Scaling TeraGrid Access: A Testbed for Identity Management and Attribute-based Authorization

Von Welch, Ian Foster, Tom Scavo, Frank Siebenlist,
Charlie Catlett, Jill Gemmill, and Dane Skow

Abstract— In this paper we describe plans for a TeraGrid testbed to evaluate identity federation and attribute-based authorization to enable scalability to larger number of user than would be possible with today's infrastructure. The presentation of the paper will include results obtained by the time of the TeraGrid conference in addition to the content in this paper.

Index Terms— Security

1 INTRODUCTION

The effectiveness of large cyberinfrastructures such as TeraGrid depends critically on (among many other things) the ease with which users can access distributed resources and the security and integrity of the services and resources (collectively and individually). Traditional site authorization mechanism on which such systems build are typically based on obtaining an "account" (login and password). From the point of view of a resource provider, creating and maintaining the many accounts and gridmap entries represents a significant cost. The communication and maintenance of many passwords also introduces a potential vulnerability. Indeed, a compromise at one single site can require thousands of individual communications to users to reissue passwords. For a single site, a user community of 1000 means 1000 accounts at that site. For even a small federation of, say, 10 sites serving the same user community, the number of accounts to be managed, protected, and communicated jumps to 10000 – a significant number.

Such traditional access control mechanisms were designed for scenarios in which a strong trust relationship exists between "users" and "resource providers." In such scenarios, it is not unreasonable to expect resource providers to know the identities of all their users ahead of time and to allow access based on authentication of the individual user. Furthermore, resource providers could draw upon this relationship to quickly contact a user in case of problems.

- *Von Welch, NCSA, vwelch@ncsa.uiuc.edu*
- *Ian Foster, U. of Chicago, foster@mcs.anl.gov*
- *Tom Scavo, NCSA, trscavo@ncsa.uiuc.edu*
- *Frank Siebenlist, Argonne National Laboratory, franks@mcs.anl.gov*
- *Charlie Catlett, U. of Chicago, catlett@mcs.anl.gov*
- *Jill Gemmill, U. of Alabam-Birmingham, JGemmill@uab.edu*
- *Dane Skow, U. of Chicago, skow@mcs.anl.gov*

In order to provide access to increasingly large communities of students and scientists, and to support systems of many federated resources, it becomes necessary to move away from a model of pre-registering each user for authorization, and towards approaches that leverage relationships with existing communities and organizations. The keys to realizing such scenarios are identity federation and attribute-based access control. Identity federation allows the resource providers to rely on identification and authentication of the user community by outsides sources, allowing users to authenticate using existing credentials at their local campus or institution. Shibboleth [4][6] has emerged from the higher-education community to allow for cross-site attribute-based access control for web applications. Recent enhancements to the Globus Toolkit version 4 (GT4), which is utilized in TeraGrid and other large Grid deployments, have introduced an attribute-based authorization framework that makes feasible the integration of these different attribute-based access control systems into large Grids. Attribute-based access control provides authorization mechanisms that allow access control decisions to be made on the basis of a variety of user attributes in addition to simple identity. The virtual organization management system (VOMS) [5] is in use by peer Grids for providing attribute-based authorization.

In this paper we describe the TeraGrid Testbed for implementing identity federation and attribute-based access control, including the key motivating scenarios and requirements.

2 MOTIVATION

As described in the introduction, our motivation is to allow the scaling of the TeraGrid in terms of users through the use of identity federation to leverage existing campus identity systems coupled

with attribute-based authorization to mitigate resource providers needing a priori relationship with every user. This will allow users to use existing authentication mechanisms at their local campus, getting TeraGrid out of the role of having to register and manage credentials for users, while providing an information-rich system for authorizing users. Thus the motivations for considering identity federation and authorization-based access control include ease of access for users, improved scalability (resulting in improved security), reduced cost and overhead for providers, and better integrating national-scale cyberinfrastructure (such as TeraGrid) with campus cyberinfrastructure, further reducing the administrative overhead faced by campus users in accessing national resources. There are three areas of interaction between users and TeraGrid, described in this section, in which we expect to evaluate these benefits.

