

Authentication & Authorization in SURAgrid: Concepts and Technologies

May 2005

NSF Middleware Initiative (NMI) Integration Testbed Case Study Series: Supplemental Documentation

This document is a technical supplement to the NMI Testbed Grid case study, "Exploring Technical and Policy Considerations for Inter-Institutional Grids", part of the NMI Integration Testbed Case Study Series. The NMI Testbed Grid began as a sub-project of the NMI Integration Testbed program in September 2003 and is continuing beyond the program as SURAgrid, a multi-institutional cooperative effort to model the use of grid technology to share distributed resources across institutional boundaries while maintaining local autonomy.

For more information on the NMI Integration Testbed Case Study Series, the NMI Testbed Grid project, or SURAgrid, contact:

Mary Fran Yafchak

Southeastern Universities Research Association (SURA)

maryfran@sura.org

The NMI Integration Testbed Case Study Series was sponsored by the NMI Enterprise and Desktop Integration Technologies (NMI-EDIT) Consortium of EDUCAUSE, Internet2, and SURA, under National Science Foundation Cooperative Agreement NSF 02-028, ANI-0123937

Contributors

Name

Adiga, Ashok
Henderson, Shelley
Jokl, Jim
McKee, Shawn
Perez, Jerry
Robinson, John-Paul
Vandenberg, Art
Yafchak, Mary Fran
Barzee, Kathleen

Institution

Texas Advanced Computing Center (TACC)
University of Southern California
University of Virginia
University of Michigan
Texas Tech University
University of Alabama at Birmingham
Georgia State University
Southeastern Universities Research Association (SURA)
Southeastern Universities Research Association (SURA)

Introduction

The NMI Testbed Grid project was formalized as a sub-activity of the NMI Integration Testbed¹ program in September 2003. The intent was to model the shared, collaborative and heterogeneous multi-institutional environment characteristic of higher-education research. The implementation of authentication (AuthN) and authorization (AuthZ) processes within a multi-participant, multi-project grid was a central focus of the collaboration, providing NMI Testbed Grid participants² with the opportunity to explore how local authentication solutions, or those within project-specific grid environments, could be extended to provide access management for grid resources within and beyond the campus. By taking this approach, the NMI Testbed Grid provided a working higher education grid environment that could be connected to the emerging enterprise middleware infrastructure.

The focus on identifying and resolving AuthN and AuthZ issues in a shared, heterogeneous grid environment is being continued through SURAgri³, a cooperative grid deployment within the SURA region⁴. This paper will review the collective and institutional experiences and knowledge gained on this topic within the NMI Testbed Grid beginning September 2003 and continuing into SURAgri through May 2005.

Motivation for AuthN/AuthZ Exploration

A primary purpose of authentication (verifying identity) and authorization (granting permission to access resources based on identity) is to implement the *policies* that organizations have created to govern and manage the use of computing resources. This is recognized by Foster et al. in describing grid technology as a “resource-sharing technology with software and services that let people access computing power, databases, and other tools securely online across corporate, institutional, and geographic boundaries without sacrificing local autonomy” (1). A researcher in the higher-education community may not only be a computer user on their campus’s primary network, they may be a user of regional, national, or international resources within grid-based projects. Each network or project typically requires its own authentication credentials, usually in the form of a digital certificate, to allow the researcher access and authorization to use network or project resources.

Unfortunately, the mechanisms used for AuthN on a grid are not typically integrated with local campus authentication so researchers must often maintain separate sign-ons for each grid they participate in, in addition to their campus network sign-on. Moreover, the IT staff responsible for implementing the policies that govern the use of these separate computing resources must maintain multiple authentication and authorization mechanisms to provide security credentials for each researcher and grid-based project on the campus network. For both users and administrators of resources, these complexities are compounded by the complexities of grid authorization mechanisms – a combination that can create a very challenging environment for practical implementation.

SURAgri

The importance of grid technology as a potential tool for research and education increased at nearly all of the participating sites throughout the duration of the NMI Testbed Grid project. As their experiences and perspectives merged, sites chose to work together within the virtual organization of the NMI Testbed program to further support and enable grid adoption as a component of campus infrastructure. Their work continues within SURAgri.

¹ NMI Integration Testbed information: <http://www1.sura.org/3000/NMI-Testbed.html>

² NMI Testbed Grid information: <http://www1.sura.org/3000/SURAgri.html>

³ SURAgri information: <http://www1.sura.org/3000/SURAgri.html>

⁴ Southeastern Universities Research Association: www.sura.org

Sites join SURAgrid as a cooperative, collaborative endeavor to investigate, develop and implement grid infrastructure. A main benefit is the opportunity for participating campus users to gain access to short bursts of significant amounts of computational resources. Sites may contribute a relatively small amount of persistent resources to SURAgrid but in return gain access to higher burst levels of computational resources than they could otherwise draw upon. (See <http://www1.sura.org/3000/SURAgrid.html> for more detail about SURAgrid.)

Addressing the environmental complexity and duplication of effort in AuthN/AuthZ for grids is a key focus of SURAgrid activity. A primary goal of SURAgrid is to create a scalable infrastructure that leverages local institutional identity (AuthN) while managing access to shared resources (AuthZ) across institutional boundaries. Through this, IT staff should have fewer user profiles and fewer processes to manage, as well as being able to insure that identity is being verified through an authoritative source. SURAgrid participants are addressing AuthN and AuthZ goals on two major fronts:

- 1) Implementing tools to allow campuses to leverage local institutional identity and authorization policies across the SURAgrid and to manage grid services similar to other campus applications;
- 2) Modeling and advancing scalable infrastructure within a working grid – developing and exploring cutting edge tools for AuthN and AuthZ to facilitate secure, grid-enabled application deployment.

