

## **Project Summary: NMI DEVELOPMENT: Policy Controlled Attribute Framework**

---

Secure authentication and authorization are growing challenges for the distributed, multi-institutional collaborative infrastructures and applications that NMI seeks to support, and without focused effort are likely to pose insurmountable obstacles to the realization of NMI goals. Will an institution allow access to expensive or sensitive resources based on information about a user obtained from an external source? If so, then inter-institutional collaboration is feasible; if not, then cyberinfrastructure breaks down into a set of mutually distrusting enclaves. Furthermore, the privacy aspects associated with the sharing of information have so far been ignored in the middleware developed for the Grid community. It is clear that as the Grid technologies mature and find broader application use, privacy considerations have to be taken into account.

The pressing need is for *robust, campus-integrated infrastructure and tools that will allow for secure verification of a user's attributes*, such as identity, institutional affiliation, and role in a collaboration. Such an infrastructure and tools do not yet exist. Fortunately, however, the building blocks do: the Grid Security Infrastructure (GSI) that is the basis for many NMI-based infrastructures and applications provides robust secure authentication, but is not yet anchored in campus infrastructures; the Shibboleth attribute service (also NMI software) allows for controlled access to attribute information in campus directories, but is not yet integrated with GSI. Both of these solutions have significant traction in their respective communities, however are disjoint in their development and deployment today. *The integration of GSI and Shibboleth, and other related technical developments, can provide the needed capabilities for a robust attribute infrastructure.* Such integration is the goal of this project.

This project will deliver a framework that allows participants in multi-organizational collaborations to control the attribute information that they want to publish, share, and reveal to other parties. Those parties will be also be able to determine whether they possess the capabilities to access a service by matching their capabilities with the service's shared policy of required attributes. Finally, pseudonymous interactions will be supported through the use of anonymous public key credentials that are mapped to the client's identity at the client's own discretion.

The project substantially leverages on and extends existing technologies, primarily Internet2's Shibboleth, the Globus Alliance's Globus Toolkit, and NCSA's GridLogon Service. The framework will use Shibboleth's Attribute Authority service (SAAS) and its attribute release policies to restrict the attributes communicated to other parties. We will enhance these Shibboleth services by enabling Web services access through integration with the Globus Toolkit. To enable pseudonymous deployment, a module will be developed for the GridLogon service to allow authenticated users to obtain public key credentials that do not reveal their identity, yet are fully compatible with the Grid Security Infrastructure. Lastly, formats and protocols will be developed and implemented to express, publish, share, and match attribute-related policies and capabilities.

Finally, the new capabilities to be developed in this project will be integrated into NMI software, and deployed and evaluated within the context of major NSF infrastructure and application projects.

**Intellectual Merit:** While the primary focus of both the NMI program and this project is development and integration, the proposed work will provide new insights into requirements, implementation techniques, policies, and usage patterns relating to the management and use of attributes in complex multi-institutional settings and advanced applications.

**Broader Impact:** This project will address what is rapidly becoming a crisis, namely the lack of robust, campus-integrated infrastructure and tools for secure verification and privacy-preserved sharing of user attributes. By so doing, it will allow existing communities and resource providers to continue their cyberinfrastructure activities and reduce barriers to the integration of new communities and resource providers.

National Science Foundation  
National Middleware Infrastructure Solicitation NSF 04-555

# **NMI DEVELOPMENT: Policy Controlled Attribute Framework**

## **A Privacy-Friendly Framework for Policy Enforced Release, Sharing, and Matching of Attribute Information**

**(Revised version based on two-year budget)**

**For the period October 1, 2004 – September 30, 2006**

**Submitted May 14, 2004**

### **Principal Investigators**

---

Von Welch  
National Center for Supercomputng Applications  
University of Illinois  
Urbana, Illinois 61801  
Tel : 217-265-7139  
Email : [vwelch@ncsa.uiuc.edu](mailto:vwelch@ncsa.uiuc.edu)

Katarzyna Keahey  
Computation Institute  
University of Chicago  
Chicago, Illinois 60637  
Tel: 630-252-1673  
Email: [keahey@mcs.anl.gov](mailto:keahey@mcs.anl.gov)

### **Co-Principal Investigators**

---

Frank Siebenlist	Argonne National Laboratory, Argonne, IL.
Tom Barton	University of Chicago, Chicago, IL.

## Project Description

---

### A. The Challenge: Attribute-base Authorization for Virtual Organization with Privacy

---

A virtual organization (VO) refers to a collection of users and resource that span multiple organizations that wish to share resource in a controlled manner in order to achieve a goal [14][32]. Example of VOs can be found through out NSF, such as the Network Earthquake Engineering Science Grid (NEESGrid)[10], Grid3[17][34], and the TeraGrid[4]. The management of access control policies in large VOs is a complex task, as users are distributed over a number of organizations and tends to be dynamic with users coming and going over the life of the project and changing roles. Since the VO users are not all part of a single organization, the VO is forced to set up an attribute service, such as the virtual organization management service (VOMS)[7], to manage user privileges. However this approach has a number of serious drawbacks, mainly that the members of the VO are typically focused on their domain science and not running security service. This lack of expertise leads to the dangerous situation of security services being run by users without experience in this area and typically in physical locations not suited to sensitive equipment.

