

**Annual Report for Period:**12/2005 - 11/2006**Submitted on:** 12/21/2006**Principal Investigator:** Welch, Von S.**Award ID:** 0438424**Organization:** U of Ill Urbana-Champaign**Title:**

SCI: Collaborative Research: NMI DEVELOPMENT: Policy Controlled Attribute Framework

**Project Participants****Senior Personnel****Name:** Welch, Von**Worked for more than 160 Hours:** Yes**Contribution to Project:****Name:** Barton, Thomas**Worked for more than 160 Hours:** Yes**Contribution to Project:****Name:** Keahey, Katarzyna**Worked for more than 160 Hours:** Yes**Contribution to Project:****Name:** Siebenlist, Frank**Worked for more than 160 Hours:** Yes**Contribution to Project:****Post-doc****Graduate Student****Undergraduate Student****Technician, Programmer****Name:** Champion, David**Worked for more than 160 Hours:** Yes**Contribution to Project:****Name:** Freeman, Timothy**Worked for more than 160 Hours:** Yes**Contribution to Project:****Name:** Scavo, Thomas**Worked for more than 160 Hours:** Yes**Contribution to Project:****Name:** Ananthakrishnan, Rachana**Worked for more than 160 Hours:** Yes**Contribution to Project:****Name:** Kettimuthu, Raj

**Worked for more than 160 Hours:** Yes

**Contribution to Project:**

Software development

## Other Participant

## Research Experience for Undergraduates

### Organizational Partners

#### **Argonne National Laboratory**

Argonne National Laboratory contributed the following to the GridShib Project: web server, CVS server, bugzilla server, mailing list

#### **Ohio State University**

Ohio State University contributed the following to the GridShib Project: wiki web server

### Other Collaborators or Contacts

Internet2 Shibboleth Project: Steven Carmody, RL 'Bob' Morgan, Scott Cantor, Wlater Hoehn, Chad La Joie, Howard Gilbert

MyProxy Project: Jim Basney, Bill Baker

eSP-grid: Mark Norman

LionShare: Mike Halm, Derek Morr, Alex Valentine, Mark Earnest, Loren Metzger

eduSource Communication Layer: Marek Hatala, Timmy Eap, Ashok Shah

Condor-Shib: Steve Moore, Arnie Miles

myVocs: Jill Gemmill, John-Paul Robinson

nanoHub: Sebastien Goasguen

SWITCH: Christoph Witzig

MAMS: Erik Vullings

### Activities and Findings

#### **Research and Education Activities: (See PDF version submitted by PI at the end of the report)**

The goal of the GridShib project is to allow computational grids to leverage existing campus identity management systems (i.e., user logins and attribute information). The Shibboleth Software, developed by Internet2, is a broadly deployed system for higher education campuses that allows intra-campus and interorganizational access (authentication and authorization) to web resources. GridShib leverages Shibboleth as a standard interface for authentication of campus users.

Achieving this goal means meeting a series of technical and policy objectives. The technical objectives are mainly concerned with the translation of security assertions as generated by Shibboleth into a format understood by Grids (our project focuses on Globus Toolkit). In technical terms, this means translating from SAML into X.509. We have been engaged in software development for performing this translation, which includes extensions to the Globus Toolkit (<http://www.globus.org>) and Shibboleth (<http://shibboleth.internet2.edu/>) as well as stand-alone gateways (e.g. the GridShib-CA).

Building on the completion of the technical objectives, a number of policy issues must be addressed. For example, what are the requirements for user authentication at a campus in order for it to be useful for grids, how should the user identity be translated from the campus to the grid,

what are grid requirements for auditing on the campus. These issues must be resolved through community engagement, driven by available solutions of the technical issues, i.e. working software, and then ultimately solved by that working software. While some of these policy issues can be decided through deep thought and agreement, some will take deployment, experimentation and research in an operational context to determine the best solution.

Our community engagements include a collaboration with the nanoHub project resulting in a trail deployment of our software in order to provide nanoHub attributes to grid Services to allow selective servicing of nanoHub users by grid resource providers.