SURAgrid participants bring a variety of backgrounds and perspectives to the common effort. Some are relative newcomers to the grid computing environment and are working to increase their own familiarity with grid tools (as typified by NMI Testbed Grid's involvement with the open source Globus Toolkit) and increase the awareness and use of grids at their institution. Some members represent High Performance Computing Centers with established user communities and production computational resources that are seeking to leverage the grid to expand access to their facilities to a larger audience. Other SURAgrid members are expert in grid research and development and, having identified critical areas for grid evolution, find that SURAgrid is uniquely positioned as an ongoing testbed for shared objectives. Common to all SURAgrid members is the interest in how grids and High Performance Computing intersect, and the need to address AuthN and AuthZ within a multi-project and dynamic grid environment.

Authentication

Authentication (AuthN) is the act of identifying an individual computer user (it does *not* include determining what resources the user can access, otherwise known as authorization, or AuthZ). AuthN is the process in which a real-world entity is verified to be who (e.g., person) or what (e.g. compute node, remote instrument) its identifier (e.g., username, certificate subject, etc.) claims. In the process, the real-world entity's authentication credentials are evaluated and verified as being trusted, or from a trusted source. Examples of credentials include a smartcard, response to a challenge question, password, public-key certificate, photo ID, fingerprint, or a biometric (2, 3, 4).

SURAgrid's principles for AuthN further specify that:

- Grid user authentication should leverage existing campus identity management processes. AuthN to various applications should be transparent to a user, seamlessly integrating with the existing campus infrastructure and user-environment.
- Local authentication is translated into grid authentication. Local identity mechanisms already in use on SURAgrid participants' campuses will be incorporated into the PKI-based process required for grid authentication.

Local authentication is a foundational IT process that all higher-education campuses routinely address and which can be implemented through a number of well-vetted mechanisms. True to their heterogeneous nature, institutions participating in SURAgrid do not necessarily implement local authentication with the same mechanism. Kerberos, LDAP (Lightweight Directory Access

Protocol), password databases, and even PKI are all used as mechanisms to establish local identity. SURAggrid participants incorporate various NMI components to integrate their local authentication processes with grid authentication (see Appendix A: Methods of Leveraging Local Authentication to Gain PKI Grid Credentials), including MyProxy and KX.509.

On SURAggrid, the GRIDS Center's Globus security component, which uses PKI, has been deployed. MyProxy and KX.509 are both used as translators between the local AuthN infrastructure of a campus and PKI, which is then used to transport the locally authenticated identity to the grid authentication infrastructure. The Globus Gridmap file is used during grid-to-local authentication translation and informs the grid resource that the user's grid-identity certificate has been verified.

Scaling the Grid through a Bridge CA

Why a Bridge CA?

When the NMI Testbed Grid began, some of the earliest discussions revolved around how authentication would be implemented in the grid. While local campus authentication solutions were known to vary widely, the PKI digital certificates used in the Globus grid security infrastructure are based on standards and could be counted on to interoperate as long as some suitable trust mechanism was established. Participants discussed a few options for building the authentication infrastructure including:

- 1) Operating a Certification Authority for the Testbed Grid and directly issue certificates for users;
- 2) Sites using their existing local CAs and agreeing to install each other's root authority certificates into their trusted Globus certificate store;
- 3) Sites exchanging cross-certificate pairs;
- 4) Implementing a Bridge CA to form the trust path between campus CAs.

The first option is difficult to implement since the different campuses use different local authentication mechanisms, making it hard to automate issuing certificates to campus end users from a central location. While the second option of installing the root certificates from all of the different campus CAs into the Globus trusted certificates directory will certainly work, it presents both management and scalability challenges and was set aside. The third option of having each site cross-certify with each other is an N^2 problem (N =number of sites; each site needing to cross-certify with all others) and was ruled out since it would not scale in a large grid. Instead, the Bridge CA was selected for implementation. In addition to other benefits (see next paragraph), this option addresses the N^2 scaling issue in that each site only needs to exchange certificates with the Bridge CA.

The Bridge CA defines the trust relationship between institutions so that when a user presents their grid identity (certificate) to a resource, that resource knows it can trust the site that issued the certificate to have properly authenticated the certificate holder. An additional strength of deploying a Bridge CA is the alignment of this concept with current effort within the Higher Education IT community to establish a Higher Education Bridge Certificate Authority⁵. As a contributor to this broader community effort, Jim Jokl, University of Virginia and SURAggrid lead for AuthN/AuthZ development, saw the potential within SURAggrid to demonstrate the immediate use of Globus in a bridged environment as a precursor to future availability of HEBCA as an underlying trust infrastructure.

Given the Globus Toolkit's use of certificates for authentication, a PKI Bridge appeared to be the most logical solution to enable the trust infrastructure for a dynamic, multi-institutional grid. The bridge could provide the technical cross-certification infrastructure to enable scalable

⁵ HEBCA information: http://www.educause.edu/content.asp?page_id=623&bhcp=1; also see: Higher Education PKI Technical Activities Group (HEPI-TAG) <http://middleware.internet2.edu/hepi-tag/>

authentication between sites and also serve as a catalyst for policy discussions and decisions required to support real-world usage. Several of the Testbed sites agreed to work with UVa to model and exercise a prototype Bridge CA in order to understand how Globus would function in a bridged environment.

