NMI-EDIT

Identity and Access Management

Case Study Series

THE UNIVERSITY OF TEXAS SYSTEM

The University of Texas System: "Extending the Reach" Case Studies

NMI

January 2006

NMI-EDIT Case Study Series

In response to calls from the higher-education community, the NMI-EDIT Consortium has developed a series of Identity Management Case Studies to explore the planning and implementation of this critical infrastructure at higher-education institutions around the country.

In the spring of 2004, NMI-EDIT released the Extending the Reach Call for Proposal with the overall vision of exploring possible models for middleware support and informing the NMI-EDIT outreach and development efforts through collaboration with a wider, more diverse group of institutions. The work outlined in this case study was supported in part by the NMI-EDIT Extending the Reach Program.

This NMI-EDIT Case Study Series is sponsored by the National Science Foundation Middleware Initiative-Enterprise and Desktop Integration Technologies (NMI-EDIT) Consortium of EDUCAUSE, Internet2, and SURA. Additional support was provided by the National Science Foundation Cooperative Agreement NSF 02-028, ANI-0123937. Thanks are extended to authors Paul Caskey, Barry Ribbeck, William Weems PhD, and Miguel Soldi.

Project Summary

In an effort to enable greater synergy and collaboration among its institutions and better serve their broad constituencies, The University of Texas System (UT System) set out in 2004 to create a core resource consortium of middleware-leveraged resource providers – a project which has matured and is now the UT System Federation. The goal of this effort was to foster the foundation for a federated collaborative infrastructure by targeting UT System institutions where collaborative efforts already existed and where core middleware technology would provide an immediate benefit. Those UT System institutions that did not have any current collaborative opportunities were mentored and provided assistance to build the middleware infrastructure they will need in the future to leverage access to resources.

When planning the approach to the exploration of a federated infrastructure model and the role of middleware in that model, UT System took advantage of existing opportunities and, importantly, used a long-term, system-wide approach in order to fully understand how these new technologies could best meet the needs of their constituents. As a result of their approach, the UT System's development work spans a variety of test environments. This document is a compilation of the five case studies, each set in a unique environment, the UT System completed as part of their exploration and implementation of middleware and the federated infrastructure model.

Four case studies were undertaken in the course of this project with partial funding from an Extending the Reach (ETR) grant from the National Science Foundation Middleware Initiative-Enterprise and Desktop Integration Technologies (NMI-EDIT) Consortium. The fifth case documents UT System's work in developing the UT System Federation.

In addition to the case studies and the consulting and consortium building, the project included significant efforts in outreach and information dissemination. This outreach focused on education and awareness building via middleware presentations as a means to build and foster a community of interest. Integral to the presentations was a discussion of the middleware implementations at the institutions participating in the case studies, and the pay-offs that were realized as a result, with a focus on the benefits of a federated model and the availability of support from with the UT System Federation

A synopsis of each of the five case studies is provided below. For more information about any of the University of Texas System Case Studies, contact Miguel Soldi at msoldi@utsystem.edu.

Case Study 1: Middleware Outsourcing Feasibility Study

There are numerous opportunities for collaboration and resource sharing throughout the UT System. However, the application technologies that allow higher-education institutions to take advantage of these opportunities are becoming increasingly sophisticated, and smaller institutions are finding it increasingly difficult to keep pace with larger institutions in this area. These new applications and back-end systems require considerable "glue" in order to deliver the promised benefits and middleware provides that glue.

Since often it is specifically these smaller institutions that stand to benefit the most from today's collaborative architectures, UT System wanted to explore alternatives their smaller institutions might use in order to reap the benefits of collaborative, resource-sharing environments. One such alternative is for the smaller institutions to outsource its core middleware services to larger institutions. What is difficult to ascertain is the proper outsource model and the difficulty in developing a working knowledge of the processes needed to implement such outsourcing.

This case study is an evaluation of the feasibility of designing, building, hosting, and operating the NMI-EDIT middleware components necessary for participation in system-wide collaborative efforts for The University of Texas at Tyler. UT System Administration hosts the servers in this pilot and used wireless network authentication as the test application.

Case Study 2: Baylor College of Medicine/UT Health Science Center at Houston Shibboleth Project

The Texas Medical Center (TMC) consists of 42 separate institutions that collaborate with each other and with hundreds of external organizations via the Internet for a multitude of purposes. A major technical challenge in these collaborations is how to securely and efficiently authenticate and authorize individuals who are not formally identified personnel of an institution in order to allow them to access to restricted resources.

Realizing that the traditional processes used for authentication and authorization are not scalable and are insufficient for UT Systems institutions' user-base, the UT System wanted to explore a federated approach as an alternative to traditional processes. For their test environment in this case, UT System tests the federated approach to enable the students, residents, fellows and faculty at Baylor College of Medicine (BCM) to securely access the information services provided by The University of Texas Health Science Center at Houston (UTHSC-H) and the TMC Library. This project is expected to serve as a proof-of-concept pilot for other TMC institutions. Such a federated approach has great potential for BCM, UTHSC-H and the TMC library.

Since developing a federated identity management and trust system not only requires the development of technical infrastructure but must also have formally defined policies, procedures and trust agreements, this case study also evaluated the approaches to the development of these critical, non-technical components as well. The BCM Shibboleth Identity Provider was successfully tested against the UTHSC-H Blackboard Learning Management System, the TMC Library Proxy Server, and the Internal BCM Shibboleth Service Provider (SP) resources

Case Study 3: Houston Academy of Medicine – Texas Medical Center (HAM-TMC) Library Shibboleth Proxy Service

At the Houston Academy of Medicine Texas-Medical Center Library (HAM-TMC) library, which provided the test scenario for this case study, access management had becomes a nearly impossible task. Remote access to the library's on-line databases is critical to research, academia and health care, and the library provides access to these resources to a group of medical research and provider institutions under a cooperative agreement and acts as the proxy for digital content providers.

However, traditional methods had proved inadequate for the library to continue to define the electronic access rights to resources that need to be given to people based upon their institutional affiliation. In this case study, using Shibboleth technology, the HAM-TMC Library was removed from the identity management and access management role. This provided an improved security and access management model by moving the responsibility of identity and access management to the Identity Providers.

The library's EZProxy server was changed to give the application the ability to utilize Shibboleth-based authentication and authorization. After completing a proof-of-concept, the library would move to integrate Shibboleth into its production systems to allow two or more of the TMC institutions running Shibboleth services to access the library's resources.

Case Study 4: Shibboleth Version of Employee Benefits Annual Enrollment Application

The goal of this case study, which is still an on-going project, was to create and pilot a federated version of the UT System Administration Office of Employee Benefits and Group Insurance (EGI) benefits annual enrollment application, UTTouch, using the Shibboleth software and the UT System Identity Management Federation. While the pilot will only be available to a small subset of UT System's user population the application will eventually be rolled out to an audience of over 80,000 and will be UT System's first large-scale, system-wide deployment of a federated application that uses Shibboleth.

Historically UTTouch, a web interface to a mainframe application, has required users to login using their social security number

(SSN). Since the use of SSNs as identifiers will be prohibited by UT System policy by 2007, UT System saw Shibboleth and the UT System Federation as a means to achieve compliance with this new policy. They also discovered that a Shibboleth solution had the potential to greatly enhance the user experience as well. The Shibboleth version of UTTouch will allow employees and retirees to simply login to their home institution's Identity Provider and let Shibboleth send appropriate attributes to UTTouch.

Although this case study is meant to serve more as description of the project, considerable progress has been made in Shibboleth-enabling this legacy application and the web agent has been modified and tested in a development environment. In the document, UT System's project team shares valuable project insights and discusses their technical and policy challenges (e.g., the policy issues surrounding the assignment, use, population, and release of identifiers in a federated environment).

UT System Case Study: The University of Texas System Identity Management Federation

This case study discusses the federation-building activities UT System undertook as an expansion of their ETR program participation. The goal of their work, which is still in progress, is to create the foundation for a federated, collaborative infrastructure that will enable greater synergy and collaboration among UT institutions, simplify application security through a common trust fabric, and provide a platform to address UT System-wide identity management initiatives. From UT System's perspective, having their own federation would provide some compelling benefits, including the ability to leverage existing inter-institution agreements, establishment of a common set of standards and attributes for UT institutions, more granular control over authentication and authorization policies, and a forum for experimentation and dialogue.

The main technical challenge to the instantiation of the UT System Federation was establishing standards and prerequisites while simultaneously leveraging each UT institution's existing infrastructure and bringing all institutions onto "the same technical page". UT System Administration developed a plan to help close any gaps between the institutions. The plan included providing various types of support and outreach, including hosting system-wide middleware events that helped participating institutions install the Shibboleth identity provider and service provider software.