We have also collaborated with the NSF TeraGrid project to develop a strategy for Grid-Campus federated identity management. This collaboration has resulted in a jointly authored white paper (<http://gridshib.globus.org/tg-paper.html>) and plans for testbed with GridShib software in early 2007. Subsequently Internet2 has joined in supporting this work, resulting in a session at the recent Internet2 Fall Member Meeting. This collaboration will be the focus of the remainder of our project (remaining funding should allow us to continue until Spring of 2007), with us leading deployment activities.

In addition we have had an ongoing technical collaboration and demonstration of interoperability with NMI-funded myVocs project based at University of Alabama-Birmingham (Gemmill, Robinson) and we led two well-attended (40+ attendees) Shibboleth-Grid developers BOFs at recent Global Grid Forum meetings.

### **Findings: (See PDF version submitted by PI at the end of the report)**

We have a major technical findings in the last year:

(1) Leverage the campus for identities and virtual organizations (VOs) for attributes. In going through use cases with TeraGrid and other communities it has become fairly clear that most user attributes of traditional interest to grids are not ones for which campuses are authoritative, but instead come from VOs — e.g. VO membership, role. In collaboration with the NMI myVocs project (Gemmill, Robinson) in demonstrating an integration of GridShib and myVocs, we have concluded a model where the campus asserts identity and then a VO augments this identity with attributes is the logical path. One possible exception to this model that we will attempt to validate in the upcoming year, is a new use case for grids involving the granting of access to classes of students (e.g. CS101) at a campus, a attribute of the students for which the campus is clearly authoritative.

(2) Pushing of attributes has significant advantages over attribute pull. We initially started with an attribute pull mode, where a relying party, on receiving a request from a user would go forth and retrieve attributes regarding that user in order to make an authorization decision. This mode of operation has a number of issues involving how the relying party determines the correct authority to contact and how it establishes a trust relationship with that authority. We also have concerns that it will require caching to address performance under heavy load, which adds complexity. In parallel with the Shibboleth project itself, we are moving to an overall push mode of operation, where the user gathers attributes themselves and presents these to the relying party at the time of their request. This doesn't preclude a pull mode of operations, but should simplify issues for the majority of use cases.

### **Training and Development:**

#### **Outreach Activities:**

Presentations made by project members in the past year:

\* Tom Barton. GridShib. Common Solutions Group, September 2006. <http://grid.ncsa.uiuc.edu/presentations/20060920-gridshib-tb.ppt>

\* Tom Barton, Jim Basney, Tim Freeman, Tom Scavo, Frank Siebenlist, Von Welch, Rachana Ananthakrishnan, Bill Baker, Monte Goode, and Kate Keahey. Identity Federation and Attribute-based Authorization through the Globus Toolkit, Shibboleth, GridShib, and MyProxy. 5th Annual PKI R&D Workshop, April 2006. <http://grid.ncsa.uiuc.edu/presentations/welch-pki-06.ppt>

\* Tom Barton, Tim Freeman, Kate Keahey, Raj Kettimuthu, Tom Scavo, Frank Siebenlist, and Von Welch. GridShib: Grid/Shibboleth Interoperability. Presentation to Naregi Visitors, November 2006. <http://grid.ncsa.uiuc.edu/presentations/gridshib-nov06.ppt>

\* Tom Barton, Tim Freeman, Kate Keahey, Raj Kettimuthu, Tom Scavo, Frank Siebenlist, and Von Welch. GridShib: Grid/Shibboleth Interoperability. GlobusWorld 2006, September 2006. <http://grid.ncsa.uiuc.edu/presentations/gw-gridshib-sep06.ppt>