Bridge CA Implementation

Jokl's work with the SURAGrid Bridge CA began with the establishment of a prototype Bridge CA within the NMI Testbed Grid project in February 2004. Jokl had experience in prototyping a similar service as part of his earlier work with HEPKI-TAG to test and understand the bridge-aware PKI path validation logic that was newly available in Windows XP. This project (<http://pkidev.internet2.edu/bridge/>) consisted of creating a few test hierarchical CAs and a test bridge CA and the use of these CAs to issue end-user certificates. These end-user certificates were then used to test applications and to understand the campus PKI infrastructure needed to enable bridge-based path validation in Windows XP. This work was done in early 2003 and, among other things, left the group with an understanding of how to build a campus PKI infrastructure to enable Windows XP to dynamically discover bridge cross certificate pairs leveraging the Authority Information Access (AIA) certificate field.

The prototype NMI Testbed Bridge CA consisted of a Unix computer system, OpenSSL, and a slightly modified version of the scripts that were initially created for the HEPKI-TAG bridge project. The prototype was initially deployed in the lab at the University of Virginia to simulate two campuses operating independent CAs that were bridged together. Certificates from the two simulated campus CAs were used in a Globus test environment to determine if Globus authentication could operate in a bridged PKI environment. This testing determined that, while the OpenSSL-based path validation logic within Globus was not "bridge-aware", Globus could be coaxed to function in a bridged PKI by preloading all of the cross certificate pairs on to the Globus computer systems.

Once the prototype phase was complete, a more secure NMI Testbed bridge CA was created to provide the production bridge service for the NMI Testbed Grid. This CA was built on a dedicated laptop computer running Linux and has been used for all subsequent cross-certifications. The laptop has never been on the network, is kept in a secure location and is only powered on to cross-certify new sites with the Testbed bridge; Certificate requests are moved to the CA laptop using a flash memory card, which is also used to transfer the signed certificates from the bridge computer back to a networked computer where they can be installed on the grid website for download.

The University of Virginia was the first site to cross-certify with the Bridge CA, as part of verifying operation over the network. The University of Alabama at Birmingham became the first external site to cross-certify in April 2004, followed by Texas Advanced Computing Center (TACC) in September 2004, Louisiana State University the following November and the University of Southern California in the first quarter of 2005. The first demonstration of inter-institutional authentication facilitated by the Bridge CA took place at the Internet2 Members' meeting in March 2004 between UAB, UVa and TACC. It was then integrated into a task farming application demonstration with LSU at SuperComputing 2004 (with LSU added to the list of those cross-certified) and shown with this same application at a joint meeting of the SURA IT Committee and regional HPC Directors in March 2005.

Cross-certification is being expanded within SURAGrid as a recommended component for sites adding their resources to the grid. The set-up process for cross-certification has generally taken two to three days of intermittent effort on the part of the site desiring to cross-certify and the University of Virginia as the host and operator of the Bridge CA. This frequently involves more than one attempt since the process is generally new to those operating the campus CAs. To help

with this, Jokl established a web page⁶ that documents the overall process for cross-certification and has continued to add detail and lessons learned to the “how to” portion, which has helped reduce the time required for sites to cross-certify.

As the first few sites went through the process of cross-certifying production campus CAs, a great deal was learned about the process. Issues ranged from incompatibilities between the Bridge CA certificate profile and those already in use at the various campus CAs, to difficulties in generating the certificate requests using the correct keys. Again, Jokl documented each of these issues as they were discovered on the project web site (<https://www.pki.virginia.edu/nmi-bridge/>) to ensure that new sites would not run into the same set of problems.

Documentation of the cross-certification process within SURAggrid is also being expanded through outreach efforts of other SURAggrid sites. This paper represents one example of that, reflecting the desire of multiple SURAggrid participants to actively disseminate knowledge and experience gained in SURAggrid to catalyze broader deployment and experimentation. In addition, Art Vandenberg, Georgia State University, and a team of students are documenting and validating an end-to-end “Twelve-Step” process – from establishing a campus grid CA to cross certifying with the Bridge CA. This “Twelve Step” process is being developed to directly assist other sites in setting up a campus grid CA and cross-certifying with the SURAggrid Bridge CA. The document aims to provide an easily followed summary for achieving a secure SURAggrid authentication model that meets SURAggrid policy guidelines with a minimum of resources.

Through the diversity of institutions cross-certifying through SURAggrid, a deeper understanding of potential issues – both technical and organizational – is emerging. This includes the variety of ways sites may have already implemented a certificate, what implications the cross-certification process has for their current certificate profile, how well the current certificate and related local processes mesh with the Bridge CA, and how the Bridge CA can scale.

Integration with a Newly Created Campus CA

The University of Alabama at Birmingham was in the process of constructing their campus grid environment at the time of their cross-certification with the Bridge CA. While cross-certifying early on in the grid building process can be easier than integrating with an existing Campus CA – it’s easier to align local CA certificate profiles with the Bridge CA, for instance – it’s also recommended that organizations in this situation first understand the basics of using certificates with SSL, since this is how Globus uses certificates.

There are many resources to assist with this, including <http://middleware.internet2.edu/hepki-tag/opensrc.html>, http://www.globus.org/grid_software/security/simple-ca.php, and "SSL and TLS" by Eric Rescorla. Reviewing such resources can help clarify many of the decisions that need to be made in setting up an application-specific PKI environment, however since much of this material covers traditional PKI environments, application-specific PKI deployments are left as an exercise for the reader. In particular, campuses starting to build a campus CA should strongly consider implementing the PKI-Lite certificate profiles. These profiles have been designed to support many campus PKI applications and are known to be compatible with for use with Globus in a bridged environment. The PKI-Lite profiles were modified based on the results of this testbed work to ensure compatibility with the bridged CA model. Thus, even if an institution isn’t ready to cross-certify in a bridged CA model, using these resources to build your CA will help ensure that a future cross-certification with a Bridge CA will go smoothly.

