



The University of Florida: Identifiers, Social Security Numbers, and Identity Management

NSF Middleware Initiative (NMI) Integration Testbed Case Study Series

Series contact: Mary Fran Yafchak, Southeastern Universities Research Association,
maryfran@sura.org.

The NMI Integration Testbed Program provided practical evaluation of NMI components within the context of real projects and application scenarios from June 2002 through November 2004. During that time, NMI Testbed sites collectively submitted over 220 evaluation reports to middleware component developers as direct feedback into the NMI development cycle. Site representatives also actively inspired, promoted and facilitated the integration of middleware throughout their institutions.

The NMI Integration Testbed Case Study Series documents the most significant influences and outcomes of NMI Testbed sites' middleware integration efforts, highlighting intersections with established projects, application contexts and influences, drivers for innovation, decision points and challenges. Through this documentation, the work of these pioneering institutions is captured to provide a breadth of insight and approaches for others to use towards successful middleware development and deployment.

This NMI Integration Testbed 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.

Copyright © 2005 The University of Florida. The University of Florida permits use of this content for noncommercial purposes with proper attribution. All rights reserved.



Executive Summary

In 1996, the University of Florida (UF) began to tackle the migration away from the use of faculty, staff and students' Social Security Numbers (SSNs) as identifiers for university business processes through their GatorLink initiative. This project provided a unique GatorLink username and password for each faculty, staff and student member. However by 1999, while the GatorLink initiative had resulted in new computer services, the use of the GatorLink as an identifier in computer system programs had not been widely adopted.

An additional concern at UF was the numerous university record systems being separately managed - each requiring users to manage separate logins. Moreover, administration of these disparate systems was decentralized. This created a security risk for UF in that it wasn't possible to insure that all of a person's authorizations were revoked when their affiliation with UF changed. In 2000, a campus team created goals for directory services at UF, including a comprehensive enterprise-level directory with appropriate levels of security and accessibility. The task of migrating from the use of SSNs as an identifier was merged with UF's new directory services initiative. UF also gained insight into how NMI middleware should be integrated into their directory services initiative through the University's participation in the NSF Middleware Initiative (NMI) Integration Testbed Program.

With assistance from business units across the campus, in January 2003 the university introduced its new directory service and issued UF identifiers (UFIDs) to its faculty, staff, students, alumni and affiliates. Over 1500 applications were modified to work with the new UFIDs, while new web interfaces were created to allow users and departmental directory coordinators to manage specific directory information. Today, the GatorLink username and password provide single credential access to a wide variety of business processes. UF's LDAP directory, which supports NMI-EDIT schema standards, provides access to public contact information through the university's white pages. Directory registry affiliations drive automated business processes, such as the provisioning of basic university services and system access.

In July 2004, UF launched its new PeopleSoft HR and Finance systems. As part of their implementation, the University has committed to role-based application access, which associates persons with services through roles. Entries into the university directory are processed by PeopleSoft to automatically provision access to HR and Finance systems based on the persons' affiliations.

For more information about this University of Florida Case Study, contact Dr. Michael Conlon at mconlon@ufl.edu.



NMI Components Highlighted in this Case Study

The NMI components discussed in this case study series encompass NMI Releases 1 through 4. Information about NMI Releases can be found at <http://nsf-middleware.org/>.

eduOrg Directory Schema

eduOrg contains organization-related attributes for higher-education institutions to deploy for enabling inter-institutional collaborations.

Home site: <http://www.educause.edu/eduperson>

eduPerson Directory Schema

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

Home site: <http://www.educause.edu/eduperson>

Pubcookie

Pubcookie is open source software for intra-institutional web initial sign-on.

Home site: <http://www.pubcookie.org/>

Conventions and Best Practices

These NMI-EDIT documents reflect current NMI research in campus core middleware. The architecture approaches and policies promulgated here are in use at several leading campuses and institutions. Discussion includes “Practices in Directory Groups”, “Metadirectories Best Practices”, and “Enterprise Directory Implementation Roadmap” documents.

Home site: <http://www.nsf-middleware.org/NMIR5/nmi-edit/bestpractices.asp>

LDAP Analyzer

The LDAP Analyzer Service determines the compliance of an LDAP directory server implementation with various object class definitions such as inetOrgPerson, eduPerson, eduOrg, H.350 and the Grid Laboratory Universal Environment (GLUE) schemas, as well as the recommendations outlined in the LDAP-recipe and other best practices.

Home site: <http://middleware.internet2.edu/dir/>



University of Florida Directory Services and Middleware

In the early days of the Internet, there was little of the concern about identity theft and individual privacy that we have seen in recent years. Universities across the United States used faculty, staff and students' Social Security Numbers (SSNs) as identifiers for registration, classes, grading, financial aid, health services, and other university business processes. The SSN was listed on the students' identification card, grade list, housing form, and on hundreds of other publicly available records.