This case study includes discussion of the technical and policy work that takes place in parallel as the UT System project team develops their identity management federation. Since UT System institutions are diverse and autonomous, there is diversity in the operating practices and standards among institutions, with no consistent

identity management trust definition or policies. The UT System Federation, which operates under the authority of the UT System Board of Regents, is comprised of all sixteen UT institutions and each comes to the project with differing motivations and varying levels of commitment and urgency. Not surprisingly, although technical challenges were certainly significant, political challenges have been the most difficult the project team has had to overcome.

Nonetheless, the UT System has implemented a technical and policy infrastructure that can both support and scale as the needs of member institutions and constituents served dictate. As the project has progressed, issues of policy, funding, and resource availability are becoming more pressing. Still ahead of for the project team is concluding policy work, implementing compelling UT System-wide applications, and reaching the critical mass necessary to elevate the UT System Federation to a fully operational status.

Participating Institutions

Four institutions formed a partnership for the implementation of this project. The University of Texas System Administration Office of Technology and Information Services, The University of Texas at Austin, The University of Texas Health Science Center at Houston, and Baylor College of Medicine. Each institution brought with it resources central to the success of this project.

The University of Texas System Administration Office of Technology and Information Services has the mission of promoting and facilitating collaboration among UT institutions including providing a state-wide telecommunications network for UT System institutions, community colleges, and public schools in collaboration with Texas A&M University; facilitating the contracts for digital content for all UT libraries; facilitating system-wide security initiatives; and providing system-wide forums to develop and promote collaboration.

The University of Texas at Austin maintains the Human Resources application for seven UT institutions and is currently implementing an enterprise directory and working to deploy a Shibboleth Identity Provider server.

The University of Texas Health Science Center at Houston operates a production Shibboleth enabled Blackboard course management system, is a charter member of InCommon and is active in many Internet 2, EDUCAUSE, and NMI middleware activities. The University of Texas Health Science Center at Houston has fostered middleware collaboration with institutions within the Texas Medical Center, UT System and other higher education domains**.**

Baylor College of Medicine has participated in middleware activities within Houston and operates a production enterprise directory. Baylor accesses resources within the University of Texas Health Science Center at Houston and shares faculty and residents collaboratively with that institution.

Institutions Served

The Houston Academy of Medicine – Texas Medical Center Library provides digital content and standard reference materials to over 40 institutions throughout the Texas Medical Center in Houston including not for profit hospital systems. Baylor College of Medicine and The University of Texas Health Science Center at Houston are the major contributors for the library.

The University of Texas at Tyler is a small general academic institution located in east Texas with 1,000 staff and faculty and with a total enrollment of 5,700 students.

NMI-EDIT Components Highlighted in this Case Study

eduPerson Directory Schema **http://www.educause.edu/eduperson**

eduPerson contains identity-related attributes for higher-education institutions to deploy to foster inter-institutional collaborations.

Enterprise Directory Implementation Roadmap
http://www.nmi-edit.org/roadmap/directories.html

The Enterprise Directory Implementation Roadmap is a web-based structure of resources that institutions can draw on to help deploy and use enterprise directories in higher education and research communities.

Enterprise Authentication Implementation Roadmap
http://www.nmi-edit.org/roadmap/auth-roadmap_200510/index.html

The Enterprise Authentication Roadmap is a collection of resources that campuses can use to assist in building enterprise authentication services.

Shibboleth System **http://shibboleth.internet2.edu**

The Shibboleth System supports inter-institutional sharing of web resources that are subject to access controls.

Case Study 1: Middleware Outsourcing Feasibility Study

The UT System is a diverse educational and research university system, consisting of nine general academic institutions and six health institutions, employing 81,000 staff and faculty, and has a total enrollment of 175,000 students. There are numerous opportunities for collaboration and resource sharing throughout the UT System. However, such collaborative endeavors require some measure of consistency among the various electronic infrastructures. Further, the infrastructure necessary to facilitate collaborative education and research is still relatively immature and implementing the necessary technology can be a difficult task that requires new hardware, software, and advanced technical skills. Understandably, smaller institutions are finding it increasingly difficult to keep pace with larger institutions in developing these newer collaborative architectures.

From a university system viewpoint, this technical imbalance becomes more than a local institutional concern. It can become a limiting factor in the development of strategic system-wide collaborative initiatives. Many would argue that it is specifically these smaller institutions that stand to benefit the most from the newly enabled collaborative architectures. Insufficient budget, limited staff availability, and lack of in-house skills can seriously impair an institution's participation in a modern collaborative infrastructure, thus preventing it from reaping the benefits of system-wide collaboration.

The objective of this case study is to evaluate the feasibility of designing, building, hosting, and operating the NSF Middleware Initiative-Enterprise and Desktop Integration Technologies (NMI-EDIT) middleware components necessary for participation in system-wide collaborative efforts for The University of Texas at Tyler (UT Tyler). UT Tyler is a small but rapidly growing general academic institution employing 1,000 staff and faculty and with a total enrollment of 5,700 students. It is also intended that this project serve as a model for helping similar institutions within the UT System to participate in the new collaborative architecture being developed both within the UT System and in higher-education worldwide.

NMI-EDIT Components Highlighted in this Case Study

U.T. made use of the research and tools provided by the NMI-EDIT Consortium.

These included the eduPerson Directory Schema[1], the Shibboleth® Software[2], and the Enterprise Directory Implementation Roadmap[3].

Problem Discovery

The technologies available in higher education are becoming increasingly sophisticated. Student portals linked into registration systems, course management systems, and a host of other back-end systems can now make the campus experience more meaningful. Higher-education institutions are also under increasing pressure to provide seamless collaboration and resource sharing between institutions.

These newer applications and back-end systems require considerable "glue" in order to deliver the promised benefits of systems integration and collaboration. Middleware provides that glue and when implemented correctly, it is invisible to end-users.

However, implementing the required middleware for today's campuses is no trivial task and requires a significant investment in time and effort. Thus, it is understandable that many small institutions cannot make the necessary commitment to developing a middleware infrastructure and are struggling to "keep the trains running" with what they have in place. This is a difficult situation to manage. The farther behind the curve these campuses get, the more difficult it will be to implement needed middleware. This is of special concern since, with minimal operating resources, it is precisely these campuses that stand to benefit the most from having a middleware infrastructure into which standard "off-the-shelf" applications and technologies can be plugged in to derive immediate benefits.

One alternative for small institutions could be to outsource core middleware services to larger institutions, the same way other information technology is now being outsourced. What is difficult to ascertain is the proper outsource model and the difficulty in developing a working knowledge of the processes needed to implement such outsourcing.

Case Study #1 was a pilot for outsourcing UT Tyler core middleware services to UT System Administration. The result of this pilot will determine the efficacy of providing these services to other UT System institutions under an outsourced-service model. UT Tyler's experience also provides a framework and documentation for middleware deployment in small institutional settings challenged by limited resources and expertise.

[1] The eduPerson Schema is located at http://www.educause.edu/eduperson
[2] The Shibboleth Software is located at http://shibboleth.internet2.edu
[3] The NMI-EDIT Enterprise Directory Implementation Roadmap is located at http://www.nmi-edit.org/roadmap/internet2-mace-dir-implementation-roadmap-200312.html

All are available from the NMI-EDIT webpage at www.nmi-edit.org

The Solution

For the pilot, it was decided to implement Lightweight Directory Access Protocol (LDAP)-based enterprise directory servers, Shibboleth Identity Provider servers, and to set up nightly provisioning of the directories that would be based on feeds from a variety of source systems of record. Such a design could handle campus-wide LDAP based authentication and authorization (which is useful for most applications) and system-wide authentication and authorization using Shibboleth. UT System Administration would host the system, while their subject matter experts would handle the design and implementation work.

The project scope included:

- Directory provisioning and de-provisioning
- LDAP-based authentication
- LDAP application integration (with both Shibboleth and the wireless network authentication system)
- Basic password management
- Implementing and maintaining a Shibboleth IdP service

The following were not included in the scope:

- Directory re-provisioning
- Detailed localPerson directory schema
- Real-time provisioning/de-provisioning
- Course management integration with the enterprise directory

Some of the administrative decisions that impacted the solution included the following:

- Student user IDs would look different than employee IDs and, for individuals that are both student and employee, a primary affiliation would be determined based on their employment status, (full-time versus part-time, with the latter being treated as students).
- Individuals who "disappear" from source feeds will be deactivated, but not deleted.
- The provisioning system would use Social Security Numbers (SSNs) to reconcile identities until other institutional identifiers are in place.

A daily identity-lifecycle report would be issued detailing the activities of the provisioning system.

Project Resources

Staff

The project was implemented primarily by two UT System staff members who both worked on the case study on a part-time basis:

- Paul Caskey, Technology Architect from UT System Administration Office of Technology and Information Services, came to the project as the technical subject-matter expert and with substantial middleware experience.
- Michael Vick, Manager, Campus Computing Services from UT Tyler, came to the project as technical-lead and with considerable knowledge of the institution's information systems,

including identity management and authentication systems, and data sources.