Beginning March 2004, UAB began investigating and building an application-specific PKI infrastructure for its campus grid, to be known as the UABGridCA. This CA is intended to be used only for supporting the UAB grid computing infrastructure on campus and not to be a general-purpose CA for other campus processes (e.g., signing formal UAB documents). The UABGridCA

⁶ <https://www.pki.virginia.edu/nmi-bridge/>

does leverage, and is integrated with, the campus AuthN infrastructure and makes assurances of user identity to grid resources based on the policies of this core AuthN framework. Since UAB was building its grid-specific CA at the same time SURAggrid was exploring the Bridge CA, UAB was able to address any Bridge CA compatibility issues during the UABGridCA roll out and adapt the campus CA configuration to the configuration requirements of the Bridge CA. The most notable benefit from this was that the campus certificate profile was developed to align at key points with profile requirements of the Bridge CA.

A valuable realization from this is that an institutional certificate infrastructure should not be an impediment to a grid implementation and that grid deployment does not need to be tied to an institutional PKI implementation. Instead, organizations should view the certificate infrastructure simply as a requirement of the grid as an application and work to leverage existing institutional identity systems to generate grid-specific identities in the form of a PKI certificate, keeping this transparent to grid users.

Integration with an Existing Campus CA

Several of the SURAggrid sites cross-certifying with the Bridge CA have integrated the process with operation of an existing campus CA. This is more difficult than parallel implementation with a newly created campus CA and often requires additional time due to the presence of existing infrastructure and the need to align with active organizational policies and processes. Though the cross-certification process itself does not require many actual labor hours, pieces of that process can be separated by days, weeks or even months, depending on the priority of the implementation, the nature of the organization's need, and the resources that can be applied, Examples below provide insights from integration with existing campus CAs at several SURAggrid sites.

University of Michigan

The University of Michigan is currently working to complete the cross-certification process with the Bridge CA. Like TACC and USC, UMich already has an organizational certificate authority in place. The UMich root certificate was issued by CREN. This umich.edu root CA has signed the University-wide Kerberos CA (KCA). Likewise, the MGRID CA, which signs host certificates for MGRID related grid compute resources, sits below the umich.edu root CA. UMich is currently working with UVA to determine what level in the UMich certificate hierarchy is best to cross-certify with the Bridge CA. For instance, cross-certification at the umich.edu root CA may be more effective since only a single cross-certification would be needed, but the certification may not work at this level due to the path validation logic used in Globus. Therefore, cross-certification may need to be done at the MGRID CA and KCA levels to cross-certify all of the certificates (user and machine) that are used in MGRID.

University of Southern California

USC's situation was similar to TACC's in that they had a pre-existing certificate profile that proved incompatible with the Bridge CA. Their original attempt at cross-certification was frustrated by this incompatibility, however, once identified, the problem was corrected without significant modification to the current USC PKI Lite CA, which is an important consideration for sites that are already supporting an institutional certificate authority. A further delay was due to technical personnel being unavailable to devote time to make the minor changes needed.

Texas Advanced Computing Center (TACC)

The TACC CA was established and deployed in a production environment prior to the development of the Bridge CA, therefore TACC certificates and signing policy had not considered the requirements of the Bridge CA. Rather than change their current certificate and environment to fit with the Bridge CA specifications, TACC used OpenSSL to generate certificates as needed to work with the Bridge CA. They maintain a special OpenSSL profile that is used for creating Bridge CA certificates. At some point, they plan to consider updating their profiles to meet the

Bridge CA requirements, but will first need to ensure that this will not impact grid users who have been issued certificates using TACC's current profile.

University of Virginia

The University of Virginia had a pre-existing campus CA in production that was designed to support a wide range of applications. Supporting the use of the Globus toolkit both locally and in an inter-institutional grid was an important goal. Given UVa's previous work with bridge CAs, they knew that their existing certificate profile was compatible with the bridge and the cross-certification process itself was relatively easy to complete since no coordination between sites was needed. One of the issues uncovered during the process is that the level in the campus PKI hierarchy where the cross-certification occurs can cause problems with the simple path validation logic used in the Globus toolkit. Cross-certifying at the lowest possible level in the campus hierarchy eliminates these problems.

It should be noted that it is not "risky" for a site to make changes to their current certificate, However, doing so can require changes in other aspects of their environment and applications that rely on the certificate. As part of the NMI Integration Testbed program, participating sites provided feedback to PKI-Lite⁷ developers regarding this and other aspects of using a Bridge CA, which resulted in several changes to PKI-Lite software that facilitated the cross-certification process. The key change was the recommendation that campuses not include Subject Name and Serial Number in the Authority Key Identifier certificate field, which causes problems for the bridged CA, and instead populate the field with only the key identifier. With the evolution of understanding and documentation to support PKI deployment and cross-certification with the Bridge CA, subsequent SURAggrid cross-certifications are proceeding more quickly.

Authorization

AuthZ typically refers to the process of determining the eligibility of a properly authenticated identifier (e.g., person) to access a networked object (application, function, or resource). Although a SURAggrid user may be authenticated as discussed in the previous section, as with other grids, that user has no rights to use SURAggrid resources until they are *authorized* – that is, granted access to particular resources based on who they are or to what group(s) they belong. Authorization can also refer to the issuing of a token that proves a subject has the right to access an object, to the right or permission that is granted to access the object, or to the token itself (e.g., a signed assertion). Signed assertions and other authorization characteristics may be kept in either system-specific or campus-wide infrastructure directories, within a file system, in an external device (e.g., a smartcard), or in a separate data system (2, 3, 6). After successfully cross-certifying through the Bridge CA, a SURAggrid user still needs some form of authorization, most likely in the form of a user account, to begin using grid resources.

