

A Proposal for UF Directory Services

Final 8/16/01

Executive Summary

Directory services provide a foundation for a modern enterprise information infrastructure. Directory services are not 'databases' but rather a range of processes, applications, systems, and data that allow an organization to provide timely, accurate, and secure information about the organization and its people. Easily accessible directory information supports the work, business practices, and inter/intra organizational communications of our faculty, staff, students, and affiliates in support of our missions.

Benefits of effective directory services include:

- Efficient implementation of a new student/person ID number to replace the use of social security number.
- User friendly processes that enter updates to individual directory information once, allowing consistent access to real time information from across the university and its affiliates. Eliminates the need for maintaining duplicate lists of directory information locally.
- Consistent policy-based access to directory information
- Support for cross-institutional work and communications
- Directory information about not only individuals, but also about colleges, departments, units, organizations, and groups
- Access to directory information remotely and regardless of location
- Enables on-line work processes including work routing.
- More effective and efficient maintenance of data

Comprehensive and effective directory services include:

- Accurate and timely online directory information. Supports the creation of consistent phonebook information in all formats
- Support for enterprise systems and business processes by providing single point of contact for a range of data and information, eliminating the need for "shadow" data systems.
- Institutional commitment and the coordination of several organizations in providing data, services, and expertise.
- Consistent use of identifiers and logins, and well defined policies and improved procedures for the use of various identifiers.
- Delivery of services using locally developed software and databases in conjunction with interfaces to other enterprise systems, including the Registrar, the Enterprise Resource Planning (ERP), Shands HealthCare, library, and local

area network systems. The result is the ability to provide consistent services and data across a broad range of systems and business processes.

- Consistent directory architecture reflecting how data is provided and received from a variety of originating sources into a commonly shared directory system. The result is a seamless "one stop shopping" experience for the user.
- Sound policies governing the management and use of directory information to allow maximum flexibility in services to users, while ensuring needed security and consistency of data.
- Subschema developed to maintain differing kinds of information for diverse groups or roles, such as credentialing or sub-specialty information for clinical and extension faculty.
- Training in the use of directory services and support for units in converting their processes and systems to allow them to use enterprise directory services, and experience maximum benefits.

An Enterprise-wide Directory Services Planning Team developed this proposal following completion of the following: 1) Identification and definition of all key elements of Directory Services, 2) Review of current UF directory services and processes, 3) Review of existing models at other universities and the work of Internet2 national initiative involving the development of directory and middleware services in higher education, 4) A collaborative planning effort involving key operational units.

The Directory Planning Team Recommends that the university:

- Establish a Directory Services Steering Committee composed of representatives from units that are data providers and IS units responsible for managing and maintaining directory services.
- Direct the committee to facilitate and oversee directory implementation activities; recommend policies and procedures; support training and communications; and ensure ongoing revisions and updates as needed to directory structures, processes, and procedures
- Assign a directory project manager to coordinate the creation and implementation of the directory system and to be the interface between the Steering Committee and the development and implementation team
- Reassign 4 FTE for 14 months to Information Systems to form the core of the system development and implementation team
- Implement policies recommended by the Directory Steering Committee, for both service providers and users, to ensure and manage Directory Services that are flexible, effective, responsive, and secure
- Direct the Steering Committee to:
 - Develop needed data models and new schema
 - Identify desired data flows, develop or refine related applications and interfaces

- Oversee the creation of a web interface to allow self-service directory access and maintenance
- Oversee development of maintenance procedures and tools for using the Directory
- Develop training and communications/information to ensure effective use of Directory services
- Continuously assess user needs and desired data elements for directory services.

The University has made progress in building components of the needed directory services. The planning team estimates comprehensive services can be in place fourteen months following commencement of a dedicated effort of the type described in this report. Three to four FTEs will be needed on average throughout the fourteen months. Maintenance of directory services will require 2 FTEs of on-going support in addition to the on-going work of a Directory Services Steering Committee.