The Project Timeline

The project did not have 100% dedicated resources from either institution, so the timeline was necessarily very flexible. It was begun in October 2004 and the initial phase was completed in September 2005 (further refinement work is ongoing). The majority of the work focused on writing and testing the code for directory provisioning and data manipulation, identity lifecycle reporting, and for the password-management web application.

The project was begun with discussions with UT Tyler stakeholders regarding the provisioning and de-provisioning of user accounts, the syntax for building user IDs, initial passwords, data items of significance in the various feeds, update cycles, and reporting. These discussions helped determine the project scope and business rules.

Milestones

In chronological order, the critical milestones for this project were:

- Initial discussions with stakeholders (scope, goals, timelines, requirements)
- Identification of systems of record and required data flows
- Discussion of proposed architecture (hardware, software, schema, provisioning)
- Initial discussion of technical implementation details (e.g., user IDs, provisioning, passwords)
- Acquisition of hardware/software (four servers)
- Ensuring directories were operational and redundant (limited records manually populated)
- Ensuring the IdP server was operational and redundant
- Ensuring basic provisioning completed
- Ensuring de-provisioning completed
- Establishment of Identity Lifecycle Reporting (daily batch)
- Completion of password-management web application
- Integration of the wireless authentication system at UT Tyler

Technical Details

Enterprise Directory

Redundant Dell model 1850 enterprise directory servers running Microsoft Windows 2003 and Sun ONE Directory Server software were implemented. These servers were deployed as active/passive nodes and negotiate a virtual IP address between them using the Microsoft Windows Network Load Balancing service. Standard directory schemas were deployed (e.g., OrgPerson and InetOrgPerson) along with the eduPerson schema, a rudimentary uttPerson schema, and an early version of the system-wide utPerson schema. Directory data synchronization is accomplished via the Sun ONE replication

engine and is backed up nightly by UT System Administration network staff.

Additionally, these servers host a password-change web page that enables UT Tyler constituents to change their password according to local complexity policy such as length or special characters. These pages are kept synchronized using a simple 'xcopy' script running on the primary server.

Shibboleth Identity Provider

Redundant Dell model 1850 Identity Provider servers running RedHat Enterprise Linux AS 3.0 were implemented. The servers were loaded with Shibboleth 1.2.1 and were configured to participate in the UT System Federation. The servers were deployed as active/passive nodes, with the virtual IP address negotiated by the heartbeat software from the Linux high-availability project[4]. Shibboleth configuration data is kept synchronized using a simple 'rsync' script on the primary server. Since UT Tyler desired a branded login page, the servers use the Pubcookie[5]system to handle the authentication (via LDAP/SSL authentication).

Directory Provisioning

Directory provisioning and de-provisioning is done with Microsoft Identity Integration Server (MIIS). This tool was selected primarily due to a low purchase cost and the institution's ready availability of MIIS skills for implementation.

MIIS is connected to the following data sources:

- Human Resources application – several UT System institutions use a proprietary system called DEFINE
- Student Information System – UT Tyler uses the POISE system
- Employee e-mail – UT Tyler uses Lotus Notes

As noted above, the provisioning system maintains information on all current students, faculty, and staff at UT Tyler.

Identities from various source systems are reconciled using SSNs. As records for new people arrive in source feeds, the system determines the person's primary affiliation based on their presence in various feeds, coupled with their employment status. This affiliation is then used to build an appropriate userID, with student userIDs taking the form of FirstName+MiddleInitial+LastName and employee userIDs taking the form of firstInitial+LastName.

In the event people have similar names, student userIDs have a unique number appended to them,, while the process for creating employee userIDs uses an increasing number of characters from the first and middle names to ensure uniqueness (and, in the event this technique fails, will append a unique number). If the

[4] See www.linux-ha.org
[5] See http://www.pubcookie.org/

person's record is present in the employee e-mail system, that userID becomes the preferred userID rather than creating an ID using the above algorithm. If a person's primary affiliation changes, their userID will change as well according to the established policy.

Once the userID has been assigned, the account is created in the directory and an entry is written into the person registry that maintains permanent data about each person ever handled by the provisioning system. If a person-record "disappears" from a source feed, their directory account is disabled and moved to an inactive container in the directory.

The daily batch cycle generates a report detailing:

- The number of entries present in the various feeds.
- The number of new entries in each feed.
- The number of deleted entries in each feed.
- The number of user records with attribute changes in each feed.
- The Identity lifecycle activities, including a list of any accounts that have been created, disabled, or renamed.

Networking

UT System Administration, which is located in Austin, TX, hosts all four servers in this pilot. The servers are accessed securely over existing Internet links using an IPSEC tunnel that was established specifically for the pilot. This setup allows the servers to appear as if they are on the UT Tyler LAN and all traffic between the two sites is encrypted using Triple Data Encryption Standard (3DES) encryption.

Application

For the pilot, the application selected to test the outsourced services was UT Tyler's wireless network authentication. UT Tyler is currently implementing a campus-wide wireless network. During this implementation, one of the first questions that had to be answered was how local users should authenticate to the network. Students, faculty, and staff at UT Tyler are currently authenticated to a variety of systems (e.g., SIS, CMS, e-mail), despite the lack of a centralized authentication system in place.

The wireless application seemed most appropriate for the middleware services being created because it would be difficult, if not impossible, to have the wireless network authenticate users from the SIS, CMS, or even the e-mail system. One alternative is to create yet another identity 'island" just for the wireless network, as it has been done for all other systems, although this means users would have yet another password to remember.

What the Future Holds

Future work with the middleware infrastructure will involve directory re-provisioning, enhanced lifecycle reporting,

and eventually course management integration with the directory. This work will lead to new applications being integrated with this infrastructure and will enable UT Tyler to more easily share resources within the UT System.

Lessons Learned

Listed below are the important lessons that were learned during the implementation of this project. :

- It is wise to spend more time up-front making business decisions and documenting procedures. Had the project team done so during this case study, some business process ambiguity may have been prevented.
- Involve all stakeholders, not just those from information technology, early in the project.
- Smaller campuses typically cannot afford extra network bandwidth, and degraded network conditions can be a killer to externally-hosted middleware services.
- It is difficult to gain buy-in and transfer knowledge over long distance.

In Conclusion

New applications, coupled with increasing pressures to collaborate and share resources, are straining conventional architectures at smaller institutions. These pressures mean that infrastructures must change to support improved identity management processes, application integration, collaborative tools and the like.

A rich set of middleware components has been emerging to respond to these challenges, although implementation of these components is no trivial task, either in terms of staff availability, in-house skills, or budget. This is especially true for smaller, resource-strained institutions that do not have the information technology resources enjoyed by larger institutions. This case study attempted to address the middleware infrastructure challenges smaller institutions face by evaluating the feasibility of outsourcing the implementation and operation of middleware components. While the hardware architecture implemented in this case is a robust design for the problem at hand, it can be concluded that undertaking a middleware outsourcing arrangement is a serious matter that should not be undertaken lightly. Middleware, when fully integrated into an infrastructure, is at the heart of information technology architecture.

Eventually, and to varying degrees, most systems will need to interact with the middleware layer. When the middleware layer is degraded or inaccessible for any reason, many things are likely to break. Further, while middleware is normally invisible to end users, in a failure scenario, it becomes painfully visible to everyone. Accordingly, in an outsourcing arrangement, the network connectivity between the main

campus and the outsourcing site becomes an incredibly critical piece to the information technology operations of the campus in question if very costly failures of the middleware layer are to be avoided. It is not only imperative that the link has the required bandwidth, but also that it be very reliable in terms of availability. For this case study, such a network connection was not available.

Since network connections can represent a significant portion of a smaller institution's information technology budget, such institutions typically cannot afford to purchase more bandwidth than they absolutely need and, likewise, cannot afford redundant connections or sophisticated fail-over designs. It is for this reason that one can conclude that while the outsourced design and implementation of middleware provides much-needed assistance to smaller institutions and is extremely valuable, externally hosting middleware services should be undertaken with extreme caution and only in environments where there is very solid network connectivity between the institutions.

More Information

For more information on the UT System's Middleware Outsourcing Study, please contact Paul Caskey at pcaskey@utsystem.edu.