- \* Tom Barton, Tim Freeman, Kate Keahey, Raj Kettimuthu, Tom Scavo, Frank Siebenlist, and Von Welch. GridShib: Grid/Shibboleth Integration. GlobusWorld 2006, September 2006. <http://grid.ncsa.uiuc.edu/events/ggf18-shib-bof/gridshib-ggf18-shib-bof-sep06.ppt>
- \* Tom Barton, Tim Freeman, Kate Keahey, Raj Kettimuthu, Tom Scavo, Frank Siebenlist, and Von Welch. Gridshib Project Update. Internet2 Fall Member Meeting, December 2006. <http://grid.ncsa.uiuc.edu/presentations/i2mm-gridshib-roadmap.ppt>
- \* Jill Gemmill, John-Paul Robinson, Tom Scavo, and Von Welch. myVocs and GridShib: Integrated VO Management. Internet2 Spring Member Meeting, April 2006. <http://grid.ncsa.uiuc.edu/presentations/i2mm-myvocs-gridshib-april06.ppt>
- \* Dane Skow and Von Welch. TeraGrid Plans for Authentication and Authorization Testbed. TeraGrid Science Gateway telecon, October 2006. <http://grid.ncsa.uiuc.edu/presentations/TG-SG-call-Oct06.ppt>
- \* Von Welch. Scaling TeraGrid Access: A Roadmap (Testbed) for Federated Identity Management for a Large Cyberinfrastructure. SuperComputing '06 TeraGrid Booth Presentation, November 2006. <http://grid.ncsa.uiuc.edu/presentations/tg-nov06.ppt>
- \* Tom Scavo. Bindings and Profiles for Attribute-based Authz in the Grid. GlobusWorld 2006, September 2006. <http://grid.ncsa.uiuc.edu/events/ggf18-shib-bof/x509-bindings-profiles-sep06.ppt>
- \* Von Welch, Jim Basney, Tom Scavo. An X.509 Binding for SAML. Poster presented at the Midwest Security Workshop, September 2006. <http://grid.ncsa.uiuc.edu/posters/Poster-MSW2006.doc>
- \* John-Paul Robinson, Jill Gemmill, Tom Scavo, and Von Welch. Building Systems with Shibboleth: Integrated VO Management with myVocs and GridShib. TERENA May 2006. [http://www.terena.nl/events/tnc2006/core/getfile.php?file\\_id=1099](http://www.terena.nl/events/tnc2006/core/getfile.php?file_id=1099)

Other contributions made by the project:

The Gridshib Project organized two BOFs for Grid developers working with Shibboleth, one at GGF 16 (Feb 14-15, 2006) and GGF 18 (Sep 11-12, 2006). Both BOFs were well attended with eight or more projects represented and over 40 attendees total. Proceedings for the two BOFs are available at the following URLs:

<http://www.ggf.org/documents/GFD.79.pdf>

<http://grid.ncsa.uiuc.edu/events/ggf18-shib-bof/>

The following whitepaper was presented and discussed at the TeraGrid Authentication, Authorization, and Account Management Workshop (Argonne National Laboratory, August 30&31, 2006):

- \* Von Welch, Ian Foster, Tom Scavo, Frank Siebenlist, Charlie Catlett. Scaling TeraGrid Access: A Roadmap for Attribute-based Authorization for a Large Cyberinfrastructure. <http://gridshib.globus.org/docs/tg-paper/TG-Attribute-Authz-Roadmap-draft-aug24.pdf>

### **Journal Publications**

#### **Books or Other One-time Publications**

Tom Scavo and Scott Cantor, "SAML Metadata Extension for a Standalone Attribute Requester", (2005). Specification, Committee Draft Bibliography: OASIS Security Services Technical Committee Draft 01, 11 April 2005.

Document ID: sstc-saml-metadata-ext-cd-01

<http://www.oasis-open.org/committees/download.php/13845/sstc-saml-m>

Von Welch, Tom Barton, Kate Keahey, and Frank Siebenlist, "Attributes, Anonymity, and Access: Shibboleth and Globus Integration to Facilitate Grid Collaboration", (2005). Conference Proceedings, Published Collection: Proceedings of the 4th Annual PKI R&D Workshop, April 2005

Bibliography: <http://middleware.internet2.edu/pki05/proceedings/welch-globus-shibboleth.pdf>

Tom Barton, Jim Basney, Tim Freeman, Tom Scavo, Frank Siebenlist, Von Welch, Rachana Ananthkrishnan, Bill Baker, and Kate Keahey, "Identity Federation and Attribute-based Authorization through the Globus Toolkit, Shibboleth, GridShib, and MyProxy", (2006). Conference Proceedings, Published  
 Collection: Proceedings of the 5th Annual PKI R&D Workshop  
 Bibliography: <http://grid.ncsa.uiuc.edu/papers/gridshib-pki05-final.pdf>

### Web/Internet Site

**URL(s):**

<http://gridshib.globus.org/>

**Description:**

The GridShib Web Site is our primary, outward-facing online resource. It contains valuable information about the project (including a brief introduction along with news, announcements, reports, and presentations) as well as important links to resources such as software distributions, archives, repositories, and various sources of end-user support. The web site is hosted by Argonne National Laboratory (whereas it was previously hosted by NCSA).

