ProjectSummary

GridShib Project Summary Report

This summary report of the *GridShib Project* spans the time period 2004-12-01 to 2005-12-06. Where possible, items in each category are listed in chronological order.

Table of Contents

Executive Summary

The NMI GridShib project is a collaboration between NCSA and the University of Chicago. The goal of the project is to allow for interoperability between the Globus Toolkit and the Shibboleth Identity Federation system. In the first year of this project basic interoperability has been achieved in that a deployment of the Globus Toolkit can query and receive attributes from a Shibboleth service regarding a Grid user, and then make an access control decision based on those attributes. In the second year, we turn our attention to refining this capability as well as addressing higher-level management issues such as management of the federation of name spaces between the Grid and campus worlds and the management of the trust configuration /metadata between the Grid and Shibboleth components.

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.

Almost completed development allows the Grid SP to use an embedded SAML `AuthenticationStatement` as the value of the `NameIdentifier` in a SAML `AttributeQuery` .

Also almost complete is the ability for the Grid SP to select from multiple !IdP endpoints and credentials from a configured SAML2 metadata instance.

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.

Project Plan

In this section we discuss our plans for work in the forthcoming year for enabling the seamless integration of Shibboleth/SAML and Grid Security/X.509. We expect there is more here then we will be able to accomplish in the remaining year, so we list these in order of priority. Many of these plans are in collaboration with the MyProxy project, an online credential management system developed at NCSA. These plans result from discussions with the MyProxy project management.

GridShib Beta

The file-based name mapping system will be augmented with a database implementation. This will not solve the maintenance problem, but it will make it easier to provide administrative tools. A database implementation will also facilitate load-balancing of IdPs (an ongoing issue in the Shibboleth Project).

One approach to the !IdP discovery problem is to include the !IdP providerId into the proxy certificate itself. Thus we are considering a modification to MyProxy that produces such certificates. For this to work, we assume initially that MyProxy resides in the same security domain as the !IdP. Further work will attempt to relax this restriction.

Metadata is an important aspect of GridShib (or any federated identity management system, for that matter). Therefore the following enhancements are being considered:

- produce !IdP metadata from the underlying !IdP configuration;
- provision attribute release policies (ARPs) from Grid SP metadata;
- consume !IdP metadata and provision Grid SP configuration; and
- produce SP metadata from the underlying Grid SP configuration.

On the !IdP side, tools to produce and consume metadata are being designed. In particular, a tool to automatically produce !IdP metadata would be very helpful. Similar tools for the Grid SP are needed.

Testing a browser-based Shibboleth deployment remains a challenge. Testing GridShib on top of Shibboleth is even more difficult. To address this problem, we provide a command-line testing tool that tests both a Shibboleth AA and a GridShib AA. A discriminating test strategy is being built around this tool.

Need for Name Binding

In the simplest case, access to a grid service is managed by providing all users with an EEC from a recognized CA, mapping the names in these EECs to another namespace local to the grid service, and using these local names in access control lists. To broaden the availability of the grid service to more users, additional naming authorities can be recognized. In particular, we wish to enable use of established naming authorities, such as those local to a user's home organization, and authentication tokens other than X.509 EECs. However, we are constrained by the requirement that an EEC must be presented to the grid service, and that only attributes correlated with the presented name in that EEC can be marshaled.

This presents two problems. One is the exchange of an original authentication token for a suitable EEC to be presented to the grid service. The other is mapping the presented name in this EEC to the name in the original authentication token, called the *principal name*, so that attributes bound to the principal name can be marshaled by the grid service. Because the principal namespace is not local to the grid service, and to support pseudonymous access scenarios, we propose to collocate this presented name to principal name mapping function with the authority for the principal namespace and the attributes that are bound to principal names.

Direct client-server use case

We see two distinct but equally important scenarios in which this name binding must take place. In the first scenario, which we discuss in this section, the client application communicates directly with the service. The second scenario, which we discuss in the next section, involves a web portal intermediary.