At the University of Florida, this was common practice, and because of the size of the institution (over 48,000 on-campus students), this was considered an efficient way to keep track of everyone. By their history, university record systems were separately managed at the University, using multiple types of hardware and software. Thus, the burden was placed on the student or staff member to navigate the various systems, using separate system usernames and manage their contact information in each of these systems individually. Some systems had no self-service and required paper forms and in some cases personal presence to update.

In 2000, an ad hoc committee was formed to develop approaches to addressing

enterprise authentication services, enterprise authorization services and the enabling of business processes using common directory information. In 2002, the provost requested a study on the feasibility of replacing SSN with a university-managed identifier. This work was merged with directory planning and also supplemented with the University's participation in the NSF Middleware Initiative (NMI) Integration Testbed Program ¹. Managed by the Southeastern Universities Research Association (SURA) on behalf of the NMI-Enterprise Desktop and Integration Technologies (EDIT) Consortium, the testbed consists of eight universities that participate in a closely coordinated effort to deploy and evaluate NMI technologies. As a Testbed Site representative, Dr. Michael Conlon, Director of Data Infrastructure for the University, evaluated NMI-EDIT² middleware components and provided valuable feedback to NMI developers while leveraging the experience to further the

¹As part of its overall effort to develop and disseminate software that lets scientists and educators share resources across the Internet, NMI began a practical deployment and evaluation effort called the NMI Integration Testbed: <http://1www1.sura.org/3000/NMI-Testbed.html>

² The primary goal of the NMI-Enterprise and Desktop Integration Technologies (NMI-EDIT) Consortium is to improve the productivity of the research and education community through development, testing, and dissemination of architectures, software, and practices in the areas of identity and access management. <http://www.nmi-edit.org>.



development of the University's middleware-enabled infrastructure. On January 19, 2003, the university introduced its new directory and issued UF identifiers (UFIDs) to over 500,000 faculty, staff, students, alumni and affiliates.

On July 1, 2004, the university launched its new PeopleSoft HR and Finance systems. The UF Directory provides identity and affiliation management services to the PeopleSoft systems. As the PeopleSoft student administration systems are implemented over the next two years, the UF Directory will be re-implemented via PeopleSoft Campus Community, providing the university with a consistent enterprise system architecture and support for both centralized and distributed university business processes.

Chronology

In 1996, the president of the university launched an initiative to develop a universal username for all users and systems. Work began on GatorLink, a common username and password for all university systems, and a set of services including email and web hosting that are provided to all university members.

UF's Cooperative Computing Initiative (CCI), an ad hoc group formed in 1999, first investigated the issue of developing an enterprise directory. In a brainstorming session, this group determined that there

were three basic problems facing university information systems:

- Maintenance of directory information – 17 core university business systems each created and maintained contact information separately creating poor customer service and duplication of effort.
- Enterprise authentication services -- the GatorLink initiatives had created new services, but had not been adopted for use by university business systems or desktop systems, leading to a large number of usernames and passwords and no association of usernames to identity.
- Enterprise authorization services -- permission information was scattered across many systems. It was not possible to answer the question "What authorizations does person x have?" As a consequence it was not possible to insure that all authorizations were removed when a person's affiliations with the university changed.

Though they were able to frame the issues and challenges, the CCI had no resources or mandate to move forward. Subsequently, the university Health Science Center received an information technology planning grant from the National Library of Medicine for Integrated Advanced Information Management Systems (IAIMS). Dr. Kenneth Berns, Vice President for Health Affairs was the principle investigator, Dr. Michael Conlon was the project manager and Marian Boyle was the project coordinator for this



grant. The Health Science Center determined that they suffered from the same three problems as identified a year earlier by the CCI. In light of this, a Directory Planning Team meeting was convened in June 2000, comprised of stakeholders and service providers involved with existing directory-related projects and services. The team settled on a set of desired outcomes for university directory services, including a comprehensive enterprise-level directory with appropriate levels of security and accessibility. This directory would need to support email addresses, multiple search terms, and business processes. After one year of close collaboration, the group generated a proposal¹ for the development of unified directory services.

Around this same time (spring 2001), the University Provost asked the Registrar to determine the best way to replace social security numbers with a university managed identifier. The Registrar convened a panel to discuss the implications of this change and the ways that it might be accomplished. The panel recommended to the Provost that they combine this work with the directory services proposal, based on the white paper written by the Directory Planning Team.

In the summer of 2001, Dr. Charles Frazier became the Vice Provost for Information Technology and his first order of business was to begin the directory project. Dr. Frazier organized a project team, led by Warren Curry, to implement the directory and a new identifier, "UFID". Units from

across campus donated manpower to assist with the implementation effort.