Background

The University has recognized the need for comprehensive directory services for some time. President Lombardi required that on-line phonebook services be created with appropriate maintenance procedures.¹ Information Services has constructed a university directory database in 1992 to coordinate access to directory information for business processes. NERDC deployed an LDAP² server in 1997 to facilitate access to email addresses. CIRCA developed on-line phonebook functions as part of the GatorLink³ initiatives to address Dr. Lombardi's requirements. In January 1999, the Cooperative Computing Initiative⁴ (CCI) identified Directory Services as a critical need facing the University IT community. UF IAIMS⁵ also identified directory services as one of the top three IT issues facing the Health Science Center in 1999. The current planning effort is a joint effort of the CCI and UF IAIMS.

Directories have been the focus of much discussion and development in the higher education community. Internet2 has developed a best practices document⁶, and Educause

¹ See the Appendix for Dr. Lombardi's original email regarding directory services.

² Lightweight Directory Access Protocol (LDAP) is an Internet standard for providing directory information to software such as email systems.

³ GatorLink authentication services provide username and password to members of the University community to access a variety of electronic resources.

⁴ The Cooperative Computing Initiative was created by Dr. Betty Capaldi and chaired by Mr. Earl Robbins in his role as University CIO to foster cross-unit application development.

⁵ Integrated Advanced Information Management Systems (IAIMS) is a grant program of the National Library of Medicine. Dr. Ken Berns is the principal investigator for UF IAIMS, NIH/NLM 1 G08 LM06821

⁶ Internet2 Middleware Initiative, "Identifiers, Authentication and Directories: Best Practices for Higher Education," May 9, 2000. <http://www.internet2/middleware/best-practices.html>

has developed a model for directory entries⁷. Dr. Ken Klingenstein of the University of Colorado and Internet2 visited UF on April 27, 2001, presenting directory concepts to the university community and discussing these concepts with the planning team.

Definitions

Directories are the operational linchpins of almost all electronic services. In future information technology environments, directories will be among the most critical services offered.

A "directory" provides a way to quickly and easily look up information for a given identifier. A common use is to find phone numbers and addresses for an individual. However, a directory can also contain critical customized information for people, processes, resources, departments, organizations, college courses, or any other entity. By placing such information in a common storage area, diverse applications from diverse locations can access a consistent and comprehensive source for current values of key data.

A directory is *not* the databases that are used to support the business of the institution. It does not replace the records maintained by the registrar, personnel, or other departments. Rather, it provides a unified view of selected subsets of these records or other information maintained by departments at the institution.

Directories are databases that are optimized for reads, and contain key institutional and personal data for use by a wide variety of applications. Directories need ways to describe the sequence of fields in the database (a schema), the names of the fields (a namespace) and the contents of the fields (attribute values). Directories also need identifiers for the entries to serve in indices to the database.

A key goal of the directory is to make the information easily available to all users and units of the institution rather than via fixed screens that provide the information in inflexible ways. Today, directories are likely to be implemented using an open protocol like LDAP rather than the relational database systems used in maintaining the institution's various databases. This provides the potential for using the information in ways that meet the unique needs of an individual department.

A directory can also be the source for information regarding university activities that may not be related to business practices. For example, URL's (website address) for an instructor's courses might be recorded in the directory and no where else.

"LDAP" is the term most often used when talking about directories. LDAP is not a directory in itself, but it is a tool that can be used to build directories. LDAP is a client-server protocol for accessing a directory service. There are several commercial and non-commercial implementations of LDAP, and many applications, including the most commonly used mail clients, can use LDAP to retrieve directory information. LDAP was originally developed at the University of Michigan and is a simplification of the international X.500 directory standard.

⁷ Educase, eduPerson object class. <http://www.educause.edu/eduperson>

Examples of "fields" in a directory may include institutional status, email aliases, personal photos, permissions, private keys and calendars.