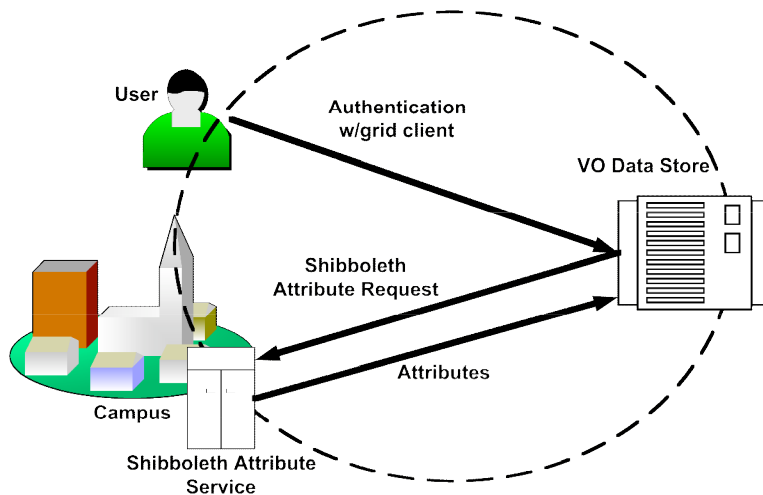
Today most VOs only share resources within their local community. In these cases, given the sensitive nature and value of the resources being shared, it is natural that authorization is based on knowing the identity of the requestor. However, as VOs begin to share their resources with larger communities, knowing individual identities of users may not only be impractical, it may be undesirable. For example, a VO focused on genomics may generate large amounts of data that they wish to make available to other genomic researchers; however if the usage pattern of any particular researcher accessing the data were to be analyzed, that analysis could give clues at to that researcher's direction and may allow someone to steal ideas and beat that researcher to valuable publications or patents. This makes the ability to have identity-hiding modes of operation beneficial, in order to protect privacy when it would encourage the use of the system while still providing assurance to the system owner that only appropriate users are being served.

### B. Our Approach: Integration of Campus and Grid Security

---

Our proposed work will produce a solution to the challenges described in the previous section through the leveraging of campus attribute technologies, namely the Shibboleth identity federation service[31][6], and Grid technologies[11][12]: specifically, the Globus Toolkit@[13] and its Grid Security Infrastructure[37] and the MyProxy-based[27] GridLogon Credential Service . These different technologies are all distributed as part of NMI, and the results of our work will be contributed to NMI. Thus, we see this proposed work as adding considerable value to the overall NMI software solution [28]. We give a detailed description of these technologies in the subsequent section.

Figure 1 shows the basic functionality we will develop through our integration. A user will use a standard grid client from the Globus Toolkit (e.g., globus-job run, globus-url-copy) to contact a grid service, as is common practice in essentially any NMI-based application or infrastructure today. A service will authenticate the user using standard grid authentication, determine their home domain from that identity and contact a Shibboleth attribute service at that domain to obtain that user's attributes. Using those attributes, the service will make an authorization decision based on the attributes, for example, only allowing members of the physics faulty department to access a data set. In this way, we deliver considerable new value (authorization based on attribute registries maintained by sites) without requiring substantial changes to applications or deployments.



**Figure 1. Our Goal: Integration of Campus Shibboleth infrastructure and Grid software.**

## C. Leveraged Technologies

---

We provide further technical details on the Globus Toolkit, GridLogon, and Shibboleth.

### C.1. Globus Toolkit

---

The Globus Toolkit (GT)[16] is a reference implementation of Grid standards for data and resource management with the goal of supporting virtual organizations. GT has been used as in production Grids worldwide, including in DOE (e.g., Particle Physics Data Grid, Earth Systems Grid, DOE Science Grid), NSF (e.g., Alliance, TeraGrid, NPACI, NEESgrid), NASA (Information Power Grid), Europe (European Data Grid), international and inter-agency efforts (e.g., Grid3, LHC Computing Grid), and elsewhere.

GT provides security mechanisms for authentication, authorization and message protection based on PKI credentials [19][35][36]. Users acquire a set of PKI credentials and use these to authenticate to Grid resources, allowing that resource to establish an identity for the user and map it to a set of local privileges. There has been some work to extend GT's identity-based authorization mechanisms with attribute-based methods that scale better to support large VOs; one example of such work is VOMS. However, these methods have not attempted to leverage the security services provided by the production information technology staff at campuses, which is the focus of our work.

GT4.0, scheduled for release in October 2004 and for inclusion in NMI releases before the end of 2004, implements a set of proposed standards collectively called the WS-Resource framework (WSRF)[41]. In its current instantiation, WSRF adopts WS-Addressing endpoint references (EPRs) as a network pointer to a stateful resource that, furthermore, can be communicated to other parties[2][20]. As we explain in the detailed description of our proposed work that follows, EPRs play an important role in the work proposed here.

### C.2. GridLogon

---

GridLogon is a credential service that is being developed at NCSA as an extension to the popular MyProxy[27] service. MyProxy, which is part of the NMI release, is a credential management service and is the de facto mechanism used to provide security to grid portals worldwide. MyProxy is being extended, through other funding, to become a full-featured credential service, with a name change to GridLogon to reflect its enhanced capabilities.

The main function of the GridLogon service is to issue X.509 credentials to clients that are used by the client to securely access Globus Toolkit enabled Grid resources. GridLogon will provide a pluggable authentication mechanism to allow for with the use of different authentication systems, such as username/password, One-Time-Password, Kerberos, Public Key, etc., so that it may be integrated with a local authentication system to provide easy access to Grid resources. GridLogon plays an important role in our work to provide anonymous/pseudonymous access.