At the time, the directory project was the largest information technology project ever undertaken at the University, eclipsing even Y2K preparations and costing in excess of \$5M. Over 1500 administrative applications were modified to work with the new UFIDs. Programmers also created web interfaces for users to perform updates on their own information, and a proxy for directory coordinators to manage contacts. Every unit of the University chose or was assigned a directory coordinator, who received training to fulfill this role.

While investigating possible options for enterprise directory systems, the University also became aware of the National Science Foundation's Middleware Initiative (NMI). Beginning in June 2002 and for two years after, the University of Florida participated as a Testbed Site in the NMI Integration Testbed Program. By evaluating relevant middleware components from the first four NMI releases (May 2002, October 2002, April 2003, December 2003), the University was able to learn more about the scope of middleware components that could be useful in the directory project and also assist others in developing directory strategies and implementations. Multiple NMI-EDIT components were explored in this context:

- [eduOrg](#), [eduPerson](#)³ – Foundational Objectclasses for identify management

³eduOrg, eduPerson information:
<http://www.educause.edu/eduperson/>



in academic institutions

- [Pubcookie](#)⁴ – An option for Web-based SSO (single sign-on)
- [Conventions and Best Practices](#)⁵ – NMI-EDIT's guides and best practices for identify management infrastructure include “[Practices in Directory Groups](#)”, “[Metadirectories Best Practices](#)”, “[Enterprise Directory Implementation Roadmap](#)”.
- [LDAP Analyzer](#)⁶ – A tool for verifying the structure and compliance of established LDAP-based directory services.

The University of Florida directory service went live on January 19, 2003. The directory addresses authentication, authorization and support of university business processes. The University currently has over 800 directory coordinators, identified on a web page that is updated daily.

Directory Architecture

The UF directory architecture consists of a directory registry, application programming interfaces (APIs) for access to directory information, message queues for supporting business processes and web and batch interfaces to the APIs for directory maintenance. The global architecture is shown in Figure 1.

Self-service and directory coordinator web interfaces are available through the university portal ([my.ufl.edu](#)). Batch interfaces are used for large-scale updates, such as those from Student Systems.

The registry contains over 140 tables of person-related directory information including contact information and affiliations. Affiliations drive automated business processes such as computer account management and provision of baseline university services and system access. Directory APIs regulate all access to the registry.

Message queues are used to track information updates of interest to specific applications. Over fifty message queues are in production to provide directory information to university systems such as PeopleSoft, LDAP, Active Directory and the Hospital Information systems. All these systems have consistent contact and affiliation information provided by the UF directory. The systems' native capabilities to manage directory information locally have been disabled and all updates to directory information are provided by directory interfaces.

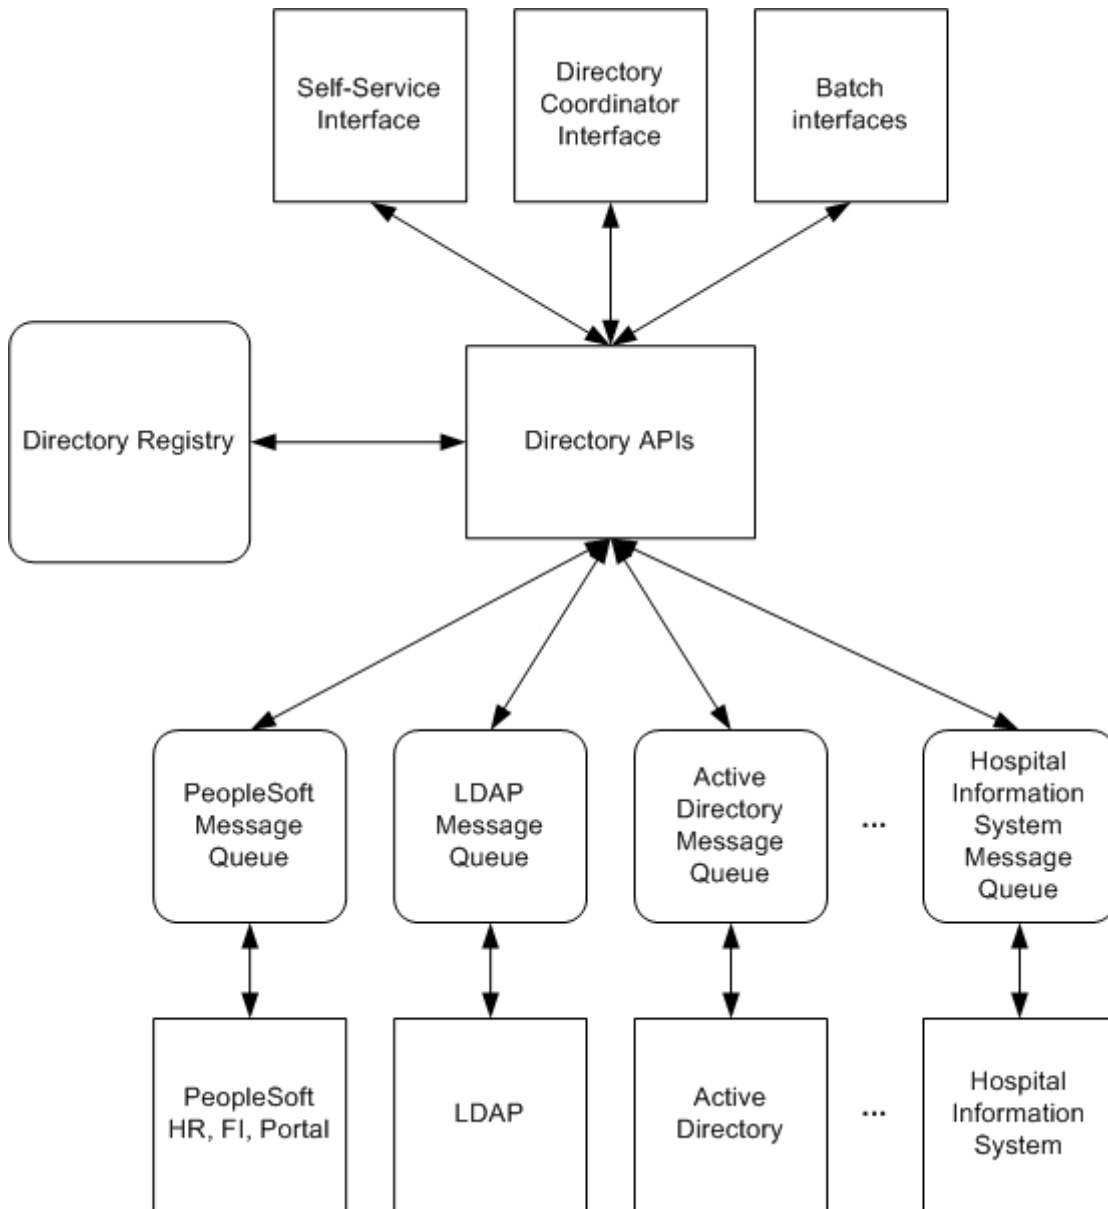
⁴Pubcookie information: <http://www.pubcookie.org/>