When the client application and service communicate directly, end-to-end X.509 authentication is performed as part of the protocol (which is either based on TLS or SOAP with message-level security based on WS-Security). The difficulty in this case is binding the identifier in the user?s X.509 credentials back to principal name so that attributes may be obtained.

In this case, we believe that the online CA functionality in MyProxy can be used to solve this problem. The user obtains short-lived X.509 credentials initially by authenticating to the MyProxy online CA using their principal name and password. The MyProxy CA would then issue the X.509 credential, embedding into it the user?s principal name. The service would then extract the principal name and use it when communicating back to the Shibboleth Attribute Authority.

We note that this approach has a distinct advantage over the current implementation in that the Shibboleth AA does not need to maintain a DN-to-principal name mapping since the principal name is in the SAML query.

An open issue is the appropriate mechanism for embedding the principal name into the X.509 certificate. Current options being considered are to use the Subject Alternate Name or the Subject Information Access extension of the certificate. One could also embed the principal name into the DN itself, however we are concerned about placing requirements on the contents of the DN.

Portal Use Case

Another use case involves the client using a web browser to access a web server, which in turn accesses Grid services on behalf of the client. This use case is becoming more common as a means to allow for easy access to Grid services with a minimal footprint installation on the client system.

The primary observation in this use case is that the portal effectively functions as a "chasm" that must be bridged. Either X.509 or Shibboleth/SAML can be used to authenticate to the portal, but there is no general method in use today to delegate authority to the portal. This is the so-called *n*-tier problem ($n > 2$), an active research area.

We note that MyProxy has been used traditionally in the Grid community to enable a portal to use a client?s username and password to obtain X.509 credentials for the client. As in the previous section, these X.509 credentials would have the principal name, taken from the `NameIdentifier` element in the SAML assertion, embedded in them. This would allow the Grid service to query the SAML Attribute Authority in an identical manner as described previously.

Publications

- Tom Scavo and Scott Cantor. *SAML Metadata Extension for a Standalone Attribute Requester.* 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-metadata-ext-cd-01.pdf

- Von Welch, Tom Barton, Kate Keahey, and Frank Siebenlist. *Attributes, Anonymity, and Access: Shibboleth and Globus Integration to Facilitate Grid Collaboration.* Proceedings of the 4th Annual PKI R&D Workshop, April 2005. http://middleware.internet2.edu/pki05/proceedings/welch-globus-shibboleth.pdf

- Tom Barton, Jim Basney, Tim Freeman, Tom Scavo, Frank Siebenlist, Von Welch, Rachana Ananthakrishnan, Bill Baker, and Kate Keahey. *Identity Federation and Attribute-based Authorization through the Globus Toolkit, Shibboleth, !GridShib, and !MyProxy.* Proceedings of the 5th Annual PKI R&D Workshop (To appear April 2006). http://grid.ncsa.uiuc.edu/papers/gridshib-pki06-draft.pdf (October 2005)

Online Resources

- *GridShib Web Site* http://gridshib.globus.org/

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).

- *GridShib Wiki* https://authdev.it.ohio-state.edu/twiki/bin/view/GridShib/WebHome

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.

- *GridShib CVS Repository* http://viewcvs.globus.org/viewcvs.cgi/playground/java/gridshib/

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.

- `gridshib-beta` *mailing list* gridshib-beta@globus.org

`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.

- *GridShib Bugzilla* http://bugzilla.globus.org/globus/buglist.cgi?product=GridShib

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

Software Components

- *GridShib Alpha* [2005-05-01]

GridShib Alpha 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.

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.

- *GridShib Beta* [2005-09-06] (included in NMI-R8)

GridShib Beta 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.

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

- *Shibboleth !IdP Tester* [2005-11-21]

The *Shibboleth !IdP Tester* 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.