### C.3. Shibboleth

Shibboleth[6][31] is a tool, developed by Internet2, for identity federation between campuses that allows resources to obtain attributes about the user (e.g. departmental affiliation, student status), while preserving the user's privacy and not having to become involved with the details of how the user is authenticated in their home domain. Shibboleth is being deployed at an increasing number of institutions in conjunction with several higher educational federations: InQueue[22], InCommon[21], and the National Science Digital Library[25] among others in the US, and a growing number of research & education federations internationally as well. Hence, Shibboleth provides a perfect base for user attributes for many Grid VOs, as mentioned in the first section, as many of those VO users are homed in such campuses and these campuses are used to supporting high-quality production secure infrastructure, something the VOs do not have the resources or expertise to accomplish.

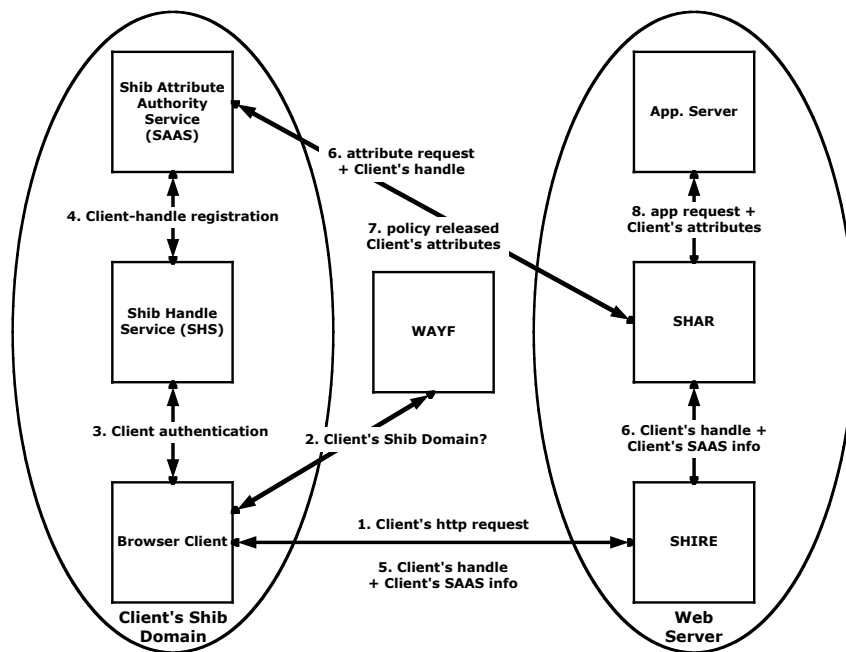


Figure 2. Shibboleth Framework

Currently is targeted at web browser-based services (we will refer to such services as *web applications* to distinguish them from more sophisticated Web Services). Shibboleth allows web applications to make authorization decisions regarding clients based on the client's attributes, while preserving the anonymity of the client and providing for policies on attribute release. One of the goals of the Shibboleth architecture team is to extend Shibboleth to support services other than web applications; our work will serve as one of the initial applications. One of our project team (Welch) has already been involved with the Shibboleth architecture team to develop the ideas presented in this proposal and provide requirements for future versions of Shibboleth.

The Shibboleth framework with its components and basic process model is depicted in Figure 2. Clients contacting a target web application are redirected to a Shibboleth Handle Service (SHS) in their local domain (a "Where Are You From" (WAYF) service may be operated by a federation as a default means of redirection). The SHS authenticates the client, using a local authentication mechanism not specific to Shibboleth, and returns to the client a handle. The handle is presented to the target web application, which presents it to the Shibboleth Attribute Authority Service (SAAS) in the client's domain and requests the set of attributes it desired in order to authorize the client's access request. The handle serves to identify the user to the SAAS without providing any identifying information to the target web applications. This is accomplished through the use of single use identifiers, the details of which are beyond the scope of this proposal.

The SAAS enforces Shibboleth's attribute release policy. The client attributes that it will communicate to the web application will be the intersection of those attributes the web application requests, those the client's SAAS has available, and those which the client's release policy allows to be revealed to the web application in question.

#### D. Developing the Attribute Framework.

We provide details of our planned integration work.

##### D.1. Globus Toolkit and Shibboleth Integration

The basic integration will use the Shibboleth framework as a sophisticated attribute service. When GT-enabled clients are to communicate with GT-enabled target web services, then the latter will have access to the Policy Released Client's attribute information as managed by the Shibboleth Attribute Authority Service (SAAS). In all these scenarios, it is assumed that the client has Globus Toolkit compatible public key credentials, which provide some form of identity authentication to the target service.