⁵Conventions and Best Practices information: <http://www.nsf-middleware.org/NMIR5/nmi-edit/bestpractices.asp>

⁶LDAP Analyzer information: <http://ldap.mtu.edu/internet2/analyzer/index.shtml>



Figure 1: UF Directory Architecture



Authentication

Today, the GatorLink username and password provide single credential access to a wide variety of university systems. Every person affiliated with the university has a UFID assigned by the directory APIs when the directory record is created. Local

directory coordinators can create new UFIDs and identity resolution processes are in place to avoid generation of duplicates and resolve rare duplicatesⁱⁱ. Once a directory entry has been created based on the person's affiliation with the university, a GatorLink username and password can be



created using self-service interfaces ⁱⁱⁱ. Establishing a directory entry, creating a new GatorLink username and password and provisioning access to the portal and university systems can be done locally and is typically completed in less than thirty minutes.

Every applicant, student, faculty and staff member must have a GatorLink username and password, also known simply as a GatorLink and associated with their UFID in the registry. The GatorLink ID must be eight characters or fewer, begin with a letter, and consist of only letters and numbers (no symbols). Future work will extend the GatorLink username to sixteen characters. GatorLink policies also require that each user have a strong password ^{iv} as a security measure to ensure that only that user can access their records and UF job-related data. While students, faculty and staff are expected to remember their GatorLinks, the UFID number exists as an opaque identifier and database key and is not displayed outside business transactions.