### Other Specific Products

**Product Type:**

**GridShib Wiki**

**Product Description:**

The GridShib Wiki includes four broad categories of web content: Installation, Deployment, Community, and Team. Currently, the Team section of the wiki, used primarily for internal collaboration, is most developed. We expect the other sections of the wiki to become more developed over time as the project and the software gain traction. The wiki is hosted by The Ohio State University.

**Sharing Information:**

The GridShib Wiki is available as a user-maintained web resource <<https://authdev.it.ohio-state.edu/twiki/bin/view/GridShib/WebHome>>

**Product Type:**

**Data or databases**

**Product Description:**

All project source code is stored in the GridShib CVS Repository, which is anonymously available to all users. GridShib software is licensed under the terms of the Globus Toolkit 4.0 Public License, which itself is based on the Apache License, Version 2.0. The repository is hosted by Argonne National Laboratory.

**Sharing Information:**

The GridShib CVS Repository <<http://viewcvs.globus.org/viewcvs.cgi/playground/java/gridshib/>> is an open web resource.

**Product Type:**

**gridshib-beta mailing list**

**Product Description:**

gridshib-beta is an archived mailing list for end-user support. Users can obtain support for installing, configuring, and deploying the GridShib software. The mailing list is hosted by Argonne National Laboratory.

**Sharing Information:**

The gridshib-beta mailing list <[gridshib-beta@globus.org](mailto:gridshib-beta@globus.org)> is an open subscription mailing list. Instructions for subscribing to the mailing list are on the GridShib Support page <<http://gridshib.globus.org/support.html>>.

**Product Type:**

**Data or databases**

**Product Description:**

Software bugs and enhancement requests (both internal and external) are posted to GridShib Bugzilla, which is hosted by Argonne National Laboratory.

**Sharing Information:**

The GridShib Bugzilla <<http://bugzilla.globus.org/globus/buglist.cgi?product=GridShib>> is an open web resource. (Only GridShib developers may modify the bugzilla database.)

**Product Type:**

**Software (or netware)**

**Product Description:**

GridShib Alpha [2005-05-01] was an initial, unreleased software component early in the project development cycle. This pre-release implementation underwent rigorous internal testing before it was repackaged and released to the general public.

**Sharing Information:**

Pre-beta versions of the software are available on the GridShib Archives page. However, users are encouraged to obtain the latest release of the software from the GridShib Downloads page.

**Product Type:**

**Software (or netware)**

**Product Description:**

GridShib Beta [2005-09-06] (included in NMI-R8) is the reference implementation of the GridShib Beta Attribute Exchange Profile. GridShib Beta consists of two separate components: GridShib for Globus Toolkit and GridShib for Shibboleth. The two components may be installed and tested separately, but both components are required for complete functionality and interoperability.

**Sharing Information:**

GridShib Beta is available on the GridShib Downloads page. The latest source code is available for download from the GridShib CVS Repository.

**Product Type:**

**Software (or netware)**

**Product Description:**

The Shibboleth IdP Tester [2005-11-21] is a standalone software tool that tests a previously installed and tested Shibboleth Identity Provider (IdP). A deployer can have confidence that an IdP that passes this test will accept the GridShib for Shibboleth plugin.

**Sharing Information:**

The Shibboleth IdP Tester is available on the GridShib Downloads page. The latest source code is available for download from the GridShib CVS Repository.