"Identifiers" are the keys that make directories usable. A directory needs some unambiguous identifier that can be used to retrieve specific information. Social Security Numbers, student id numbers, computer accounts, email addresses, Internet addresses, can all be identifiers that uniquely identify some person or thing. It is crucial to the implementation of a directory that the relevant identifiers and their relationships to each other are identified and incorporated into the directory structure.

"Middleware" is a new term used to designate the protocols, tools, and software that can make directories easier to use, and reduce the complexity of related application development. Directories should not just provide an alternative way to access the raw data stored in the institution's databases. Middleware should be available to make the information more accessible, and to make more sense out of complex information. When it is difficult to answer simple questions using the institutional data, it is probably because the middleware is missing.

"Authentication" and "authorization" are necessary characteristics of any directory as well as a service provided by a directory.

"Authentication" is the process of determining who you are. Fingerprints, identification cards, DNA, or the testimonials of others can all establish your identity, but in the computer world, it is usually established by the correct entry of a computer username and password.

"Authorization" is the process of determining what you can do, what you are permitted to access. Can you transfer money from an account? Can you submit grades for a class?

Clearly, both authentication and authorization are essential components of a directory. They are the tools that are used to implement the policies that dictate who can view and modify the information in a directory. If you "sign in" as required by the directory, it might be possible to update/correct your individual data. In addition, an authorized person may also be able to sign in and update your data for you.

A directory can also provide authentication and authorization as a service. A directory can provide a common and standard way to answer both the "who am I" question as well as "what can I do". At UF, the Kerberos⁸ server that supports GatorLink usernames and passwords provides an authentication directory. However, just answering the "who am I" question is rarely sufficient. Many services that could use GatorLink for authentication also need to authorize a person, sometimes with the answer to a question as simple as "is this person a student". Many authentication questions can be answered using "role" information – what roles does a person have in the university community? Common roles are student, faculty member and staff member, but there are many others. Many people have more than one role.

In addition to campus-wide directory services, Local Area Network (LAN) systems may also have associated directory services. Novell Netware Directory Services (NDS), Microsoft Active Directory (AD) and Sun Microsystems Network Information Services

⁸ Kerberos is a standard method for providing authentication services developed at MIT.

(NIS) are all used at UF and each has a purpose in supplying information to users of these systems. Proposed relationships between these systems and university directory services are presented in this document.

Policy issues

Challenges to building enterprise-wide directory service in higher education lie in conducting consensus processes for policy and funding. Given our ad hoc data administration environments, we need to clearly define issues of data ownership and access. A directory project brings all these issues to the fore. Similarly, there are many issues that need to be resolved regarding directory data including privacy, Open Records, Family Educational Rights and Privacy Act (FERPA), etc.

Directory services and systems must operate in service to other systems such as library, Enterprise Resource Planning (ERP), departmental, student information, university business, health care, extension and others. In each case, the directory provides a consistent means of managing and providing information regarding people and organizations to each of the other systems.

The funding of a directory initiative is also a challenge. The cost of a central infrastructure includes multiple servers, development of interfaces to legacy systems, and management of the schema. Someone must pay to retrofit key applications to use directories. Departments must be trained in new tools and redeploy their environments to reflect the enterprise service.

Advocacy of a directory initiative needs to draw senior leadership as befits data infrastructure issues. At the same time, there are some external drivers that reflect the urgency and importance of directory services. Federal initiatives in digital signatures for students, scientific communities' need to exchange management data, and licensing of scholarly materials are some of the national issues that directory services can help to address.

Charge to team and work of team

Members of the Cooperative Computing Initiative (CCI) and the UF IAIMS Steering Committee identified the need for comprehensive directory systems and services for the university and healthcare enterprise. Members of these two groups developed a charge to a collaborative planning team. The planning team includes representatives from related operational and planning areas. The team began work in June 2000. A website (<http://www.cci.ufl.edu/dir.htm>) and listserv (CCI-DIRECTORY-L@lists.ufl.edu) were established to support the effort. The charge to the planning team was:

Recommend processes and systems resulting in enterprise-level support of directory services, including a common, centralized address book for all UF and Shands users and affiliates.