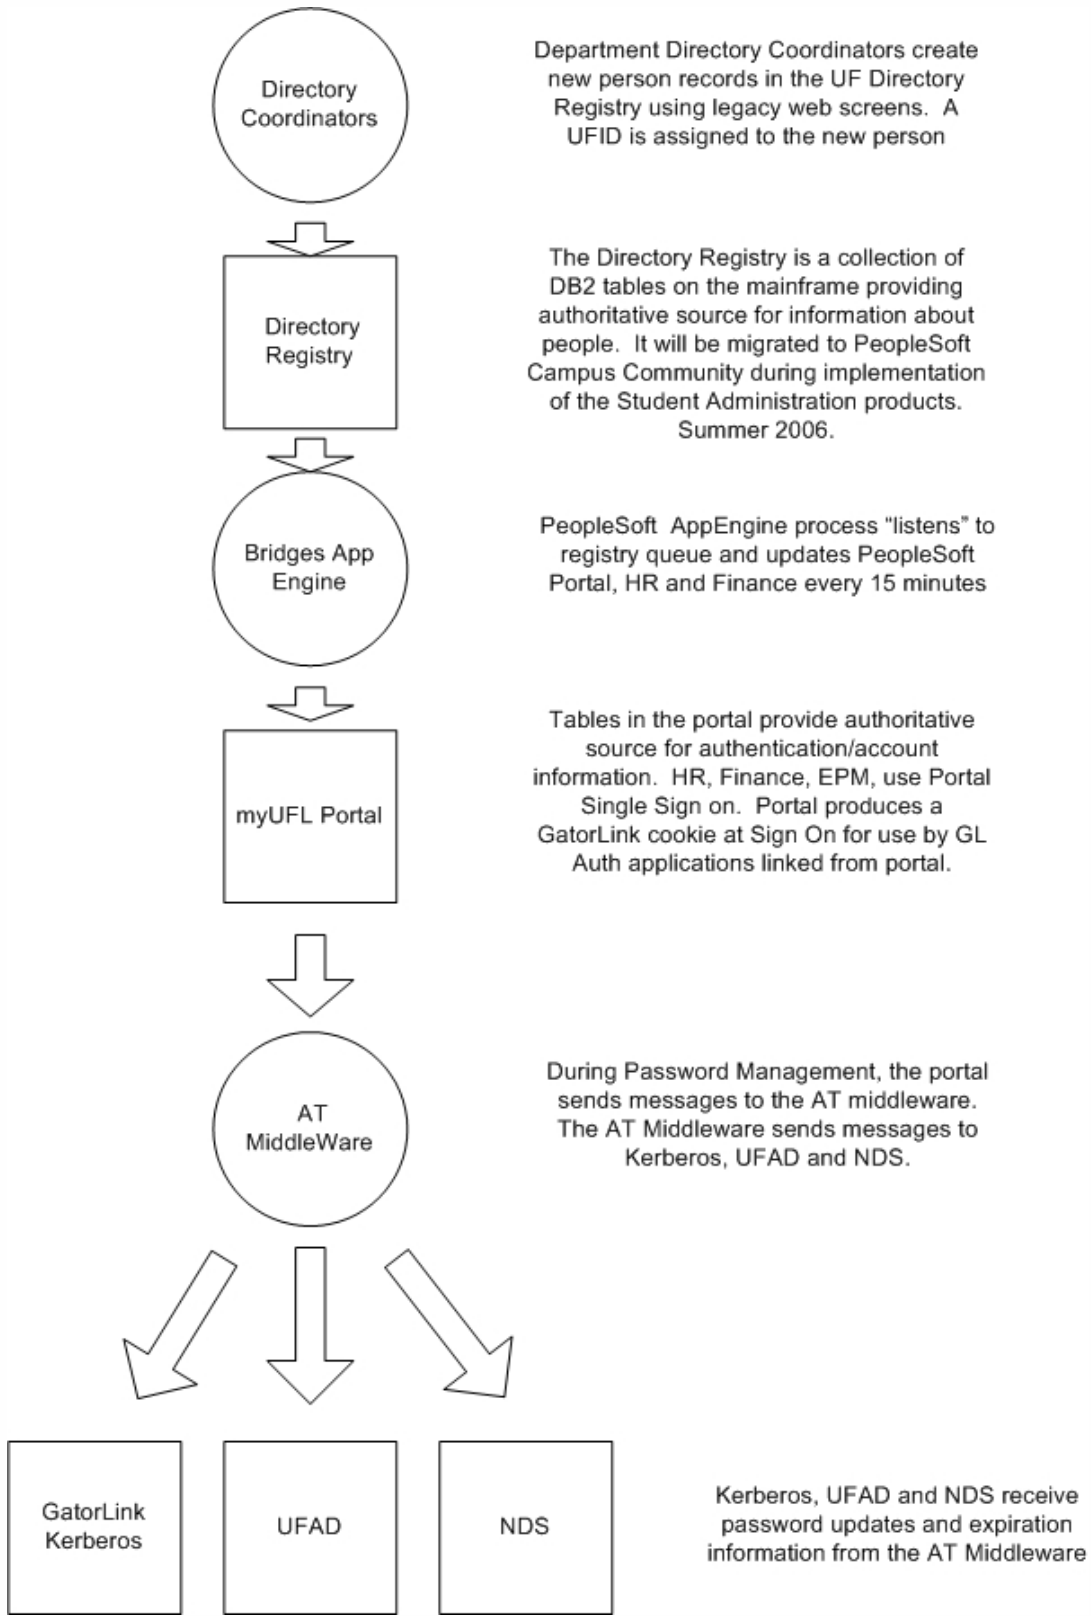
Many campus services use GatorLink authentication, including the myUFL portal,

access to PeopleSoft services, ISIS Student registration system, Library system, WebCT course management system, computer lab and network connections, print services, online training and the download of university-licensed software. Authentication services are synchronized to provide single credential access regardless of the authentication technology in use. See Figure 2.