AuthZ is used on grids to enforce conditions of use on a grid resource as specified by the resource owner. Organizations develop policies that specify these conditions and use authorization tools and technologies to implement those policies. While many authorization technologies have been developed and are in active use on various grids, there is no universal consensus on which technologies are or will prove to be most effective in an inter-institutional grid environment. In Globus-based grids, the Gridmap file can be viewed as the mechanism for AuthZ. Authentication occurs when the users local identity (not necessarily mapped to a local institutional identity, though that is the case in SURAggrid) is expressed as a certificate that is understandable within the grid PKI infrastructure. On the resource side, authentication is completed when this certificate is verified. From that point on, use of the certificate to obtain access to grid resources can be considered authorization. The Bridge CA further facilitates AuthZ by defining a trust relationship between institutions such that, when a user presents their grid identity (certificate), the resource being accessed knows it can trust the site that issued the certificate to have properly authenticated the certificate holder.

⁷ The PKI-Lite Framework: http://middleware.internet2.edu/hepki-tag#PKI_Lite

Implementation of AuthZ within SURAggrid will evolve using the following principles:

- *Model a heterogeneous, real-world environment* – As separate and autonomous institutions, SURAggrid participants are not required to use a specific and common CA provider. While this increases the complexity of implementing AuthN and AuthZ, it will ultimately result in a more robust and scalable solution. The BridgeCA is consistent with concepts of federated, autonomous regions emerging within the national Higher Education community.
- *Phased deployment* - In the initial months of the NMI Testbed Grid, sites had few sharing policies in place regarding contributed resources and didn't feel a need to implement technology-based authorization tools. To gain access to grid-based resources, site representatives simply phoned each other to work out the details. This can be an easy and effective model but not one that can scale as the number of resources and users grows. For the future, SURAggrid will address authorization incrementally, creating and implementing mechanisms as needed to support current usage but keeping the long-term goal in mind - scalable and dynamic authorization that respects the administrative policy boundaries of each resource provider and empowers SURAggrid site representatives to serve their users without having to coordinate with other site representatives.

Common Authorization Methodologies

SURAggrid participants have integrated their local campus authorization with SURAggrid authorization in various ways, as illustrated below. Common methods include the use of the Globus Gridmap file, interaction with schedulers, and interaction with file or operating systems. The distinction between methods lies in where authorization decisions are made within the system rather than the inherent complexity of making them. Generally, a single institution will have some form of integration with all three of these methods, depending on the demands and complexity of their local system environment.

In Globus-based grids in particular, the Globus Gridmap file provides basic authentication but authorization is not directly addressed. Instead, information is provided for a site to use, if and how, they choose to implement authorization policies. Upon presentation of a PKI certificate, Globus provides mapping from that certificate to local accounts listed in the Gridmap file via a "yes/no" policy. A "yes" is generated if the users' listing is found in the Gridmap file, then the user is allowed access and their job runs locally under whatever environment pertains to their local account. A "no" is generated if the users' listing is not found in the Gridmap file and the user is refused access to the resource. The Globus Gridmap file can be configured to be more granular and provide a somewhat dynamic yes/no policy (e.g., certain users are allowed daytime access, others nighttime only). The Gridmap file can also be modified to generate a request to map or unmap a specified Distinguished Name to a user's account. A callout feature in version 3.2 of the Globus Toolkit also provides the ability to customize a Gridmap file and a mechanism for doing fine-grained authorization in the GRAM Jobmanager (7).

University of Michigan

The MGRID authorization process involves the use of several currently available tools (e.g., PBS Pro, Condor) depending on the specific MGRID resources a user needs to access. MGRID would prefer that the resources they've dedicated to the SURAggrid be mapped to person's DN (distinguished name) temporary accounts, thus giving them a priority in Condor. MGRID uses the Globus software to make a callout to the XACML language that provides a rich set of features to support fine-grained control in the form of a policy. Based on the specifications of the resource owner, an MGRID XACML policy can use any of the following criteria (which can also be combined as a set of criteria) as the basis for an authorization decision:

- name of executable to be run
- required lifetime of output data

- amount of data to be stored
- required duration of execution
- job manager (for example: fork, PBS, Condor, etc.)
- membership in a group
- number of compute nodes requested
- name of queue (PBS, Condor, and SGE support named queues)
- resource requested
- start time of job
- role in which user is acting

Using this infrastructure, policies such as the following can be instantiated:

- Allow any user to submit single-node jobs that run for 5 minutes or less on the head node.
- Jobs that take longer than 12 hours or using more than 20 compute nodes can only be run on the LOW_PRIORITY queue.
- Users running the Atlas executable can only run on the atlas queue, and must be a member of the ATLAS group.
- If a user is a member of the ADMINISTRATOR group then they can run jobs with no restrictions.
- If a user is a member of the BAD_PEOPLE group they can't run any jobs.

Texas Advanced Computing Center (TACC)

TACC is participating in several grid projects in addition to SURAGrid - a campus grid (UT Grid), a statewide grid (TIGRE) and a national grid (TeraGrid). In the absence of a globally accepted set of tools for authorization, TACC is currently using different tools for UT Grid and TeraGrid to process requests for new user accounts, creation of accounts across their local resources, and mapping of the accounts to the requesting user's distinguished name in their X.509 certificate.

For UT Grid, a user requests an account via the UT Grid User Portal, which is then forwarded by the portal to the TACC CA administrator for authorization. (If the user already has a valid account on a TACC resource, this step can be bypassed.) There is a formal process for the TeraGrid in which users request accounts on-line via a central authorization service, followed by a manual step to complete user authorization. The Principal Investigator of the project is contacted for verification and, once manual verification is complete, the user information is entered into a TeraGrid-wide database. Requests for new user accounts are automatically forwarded to each TeraGrid resource provider site. Most, if not all, of TACC's current users prefer shell-logins, therefore TACC supports both shell-logins and browser based user portals for people to submit and manage jobs.