2.1 Community Access for Science Gateways

Systems such as TeraGrid are increasingly focusing attention on enabling access via "science gateways" [3]. In such systems, users access a "gateway" (e.g., via a portal) which then performs operations on their behalf. As compared to traditional Grid systems, gateways may use more "lightweight" approaches to authentication and authorization in which (for example) the gateway authenticates the user (e.g., using portal-specific username and password), vets the user, and then submits requests to TeraGrid resources on their behalf, perhaps executing the requests using a "community account." Several TeraGrid Science Gateways are already pursuing first steps in the direction laid out in this paper, but are hampered by TeraGrid's lack of attribute-based authorization. Since authorization on TeraGrid is currently identity-based, community credentials must be deployed by the Gateways, with audit mechanisms used to determine specific user identities outside the course of normal authorization.

Support for attributes will allow Science Gateways to use identity credentials for their users, coupling these identify credentials with attribute assertions regarding the user's community membership. Thus, resource providers will be able to authenticate community users, even those previously unknown to the provider, and authorize them based on attributes. Attribute-based authorization would allow such gateways to either obtain light-weight identity credentials for their users (perhaps created on the portal without the user needing even being aware of it) or use identity credentials from outside sources (e.g., campus CAs) and couple those credentials with an attribute indicating their community membership. TeraGrid resources would recognize the user as a community member based on the attribute and map their request to the community account, removing the need for individual accounts, while also obtaining an identifier for the user from the community credential, allowing for strong auditing.

An example use scenario would be a user registering for and being granted access to a Science Gateway (we skip the details here, presumably the user is a member of the community served by the Science Gateway). The user then subsequently authenticates, and the Gateway creates a credential for the user which conveys not only the user's identity, but the fact they are a community member as well as potentially other information, such as the IP address currently in use by the user, a mapping of that IP address to a physical location (e.g. country), etc. If the user were using Shibboleth to authenticate to the portal, attributes from the user's home institution could also be collected and conveyed. A request on behalf of the user is then presented to the resource provider, which parses the attribute information. The request would be authorized based on the user's community membership, but other attributes could also be considered before granting the request, e.g. the user's identity could be checked against a "blacklist" of unacceptable users, if the service being access falls under export control, the user's current physical location could be checked against unauthorized countries, etc.

2.2 Grid Interoperation

Grid interoperation has proved to be important in situations where a user or community operating on one Grid needs access to capabilities or capacity provided by another Grid. However, again, we see a variety of ad hoc mechanisms being used for authentication and authorization. For example, the bridge that currently exists between TeraGrid and Open Science Grid (OSG) is achieved by exchanging user lists between the two Grids, in conjunction with each Grid mapping the list of users from the other Grid to a community account. This approach is fragile and requires constant user list exchange between the two Grids. It will also grow more complicated if the Grids want to distinguish between groups of users in the other Grid and run them in separate accounts. VOs in OSG already have the ability to define attributes for their users via VOMS [5]. If TeraGrid resources were able to consume and utilize these attributes, it would remove the need for TeraGrid to know all OSG users by individual identity (although those identities would still be available for auditing purposes). Likewise, if TeraGrid were to assert attributes for its users (either allocated or community), OSG could consume those attributes and thus need not know all TeraGrid users by identity.

An example use scenario here would be a user obtaining their credential for the peer grid in the usual manner, which includes attribute information about their community membership in that grid.

Similar to the Science Gateway scenario in the previous section, when they present their request to the TeraGrid resource provider, they would be authorized based on their community membership, but other factors, such as their identity could be used in the authorization process.

2.3 Campus Access to TeraGrid

Campus researchers and students make up a large portion of both the current TeraGrid user community and potential future users. In general, campuses are making (many already have made) large investments in identity management infrastructure that is used locally for trusted operations such as controlling access to grades and other private information. The work of the Internet2 community in developing Shibboleth allows other organizations outside of those campuses to leverage that identity management infrastructure when handling requests from a campus's users.

An example use scenario here would be that of a normal, allocated TeraGrid user, who uses their local campus credentials to authenticate to TeraGrid, much in the same way they might use their TeraGrid MyProxy credentials today. A more complex scenario would be a user who accesses TeraGrid as part of their participation in a class or workshop, being authorized based on this attribute rather than their identity.