The university has implemented a cookie-based authentication mechanism, GLAuth, to simplify the use of GatorLink authentication for web-based applications. GLAuth uses Kerberos as its authentication service and provides an Apache authentication module, which can be used on Windows or Linux web servers to require GatorLink authentication for access to specific web resources. Local web administrators can easily implement GatorLink authentication on their sites without interaction with central IT.



Figure 2: Authentication architecture



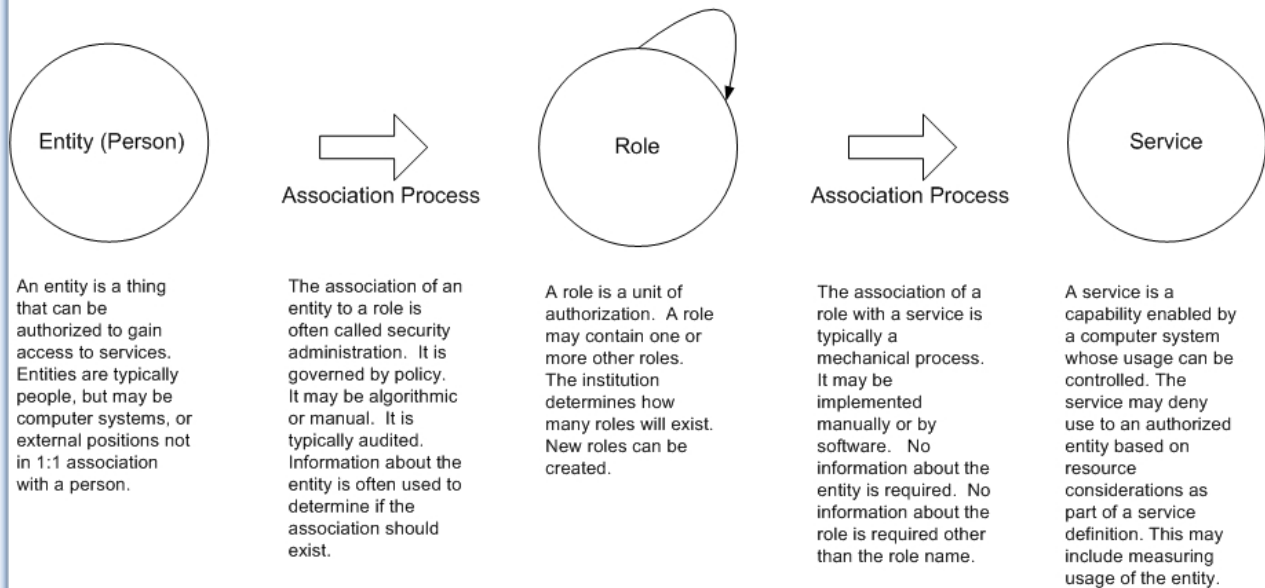
Authorization

As part of the implementation of PeopleSoft HR, Finance, Student, Portal and Data Warehouse systems, the University has committed to role-based application access.

Role-based access associates persons with services through roles that can be assigned by algorithm or through a manual authorization processes. See Figure 3.

Figure 3 Role-based authorization concept

Entities (typically people) are associated with roles. The roles provide the authorization to use services. Roles may contain other roles. The association of entities to roles is a policy-based process. The association of roles to services is a mechanical process.