Case Study 2: Baylor College of Medicine/UT Health Science Center at Houston Shibboleth Project

Baylor College of Medicine (BCM) is located in the Texas Medical Center (TMC) in Houston, TX, and its' computer users need to share restricted resources with numerous external organizations. BCM students, residents, fellows and faculty have critical requirements for appropriately utilizing the secure information services provided by The University of Texas Health Science Center at Houston (UTHSC-H) and the TMC Library. This case study was designed to assist BCM with the following:

- Develop a Shibboleth®[6] Identity Provider (IdP) that is coupled with the BCM Enterprise Directory Service.
- Test the ability of the BCM IdP to assert identity information for authentication and authorization purposes to the UTHSC-H Shibboleth-enabled Blackboard Learning Management System, the TMC Library proxy server and to internal, restricted BCM information resources.
- Consider additional BCM and UTHSC-H service providers (SPs) that should be Shibboleth-enabled in order to facilitate collaboration between the two institutions.
- Evaluate federated approaches that permit SPs at each institution to trust the IdPs of the other two institutions to authenticate the identity of their institutional affiliates and provide appropriate identity information to allow SPs to make authorization decisions.

This project is also expected to serve as a proof-of-concept pilot for other TMC institutions.

Problem Discovery

The TMC consists of 42 separate institutions that collaborate with each other to varying degrees, and also with hundreds of external organizations via the Internet for a multitude of purposes, such as education, research, health care and business activities. Currently, a major problem that greatly inhibits such collaborative activities in cyberspace is how to securely and efficiently authenticate and authorize many individuals who are not formally identified personnel of an institution to access that institution's restricted resources.

[6] The Shibboleth Software is located at http://shibboleth.internet2.edu and at the NMI-EDIT web page at www.nmi-edit.org.

The traditional authentication and authorization process is for each institution to identify all individuals that need to access their organization's information services and then somehow securely assign each person a username and password for every application. Not only is this traditional process not scalable in today's digital world, it also has major security problems and is extremely frustrating to users who must have their identity vetted multiple times and somehow keep track of an increasing number of usernames and passwords.

The Solution

A federated approach to identity management allows institutions that have formally agreed upon a set of policies, procedures and technologies to each function as an IdP. Thus, an institution's IdP can authenticate the identity of a person affiliated with it to access restricted information services provided by other member institutions. The IdP can also provide the relying SP with identity information about the authenticated person that may be required for authorization and/or provisioning purposes.

Such a federated approach has great potential for BCM, UTHSC-H and the TMC library, where faculty, students and other personnel of each organization must access restricted resources of the other institutions. The UTHSC-H already had an identity management system in place along with a Shibboleth IdP service and service provider resources. Therefore, this project was undertaken to develop a BCM Shibboleth IdP along with intra-institutional Shibboleth service providers. This in turn would allow BCM personnel, using their BMC identity credentials, access to Shibboleth-enabled resources at both UTHSC-H and the TMC library. Likewise, UTHSC-H personnel would gain access, when appropriate, to BMC Shibboleth-enabled SP resources.

Developing a federated identity management and trust system not only requires the development of technical infrastructure but must also have formally defined policies, procedures and trust agreements. Since they are indispensable to the development of the required trust fabric among federation members, approaches to the development of these critical, non-technical components were evaluated during this project as well.

The project included the following activities:

- Installation of a Shibboleth IdP at BCM using Solaris 8, Tomcat and Apache.
- Installation of a BCM Apache Service Provider.
- Enrollment of BCM in the InQueue[7] Federation for testing the Shibboleth components.

In order to provide them with base level federation policy and operating practice agreements, the decision was made for BMC, UTHSC-H, and TMC to join the

[7] See http://inqueue.internet2.edu/

InCommon Federation[8]. This would allow all three institutions to be able to use the InCommon WAYF ("where are you from") Shibboleth resource. The federation policies will be appropriate for the TMC Library proxy server to trust the BCM and UTHSC-H IdPs.

In order to address the greater security requirements associated with many of the BCM and UTHSC-H Service Provider (SP) resources, BCM and UTHSC-H will enter into a two-party agreement that is expected to expand the InCommon policies to include operating practices similar to those being considered for final approval by the UT System Identity Management Federation.

Project Resources

Staff

The project was implemented primarily by three UT System staff and all worked on the case study on a part-time basis. Programming services were obtained on a contract-basis for the testing of the Java implementation of the Shibboleth Service Provider software.

The Project Timeline

Implementation took approximately twelve weeks for completion of configuration and testing. The BCM Shibboleth IdP was successfully tested against the UTHSC-H Blackboard Learning Management System, the TMC Library Proxy Server, and the Internal BCM Shibboleth SP resources.

Technical Factors

A technical snag occurred when the service provider software on the project's Sun Web server did not install as expected. One alternative for getting around the snag that was discussed was to use a session-saving technique that would add a Shibboleth method to the UTHSC-H Java authentication method. A second alternative considered was to test the new 100% Java version of the Shibboleth Service Provider software. This packaged version works with both Java 1.4 and Java 1.5 and has been tested with Tomcat. Functionally, it is close to Shibboleth v1.3.

The second alternative was selected and once it was understood, it would be used to offer web-based applications to BCM users that have been authenticated and authorized via Shibboleth. The Shibboleth Service Provider Java development Setup can be found at: http://tpappsrv.its.yale.edu/tp/shibboleth_eclipse.htm. Code provided by The Ohio State University was used for the Java authentication against a Shibboleth server and worked as advertised.

The milestones accomplished in this project were:

- Solaris 9 implementation of a Shibboleth Service Provider server, integrated with Apache 1.3.x, as a proof-of-concept and base reference implementation.

[8] See http://www.incommonfederation.org/

- Integration of this reference implementation with a Shibboleth IdP server at UTHSC-H.
- Solaris 9 implementation of the Sun Web Server with Java Authentication model code provided by UTHSC-H.
- Tested the Jakarta Tomcat application server, the Sun Java Enterprise Application Server, and the Sun Java Enterprise Web Server against a Java Shibboleth Service Provider server at UTHSC-H.

Lessons Learned

There were some important lessons learned throughout the implementation of this project, including:

- Policy work is slow, and inter-institutional policy work is very slow, so pick your projects carefully.
- Smart collaboration dictates teams should work to solve real problems, find very low hanging fruit, and make people's lives better.
- No matter how compelling the need behind a project, that project must have a champion to drive it forward.

More Information

For more information on the UT System's Baylor College of Medicine/UT Health Science Center at Houston Shibboleth Project, contact:

William A. Weems, Assistant Vice President
Academic Technology University of Texas Health Science Center at Houston,
William.A.Weems@uth.tmc.edu
Or:
Stephen R. Ford, Assistant Director
IT Security and Compliance
Information Technology Program
Baylor College of Medicine,
Stephenf@bcm.edu

Case Study 3:
Houston Academy of Medicine – Texas Medical Center (HAM-TMC) Library Shibboleth Proxy Service

Controlling access to electronic resources over the Internet is easy when all of the information to be shared is public. However, when access must be restricted to specific individuals or groups, while maintaining a modicum of control over assets that must be protected, the job becomes much more difficult. This is especially true when one really has very little control over the population requesting access.

The Houston Academy of Medicine Texas-Medical Center Library (HAM-TMC), an institution supported by a consortium of more than forty research and health care institutions, is located in the heart of the Texas Medical Center (TMC). This library, which is supported by four universities (two at the University of Texas, plus one each at Baylor College of Medicine and Texas Woman's University), provided the test scenario for this case study.

The library provides resources to a group of medical research and provider institutions under a cooperative agreement and acts as the proxy for digital content providers. Since the library's environment includes a relatively small IT staff and there are more than 100,000 institutions accessing the library's data, access management had become a nearly impossible task. In such an environment, one may rightly ask how the Library is to keep up with defining the electronic-resource access rights that need to be given to people based upon their institutional affiliation when the library has no means of tracking or verifying such affiliations. The obvious answer is it can't be done in the traditional manner.

In this case study, using Shibboleth®[9] technology, the HAM-TMC Library was removed from the identity management and access management role. This provided an improved security and access management model for the library, for the content providers, and for the participating institutions by putting the burden of identity and access management back on the Identity Providers (who have a larger stake in keeping this information current). In this case, the stakeholders are the participating institutions.

Problem Discovery

Remote access to the Library's on-line databases is considered critical to research,

[9] The Shibboleth Software is located at http://shibboleth.internet2.edu and at the NMI-EDIT web page at www.nmi-edit.org.

academia and health care. Until recently, all web access for these areas has been available through a proxy-based access (via IP authentication) authentication and authorization model or through a locally-hosted, account-based model. While these models serve the community, they do not scale and it is difficult, if not impossible, to control the scale of account management as more users request access from home and off-site facilities. A better, more scalable and manageable mechanism is required for long-term use.

NMI-EDIT Components Highlighted in this Case Study

UT made use of the research and tools provided by the NSF Middleware Initiative-Enterprise and Desktop Integration Technologies (NMI-EDIT) Consortium, including the Enterprise Authentication Implementation Roadmap[10], the Shibboleth Software, and the Enterprise Directory Implementation Roadmap[11]. Discussions with colleagues at NMI-EDIT and EDUCAUSE also proved valuable.