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

Presentations

- Tom Barton. *NSF Middleware Initiative: !GridShib.* Internet2 Fall Member Meeting, October 2004. http://www.internet2.edu/presentations/fall04/20040930-gnsfmi-barton.ppt

- Tom Barton. *Shibboleth for Non-Web-Based Applications: !GridShib.* Internet2 Fall Member Meeting, October 2004. http://www.internet2.edu/presentations/fall04/20040929-gridshib-barton.ppt

- Von Welch. *GridShib: Grid-Shibboleth Integration Overview.* NMI Management Call, November 2004. http://grid.ncsa.uiuc.edu/presentations/GridShib-Overview.ppt

- Von Welch. *GridShib: Grid-Shibboleth Integration.* GlobusWORLD 2005, February 2005. http://grid.ncsa.uiuc.edu/GridShib/presentations/GridShib-GW05.ppt
- Von Welch. *GridShib: Grid-Shibboleth Integration.* caBIG Telecon, March 2005. http://grid.ncsa.uiuc.edu/GridShib/presentations/GridShib-caBIG-March05.ppt
- Von Welch. *GridShib: Grid-Shibboleth Integration (Identity Federation and Grids).* UK eScience Security Workshop, April 2005. http://grid.ncsa.uiuc.edu/GridShib/presentations/GridShib-uk-april05.ppt
- Von Welch. *Attributes, Anonymity, and Access: Shibboleth and Globus Integration to Facilitate Grid Collaboration.* 4th Annual PKI R&D Workshop, April 2005. http://middleware.internet2.edu/pki05/proceedings/welch-globus-shibboleth.ppt
- Tom Barton. *NSF Middleware Initiative: !GridShib.* Internet2 Spring Member Meeting, May 2005. http://home.uchicago.edu/~tbarton/gridshib/20050503-gridshib.ppt
- Frank Siebenlist. *Globus Security with SAML, Shibboleth, and !GridShib.* GT4 Tutorial, May 2005. http://www.mcs.anl.gov/~franks/GT-Security-May4-2005.ppt
- Von Welch. *Tools for Grid/Campus Integration: GridShib and !MyProxy.* Internet2 Advanced CAMP, July 2005. http://grid.ncsa.uiuc.edu/presentations/welch-adv-camp-july05.ppt
- Von Welch. *GridShib: Campus/Grid RBAC Integration.* !GridWorld/GGF-15, October 2005. http://www.ggf.org/GGF15/presentations/gridshib-welch-ggf-oct05.ppt
- Von Welch. *Authentication for Virtual Organizations: From Passwords to X509, Identity Federation and !GridShib.* BRIITE Meeting, November 2005. http://grid.ncsa.uiuc.edu/presentations/welch-brite-nov3.ppt
- Tom Scavo and Von Welch. *GridShib: An Attribute-Based Authorization Framework.* Cyberinfrastructure Seminar Series. December 2005. http://grid.ncsa.uiuc.edu/presentations/gridshib-cip-seminar-dec05.ppt

Educational Materials

- *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
- *Security Assertion Markup Language: An Introduction to SAML 2.0* 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

Contributions

The following handbook (included in NMI-R7 and NMI-R8) was developed in collaboration with the Shibboleth Project:

- *Shibboleth Architecture: A Technical Overview* http://shibboleth.internet2.edu/docs/draft-mace-shibboleth-tech-overview-latest.pdf

Significant contributions have been made to the Shibboleth Wiki:

- *Shibboleth Wiki* https://authdev.it.ohio-state.edu/twiki/bin/view/Shibboleth/WebHome

GridShib project members also organized a Workshop at GGF 15 in Boston entitled "Leveraging Site Infrastructure for Multi-Site Grids". This workshop had nine different project presentations, including a GridShib presentation and demonstration, and was well attended by the GGF community.

- *Leveraging Site Infrastructure for Multi-Site Grids* http://www.ggf.org/GGF15/ggf_events_schedule_MultiSite.htm