**Product Type:**

**Teaching aids**

**Product Description:**

PowerPoint presentations:

Security Assertion Markup Language: A Brief Introduction to SAML <http://grid.ncsa.uiuc.edu/presentations/saml-intro-dec05.ppt>

Security Assertion Markup Language: SAML 1.x Technical Overview

[http://grid.ncsa.uiuc.edu/presentations/saml-v1\\_x-tech-overview-dec05.ppt](http://grid.ncsa.uiuc.edu/presentations/saml-v1_x-tech-overview-dec05.ppt)

Security Assertion Markup Language: An Introduction to SAML 2.0 [http://grid.ncsa.uiuc.edu/presentations/saml-v2\\_0-intro-dec05.ppt](http://grid.ncsa.uiuc.edu/presentations/saml-v2_0-intro-dec05.ppt)

Shibboleth: An Introduction <http://grid.ncsa.uiuc.edu/presentations/shibboleth-intro-dec05.ppt>

Shibboleth: A Technical Overview <http://grid.ncsa.uiuc.edu/presentations/shibboleth-tech-overview-dec05.ppt>

GridShib: An Introduction <http://grid.ncsa.uiuc.edu/presentations/gridshib-intro-dec05.ppt>

GridShib: A Technical Overview <http://grid.ncsa.uiuc.edu/presentations/gridshib-tech-overview-dec05.ppt>

**Sharing Information:**

All PowerPoint presentations are open web resources.

**Contributions****Contributions within Discipline:**

The GridShib project is doing applied research and software development in the field of distributed identity management and access control. Identity management is a term for the act of managing a group of users in the context of computer, i.e. creating accounts for those users,

creating passwords (or other forms of authentication) for them, and managing their roles (e.g. normal user, system administrator, principal investigator). Access control is the practice of making a decision whether a request by a particular user should be serviced, generally based on information provided by a identity management system.

As the Internet and the web have fostered greater amounts of distributed collaboration, identity management and access control have become increasingly distributed, with identity management being done by one organization and access control by another organization. This has lead to a field of study on how to cleanly separate these concerns and how they should interact. The GridShib project is studying this distribution in the context of higher education campuses using Shibboleth, interacting with computational grids, all with the end goal of supporting virtual organizations.

The major contributions of GridShib have been the definition of standards and techniques to enable this distributions, production of software to create interoperability between Shibboleth and the Globus Toolkit to enable this distribution, and now the experimental validation through deployment to validate our software and develop real-world policies to make this viable in practice.

**Contributions to Other Disciplines:**

**Contributions to Human Resource Development:**

**Contributions to Resources for Research and Education:**

**Contributions Beyond Science and Engineering:**

**Special Requirements**

**Special reporting requirements:** None

**Change in Objectives or Scope:** None

**Unobligated funds:** \$ 0.00

**Animal, Human Subjects, Biohazards:** None

**Categories for which nothing is reported:**

Activities and Findings: Any Training and Development

Any Journal

Contributions: To Any Other Disciplines

Contributions: To Any Human Resource Development

Contributions: To Any Resources for Research and Education

Contributions: To Any Beyond Science and Engineering

## Project Activities

As computational Grids have grown, there has been increasing interest in leveraging existing site infrastructure to support Grid authentication and authorization. For example, Shibboleth has been developed by the Internet2 community and increasingly deployed both in the U.S. and abroad as a mechanism for cross-site access control for web-based resources. Shibboleth utilizes OASIS SAML standards for authentication and attribute assertions to achieve its purpose.

### GridShib: X.509 and SAML integration

GridShib is a software product that allows for interoperability between the Globus Toolkit and Shibboleth. The complete software package consists of two plugins, one for the Globus Toolkit (GT) and another for Shibboleth. With both plugins installed and configured, a GT Grid Service Provider may securely request user attributes from a Shibboleth Identity Provider.

#### GridShib for Globus Toolkit

GridShib for Globus Toolkit is a plugin for Globus Toolkit 4.0. The plugin implements a policy decision point (PDP) based on attributes obtained from a Shibboleth attribute authority. A policy information point (PIP) does the actual work of requesting attributes. The separation between PIP and PDP allows the plugin to be used in flexible ways within the toolkit's authorization framework.

The Grid Client obtains and uses a proxy certificate to authenticate to a Grid Service Provider (SP). The Grid SP extracts the DN from the proxy certificate and uses the DN as the value of the NameIdentifier in a SAML AttributeQuery.