University of Southern California

Local authorization considerations should prevail and USC policy requires that individual accounts be created for grid access, in contrast to, for example, mapping grid IDs to a pool of temporary accounts. Consideration is also being given to how to represent "grid guest" accounts in the enterprise directory, and how to insure these accounts are deleted in an acceptable manner once they are no longer in use. These "grid guest" accounts, when originally created, will be authorized to use HPCC's Sunfire 15k⁸ and the USC Condor flock (which runs on Solaris workstations in public user-rooms). Access to the Linux cluster will not be automatic upon account creation; an application for an allocation will be needed. The SURAGrid accounts will have their own disk area; the pathname is under discussion. USC is working to develop a long-term solution that will allow automatic account creation on persistent resources (but not dedicated solely to SURAGrid) and be compatible with other SURAGrid sites, while retaining local control over methods, mechanisms and policy.

⁸ almaak.usc.edu for timesharing; rcf.usc.edu for short-term execution

Texas Tech University

TTU provides grid user accounts for HiPCAT⁹ (High Performance Computing Across Texas) user groups. These accounts initially allow resource use up to a certain level, and attributes can be changed if a user needs more resources. Though this can require some additional effort, the process ultimately provides more resource control.

University of Alabama at Birmingham

Authorization for SURAggrid members to access resources on UABGrid (currently under construction) will reflect a desire to respect inter-departmental autonomy in UAB campus resource management and therefore mirrors many of SURAggrid's inter-institutional design considerations. The UABGrid is composed primarily of central IT resources, including those that will be made available to SURAggrid.

The UABGrid is envisioned to be an integrated interface for SURAggrid in that the UABGrid authorization model will enable the resource provider to set authorization policies for uses based on attributes supplied by the users identity provider (typically their home institution or project). These user identity attributes will be collected by the UABGrid portal and made available to grid resources in a number of ways. UABGrid is currently exploring the use of a VOMS (Virtual Organization Membership Service)¹⁰ database system to package attributes received via Shibboleth and supply them to grid resources. In order to facilitate participation in SURAggrid, UAB has provided a dedicated cluster that uses the simple Gridmap file authorization model; if a SURAggrid user is in this Gridmap file, they will be authorized to use the cluster resource. As load increases, a scheduler will be introduced to manage resource contention.

UABGrid will move away from a homogenous system configuration approach and toward a heterogeneous environment supporting local autonomy. UAB is currently pursuing an approach similar to the University of Michigan's Walden¹¹ project, to give resource providers full control over the allocation of resources to others. Ultimately, compute center autonomy needs to be respected and the ideas embodied in Walden can lead to an effective way to implement this. Walden replaces the Globus Toolkit's Gridmap file with a more extensible architecture, designed to overcome the Gridmap file maintenance problem, allow the separation of a grid users identity from their unix account identity, and give resource providers control over the account definition conventions. Through this, SURAggrid users at UAB can be allocated resources based on the specific resource provider's policies.

University of Virginia

UVa is testing Globus groups along with other groups on campus. Owners of resources make decisions about who will be in a group to use their resource. UVa also plans to leverage policies from Condor clusters - resource owners will donate some (e.g., 20%) of their own resources in order to join and use others' clusters.

Georgia State University

Georgia State is looking at scheduling as an important component of implementing the SURAggrid authorization policy. Through SURAggrid participation, they are learning about various scheduling solutions and how that experience can be applied to a campus grid. It is possible that Georgia State's "campus grid" might become just a virtual piece of the SURAggrid model as there are definitely interesting benefits in a cooperative model – for example, if a contribution of a relatively small number of CPUs to SURAggrid could result in potential access to the total of CPUs contributed by other sites.

⁹ www.hipcat.net

¹⁰ <http://www-unix.grids-center.org/r6/ecosystem/security/voms.php>

¹¹ <http://www.mgrid.umich.edu/projects/walden.html>

Future Phases

Short-term

The activities that SURAggrid participants will undertake to evolve authentication and authorization processes in the coming year include:

- *“Automating” the Gridmap file.* This will be done through either LDAP callouts or by generating a group Gridmap file. A "call-out" could be utilized from the Globus code at the same point where the Gridmap file would normally be read. This allows a more scalable and timely system to determine who is authorized to utilize resources, particularly for a dynamic multi-institutional setting. The service that is accessed to make this determination can return much the same system information as would normally be encoded in the Gridmap file, but new possibilities present themselves. For example, the response to the call-out could be source dependent as well as time dependent. A different list could be provided depending upon when and where the calls are made to the service.
- *Increasing the application set that can benefit from and demonstrate the ease of cross-certification.* With the Bridge CA in place, the SURAggrid trust framework has been simplified and SURAggrid users are able to use local campus grid certificates to access remote grid resources. This functionality enables transparent access that can readily adapt to dynamic changes in grid resources or the SURAggrid userbase. SURAggrid, led by Application Development lead, Art Vandenberg (Georgia State University) is working to identify and promote applications with specific need for this capability, to affirm the efficacy of the Bridge CA concept and further refine the architecture.
- *Enabling each site to police their usage.* SURAggrid participants are responsible for their site's usage of SURAggrid resources and are working to develop mechanisms to monitor this. This capability will be provided, in part, via the SURAggrid portal, under development by TACC. For jobs submitted through the portal, sites will be able to track usage through the portal interface. To track jobs not submitted through the portal, sites could use a report generator from the job manager or scheduler. For this latter task, SURAggrid may implement a solution developed by the University of Michigan's MGRID project. While MGRID staff used both Condor and Condor-G extensively on MGRID for scheduling (the latter as the front-end scheduler), their applications needed the more detailed information that would result if the accounting information from their Condor processed jobs output was in a GGF Usage Record Working Group-compatible format. UMich patched the Globus job manager for Condor and PBS, creating a modified Condor¹² scheduling package solution that outputs accounting information in a format compatible with the GRIDS Center's PBS Accounting Toolkit¹³. This modified output data can in turn be processed by the MGRID Accounting package¹⁴ (an extension of the PBS Accounting Toolkit).