3 Testbed Overview

In order to instantiate the motivating scenarios described in Section 2, TeraGrid needs to deploy an enhanced infrastructure capable of attribute-based access control. This infrastructure must be able to enforce a consistent policy across its member resource providers regarding what attributes are meaningful and who is able to assert those attributes, provide means for delivering attributes to resource providers, and for them to enforce the policy. Multiple sources of attributes and identity must be supported, including other Grids (e.g., OSG), science communities, and campuses as well as other organizations hosting TeraGrid user communities. Other Grid deployments have accomplished portions of such an attribute deployment, but deployment of this complete set of functionality to enable all the given scenarios has yet to be achieved.

We propose that TeraGrid achieve this functionality by deploying a testbed amongst a small number of participants, using an enhanced version of the Common TeraGrid Software and Services (CTSS) to prototype the needed infrastructure. The goal of the testbed is to validate a future version of CTSS for full deployment to TeraGrid resource providers to provide needed functionality, identify additional services that will be required for attribute and policy distribution and maintenance, and develop policies to meet TeraGrid needs in managing the motivating scenarios. Upon successful validation, a plan will be drafted to migrate or replicate these services into production so that TeraGrid can provide this functionality on a daily basis to a large user community.

The following set of software components would provide a minimum set of functionality for a useful testbed:

- To allow resource providers to make attribute-based authorization decisions, an enhanced CTSS software stack capable of parsing attributes from both Shibboleth and VOMS. This software stack would include software developed by the GridShib project [1][2] for supporting Shibboleth attributes and from the Virtual Workspaces project to support VOMS attributes [7].
- To allow users to use their existing campus credentials to access the Grid, a GridShib-CA deployment to convert from Shibboleth-based authentication to X.509 authentication used by Grids.
- Software tools for Science Gateways that allow them to manage attributes, created locally as well as from the user's home institution.

Integration and deployment of these software components will occur in January through March, with several months of expanded deployment and evaluation following.

4 System Requirements

There exists a set of functional requirements that may not be apparent from the functional description of the testbed. In this section we discuss those requirements.

4.1 Site Autonomy

While wanting to outsource user management, in order to be comfortable with the system, resource-providing sites must maintain ultimate control of authorization, ideally at the finest granularity possible (e.g. the ability to deny individual users in large communities). This requirement can become complicated as the site's method of authenticating and authorizing users become more distributed. This requirement is primarily driven by incident response, in the event of (even suspected) misbehavior by a specific user the site will require fine-grained policy control to block the specific user, but other motivations exist, such as export control, which would lead a site to deny users based on attributes that may not be meaningful to communities.

4.2 Campus Service Agreements

As TeraGrid resource providers will rely on attributes provided by campuses and other organizations, what sort of relationship will need to exist between TeraGrid and these campuses to ensure reliable security and incident handling? An approach that requires the establishment of legal obligations is likely to be a major (even fatal) hurdle. However, as the number of organizations involves expands, some form of explicit understanding will probably be needed. In a similar vein, what standardization or conventions will be needed for identity management systems at campuses providing attributes used for TeraGrid access?

Currently the agreement is that campuses will provide assistant in the event of an incident, but the details remain vague. Is this assistance 7x24? Business hours? Does it take the form of providing contact information for users in questions or assuming responsibility for investigating those users? In the author's experience, different campuses have very different policies and procedures for similar situations today. It is probable that a range of acceptable agreements is needed.

4.3 Auditing for Incident Response

While attributes are valuable for making access control more efficient, identification of the user still has a role in audit logs. There are two common situations. First, actions need to be correlated to the same requestor in order to investigate suspicious behavior. Second, one needs to contact the user in order to obtain assistance. In the former, a unique, but opaque, identifier would suffice, whereas for the latter, one needs to map the identifier to a email address or real-world contact information. Being able to track down a user after the fact does not require that a resource provider know details regarding the user before their request, just that they have enough information such that when a situation arises, the individual can be located in a timely fashion. Logging a unique identifier for the user, even if that identifier is not used in authorization, along with the source of that identifier (e.g., campus, science gateway), allows this form of after the fact location of a user.