Desired outcomes include:

- Design, develop and implement comprehensive enterprise-level directory processes and systems.

- Provide accurate data about people and organizations, accessible by secure processes and by all people and applications.
- The process represents the consensus of the units with operational responsibility for constructing and maintaining the directory.
- The process is documented.
- Executing the process will result in a directory having the desired attributes below.

Desired attributes for the directory include the following:

- Individuals can update their directory entries.
- Allows individuals to search by each variable; i.e. name, position, dept., etc.
- Individuals can easily search for information which they are entitled to access.
- The directory can cover the entire enterprise – UF faculty, staff, students and alumni, Shands HealthCare (all locations) physicians, staff, patients.
- The directory supports email addresses and other information needed for business processes.
- Appropriate security is maintained.
- Responsibility for technical support and future development of the directory is clearly delineated.
- The directory positions the enterprise for new opportunities.

Benefits of Directory Services

The following scenarios demonstrate the potential benefits of directory services and directory-enabled business processes. Directory-based services are the central components of a streamlined experience for students, faculty and staff.

A New UF Student ID Number

Maria is a new student at UF. As part of Preview, she is assigned a UF student ID number. Her student ID number is stored for her in the directory. Maria uses her student ID number for all her interactions with the University. She won't need to provide other IDs. Her information is protected and her other identifiers such as Social Security Number and GatorLink ID are provided only as needed by policy and business process.

To insure her privacy in public communications, Maria uses only her student ID number. An enterprise ID number is assigned to her records for use in securing existing UF data systems. This number is never public and is used internally for record linkage. Should Maria need a new student ID number, one can be provided and linked to the underlying enterprise ID without having to make multiple changes across UF systems. Should UF adopt an Enterprise Resource Planning (ERP) system, identifiers from that system would be stored in the directory as well.

Update Data Once

Bill is a physician faculty member in the College of Medicine. He and/or his department administrator can update his contact information using a web page. This information automatically populates/updates the personnel system, the Shands Communications system (CHRIS), the Shands HealthCare on-line directory of physicians, Bill's entry in Netware Directory Services, Active Directory, the on-line phone book, the UF enterprise directory database, and the UF LDAP directory. People using email programs and their address books always automatically access Bill's current email address. UF business processes have access to Bill's current information. Bill's information is updated once and is used and accessed consistently across the enterprise.

Obtain Appropriate Services Regardless of Location

Jake is an IFAS researcher in Lake Alfred. Jake can access UF Library materials because he is a UF faculty member. Jake can present his GatorLink username and password to a Library web site. The web software accesses the directory to confirm that Jake is a faculty member. Jake gets access to the Library's electronic resources remotely based on his role with the University. When Jake leaves the research facility and visits extension offices, his Internet access may be provided in many ways. No matter what method Jake uses to access the Internet and regardless of his location, he is granted access to Library services based on his affiliation with the University as identified by the directory.

Participate in Cross-Institutional Work

Sapna is a graduate student in Physics. She uses computational facilities located at laboratories at Los Alamos, Lawrence Livermore, Argonne and CERN. The UF directory stores her Public Key Infrastructure (PKI) identifiers and provides credential information to remote systems that grant her access. Sapna logs on to those systems using her GatorLink username and password as managed by UF. She does not need multiple usernames and passwords.

Her directory information is automatically made available to the Higher Education Directory of Directories, a project of Internet2⁹. Her advisor's information is automatically updated in the Community of Science database¹⁰.