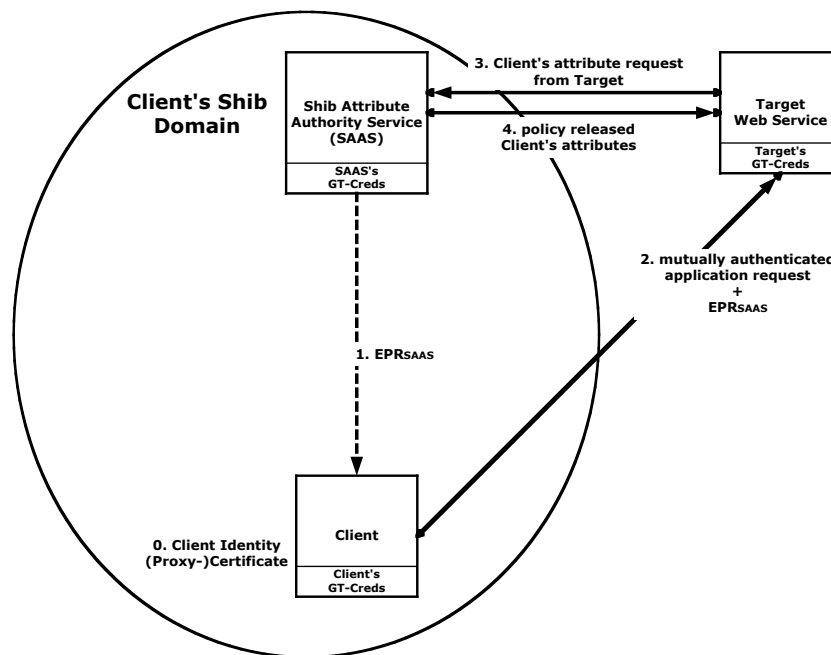
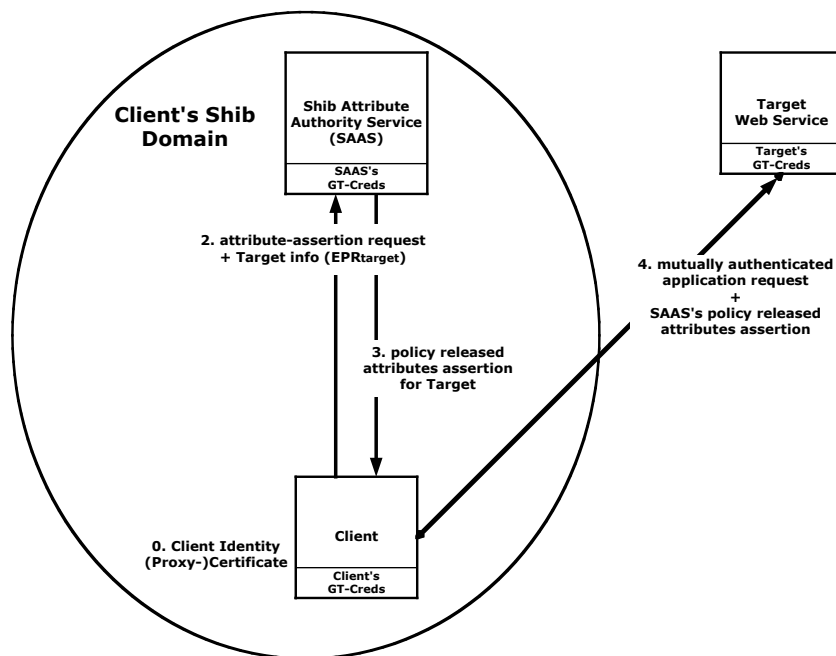


Figure 3. Basic Globus-Shibboleth Integration without anonymity

We can distinguish two modes of operation. The first mode, illustrated in Figure 3, resembles the normal web browser mode of operation, where the target service pulls the client's attributes from the client's Shib Attribute Authority Service (SAAS). In Figure 3, one can see that the client communicates the end pointer reference (EPR) (as described in Section C.1) of its SAAS ( $EPR_{SAAS}$ ) with the application request to the target service (step 2). It is assumed the client has obtained  $EPR_{SAAS}$  through some means (step 1), like a directory service. It is also assumed that the client has identity credentials in the form of a X.509 certificate, which is used to authenticate to the target service. With the  $EPR_{SAAS}$  and the authenticated client's identity, the target can ask the client's SAAS for the attribute information it needs to allow the client access to its resources.

Note that there is no anonymity in this scheme as the client uses its normal identity certificate for authentication, but the client's Attribute Release Policy (ARP) is used to filter the attribute information to the target to provide a controlled sharing of the client's attributes.



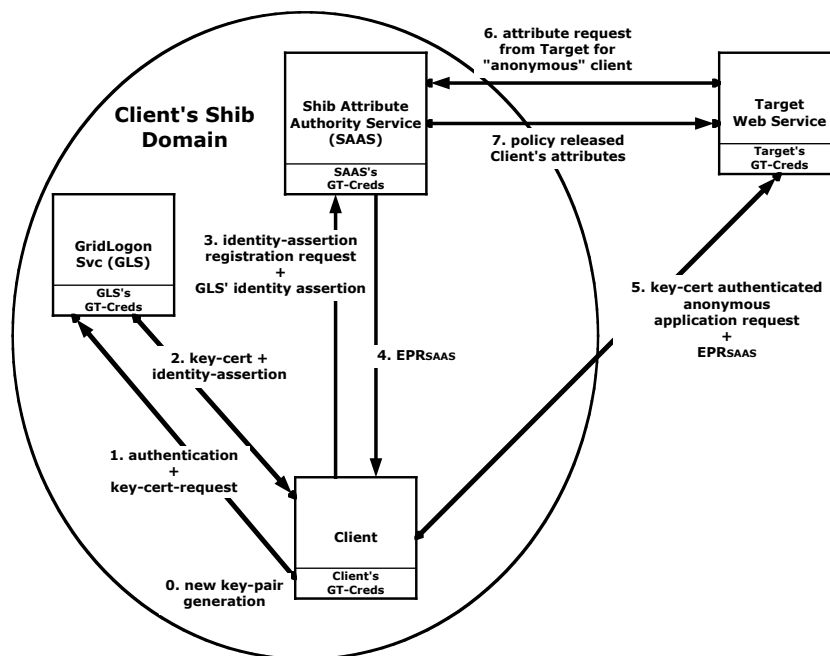
**Figure 4. Basic Globus-Shibboleth integration without anonymity plus pushing of attributes**