The Solution

One of the factors in the HAM-TMC Library environment that exacerbated the need to move to a new authentication and authorization was that, generally, the library's IT staff had no way of knowing if a user still had a relationship with a particular supporting institution. The environment also posed significant accounting challenges as well, since the library services approximately 5,000 people per week and access to resources is user-based.

Given the need to move to a new model, this case study was adopted as a way of solving the problems of scale and management of library resources in the two existing authentication/authorization models. By way of a distributed administration model, the Shibboleth technology provides both the scale and manageability the HAM-TMC Library needed. Using Shibboleth, the user's institution is responsible for the creation, management and removal of account information and user status. Implementing this model would remove those responsibilities from the HAM-TMC Library's IT staff. Long-term, as more of the institutions have the ability to implement Shibboleth technology, the library would be able to move over 98% of its identity management responsibilities back to the supporting institutions and need only provide identity services to non-institutional members.

[10] The Enterprise Authentication Implementation Roadmap can be located at http://www.nmi-edit.org/roadmap/auth-roadmap_200510/index.html

[11] The NMI-EDIT Enterprise Directory Implementation Roadmap is located at http://www.nmi-edit.org/roadmap/internet2-mace-dir-implementation-roadmap-200312.html.

Both are also available from the NMI-EDIT web page at www.nmi-edit.org.

The proposed project included changing the current library proxy server model (EZProxy) to incorporate the application's ability to utilize Internet2, Shibboleth-based authentication and authorization. After completing a proof-of-concept, the library would then undertake a pilot to integrate Shibboleth into its production systems to allow two or more of the TMC institutions running Shibboleth Identity Provider (IdP) services to access library resources.

Project Resources

Staff

The project was implemented primarily by three UT System staff, two HAM-TMC Library staff and one Baylor College of Medicine staff member. All staff worked on the case study on a part-time basis.

The Project Timeline

The implementation required approximately eight weeks to order, configure and implement the applications. The testing phase required an additional two weeks in order to confirm configurations, connectivity and access management functionality.

Technical Factors

Project staff created a Shibboleth-enabled proxy service for the HAM-TMC Library, including a local Shibboleth-based IdP service and a local LDAP directory. Remote access user accounts created within the EZProxy server were migrated to a single LDAP entry for each user. This provides additional meta-data collection to allow better management of non-institutional user accounts and, in the long-term, simple migration to the InCommon Federation[12]. A Shibboleth IdP service was instantiated to leverage its directory for integration with a federation. The EZProxy server was brought up to allow remote users to authenticate via Shibboleth or their previous local library account. The new server required Secure Socket Layer (SSL) for access to authentication content.

The case study's technical requirements, procedures followed and project "cookbook" are available at: https://is.rice.edu/~bribbeck/Projects/NMI/NMI-EDITHAM-TMCCaseStudy.html

Lessons Learned

The migration from a production system needs to be considered in such projects. Fortunately, the EZProxy service could operate in both the traditional local database model as well as by utilizing a Shibboleth resource provider service. This made the transition easy to manage from the library's perspective. Due to thin information technology staffing at the library, a great deal of the work was done by information technology staff from other institutions. The process for coordinating and requesting information technology resources between institutions should be agreed upon and well understood before any such work is undertaken. It was sometimes difficult to coordinate testing between external components due to competing local projects.

[12] See http://www.incommonfederation.org/

This process did provide for some insight into what might be entailed in troubleshooting with inter-institutional information technology staff. Other elements to consider are the trouble reporting structure for inter-institutional applications, including how problems should be worked out and reported, who should take ownership of problems, and what tools are available for problem diagnosis.

In Conclusion

A great deal of technical and logistical insight was gleaned in this case study. Of specific note was the tradeoff between embracing a new technological model while attempting to scale the existing systems. The federated model proposed in this case study presented many new and untested challenges in the realm of inter-institutional support mechanisms. The challenges were made easier by the close collaborative nature of the institutions engaged in the project and the desire by all to ensure that the collaborative library resources could be maintained.

There are potential monetary factors to ponder when considering the security aspects of maintaining the status quo. One such factor is the prospect of increased fees, forwarded from the library to the participants, for compensation due to data spills or improper use by the institutional member. Other factors include lean staffing, the lack of scalability in existing models and the need to embrace the changing world of data access.

Clearly, the new federated access model will require a great deal of change in the way that institutions operate and *inter*operate. In order to change policy, operational, identity management and inter-institutional access control infrastructures, time, effort and the institutional will to change are all required. Nevertheless, as the federated access model has no requirement for over arching management or control, it will be easier to implement this model in collaborative projects, since maintaining the status quo would be more difficult.

The most viable means of providing the leverage required to build federated infrastructures seems to be existing collaborative needs or top-down control. Once the infrastructures are in place, the operational challenges may seem easy in hind site, and the overall benefits of ease of use and scalability should balance the effort required to make the change.

More Information

For more information on this HAM-TMC Library Shibboleth EZ Proxy Service contact

Barry R. Ribbeck
Director of Systems, Architecture and Infrastructure
Information Technology
Rice University
Barry.R.Ribbeck@rice.edu

Case Study 4: Shibboleth Version of Employee Benefits Annual Enrollment Application

The UT System Administration Office of Employee Benefits and Group Insurance (EGI) manages the insurance benefits of all employees and retirees of the UT System. Every year, the EGI office requires all employees and retirees to participate in the annual benefits open enrollment and to use the benefits annual enrollment application to select and/or update their insurance coverage for the upcoming fiscal year. During the rest of the year, employees can use the benefits application to view their current coverage.

The goal of this case study was to create and pilot a federated version of the UT System Administration Office of EGI benefits annual enrollment application, UTTouch, using the Shibboleth®[13] software and the UT System Identity Management Federation.

While the pilot will only be available to a small subset of UT System's user population, this application will eventually be rolled out to over 80,000 benefit-eligible UT System employees and retirees.. This rollout will be UT System's first large-scale, system-wide deployment of a federated application that uses Shibboleth.

This is still an on-going project.

Problem Discovery

Historically, UTTouch has required employees and retirees from all UT institutions to login using their social security number (SSN). Since the use of SSNs as identifiers will be prohibited by UT System policy by 2007,Shibboleth and the UT System Federation are seen as a means to achieve compliance with this new policy. In addition to requiring SSNs for user-login, an additional drawback of the existing UTTouch application was the complexity in administering the application. During their exploration of a Shibbolized version of UTTouch, UT System discovered that a Shibboleth solution could not only help them make great strides in the elimination of the use of SSNs as credentials and simplify application administration, it had the potential to greatly enhance the user experience as well.

The Solution

UTTouch is a web interface to a mainframe application hosted by UT at Austin. The application is written in Natural using an ADABAS database. Web agent code is homegrown. The Shibboleth version of

[13] The Shibboleth Software is located at http://shibboleth.internet2.edu and at the NMI-EDIT web page at www.nmi-edit.org.

UTTouch will allow employees and retirees to simply login to their home institution's IdP server and let Shibboleth send appropriate attributes to UTTouch. In addition, the Shibboleth version of UTTouch will provide year-round access to UTTouch to new and current employees plus retirees for their use in making an initial insurance selection or viewing their current coverage, respectively.

During the first week of July 2006, two or three small UT institutions (UT System Administration, UT Tyler, and/or UT Permian Basin) will participate in the pilot in parallel with regular annual open enrollment participants.

The size of the pilot institutions, along with the project staff's familiarity with their infrastructure and its managing staff, should help mitigate any potential data, infrastructure, and scalability challenges during the rollout of the pilot application.

Project Resources

Staff

The project is being implemented primarily by three UT System staff:

- Ben Armintor, Operation Systems Specialist at UT at Austin. UT at Austin hosts UTTouch and the administrative systems used by EGI.
- John Cotton, Manager of Information Systems at UT System Administration EGI, who comes to the project as process and data owner with substantial UTTouch experience.
- Paul Caskey, Technology Architect from the UT System Administration Office of Technology and Information Services, who comes to the project as the technical subject matter expert and with substantial middleware experience.

All staff members are working on the case study on a part-time basis.

NMI-EDIT Components Highlighted in this Case Study

UT made use of the research and tools provided by the NSF Middleware Initiative-Enterprise and Desktop Integration Technologies (NMI-EDIT) Consortium, including the Shibboleth software and the eduPerson Directory Schema[14]. Outreach, workshops and discussions with colleagues in NMI-EDIT and EDUCAUSE about federated identity management also proved valuable.

The Project Timeline

The pilot of the Shibboleth version of the UT System Benefits Annual Enrollment Application began in July of 2005.

Milestones

In chronological order, the critical milestones of this project were:

- Modification of the web agent.

[14] The eduPerson Schema is located at http://www.educause.edu/eduperson and at the NMI-EDIT web page at www.nmi-edit.org.

