," In Proceedings of the 3rd Annual PKI R&D Workshop, April 2004.

[31] Welch, V., R. Ananthakrishnan, F. Siebenlist, D. Chadwick, S. Meder, and L. Pearlman, "Use of SAML for OGSI Authorization," Global Grid Forum Experimental Document GFD-E.66, March 2006.

[32] Welch, V., T. Barton, K. Keahey, and F. Siebenlist, "Attributes, Anonymity, and Access: Shibboleth and Globus Integration to Facilitate Grid Collaboration," In 4th Annual PKI R&D Workshop, Apr. 2005.

[33] Welch, V., 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, Volume 19, Issue 6, March 2007.

[34] Ylonen, T., and C. Lonvick, "The Secure Shell (SSH) Protocol Architecture," IETF RFC 4251 (Standards Track), January 2006.