The second mode of operation is the push model shown in Figure 4, which facilitates deployment scenarios where the target services may not be able to communicate directly with the client's SAAS because of firewalls and other network hurdles. In this mode, the client knows what target service it wants to access as it has obtained the target's  $EPR_{target}$  plus the target's identity info, which was possibly embedded in the EPR. The client can now ask the SAAS (step 2) for an attribute assertion that contains all the attributes appropriate to reveal to the given target. The SAAS will package those attributes up in a signed attribute assertions (step 3), which the client will forward to the target together with the application service request (step 4). The target will authenticate the client, and will verify the SAAS' attribute assertion that is bound to that client's identity. Assuming successful verification, the target service will then use those asserted attributes to determine the access right for that client to its resource.

## D.2. Pseudonymity Support through GridLogon

In order to provide anonymity, we propose to integrate the GridLogon service with Shibboleth and the Globus Toolkit. As we described previously in Section C.2, the GridLogon service issues X.509

certificates to authenticated clients. This activity proposes the development of an additional GridLogon module that issues an authenticated client a set of credentials that hide the client's identity.



**Figure 5. Globus-Shibboleth integration with anonymity.**

The details of this integration are shown in Figure 5. The GridLogon service issues certificates that include a Subject name derived from a public key, newly generated by the client for this purpose (steps 0, 1 and 2). Such a certificate would not reveal anything about the client's.

Besides this "key-certificate," the GridLogon service will issue a separate identity assertion that binds this key-certificate to an identity attribute used for that client in its domain. The client communicates this identity mapping assertion to the local attribute service, allowing the attribute service to bind the key-derived identity in the client's key-certificate to the user's permanent identifier (steps 3 and 4).

When the client uses the key-certificate to access a foreign target service (step 5), the target service will contact the client's attribute service with the certificate's key-derived subject information as the lookup key (step 6), and will receive only the attributes it is allowed to see without ever knowing the client's identity (step 7).

One other implementation option we will explore, that does not force the attribute service to maintain state about the mapping, is for the attribute service to encrypt the mapping information for its own use, and to add it to its EPR resource properties. If the attribute service's EPR for the client's attribute information is handed to the target service, then this target will not be able to decrypt the embedded mapping information and will send it automatically with the subsequent request for the client's attribute information to the attribute service. On receipt, the attribute service will look for the encrypted mapping information, decrypt it, and know the client's identity to use for the attribute lookup. Another option would have the client encrypt GLS' identity assertion with SAAS's encryption key, and push that token directly to the target service. This would eliminate the initial roundtrip from the client to SAAS at the expense of having to make the target explicitly aware of this encrypted mapping token as this scheme will not allow that token to be hidden in the reference property of SAAS' EPR. Further investigation and experience should result in the most practical solution for deployment.

The push model is shown in Figure 6 and will work in a similar way: by having the attribute service issue the client's attributes in an attribute assertion bound to the key-certificate.

Note that in these schemes, the client could generate new key-pairs at will, and obtain as many anonymous key-certificates as allowed by the GridLogon policy.

In all cases we have described, the client's true identity is hidden from the target service, but an audit trail can still be maintained by the client's GridLogon or attribute service to allow for audit and reconciliation purposes. Note that the reconstruction of these audit trails from distributed logs is addressed in the next sections.

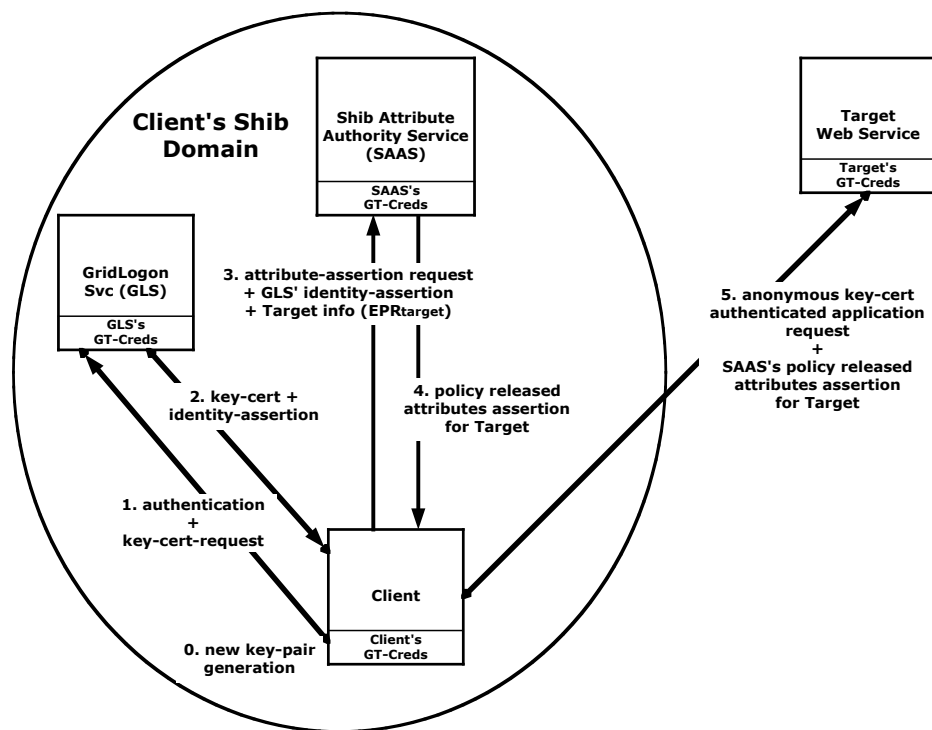


Figure 6. Globus-Shibboleth integration with anonymity and pushing of attributes.