- Determining how to access HTTP assertion headers in Natural.
- Configuration of environment variables with clients.
- Configuration and implementation of Shibboleth SP server.
- Base installation of the UTTouch application in the SP server.
- Implementation of application changes for receiving user data via Shibboleth target/http headers.
- Definition of application indexes to access the ADABAS database.
- Modification of login verification modules.
- Addressing policy issues.
- Creation of proper authorizations for the new Shibboleth SP server.
- Migration of year-round and annual enrollment applications to the new Shibboleth server.
- Initial testing of the new application version for authorization and proper functionality.
- Implementation of the pilot with completion, based on feedback, iterative improvements/fixes.
- Assessment of application service level considerations.
- Addressing help desk and support issues.
- Moving to production mode.

Federating legacy systems can present tremendous challenges. Over time, these legacy environments are often adapted to be very specific to the applications they serve. Accordingly, problems may be encountered when applying modern technologies to these systems.

Policy Factors

As expected, this project has presented significant challenges in both the technical and policy areas. Resolving just this project's policy issues could take more time than would likely be have allotted for completing the entire case study pilot.

The first policy issue to be resolved revolved around the population, release, transmittal, and use of SSN data. A significant, though short-term, change that could have to be made while implementing this pilot must be considered for the possible scenario in which UTTouch no longer asks for the SSN and then uses that SSN as the user's identifier, but instead, the user authenticates with their local institutional credentials and the Shibboleth Attribute Authority asserts the SSN in the background over the encrypted connection.

The concern with such a scenario is whether it is appropriate to assert an SSN. So far during this pilot, the answer received from UT institutions indicates that this is acceptable, provided that assertion occurs over an SSL connection - something that is configurable at the IdP. In addition, if the SSN is the only attribute being asserted, then one can argue that the information in the assertion is not necessarily personally

identifiable since there would be no other data except for the SSN.

Still, there is always the possibility that a UT institution may decide it does not want to assert SSNs, even in cases where it is clear it will only be used for reconciliation purposes, because the institution may not want to put users in possible jeopardy of identity theft.

Another important consideration with this policy issue of SSN assertion is the fact that not all UT institutions have SSNs in a directory. Shibboleth does not insist that data reside in an LDAP directory and can just as easily pull SSNs from a SQL database. Though this is not a very significant technical issue, it may present a considerable amount of effort to identify authoritative sources at each UT institution and devise rules and policies for the release and use of SSNs.

The objective for addressing this policy issue in the medium or longer-term horizon of this pilot is to devise a way to modify UTTouch so that it would not require assertion of SSNs but instead require some type of unique local identifier. One option being considered is to use the institution code and the institution local UUID to identify users and to then cross-reference that data with the SSN in tables already maintained by the application itself.

Lastly, UTTouch is currently used by employees and retirees during the annual benefits enrollment. While UT institutions gather and maintain data about all their employees and maintain identity credentials for them, the same cannot be said for retirees. Most UT institutions do not maintain data about, much less credentials for, their retirees. Providing retirees with alternate ways to authenticate to UTTouch will require consideration of various business processes.

Technical Factors

UT at Austin's web agent is a custom web server plug-in that is written in Natural and provides web applications with a convenient, script-based interface to legacy mainframe systems.

The main significant technical hurdle that must be overcome during the pilot involves modifying UTTouch's web agent plug-in to recognize the HTTP headers that comprise the Shibboleth assertions and to then successfully relay them to the web-agent-based applications. Applications running in the modified web agent environment will have access to assertions from Shibboleth IdP servers within the UT System.

Lessons Learned

Even though there is considerable pilot work left to do, some lessons have already been learned, including:

- Not all web applications are easily Shibbolized, especially those that involved legacy systems.
- There are considerable policy issues surrounding the assignment, use, population, and release of identifiers in a federated environment, especially when dealing with sensitive data, such as the employee benefit information being worked with in this pilot.
- Despite the substantial benefits, there are many hidden costs and hurdles in creating a Shibboleth version of an application.

In Conclusion

The Shibboleth version of the UT System Annual Benefits Enrollment application is still a work-in-progress. This report is meant to serve more as description of the project rather than as documentation of a successfully implemented pilot application.

Considerable progress has been made. Most importantly, the UTTouch's web agent has been modified and tested in a development environment. But moving the project forward, and coming to grips with the most difficult issues (as described in the Policy Factors section) has brought to the fore the need for permanent, unique identifiers while addressing the problems and issues associated with them.

What this project is providing evidence for is that application owners usually require their applications consume only their own local identifiers yet, in a collaborative environment where federated institutions require access to remote applications, not all institutions requiring access may have the capability of asserting these identifiers.

Likewise, it is very challenging to gain consensus on a permanent unique identifier. Further, even if such consensus is gained, application owners may still be unable to modify their applications to consume the new identifier.

This is still an on-going project - stay tuned.

More Information

For more information on the Shibboleth Version of Employee Benefits Annual Enrollment Application case study, please contact:
Miguel Soldi at msoldi@utsystem.edu.

Case Study 5: The University of Texas System Identity Management Federation

A significant by-product of the UT System's participation in the NSF Middleware Initiative-Enterprise and Desktop Integration Technologies (NMI-EDIT) Extending The Reach (ETR) project is the System's creation of the UT System Identity Management Federation (UT System Federation). The UT System has prepared this case study to educate others about the federation-building activities they undertook as an expansion of their ETR program participation.

.

The UT System is a diverse educational and research university system, consisting of nine general academic institutions and six health institutions, employing 81,000 staff and faculty, and has a total enrollment of 175,000 students. In an effort to enable greater synergy and collaboration among its institutions and better serve their broad constituencies, the UT System set out in 2004 to create a core resource consortium of middleware-leveraged resource providers. The goal of this effort is to create the foundation for a federated, collaborative infrastructure that will enable greater synergy and collaboration among UT institutions, simplify application security through a common trust fabric, and provide a platform to address UT System-wide identity management initiatives.

Though still a work-in-progress, the UT System Federation is comprised of all sixteen UT institutions and operates under the authority of the UT System Board of Regents and the guidance of a six-member Executive Committee. The project work to date has been targeted at moving the project from a conception and planning stage to an instantiation and implementation stage.

Problem Discovery

Shaping up the business problem behind the UT System Federation was akin to pulling on a string of yarn and finding out that there's a big ball of it around a corner. Early in 2004, UT System started looking for a way to solve some library authorization issues and, in addition, look for an alternative to using Public Key Infrastructure (PKI) to secure electronic communications. The ensuing dialogue brought forth the fact that application security is becoming increasingly onerous. In institutions of higher education like the UT System, it is not uncommon to find staff, faculty and students requiring access to a variety of applications that may be dispersed among multiple institutions and that may be accessed under numerous user roles in multiple contexts.

Under these conditions, traditional forms of authentication and authorization are no longer sufficient for providing the level of assurance needed by current Internet-based applications. In addition, new laws and

regulations continue to dictate more stringent identity management processes.

The Solution

The UT System Information Technology Strategic Leadership Council, a UT System-wide governance body charged with providing leadership for system-wide initiatives, agreed that deployment of a robust, secure, interoperable infrastructure for identity management in support of inter-institutional collaboration was a UT System strategic goal. This infrastructure is to be based upon the following available standards and best practices:

- Lightweight Directory Access Protocol (LDAP) compliant directory services
- eduPerson schema as promulgated by EDUCAUSE and Internet2
- utPerson schema (yet to be developed),
- Inter-institutional access control that utilizes Internet 2 Shibboleth®[15]
- Consistent institutional definitions and identity management trust policies for students, faculty, and staff, as well as sponsored affiliates.

The infrastructure must provide two sets of services:

- Metadata management – This includes the aggregation, distribution and maintenance of federation members' attribute data, syntax and semantics.
- Trust management – This includes federation and member operation practices, along with policies for control, privacy, and security.

To provide these services, participating institutions *must:*

- Build the architecture so that it is manageable and reliable.
- Identify and promote applications to showcase and demonstrate the benefits of middleware initiatives.
- Coordinate the initiatives of multiple institutions.

Despite the fact that all UT institutions share a common vision and have the same legal requirements, they enjoy considerable autonomy and exhibit significant differences in size and budget. In a sense, they are sixteen stovepipes. Thus, the primary political challenge facing the UT Federation involved having all institutions *agree* on technical specifications (e.g., the exchange of data attributes, interoperable software) and policy specifications (e.g., privacy, establishing trust and trustworthy data) to successfully provide the services mentioned above. In addition, each institution has it own portfolio of information technology projects and corresponding agendas, so coordinating the initiatives of, and with, multiple institutions has been (and will continue to be) a challenge.