Manage Email Lists

Frank and Kathy manage email lists for their respective student organizations. They indicate in the directory that they have a group of students they would like to email. They provide a list of directory entries/names for the people on their lists. When they send email, the directory produces a list of the current email addresses for these people and that list is used to send the email. Such dynamic lists are always up to date. Frank and Kathy do not need to make changes to their lists as email addresses change – which

⁹ Olsen, F. "Internet2 Plans an Electronic Directory of Millions of People in Higher Education,": Chronicle of Higher Education, June 20, 2001. <http://chronicle.com/free/2001/06/2001062001t.htm>

¹⁰ Community of Science maintains a national directory of faculty research interests. See <http://www.cos.org>

is handled automatically by directory-based services. Frank and Kathy only need to manage “who” is on their lists – email addresses come from the directory automatically.

Contact UF Departments and Organizations

Kay is considering a position at the University in the School of Accounting. She goes to the UF web site and uses the directory to locate the email address and phone number of the UF Fisher School of Accounting. As part of this process she discovers that the School has a web site and she accesses the site to learn more about the school.

Simplify Department Processes

Jared maintains a local database of names and addresses for his UF department. After meeting with UF directory services staff and learning about the capabilities of the new directory system, Jared discovers he can produce the information his department needs using the UF directory without having a locally maintained copy/database. Jared needs to collect a couple of pieces of information that are not currently in the directory. After working with the staff, the directory is amended to provide Jake an opportunity to collect and store his data directly. Jared abandons his local database and the effort required to keep it up to date, and begins to use the UF directory directly.

Current UF Directory Status

UF has several of the major components of a modern directory service. Each is described below. Missing are interfaces, application programming interfaces¹¹ (APIs), policies, procedures and connecting software. UF has made a good start on individual pieces of directory systems and is in a strong position to create the needed pieces to have modern directory services.

Enterprise Directory Database

This current directory database was begun in 1992 by UF’s Information Systems department and has gradually become more and more robust and accepted as a university resource. Known as WCGOLD, the current database and update mechanisms are close to providing the single source of authoritative data needed for the proposed directory services infrastructure at UF. These initial directory processes were used to automate and simplify related applications, the Gator 1 ID card, and was relied heavily upon during the implementation of PIN security, RACF¹² account management and more recently GatorLink account management. This database currently accumulates directory information into a DB2¹³ database for the entire university community including: UF

¹¹ APIs are software interfaces that allow developers to write software to interface their systems to other systems. Directory APIs are critical components which enable use of directory information across systems.

¹² RACF (Resource Access Control Facility) is the authentication and authorization system used on IBM mainframes running the MVS operating system. MVS and RACF are used by NERDC at UF.

¹³ DB2 is IBM’s relational database product. DB2 is used for UF mainframe applications storing their data in relational tables.

personnel, UF students, Shands personnel, University Athletic Association, University Florida Foundation, retired and emeritus faculty, courtesy faculty and other UF affiliated individuals as entered by UF's department phonebook coordinators. In addition, listings for UF's Authority code groups and departments are included in this database. Upgrades are needed to the database structures to comply with evolving directory standards developed in the higher education and Internet 2 communities. With the addition of enhancements to the update mechanisms and development of policies relating to the management and use directory information, this database would provide UF with a source repository from which to build the needed middleware components. In turn, UF would reap improvements based on these fundamental enabling technologies.

UF LDAP

This service grew out of the early work involving UF PH¹⁴ Directory. NERDC installed an LDAP service in 1997. Efforts continue to enhance and move toward the standards based LDAP directory protocol. The UF LDAP implementation gathers data from multiple sources to populate itself on a nightly basis. The major sources of these data are the UF registrar system and WCGOLD. LDAP works closely with the Kerberos server that implements GatorLink accounts and is the primary repository of email address information for the entire UF community. LDAP has been integrated with several UF applications. As with the enterprise directory database, these initial efforts have brought us quite far. By working to standardize the data schemas housed inside UF's LDAP including implementing the evolving eduPerson standard and extending it to a UF person standard, UF's LDAP can be developed to provide high quality service. The effort to improve directory services will require technical expertise in several related areas including core directory technologies (such as LDAP, X509.3, HTTP), distributed directory technologies (such as Novell's Netware Directory Services and Microsoft's Active Directory).