Long-term

For longer-term functionality and scalability of AuthN/AuthZ, SURAggrid is exploring:

- *Metascheduler for SURA Grid.* A Metascheduler is a tool for scheduling jobs, given numerous possible resources and accompanying schedulers. For an entity like SURAggrid, which will include many and various resources, a Metascheduler is a very important tool to effectively utilize group resources and provide a simplified interface for users to run their jobs. One option under consideration is the use of MGRID's NSF-funded development, MARS

¹² Condor's strength is in its ability to make "worker node" machines work together as a computing resource. Condor-G, on the other hand, is the marriage of technologies from the Condor and the Globus projects and makes Condor more "grid-aware". Condor-G is a computation management agent that provides front-end scheduling functionality for Condor.

¹³ PBS Accounting Toolkit information: <http://pbsaccounting.sourceforge.net/>

¹⁴ MGRID Accounting package information: <http://www.mgrid.umich.edu/projects/accounting.html>

(<http://www.mgrid.umich.edu/projects/mars.html>), which is being designed to address the problem of scheduling in the context of many resources, some of which may need to be co-scheduled (e.g., the network and a set of computers and a storage area).

- *Promoting increased sharing of resources.* SURAggrid participants are considering ways to encourage resource owners to add their resources to SURAggrid, including unique, high performance or production resources. One possible approach is that a resource owner would earn “credit” when their resource is used, based on a “market value” of the resource (e.g., number of cycles, power, quantity, and availability). Market value could be determined by creating a “SURAggrid economy” that reflects the value of various resources to the user community. Resource owners could make their resources accessible for bidding (either manually or automated) and the value of specific resources would float up or down depending upon the needs of users – users will place higher bids and drive up the value of certain resources. In this context, a cluster may be very powerful, but if unreliable, it may be valued lower (at least for some applications) than a less powerful but stable, system elsewhere. Though such a “grid economy” may be complex to develop, prototype and implement, it could provide a means to fairly compensate resource contributors and motivate others to share.

References:

- (1) Foster, *The Physiology of the Grid: An Open Grid Services Architecture for Distributed Systems Integration*, 2002.
- (2) <http://www.nmi-edit.org/glossary/index.cfm>
- (3) <http://www.gridforum.org/documents/GFD.42.pdf>
- (4) <http://middleware.internet2.edu/core/authentication.html>
- (5) http://www.nsf-middleware.org/testbed/testbed_status.asp#studies
- (6) <http://middleware.internet2.edu/core/authentication.html>
- (7) <http://www-unix.globus.org/security/callouts/>

Appendix A

Methods in Leverage Local Authentication to Gain PKI Grid Credentials

1) Password Database

Georgia State University

Georgia State is initially using certificates provided by Globus for testing of the Globus grid security infrastructure. For deployment of a campus grid, however, they are planning a more rigorous grid security infrastructure through the implementation of a campus GridCA that will follow PKI-Lite guidelines and cross-certify with the SURAGrid Bridge CA. They will use campus enterprise logins based on the existing campus enterprise authentication infrastructure - a metadirectory architecture and comprised of an Oracle based *Person Registry* password table - to authenticate users and provide certificates for SURAGrid access.

2) LDAP (Lightweight Directory Access Protocol)

University of Alabama at Birmingham

UAB leverages the human resource and student databases to enable members of the university to self-register for a campus identifier and associate a password with this identifier. These credentials are verifiable via LDAP and form the backbone of the campus authentication infrastructure that is leveraged by many applications. A campus grid CA application was built using PHPki (<http://phpki.sf.net>), an open source web-based interface based on OpenSSL. In order to leverage the campus credentials and allow anyone at UAB automatically generate their working grid certificate, a WebISO system using Pubcookie (www.pubcookie.org) to verify the credentials via LDAP was introduced and integrated with the web interface of the CA software. Anyone who has followed the standard identity registration processes of the University can now generate their own campus grid certificate and use that certificate across the UABGrid and SURAGrid.

3) Kerberos

University of Michigan

The University of Michigan has a unique-name database (unique names, permanently assigned) with 100,000's of entries. Replicating or recreating this information in another form for use, for instance, in grid AuthN/Z, is not tenable. The University had previously chosen Kerberos as the local authentication methodology and the challenge for grids and authenticated Web access was to provide a translation from Kerberos into the PKI space. To meet this challenge, University of Michigan researcher, Bill Doster, developed the KX.509 security protocol, which the University of Michigan subsequently contributed to the NMI. The protocol, now known as NMI-EDIT's¹⁵ KX.509/KCA (Kerberized Certificate Authority¹⁶) or simply as "KX.509", bridges the existing standard for authenticating *within* a community (Kerberos) to the existing standard for authenticating between communities (X.509) (5).