On a technical level, building the architecture so that it is manageable and

[15] The Shibboleth Software is located at http://shibboleth.internet2.edu and at the NMI-EDIT web page at www.nmi-edit.org.

reliable, as well and identifying and promoting applications to showcase and demonstrate the benefits of middleware initiatives, are increasingly important tasks in the successful implementation of the solution.

Project Resources

Staff

Paul Caskey, Technology architect at UT System Administration, is the staff member dedicated to the UT System Federation. All other staff involved in the governance and operation of the UT System Federation work on a part-time basis and in addition to their respective institutional projects.

Funding

To date, the UT System Administration Office of Technology and Information Services (OTIS) has funded the UT System Federation project. The project has had no specific resources allocated to it and was simply added to OTIS' project list. This funding model should change in the future, since a schedule of federation membership fees has been approved to provide continued funding for the project.

Maintenance of the UT System Federation WAYF ("where are you from") resource requires about a third of one full-time employee (FTE). This estimate does not account for deploying new applications, integration, or the time required for ongoing collaboration within, and outside of, UT System Administration. Administration and ongoing policy work should require about a half of one FTE. The hardware and software required to support the UT System Federation cost under $10,000.

NMI-EDIT Components Highlighted in this Case Study

UT made use of the research and tools provided by the NMI-EDIT Consortium. These included the Shibboleth Software, and the eduPerson Directory Schema[16]. Outreach, workshops and discussions with colleagues in NMI-EDIT and EDUCAUSE about federated identity management also proved valuable

Project Implementation

Conception and Planning

Instead of being application-driven in nature, the business drivers behind the UT Federation are long-term and high-level. The drivers are:

- A need to encourage greater synergy and collaboration among UT institutions.
- A need to simplify application security through a common trust fabric that allows a UT System-wide secure exchange of authentication and authorization attributes
- A need to provide a platform to address UT System-wide identity management initiatives.

[16] The eduPerson Schema is located at http://www.educause.edu/eduperson and at the NMI-EDIT web page at www.nmi-edit.org.

Though creating and supporting a customized federation is considerably more complex and time-consuming than joining an existing federation such as InCommon, from UT System's perspective, having their own federation would provide some compelling benefits, including:

- The ability to leverage existing inter-institution agreements.
- Establishment of a common set of standards and attributes for UT institutions.
- More granular control over authentication and authorization policies.
- A forum for experimentation and dialogue.

In Phase I of the project, membership in the UT System Federation will be restricted to the sixteen UT System institutions and will support their intra-institutional collaborative efforts. This will level the field among the UT System institutions and allow the System to get "the house in order" before opening it to other organizations.

Three strategies were used for planning the implementation of the UT System Federation. The first was to obtain agreement from UT leadership that the creation of a federation was a strategic goal and then to translate that agreement into a Statement of Direction that all UT institutions could understand and follow. The second strategy was to communicate and inform stakeholders, UT leadership, and information technology steering committees about the project's direction, potential benefits and high-level technology and policy requirements as a means of gathering feedback and fostering buy-in from these institutions. The third strategy was for the project team to be especially careful when addressing institutions' concerns about privacy, security, and trust.

The availability of resources to dedicate to the project was, and continues to be, a significant challenge. Since each institution has it own portfolio of information technology projects and corresponding agendas, finding value-adding applications and drivers that can justify the commitment of time and effort by all proposed member institutions is an ongoing activity.

Early in the project, the following three high-level items were identified as critical conditions for project success:

- Trust between all UT institutions
- Policies that implement the common trust fabric across the system
- Common technology framework, standards and protocols

Two parallel tracks were pursued in order to ensure these project conditions were met. The first track addresses the organizational and policy considerations of federation by establishing governance, drafting policy documents, and negotiating the operational procedures that, together, are the basis for a common trust fabric. The second track addresses technical considerations such as

infrastructure building, authentication and authorization mechanisms, and standards and protocols for attribute naming.

Federation Instantiation

Organizational and Policy Factors

When creating an identity management federation, technical and policy work take place in parallel. The problem with this is that policy work is slow and inter-institutional policy work is very slow so, inevitably, policy will lag behind technology. So while all the technical pieces may be in place and working properly, the agreements and rules that establish how things should happen will remain a work-in-progress.

Policy considerations are a function of the broader relationship that federation members have with each other. The closer that relationship is, the fewer the policy reconciliation conflicts. Conversely, the more distant the relationship, the greater is the possibility for conflict. As mentioned before, UT System institutions are diverse and autonomous, and consequently, there is also diversity in the operating practices and standards among institutions, with no consistent identity management trust definition or policies.

Policy documents that address these factors have been drafted and, at the time of this report, are about to be approved and implemented. The documents, to a great extent, use the InCommon policy documents as a model. There are however, two exceptions to this – the Members Operating Procedures and the Membership Agreement. The Member Operating Procedures document was drafted using The University of California UCTrust Federation documents. This was done with the intention of creating a document that explicitly defines the standards, policies and procedures that UT System Federation members must put in place in order to be able to make informed relying-party decisions. The Membership Agreement was drafted to leverage the existing UT System Inter-institution Agreements in order to simplify membership contract. These draft documents of the UT System Federation are available at: https://idm.utsystem.edu/utfed/index.html

Governance Structure

The UT System Federation, which is sanctioned by the UT System Strategic Leadership Council Statement of Direction, is comprised of all sixteen UT institutions and operates under the authority of the UT System Board of Regents (see Figure 1 below). The affairs of the UT System Federation will be managed by the UT System Chief Information Officer and governed by a six-member Executive Committee charged with the following:

- Provision of oversight and conflict resolution.
- Establishing and managing trust agreements.
- Determining direction and formulating policy.

- Ensuring services meet business needs, while maintaining appropriate security and compliance with legal requirements.
- Establishing and communicating operational standards and processes.

Managing trust relationships has proven to be complex, even when dealing with institutions within the same system (among "family"). This is especially true when the governance work is entangled with power and/or autonomy conflicts. Priorities vary from institution to institution and, at times, the standards and conventions proposed in policy documents are perceived as dictates. Maintaining open communication channels while trying to reach some convergence on definitions and practices has been an exercise that requires lots of patience, active listening and a willingness to explore different ideas.

As the project moves forward, the project team will need to address issues related to indemnification, such as what happens when something goes wrong, who is liable, and what is the impact on intra-institutional trust? These Membership Agreement questions require addressing without turning the document into an overbearing legal contract.

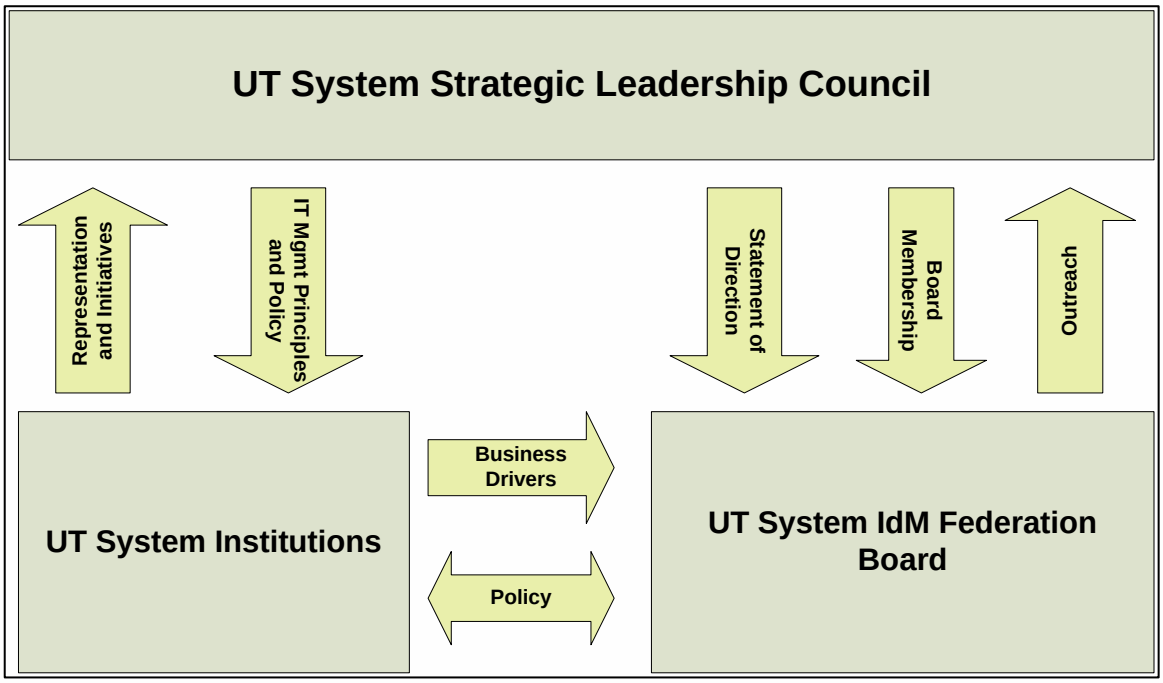

Figure 1: UT System Federation Governance Structure

Technical Factors

The main technical challenge to the instantiation of the UT System Federation was establishing standards and prerequisites while simultaneously leveraging each UT institution's existing infrastructure and bringing all institutions onto "the same page."