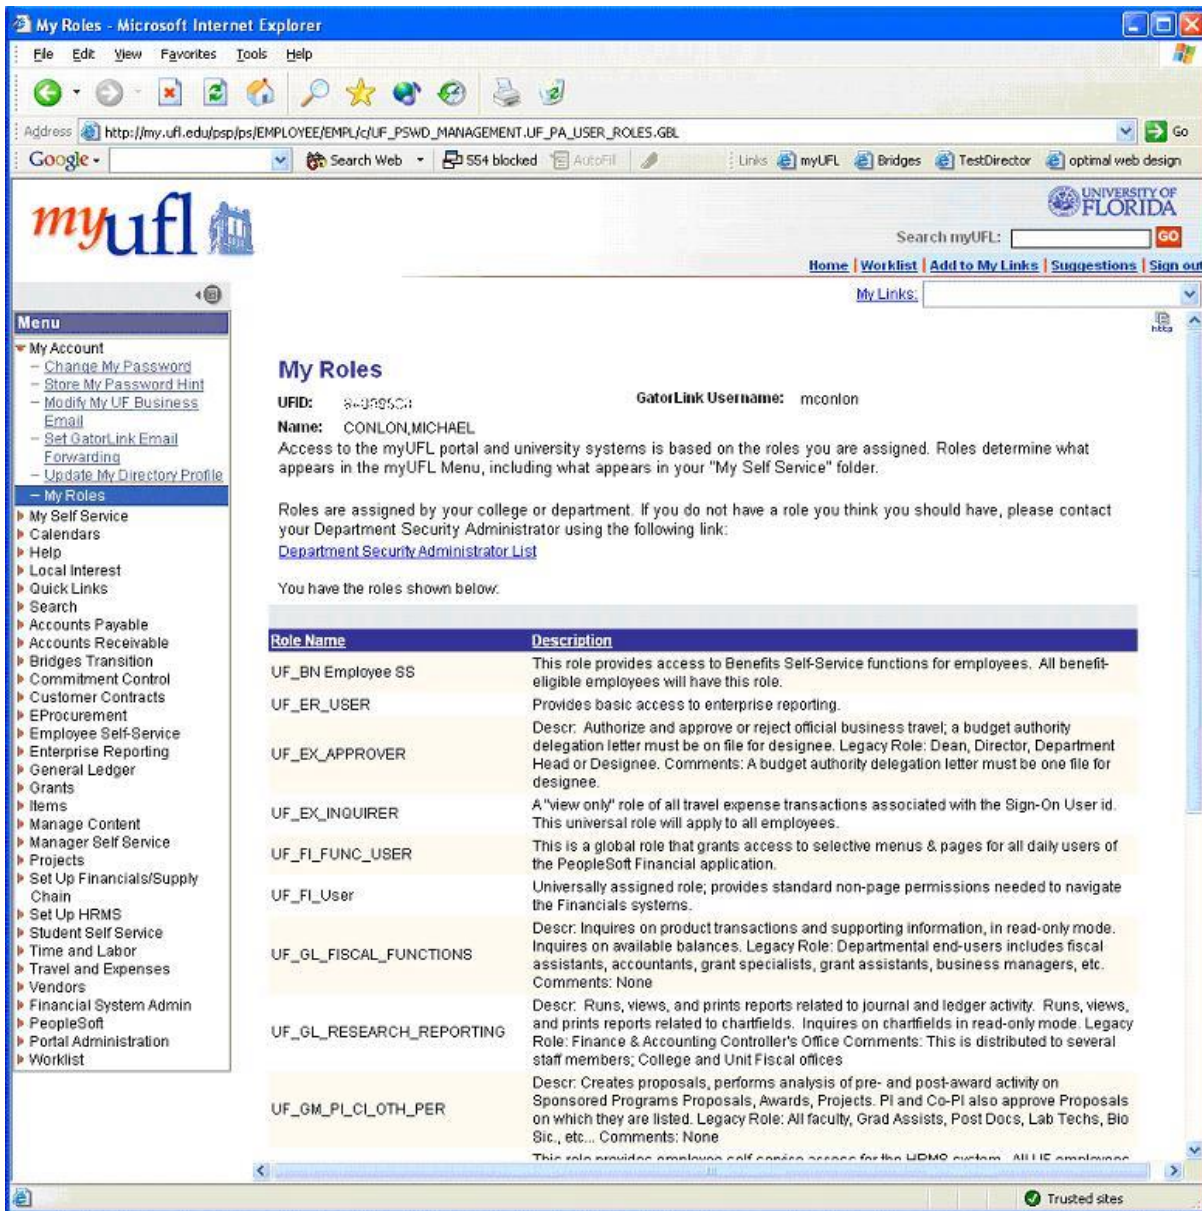


Over 250 roles are in place to govern access to PeopleSoft supported services in HR, Finance, the portal and the data warehouse. A University-developed application, the Access Request System (ARS), provides a user interface for department security administrators to request roles for people in their

departments. A second application, “My Roles,” has been developed as part of the university portal to provide self-service access to personal role information. See Figure 4.



Figure 4: My Roles in the myUFL portal provides self-service access to personal role information



Supporting University Business Processes

The University of Florida directory service supports a wide variety of business processes. The following five examples demonstrate the breadth of this use.

Distance, Continuing and Executive Education

UF has over 50,000 distance, continuing and executive education (DCE) students in a wide range of programs that are administered by the colleges and



departments. DCE students have a “DCE Student” affiliation in the directory, which can be assigned by local directory coordinators. DCE Student affiliation is used by GatorLink authentication services and portal services to insure that these students have GatorLinks and can access the portal.

UF Active Directory

Begun in 2003, the UF Active Directory (UFAD) provides desktop authentication services to departments using Microsoft operating systems. UFAD is fully integrated with directory services and all accounts in UFAD are populated from the UF Directory. All contact information is populated from the UF Directory using Microsoft Identity Integration Server (MIIS) and no local accounts are created.

Organizational Units (OUs) are populated based on a person’s “Network Managed By” relationship in the directory. The person’s account information in UFAD is located in the OU determined by the Network Managed By relationship and allows each person in the directory to have a specific department assigned to provide their network support. Directory Coordinators can update the user’s Network Managed By relationship.

UFAD authenticates using GatorLink, simplifying access for its users. GatorLink passwords are managed by processes that are determined by the user’s role. See below.