KX.509 allows any locally authenticated user to have a PKI identity on-demand. This service is provided by a KX.509 software component called the KCA (Kerberos Certificate Authority). By running KX.509 client software, any user with an existing local Kerberos TGT (ticket-granting ticket) for their local institution can create a short-lived, PKI equivalent identity for use on grids

¹⁵ NSF Middleware Initiative-Enterprise and Desktop Integration Technologies (NMI-EDIT): <http://www.nmi-edit.org/>

¹⁶ KX.509/KCA information: http://www.citi.umich.edu/projects/kerb_pki/

and authenticated web services. The KCA component of KX.509 creates the short-term PKI identity when the user runs KX.509 client software. The client software uses the users existing Kerberos tickets to authenticate their request to their local KCA (the KX.509 server) to have a new X.509 client certificate signed. When a local KCA is cross-certified by one or more Bridge CAs, users with a short-term PKI identity from that KCA can access grid-based resources at other sites that have been signed by (cross-certified by) the same Bridge CA(s).

The University of Michigan has implemented KX.509 for their Kerberos users. This allows UMich users with a uniqname to interact with grid-based resources offered by UMich, such as those on the MGRID (Michigan Grid Research and Infrastructure Development)¹⁷. KX.509 has become a key technology for authentication in grids at the University of Michigan, as well at numerous other institutions, including SURAgrid member institutions.

University of Southern California

USC has a two-pronged approach to local authentication, split between users and non-users (hosts and persistent services).

If a host or persistent service requires a PKI certificate, a responsible party (human) creates a CSR and emails it to the group in charge of the USC PKI Lite CA. One of the CA verifiers then checks that the CSR meets posted criteria, and then arranges for in-person identity verification. Once identity is verified, the CSR is dispatched to the group responsible for signing certificates. Once the CSR is signed, it is returned to the responsible party.

Users requiring a PKI certificate first login to a USC-supported host (either a desktop or one of the time-sharing systems). The user then authenticates to Kerberos, receiving in return a Kerberos TGT. When the USC user runs KX.509 client software, the USC KCA issues the user-proxy certificate (the newly signed X.509 client certificate signed that provides the short-term PKI identity) based on the user's Kerberos TGT. USC's KCA is subordinated (i.e. its certificate is signed by) the USC PKI Lite CA. Thus the Bridge CA only needs to cross-certify with the USC PKI Lite CA, though operational efficiencies may suggest the Bridge CA should cross-certify to the USC KCA as well.

4) PKI

University of Virginia

UVa maintains a central user identification and account management system. This system receives data from all of the various sources of payroll and student information and generates a unique identifier for each individual affiliated with the university. This system is also a core component of UVa's central directory infrastructure. Once a person is in the directory, they can activate their accounts on any of UVa's central systems (e.g., electronic mail, network file storage, Unix, etc) and thus obtain a password.

The university operates a central campus PKI Certification Authority that is also used to support campus Globus activities. In order to obtain their certificate, a user goes to a web site, enters a UVa central system password and other private information about themselves, and presses the submit button on their browser. The browser then generates the user's key pair and sends the certificate request to the campus CA. The CA generates and signs the certificate and automatically downloads the result and places the certificate into the user's certificate store.

The user obtains their key and certificate in the format that Globus requires via exporting their certificate and key into a PKCS-12 formatted file and using OpenSSL to obtain the separate

¹⁷ See *University of Michigan: How Grid Science and MGRID are Changing Research and Education* at: <http://www1.sura.org/3000/NMI-Testbed/UMich-MGRID.pdf>

files for the certificate and key. In the future, UVa intends to provide a simple web interface for splitting the certificate and key into individual files.

Texas Tech University

TTU uses a single sign-on (SSO) mechanism as part of their campus “eRaider” account management system. Students, faculty, and staff need only use their own eRaider username and password to access various electronic resources at Texas Tech. An eRaider account allows Texas Tech users to send and receive email, update online directory information, create and manage an email alias, access the Internet, access TechSIS (the Student Information System), register, enroll in computing short courses, take advantage of online training, create a personal website, and download free software.

The TTU High Performance Computing Center (HPCC) currently uses a PKI Certificate Authorization with the help of a simple CA that came bundled in the Rocks¹⁸ clustering software stack. Users may apply for certificates if they have existing HPCC accounts on the compute clusters via *globus-cert-request*. A file is installed in the */etc/grid-security/* directory to specify whose authority the host will recognize. TTU users generate their certificates after setting environment variables. The users then supply a pass phrase to be used each time they run Globus. The Globus certificate request program saves two files, *userkey.pem* and *usercert_request.pem*, which are mailed to the certificate authority manager at the TTU HPCC. The user copies these files to the same the directory on each machine on which they will run Globus.

Users need a host certificate for each of their hosts running Globus. They run the host certificate request on each machine, and then mail the *hostcert.req* file to the campus CA administrator. The host certificates and keys are installed in the */etc/grid-security/* directory, and owned by *root*. Since some hosts do not have email clients, the email need not be traceable to the host. It is easiest to use the same email account at the same time as for the user certificates.

Texas Advanced Computing Center

TACC provides a user portal for its campus grid (UT Grid) users as well as external partners who are involved in collaborative activities with TACC. By creating portal accounts, users also get a grid certificate signed by the TACC CA that allows them to interact with grid resources through the portal. Once a user has logged in to the user portal, a proxy certificate is placed in their session which allows the user to perform typical grid tasks such as monitoring the available resources on the grid, accessing and manipulating their files and data, and submitting, monitoring, and deleting their jobs.

The authentication process currently requires a valid system account on a TACC compute resource. TACC is working on an enhanced account management system, which will remove this dependency on an existing user account and allow users to more easily manage their system accounts, allocations, and grid credentials from within the portal interface.

For their CA, TACC is using CACL (<http://www.npaci.edu/CA/cacl.1>), an OpenSSL based CA implementation from the San Diego Super Computing Center that provides client software to automate the creation of user certificates from local machines.

¹⁸ <http://www.rocksclusters.org/Rocks/>