Note that the GridLogon service is also able to issue assertions about the strength of the client authentication to the GridLogon service, which the target service can use to enforce its own access policy. This could be used, for example, to ensure that the client was authenticated through a one-time-password mechanism.

## E. Technology Evaluation and Transfer

Engagement with users is central to our work program. We will work with both application scientists and infrastructure projects to deploy and evaluate our technologies in realistic scientific collaboration scenarios. We describe in the following the specific tasks that we are considering in this area.

### E.1. Technology Evaluation

We have identified a group of infrastructure and application projects with which we will work closely to identify requirements, design deployments and evaluation, and obtain feedback. We will invite representatives of these projects to participate in an Advisory Committee with whom we will consult via email and teleconferences.

**U.Chicago.** We have teamed with the University of Chicago’s Networking Services and Information Technologies (NSIT [26]) organization to obtain a local testbed for this integrated system. NSIT is deploying Shibboleth, and U. Chicago is an advanced user of NMI Grid technologies due to its engagement in projects such as GriPhyN, iVDGL, PPDG, Access Grid, and NEESgrid. Thus, U. Chicago is a natural testbed for early rollouts of our integrated system.

**NEESgrid [10][24].** We have identified the Network for Earthquake Engineering Simulation (NEESgrid) as an important early driver for our functionality. The NEESgrid system developed by a multi-institutional group including NCSA and U. Chicago makes heavy use of GSI security mechanisms, but also has authorization requirements in which attributes can play an important role as a means of increasing flexibility and scalability. For example, in educational settings, observer-status access to an experiment may be granted to any accredited student. We will work with the NEESgrid team to design and integrate an authorization system that exploits the techniques and software to be developed in this project.

**Earth System Grid (ESG)[5][15].** ESG is an important driver as they have the need to distribute data to a large community of users, composed to a large degree of university faculty. These users also tend to only be occasional users of data, making the acquisition of permanent identity credentials overly burdensome.

**TeraGrid [4].** The TeraGrid system is a production grid system currently encompassing resources at four sites which is currently expanding to nine sites (NCSA, SDSC, Argonne National Laboratory, Cal Tech, and growing to include Indiana University, Purdue, ORNL, Pittsburg Supercomputing Center, and Texas Advanced Computing Center). Managing access control between nine sites threatens using conventional identity-based access control lists threatens to be a monumental burden. NCSA is part of the TeraGrid security efforts and will be funded by TeraGrid to specifically address security requirements. We will leverage this relationship to gather requirements and ensure our results meet those requirements.

**Grid3 [17][34].** The Grid3 international infrastructure created by the “Trillium” collaboration of GriPhyN, iVDGL, and PPDG, spans some 28 sites with 3000+ CPUs and supports a dozen different application groups from half a dozen different virtual organizations. It is thus a wonderful testbed for our technology, as the different virtual organizations have a variety of different policies, many relating to roles that people play in various organizations. We will work with the Grid3 security team to identify opportunities for the deployment and application of our technology.

## **E.2. Technology Transfer**

---

We will do the following to ensure transfer of our technology to a broad base of user communities:

- **Contribution to NMI:** All of our work is based on current NMI components and our enhancements will be contributed back to NMI. All of our development will be open source and available to the public through anonymous FTP and/or CVS.
- **Open Standards:** We will work in the Global Grid Forum to standardize our developed methods.

## **F. Research Approach and Plans**

---

The project’s research program will be structured in terms of an integrated set of problem-oriented *research tasks* situated within a user-oriented context as defined by a set of integrating *development and application milestones*.

All technology will be rendered on the emerging Web Services Resource Framework (WSRF), web service security standards, and standard web services protocols. Furthermore, all software deliverables will be implemented on top and integrated with the open source Globus Toolkit, which will also facilitate technology transfer to commercial vendors. The policy language, attributes, and policy assertions will be expressed and defined in XML, and in a form compatible with Web services standards where appropriate.

These deliverables will be produced via the completion of the following tasks and milestones.

**Year 1:** Basic Integration.

- *Research.* Investigation of policy description formats and protocols associated with attributes. Investigate how to embed policies in EPR.
- *Technology.* Providing a basic integration of existing Shibboleth and the Globus Toolkit. Enable GT-clients and services to use Shibboleth's Attribute Authority service (SAAS) in the simple pull-configuration.