Following the Statement of Direction provided by the UT System Strategic Leadership Council, the technology required for Identity Provider institutions was:

- Shibboleth v 1.3
- RedHat Enterprise Linux 3.0
- Apache 2.0
- Tomcat 5
- LDAP infrastructure for authorization
  - → LDAP Server (Sun ONE Directory Server, OpenLDAP, etc.)
  - → eduPerson schema
- Verisign Server Digital Certificate

Based on the established standards and prerequisites, UT System Administration OTIS took on the responsibility of assessing the readiness status of all UT institutions and coming up with the best way to bring all institutions to a common readiness point. This assessment was completed by surveying the staff responsible for middleware at each UT institution. Based on survey results, a plan was developed that would help to close any gaps between the institutions. The plan included providing help-desk type support, maintaining a discussion group and a middleware-related web page, providing on-site consulting services, and hosting system-wide middleware events.

Shibboleth Identity Provider Install Fest

The first system-wide middleware event, a one-day Shibboleth Identity Provider (IdP) Install Fest, was held in September 2004 and hosted by UT System Administration. Representatives from twelve of the sixteen UT institutions attended the workshop that was designed as a hands-on exercise to assist each institution in implementing a Shibboleth IdP server. The goal of the Shibboleth IdP Install Fest was to have an operational test UT System identity management federation by the end of the day. Emphasis was placed on the idea that the install fest was not merely a practice exercise or discussion of Shibboleth theory, design and operations, but rather was a carefully scripted, technical consultation on how to implement Shibboleth on a live system. What this meant for a lot of folks was that, in fairly short notice, they had to get a directory up, provisioned, and synchronized while also putting an authentication infrastructure in place.

Topics discussed included how to configure Apache for Shibboleth, Shibboleth for certificates, Origin.xml, sites.xml, resolver.xml, and ARP. The following material was made available to all UT institutions at the time the Install Fest was announced:

- Readiness Assessment Survey

- Install Fest Prerequisites
- Install Fest Agenda
- Shibboleth Identity Provider Installation Checklist
- Build Instructions for Shibboleth Identity Providers
- List of useful IdM downloads

These documents are available at: https://idm.utsystem.edu/etrnmi/ShibbFest/index.html

The Shibboleth Identity Provider Install Fest met its main objectives and, by the end of the day, nine of the twelve attending institutions were successful in getting their IdP server up and running (albeit, to varying degrees of production-ready status). Follow-up on-site visits helped bring the other three institutions to the same production-ready status as their peer institutions. All sixteen UT institutions now have operational Shibboleth IdP servers.

Shibboleth Service Provider Install Fest

As a follow-up to the September 2004 Shibboleth IdP Install Fest, UT System Administration hosted a two-day Shibboleth Service Provider (SP) Install Fest in June of 2005. Representatives from eleven of the sixteen UT institutions attended the workshop. Like its predecessor, the Shibboleth SP Install Fest was also designed as a hands-on exercise. The goals for the SP Install Fest were, by the end of the day, to have assisted each institution in learning the basics of protecting web content using Shibboleth (as implemented in the UT System Federation) and to have Shibboleth-protected web content running at each participating institution - even if the content was just a temporary trial page.

The event was targeted primarily to the webmasters and information technology infrastructure staff responsible for setup and operations of Shibboleth-protected web applications. Rather than webmasters or content providers, most of the attendees were more technical staff than, which was a slight disappointment.

As with the Identity Provider Install Fest, there were specific requirements that needed to met prior to an institution's participation in the install fest, including:

- Having an accessible installation of IIS or Apache 2.0, configured with SSL support.
- Having a Verisign server digital certificate (signed by a commonly trusted CA) configured for the web server.
- Having a firewall configured to allow tcp/443 and tcp/8443 outbound and tcp/443 inbound to/from the web server.
- Having an application or web-based content to make available via Shibboleth (a 'Hello World' page or PowerPoint HTML was sufficient).
- Having Shibboleth SP Software built and installed.
- Having an AnyTarget ARP to test the installation's new authorization features

Topics of discussion included registering applications with the federation, installation and configuration of Shibboleth SP, and configuring attribute acceptance policies. The following material was made available to all UT institutions at the time the Install Fest was announced:

- Install Fest Prerequisites
- Install Fest Agenda
- Shibboleth Service Provider Workshop Script
  - → Initial Shibboleth configuration for Apache
  - → Initial Shibboleth Configuration for Microsoft IIS
  - → Instructions on the SHAR, SHIRE, Applications Element, CredentialsProvider Element, and the Application Element
  - → Handling Authorization
  - → Production Considerations
  - → Advance Issues
- Shibboleth Service Provider Setup Guide

These documents are available at: https://idm.utsystem.edu/SPfest/

The Service Provider Install Fest met its main objectives and by the end of the day, everybody had a working application up (although most of these applications were a "Hello World" page or the like).

Applications

Thus far, Shibboleth-enabled access to the UT System Administration wireless network has been only application added to the UT System Federation. This application is typically used several times a year when staff employees and executive officers from UT institutions come to meetings at System Administration. During these meetings, users connect to the network via the wireless access available in the conference rooms. Individuals connecting as “Guest Users” are given access via a slower, restricted connection, while System Administration employees are given access via a considerably faster and unrestricted connection. Shibboleth-enabled access to the wireless network allows employees from UT System Federation member institutions to obtain faster and unrestricted access to the wireless network when they login with their local institution’s digital credentials. This very simple application has proven highly effective in illustrating the use and value of the federation. Some of the peer pressure this generated has helped facilitate a portion of the work needed to encourage all UT institutions to participate in the federation.

Other Shibboleth applications that can be made accessible via the UT System Federation include the Employee Benefits Annual Enrollment application (currently in progress), access to UT Blackboard applications for the University of Houston/UT Health Science Center school of Nursing, and EZ Proxy services for the UT Health Science Center at Houston School of Public Health.

What the Future Holds

The priorities for the UT System Federation project's immediate future are to conclude the review and approval of the policy documents, formalize the federation membership of all UT institutions, and increase the number of Shibboleth-enabled applications that promote synergy and collaboration among UT institutions.

Federation Operations

Progress thus far notwithstanding, the UT System Federation is still in its infancy. Eventually, as the federation evolves, the consortium will need to consider and address the following:

- How the creation, provisioning, authentication, and termination of user accounts should be managed.
- Whether user accounts should be reused or recycled by member institutions.
- The integration of Shibboleth with legacy systems.
- How to measure and document the benefit and performance of the UT System Federation.
- Continual assessment of the risks to UT System stemming from malicious exploitation of the Federation and/or inconsistent practices among UT institutions. These can result in unauthorized disclosure, fraud, and liability
- Preparation of a Security Plan and a Disaster Recovery Plan.
- The approvals and processes that are needed to enter into trust agreements with external institutions and other federations.

Federation Maintenance

In addition to operational considerations, as technology evolves and the requirements of the constituencies served change, the consortium will need to consider and address the following maintenance issues in order to keep the UT System Federation "fresh", relevant, and compliant with the increasing number of federal and state regulations:

- Establishing an Audit function and related procedures.
- Establishing a plan and gathering the resources needed to keep up with technology changes and with federal and state compliance requirements.
- How should the Federation be scaled? What opportunities exist to leverage the infrastructure?
- How to keep the Federation application relevant.

Lessons Learned

When implementing the UT System Federation, the following lessons were learned:

- The technical implementation aspects of Federation can get way ahead of policy and governance.
- The integration and Federation process is entangled with power and autonomy conflicts.

- Managing trust relationships is complex, even when dealing with institutions within the same university system.
- One must understand the importance of communicating business needs and the importance of the federation project.
- One shouldn't underestimate political issues; take political realities into consideration during implementation.
- Set standards, but be prepared to adjust to local requirements.

In Conclusion

Differing motivations and varying levels of commitment and urgency underlie the participation of UT System institutions in the UT System Federation. Nonetheless, the UT System has implemented a technical and policy infrastructure that can both support and scale as the needs of member institutions and constituents served dictate.

Although technical challenges were certainly significant, the most difficult challenges that have had to be overcome during the UT System Federation project have been political in nature. As the project moves forward, issues of policy, funding, and resource availability are becoming more pressing.

Still ahead of for the project team is concluding policy work, implementing compelling UT System-wide applications, and reaching the critical mass necessary to elevate the UT System Federation to a fully operational status.

More Information

For more information on the UT System Identity Management Federation, contact Miguel Soldi at msoldi@utsystem.edu