UF Kerberos

NERDC operates the UF Kerberos server for the purpose of authenticating users via GatorLink usernames and passwords. The UF Kerberos server has been operational since 1997 and contains over 100,000 GatorLink usernames.

Security

UF currently uses a variety of means to assign identifiers to individuals and organizations and authenticate and authorize access to systems and data. Security mechanisms and accounts are directly related to and can be created from the identity that is given to an individual within the directory. The UF GatorLink account, PIN, Gator 1 Card, and NERDC accounts all depend on directory information in their establishment and management. These services are implemented using technologies such as Kerberos,

¹⁴ PH (Phonebook) developed at the University of Illinois, was an early system for making contact information accessible over the Internet.

RACF, DES¹⁵, etc. Particular attention needs to be paid to the interactions of the directory and security system. The security system is essential for controlling access to the directory. In turn the directory will be storing authoritative identity information that is critical for establishing and monitoring security for UF's resources and services. We are blending areas of computer security with identity management in ways not experienced before at UF. This will require us to address the confusion and concerns related to boundaries and responsibilities of data managers/owners and the security managers.

In addition, copyrighted materials and other intellectual property needs protection using means beyond those currently available. Securing access to material by use of IP (Internet Protocol) numbers rather than personal identifiers is a common technique, and replete with problems. IP access control fails to satisfy the needs of copyright holders and often blocks access to legitimate members of the community.

NDS, Active Directory and independent directories

Novell's Netware Directory Services (NDS) and Microsoft's Active Directory (AD) systems represent technologies that may be able to be supported from the primary LDAP and enterprise database. These services can be developed to provide college/department level support as needed. Tools to populate and manage these and other technologies of this type and scope should be provided to support the more focused and local needs of UF's community. However, these distributed directories need to be fed information and updated on a routine basis from the master enterprise wide UF directory system. Policies need to regulate the updates and maintenance of the directory data received at the enterprise level and then distributed throughout the various directories as needed. This change in approach will be a challenge for UF culture.

Identifiers

UF uses many different identifiers for individuals. Different processes, vendors, agencies and partners require these different identifiers. The appendix contains a chart compiling the list of commonly used identification mechanisms used at UF. Directories can shield this complexity from users by translating between identifiers as needed to conduct work. See Attachment -UF Identifier Mapping.

Numerous independent business processes for maintaining information

Current mechanisms for updating directories are redundant, confusing and in some cases conflict with one another. This situation exists for many reasons. The primary reason is a lack of clearly defined policies and responsibilities for directory middleware at UF. Current practices have resulted in a "catch as catch can" integration between various providers of information.

Some of these conflicts are due to the timing and content of updates. With the creation of effective policies to support a shared enterprise directory, and the design of a single

¹⁵ DES is the Data Encryption Standard, a strong security encryption algorithm originally created by the National Bureau of Standards, now known as the National Institute of Science and Technology.

system of management tools to support staff charged with maintaining the directory, a true single source database can be produced that can provide the warehouse of data for the UF enterprise. Tools to support the use of the warehouse are equally as important if independent directories are to use the enterprise directory as a source for their local and more unique needs. A single reliable and robust set of tools must be built to enhance ease of use and support policies for managing and operating directory services at UF.

Proposed UF Directory Services

The implementation plan, which follows, is designed to create the following services:

- An enterprise directory strategy. The proposed strategy developed by the planning team will allow for the delivery of enterprise directory services using locally developed software and databases in conjunction with interfaces to other enterprise systems. These locally developed systems will interface to the Registrar's systems, the Enterprise Resource Planning System, Shands HealthCare Systems, library systems, and local area network systems including NDS and AD. By constructing and maintaining our directory solution, we can provide services across a broad range of systems and business processes.
- Identity resolution. To provide student IDs and services based on the directory, we must ensure that each person is represented in the directory just once. As new people present for service, we must create the needed directory entries and have processes in place that prevent redundant or multiple entries.
- New UF Student/Person ID number. The use of Social Security Number (SSN) for identification purposes has been deemed unacceptable at UF. A provost's task force chaired by Steve Pritz recommended the adoption of a new identifier to be issued to all members of the UF community. The new identifier can be mapped to existing identifiers for University purposes using directory services. This yields several benefits – effort to implement the new identifier will be reduced, an underlying system-level identifier can be used and mapped to the public identifier to improve security, the number of identifiers needed for public transactions can be reduced, assigned identifiers can be changed without significant impact on UF systems.
- Institutional commitment. The benefits, scope, and complexity relating to the development and operation of effective enterprise directory services will require on-going institutional commitment as described in the Implementation Plan. A commitment to commonly agreed upon policies, of resources, and to maximizing the benefits to the community is needed.
- Coordinated effort. To produce effective directory services, several organizations will need to work together to provide data, services and expertise. We have experienced very good coordination during the planning process and anticipate continued cooperation during the implementation phase.
- Consistent directory architecture. The proposed approach is based on an enterprise directory architecture pictured below. The architecture describes how various data systems will receive and provide (data feeds) directory information

so that directory services will appear as a seamless “one stop shopping” experience for the university community.

- Consistent directory policies. Access to directory information is governed by policy. The planning team proposes that the Data Infrastructure Advisory Group described in the IT Review Report be responsible for directory policy development and recommendations¹⁶.
- Consistent use of identifiers. UF uses many identifiers – student ID, SSN, PIN, GatorLink, and various other logins. A UF Identifier Map (See Attachments) has been developed by the planning team to assist UF in defining policy and improving procedures for using various identifiers.
- Support for organizations. Directory services will provide entries in the UF directory for use by units, organizations and other groups. This data can eventually be used to produce the “blue pages” of the UF phone book, as well as provide on-line directory information. In addition, email to members of groups can be facilitated by directory services. Email sent to the group can be sent to each of the groups' members using directory look-ups to ensure accurate addresses for each group member.
- Support for eduPerson. EduCause has developed a standard for the representation of directory information for persons in higher education. The eduPerson object class¹⁷ is a standard for LDAP directories to enable cross-institutional searches.
- Support for subschema. Different individuals in the directory require different information to be held depending on the attributes and role of the individuals. For example, directory information for clinical faculty should provide information regarding sub-specialties and board certifications. This information is held in subschema that can be developed as needed to meet the needs of various and diverse groups of people.
- Support for enterprise systems including ERP. Most enterprise systems need access to contact information regarding people and organizations. Directory services provide a single source for such data. Enterprise systems can then make use of consistent, accurate and timely data.
- Support for LAN systems. Directory services will supply updates to information stored in NDS and Active Directory.
- Improved maintenance procedures. By standardizing the architecture, software and processes used to create, maintain and provide directory information, we will improve the maintenance and quality of such data.
- Training. To achieve the benefits of effective directory services, the institution must commit to providing training to users and producers of directory information. Individuals will need to know how to use the directory and provide updates. Business processes will change to better use directory-based services.

¹⁶ See Report of the IT Review Committee, <http://www.aa.ufl.edu/itr>

¹⁷ See eduPerson object class <http://www.educause.edu/eduperson>

Staff members involved in those processes will need to be retrained regarding simplified processes resulting from directory services.

- Participation of business and academic units. Outreach efforts will be needed to involve business units and academic units in the use of new directory processes and services. Dedicated effort should be made available to assist units converting their processes and systems to enable them to use enterprise directory services.
- Oversight. A Directory Services Steering Committee reporting to the Data Infrastructure and Administration subcommittee of ITAC is recommended to provide on-going oversight, policy formulation, needs assessment and review.