GatorLink Password Management

On May 5, 2004, the university implemented a GatorLink password management system. Previous passwords were either too weak (easily “crack-able”) or too strong (too many characters to remember so often written down). The new approach solves these issues by assigning a password policy to each individual based on their level of responsibility in authorized university processes. Each role has an associated password policy, from P1 to P5, developed by a university security group^v and designed to reflect various levels of access to university information and business processes:

- P1 – self-service/e-biz/vendors
- P2 – information only about selves, includes most students, some faculty, and self-service staff.
- P3 – can access and manipulate limited info about others – faculty, department administrators.
- P4 – can access and manipulate information at the university or college level – college administrative staff and directors.
- P5 – systems administrators for university systems that can access systems and data outside normal application security.

Approximately 100,000 UF users have either P1 or P2 password policy. Individuals at the P4 and P5 level must pass security checks. They attend a password certification course and are certified for the year following. The security group developed a matrix to simplify



the password policy rules, including length of password in characters, length of time before password expiration, and availability of password hints (e.g., none for P4 and P5 users).

PeopleSoft

Directory coordinators enter people into the university directory and create UFIDs for them. PeopleSoft Application Engine (AE) programs process message queues to automatically provision access to HR and Finance systems as appropriate based on the persons' affiliations. When a person is an employee, the HR system provides additional information to the directory and assigns employee affiliations. Non-employees often participate in university business processes and the directory can record appropriate affiliations that lead to provisioning access. The ARS can then be used to provide specific roles needed to handle special cases.

LDAP

The university LDAP service is populated from a message queue from the UF directory. LDAP provides access to public contact information and is used by the university white pages as a data source. LDAP is also used by university applications requiring current contact information for university members. UF LDAP supports the NMI-EDIT [eduPerson](#) schema standards.

Future Work

The UF Integrated Student information System (ISIS) is currently running on a

legacy system. The PeopleSoft Student Administration system will be brought online as a replacement for this in the summer of 2006. In conjunction with this, the UF Directory will be updated to use the PeopleSoft Campus Community and emerging PeopleSoft "person model."

GatorLink services such as email and web hosting are currently provisioned automatically using directory affiliation information. Legacy systems for GatorLink service authorization will be revised using roles that are automatically populated in PeopleSoft using directory message queues.

Additional college and department applications will be integrated with directory services, including building access control and VoIP (Voice over IP). Direct access to the directory database via API stored procedures will be replaced with a messaging architecture.

Reports on the testing and implementation of NMI components performed by the University of Florida were submitted to the NMI to provide additional architectural guidance and ownership perspectives on the development and deployment of middleware in complex organizations.

More Information

For more information about this University of Florida Case Study, contact Dr. Michael Conlon at mconlon@ufl.edu.



References

Information in this paper is based on interviews with the following individuals:

- Dr. Michael Conlon, Director of Data Infrastructure, NMI Project Lead, mconlon@ufl.edu
- Dr. Charles Frazier, Vice Provost for Information Technology, frazier@ufl.edu
- Mei-Li Cheng, Coordinator, Computer Applications, mlpssk@ufl.edu
- Warren Curry, Assoc Dir / ERP Integration Management, whcurry@ufl.edu
- Marian Boyle, Associate Director, Integrated Advanced Information Management Systems, mboyle@vpha.health.ufl.edu
- Steve Pritz, University Registrar, spritz@ufl.edu



Links of Interest

University of Florida <http://www.ufl.edu/>

GatorLink Password Management Policy <http://www.it.ufl.edu/policies/passwords.html>

GRIDS Center <http://www.grids-center.org/>

LDAP service at UF <ldap://dir.ufl.edu>

NMI-EDIT <http://www.nmi-edit.org/>

NMI Integration Testbed Program <http://www.nsf-middleware.org/testbed/>

NSF Middleware Initiative <http://www.nsf-middleware.org/>

UFID information <http://ufid.ufl.edu> and <http://www.it.ufl.edu/ufid>

White Pages at UF <http://phonebook.ufl.edu/>

ⁱ Hard copy of the UF Directory Planning Team's white paper is available at

<http://www.it.ufl.edu/projects/directory/planteamdocs/02-recommendations-20010816mc01.doc>

ⁱⁱ A university audit of UFID after 6 months of operation found that of over 500,000 UFIDs assigned, approximately 160 were duplicates. Additional effort was allocated to reconciliation of these duplicates.

ⁱⁱⁱ All university faculty, staff and students are required to have GatorLinks. Alumni currently are not assigned GatorLinks. The university currently has over 109,000 active GatorLinks.

^{iv} A GatorLink password must be at least eight characters long and contain a combination of uppercase letters, lowercase letters, numbers, and punctuation (three of those four elements). The password-checking program is linked to eight dictionaries, so that some users find difficulty choosing a password that is allowed by the system.

^v Information Technology Advisory Committee (ITAC) Subcommittee on Information Security Management (ISM). See <http://www.it.ufl.edu/committees/ism/>