4.4 Revocation

Identifiers and attributes issued by Shibboleth are extremely short-lived, on the order of minutes, and are consumed relatively quickly, a practice which has, to date, not required any revocation process for these identifiers and attributes. When we convert these identifiers and attributes into Grid credentials, their lifetime will be extended to roughly twelve hours. This is similar to lifetime of credentials from on-line certificate authorities such as MyProxy and the Kerberos CA, which has similarly managed to avoid having a revocation process to date. It is unclear what affect adding attributes to identity credentials will have on the requirement to revoke such credentials. Assuming attributes are fairly static and the authorities issuing such attributes are reliable, there should be little effect, but more production experience is needed to validate this claim.

5 Related Work

5.1 Virtual Organization Memebership Service (VOMS)

The Virtual Organization Membership Service (VOMS) [5] is used in Grid deployments to generate X.509 attribute certificates that assert a particular user is a member of a particular virtual organization. The VOMS attributes include as well fields to describe membership in groups within the organization and roles within each group. VOMS is used by the Open Science Grid, EGEE, and others projects with which TeraGrid is interested in interoperating.

Arguably, VOMS combined with a online certification authority such as a Kerberos CA [8] or a MyProxy CA [2] would provide similar functionality as Shibboleth and Gridshib. While both approaches have technical merit, we believe Shibboleth is more likely to be be widely deployed in U.S. institutions of higher education in the foreseeable future, reducing the barrier to use by TeraGrid.

5.2 Other Shibboleth-based Grid Deployments

Aside from GridShib, there are several other projects developing technologies for leveraging Shibboleth to support Grids. While their approaches differ, there is some overlap between the goals all these projects. As these technologies mature, we expect the most successful components from each could be harvested.

Two projects are the U.K.-based Shibboleth Enabled Bridge to Access the National Grid Service (SHEBANGS) project [9] and the ShibGrid project [10]. Both of these projects are developing prototypes for access to the U.K. Natioanl Grid Services via Shibboleth, which is being heavily deployed in the U.K. SHEBANGS uses a trusted intermediary service, known as the Credential Translation Service (CTS), to create a X.509 credential for the user in MyProxy, while ShibGrid is developing support for processing Shibboleth authentication in MyProxy itself. SHEBANGS requires less modification to other software, but requires the user be more aware of the process since they must user a username:password:server triplet from the CTS to authenticate to the application portal.

A third project is the Shibboleth-based short-lived certificate service (SLCS) being developed by Swiss Education and Research Network (SWITCH) [11]. This service takes advantage of the relatively small number of Swiss higher education sites in order to have a non-web browser client that still uses Shibboleth (something not practical in larger federations since it requires parsing each site's Shibboleth authentication web pages).

ACKNOWLEDGEMENTS

Additional contributors to the ideas presented in this paper include Rachana Ananthakrishnan, Tom Barton, Sebastien Goasguen, Michael Grobe, Jim Rome, and Tim Freeman.

REFERENCES

[1] GridShib Project. http://gridshib.globus.org/

[2] 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. In 5th Annual PKI R&D Workshop (To appear), April 2006. http://grid.ncsa.uiuc.edu/papers/gridshib-pki06-final.pdf

[3] TeraGrid Science Gateways Program. http://www.teragrid.org/programs/sci_gateways/

[4] The Shibboleth Project http://shibboleth.internet2.edu/

[5] EU DataGrid VOMS Architecture v1.1, 2003. http://grid-auth.infn.it/docs/VOMS-v1_1.pdf.

[6] S. Cantor et al., Shibboleth Architecture: Protocols and Profiles. Internet2-MACE, 10 September 2005. Document ID internet2-mace-shibboleth-arch-protocols-200509 http://shibboleth.internet2.edu/docs/internet2-mace-shibboleth-arch-protocols-latest.pdf

[7] Globus VOMS Support. http://dev.globus.org/wiki/VOMS

[8] Kerberos Leveraged PKI http://www.citi.umich.edu/projects/kerb_pki/

[9] Shibboleth Enabled Bridge to Access the National Grid Service (SHEBANGS) http://www.mc.manchester.ac.uk/research/shebangs

[10] ShibGrid Project http://www.oesc.ox.ac.uk/activities/projects/index.xml?ID=ShibGrid

[11] SWITCH Short Lived Credential Service (SLCS) http://www.switch.ch/grid/slcs/