#### GridShib for Shibboleth

GridShib for Shibboleth is a name mapping plugin for a Shibboleth 1.3 Identity Provider (IdP). The plugin allows the attribute authority to map the DN of the user's X.509 proxy certificate to a local principal name. Upon receiving an attribute query, the Shibboleth attribute authority maps the DN and utilizes the resulting principal name to resolve attributes.

The name mapping is a memory-bound collection of name-value pairs. The name (key) is a canonicalized DN that conforms to RFC 2253. The value is the local principal name.

The collection is initialized when the IdP starts up. The current implementation of the name mapping construct is file-based, that is, the name-value pairs are read from an ordinary text file. This text file is similar to the grid-mapfile used by Globus Toolkit.

#### GridShib Attribute Exchange Profile

The GridShib Attribute Exchange Profile is an extension of the Shibboleth Attribute Exchange Profile. The primary difference is the use of X.500 distinguished names (DNs) to identify principals.

The GridShib Attribute Exchange Profile is designed for a standalone attribute requester, that is, an attribute requester that does not participate in a Shibboleth browser profile. As a result, the Grid SP does not have access to an opaque, transient handle typically issued by the IdP on the front end of the

browser profile. In lieu of a handle, the Grid SP uses the DN obtained from the client's proxy certificate.

Our use case involves a Grid Client that already possesses an X.509 end-entity certificate (EEC). As is often the case in grid-based scenarios, the established user uses this EEC to generate a proxy certificate as part of single sign-on. The proxy certificate is then used to authenticate to Grid SPs as part of the act of requesting service.

Beta software that implements the GridShib Attribute Exchange Profile may be downloaded from the GridShib web site. A technical overview of the GridShib Attribute Exchange Profile is also available.

#### Globus Toolkit Authorization Framework

As the Globus Toolkit (GT) is used by many different projects and by many different Grid communities, it is clear that GT cannot mandate the use of particular technologies and mechanisms. Specifically in the area of attributes and authorization policies, the toolkit has to be flexible enough to accommodate locally preferred assertion formats and usage patterns. The Globus Toolkit Authorization Framework is designed to handle these different mechanisms in a consistent manner and to combine authorization decisions from many different sources to yield a single access decision for each invocation request.

#### Current and Future GT Support

The currently shipping GT 4.0 implementation includes a simplified version of the described attribute collection and authorization framework, but does not fully support attribute-based authorization and has no support for fine-grained delegation of rights. It includes support for proxy certificate delegation, call-out support to SAML 1.1-compliant authorization services, grid-mapfile authorization, and an XACML evaluator.

Enhancements to support Shibboleth and SAML attribute assertions have been added as part of the GridShib effort, and are part of the GridShib beta release.

The full-featured authorization framework is under active development, has produced a number of prototypes, and will ship with our next major release GT 4.2.

## Project Findings

At the architectural level, the project has identified a number of challenges that lie before us in regards to managing different namespaces. These challenges are described in more detail in the subsequent project plan as well as our most recent publication to the 5ht Annual PKI workshop, but briefly how does one establish that a user's identity at a University should be linked to their Grid identity. We currently have manual administration to solve this problem, but are concerned with scaling issues with this simple approach.

The following limitations of the current beta software implementation have been identified:

- \* The file-based name mapping doesn't scale.
- \* IdP discovery must be generalized.
- \* Metadata production and distribution needs to be automated or simplified.

The fact that the DN-principal name pairs are read from a file is a major concern. Even if we were to provide administrative tools to manage the name mapping files, the administrative overhead associated with this maintenance would be prohibitive. Clearly, this overhead must be eliminated or at least reduced.

In step 1 of the GridShib Attribute Exchange Profile, we assume that the Grid Client somehow includes the IdP providerId in the request. Unfortunately, the current implementation of the software does not satisfy this condition. Instead the providerId is configured into the Grid SP, which essentially forces both the IdP and the Grid SP to reside in the same security domain.

Trust in a GridShib deployment is based on a bilateral arrangement between the IdP and the Grid SP. By virtue of the fact that the two entities exchange and consume each other's metadata, a trust relationship is established. The problem is that n entities give rise to  $O(n^2)$  bilateral relationships, which of course does not scale.