**Year 2:** Advanced features and integration.

- *Research.* Investigate applicability of advanced privacy preserving negotiation protocols and privacy specification languages.
- *Technology.* Develop the push-model interaction where attribute assertions are issued by SAAS, such that the target service does not have to call back to SAAS. Develop the pseudonymous module for GridLogon. Implement standard formats and processing of policy information

---

## G. Relationship to Other Projects and Activities

---

The Shibboleth core architects and developers are discussing on mailing lists and weekly telephone conferences the requirements, design and development of the next generation of the Shibboleth framework. As mentioned, members of our team participate actively in this effort with a focus on the non-browser requirements and application of that framework. The intent is to leverage the knowledge and expertise of the Shibboleth architects, and to feedback any results from our work. Furthermore, we realize the relevance and importance of the standardization of useful and common schemas that can be used for the communication of attribute information, and we already initiated close collaboration with the Shibboleth team on that subject.

The PIs are aware of numerous research results about sophisticated negotiation techniques where parties will gradually reveal more and more to each other as the trust level is raised with each negotiation step[38][39][40]. There is also an interesting effort EPAL [1] that addresses the privacy policy and enforcement through a language based on the authorization policy language core of XACML [9]. However, we propose first to take a pragmatic approach that will address the complete lack of any standardized method to share and match policies with capabilities, with the most basic and simple solutions, and extend the sophistication over time driven by our user requirements.

The Condor Project at the University of Wisconsin has developed very interesting languages and methods for the matchmaking of customers to resources, and is extending this to multilateral model [33][29]. We plan to leverage on the results from this work to address our policy requirements.

From the many foreign groups that work in similar areas, we would like to point out a number of British research groups that recently have received funding through The Joint Information Systems Committee (JISC)[23] for projects that will extend and enhance the Shibboleth framework. It is our intent to work closely with these groups, such that we can leverage each other's work and experiences. In particular we have already been in contact with the lead of the ESP-Grid Project to ensure complementary efforts and collaboration.

Finally, while our team has unequalled expertise in security, middleware, and applications, we will be dependent on close contacts with other researchers to ensure continued relevance. We will achieve this information flow via engagement with the appropriate research communities; and participation in the work of standards bodies such as GGF, IETF, OASIS, and W3C.

---

## H. Management Plan

---

Von Welch will have primary responsibility for project direction. On funding, we will have an initial all-hands meeting to plan year one development and research activities. We will communicate weekly through teleconferences, access grid meetings and in-person meetings. We will have all-hands in-person

meetings at least annually for planning purposes. Each annual meeting will result in a detailed, month-by-month work plan for that following year.

The specific focus of each of the participants is given in the following sections.

### **H.1. NCSA**

---

NCSA's efforts will be focused on modifications to MyProxy/GridLogon to support anonymous credentials and modification to support Shibboleth attribute services with the Globus Toolkit security.

### **H.2. University of Chicago**

---

The University of Chicago will deploy our results repeatedly during the project in their testbed environment to validate and guide our research and development through user feedback.

### **H.3. Argonne National Laboratory (ANL)**

---

ANL's efforts will be focused on modifications to the Globus Toolkit to support integration of Shibboleth attribute services as well as extending the Shibboleth service to enable access over the web service's protocols.

## **I. Results from Prior NSF Support**

---

---

**None of the project PIs have** served as a PI on prior NSF-funded research.

## References

---

- [1] Ashley, P., Hada, S., Karjoth, G., Powers, C., Schunter, M. Enterprise Privacy Authorization Language (EPAL 1.1), IBM Research Report, 2003, <http://www.zurich.ibm.com/security/enterprise-privacy/epal>
- [2] Bosworth, A., Box, D., Christensen, E., Curbera, F., Ferguson, D., Frey, J., Kaler, C., Langworthy, D., Leymann, F., Lovering, B., Lucco, S., Millet, S., Mukhi, N., Nottingham, M., Orchard, D., Shewchuk, J., Storey, T., Weerawarana, S. Web Services Addressing (WSAddressing), <ftp://www6.software.ibm.com/software/developer/library/ws-add200403.pdf>, 2004
- [3] Box, D., Curbera, F., Langworthy, D., Nadalin, A., Nagaratnam, N., Nottingham, M., von Riegen, C., Shewchuk, J., Hondo, M., Kaler, C. Web Services Policy Framework (WSPolicy), 2003, <http://www.ibm.com/developerworks/library/ws-policy>
- [4] Catlett, C. The TeraGrid: A Primer, 2002. [www.teragrid.org](http://www.teragrid.org).
- [5] Earth System Grid, <https://www.earthsystemgrid.org/>
- [6] Erdos, M. and Cantor, S., Shibboleth Architecture. Internet 2. 2002. <http://shibboleth.internet2.edu/docs/draft-internet2-shibboleth-arch-v05.pdf>.
- [7] EU DataGrid, VOMS Architecture v1.1. 2003. [http://grid-auth.infn.it/docs/VOMS-v1\\_1.pdf](http://grid-auth.infn.it/docs/VOMS-v1_1.pdf).
- [8] eXtensible Access Control Markup Language (XACML) 1.0 Specification, OASIS, February 2003. <http://www.oasis-open.org/committees/xacml/>
- [9] eXtensible rights Markup Language, <http://www.xrml.org>, 2004.
- [10] Finholt, T.A., Wierba, E.E., Birnholtz, J.P. and Hofer, E., NEESgrid User Requirements. Technical Report NEESgrid-2002-xx. 2002. [www.neesgrid.org](http://www.neesgrid.org).
- [11] Foster, I. and Kesselman, C. (eds.). *The Grid 2: Blueprint for a New Computing Infrastructure*. Morgan Kaufmann, 2004.
- [12] Foster, I. and Kesselman, C. (eds.). *The Grid: Blueprint for a New Computing Infrastructure (2nd Edition)*. Morgan Kaufmann, 2004.
- [13] Foster, I. and Kesselman, C. Globus: A Toolkit-Based Grid Architecture. Foster, I. and Kesselman, C. eds. *The Grid: Blueprint for a New Computing Infrastructure*, Morgan Kaufmann, 1999, 259-278.
- [14] Foster, I. Kesselman, C., and Tuecke, S. The Anatomy of the Grid: Enabling Scalable Virtual Organizations. *International J. Supercomputer Applications*, 15(3), 2001.
- [15] Foster, I., Alpert, E., Chervenak, A., Drach, B., Kesselman, C., Nefedova, V., Middleton, D., Shoshani, A., Sim, A. and Williams, D., The Earth System Grid II: Turning Climate Datasets Into Community Resources. In *Annual Meeting of the American Meteorological Society*, (2002).
- [16] Globus Toolkit. <http://www.globus.org/>, 2004.
- [17] Grid2003, <http://www.ivdgl.org/grid2003/>
- [18] Hondo, M., Kaler, C. (eds.) Web Services Policy Framework. <http://www-106.ibm.com/developerworks/webservices/library/ws-polfram/>, 2003
- [19] Housley, R., Polk, W., Ford, W., and Solo, D., Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile. *RFC 3280*, IETF, April 2002
- [20] IBM, BEA, Microsoft, TIBCO Software. Web Services Addressing. <http://www-106.ibm.com/developerworks/webservices/library/ws-add/>, 2003.
- [21] InCommon Home. <http://www.incommonfederation.org/>, 2004.
- [22] InQueue Home. <http://inqueue.internet2.edu/>, 2004.
- [23] JISC: The Joint Information Systems Committee, <http://www.jisc.ac.uk/>
- [24] Kesselman, C., Foster, I. and Prudhomme, T. Distributed Telepresence: The NEESgrid Earthquake Engineering Collaboratory. In *The Grid: Blueprint for a New Computing Infrastructure (2nd Edition)*, Morgan Kaufmann, 2004.
- [25] National Science Digital Library home. <http://www.nsdlib.org/>, 2004.
- [26] Networking Services and Information Technologies (NIST), University of Chicago, <http://nsit.uchicago.edu/>

- [27] Novotny, J., Tuecke, S., and Welch, V., An Online Credential Repository for the Grid: MyProxy. *Proceedings of the Tenth International Symposium on High Performance Distributed Computing (HPDC-10)*, IEEE Press, August 2001
- [28] NSF Middleware Initiative (NMI). 2004. [www.nsf-middleware.org](http://www.nsf-middleware.org).
- [29] Raman, R., Livny, M., Solomon, M., Policy Driven Heterogeneous Resource Co-Allocation with Gangmatching. *Proceedings of the Twelfth IEEE International Symposium on High-Performance Distributed Computing*, June, 2003
- [30] Security Assertion Markup Language (SAML) 1.1 Specification, OASIS, November 2003.
- [31] Shibboleth Project, Internet2, <http://shibboleth.internet2.edu/>
- [32] Siebenlist, F., Nagaratnam, N., Welch, V., Neuman, B.C. Security for Virtual Organizations: Federating Trust and Policy Domains. In *The Grid: Blueprint for a New Computing Infrastructure (2nd Edition)*, 2004.
- [33] The Condor Project, University of Wisconsin, <http://www.cs.wisc.edu/condor/>.
- [34] The Grid2003 Project., The Grid2003 Production Grid: Principles and Practice. In *IEEE International Symposium on High Performance Distributed Computing*, (2004), IEEE Computer Science Press.
- [35] Tuecke, S., Welch, V. Engert, D., Thompson, M., and Pearlman, L., Internet X.509 Public Key Infrastructure Proxy Certificate Profile, *draft-ietf-pkix-proxy-10 (work in progress)*, IETF, 2003.
- [36] Welch, V., Foster, I., Kesselman, C., Mulmo, O., Pearlman, L., Tuecke, S., Gawor, J., Meder, S., Siebenlist, F. X.509 Proxy Certificates for Dynamic Delegation. *Submitted to 3rd Annual PKI R&D Workshop*.
- [37] Welch, V., Siebenlist, F., Foster, I., Bresnahan, J., Czajkowski, K., Gawor, J., Kesselman, C., Meder, S., Pearlman, L. and Tuecke, S., Security for Grid Services. In *12th IEEE International Symposium on High Performance Distributed Computing*, (2003).
- [38] Winsborough, W. and Li, N., Safety in automated trust negotiation. In *2004 IEEE Symposium on Security and Privacy*, (2004), IEEE Computer Society Press.
- [39] Winsborough, W. and Li, N., Towards Practical Automated Trust Negotiation. In *IEEE 3rd International Workshop on Policies for Distributed Systems and Networks*, (2002).
- [40] Winslett, M., Yu, T., K. E. Seamons, Hess, A., J. Jacobson, Jarvis, R., B. Smith and Yu, L. Negotiating Trust on the Web. *IEEE Internet Computing*, 6 (6). 30-37. 2002.
- [41] WS-Resource Framework. <http://www.globus.org/wsrfl/>, [http://www.oasis-open.org/committees/tc\\_home.php?wg\\_abbrev=wsrf](http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=wsrf), 2004.