Proposed Approach

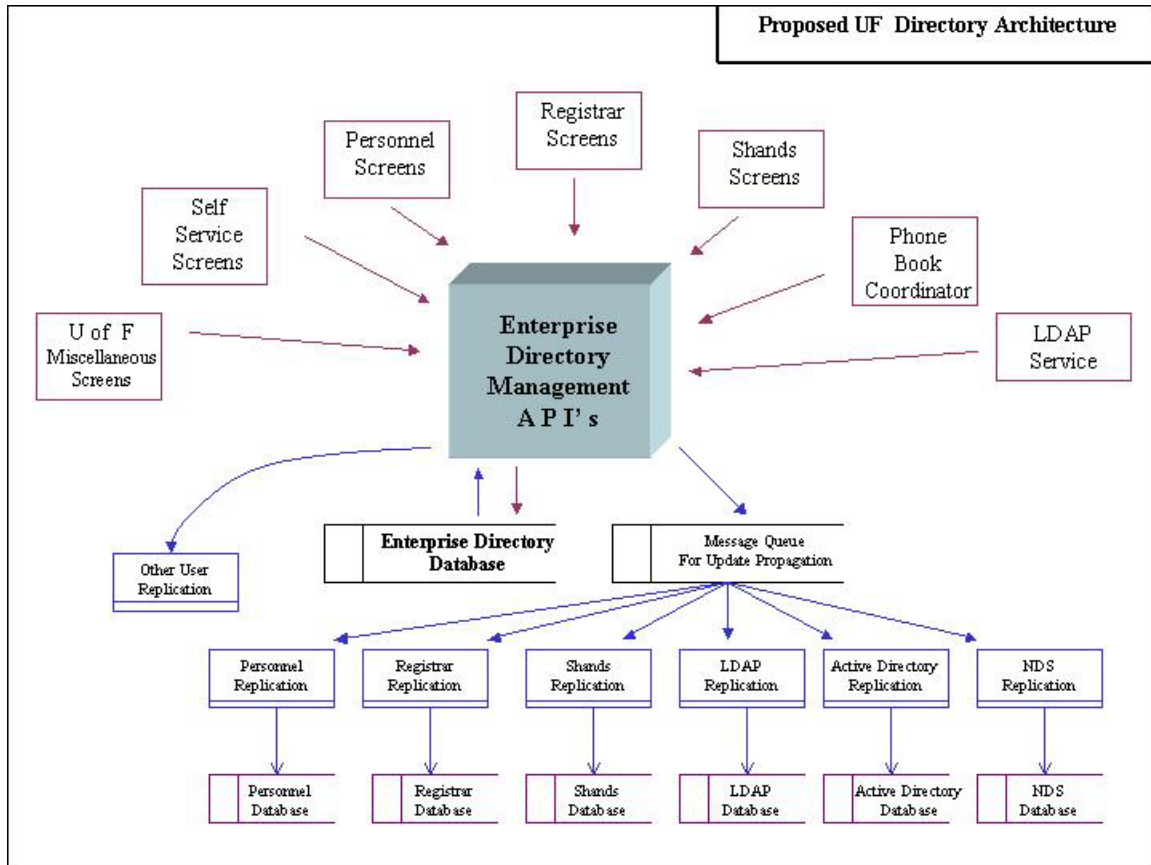
The Directory Planning team considered several approaches to implementation of directory services at UF:

1. Rely on LAN-based systems such as NDS or Active Directory. This approach does not scale to the multiple systems and complex environment of UF. These systems need to be integrated into the set of directory services, but cannot be the primary source of authoritative directory information.
2. Rely on Enterprise Resource Planning (ERP) directory services. UF may choose to implement an ERP solution. This would take several years. UF's IT services span activities far beyond the business domain of the ERP. Library, health care, extension, research and other processes all must benefit from directory services. Directory services should provide information to the ERP as well as these other systems.
3. Use consulting services to build a new directory solution. Many firms are capable of offering consulting services to determine requirements and build custom solutions for directory services. Work could be done in a timely manner assuming adequate resources were provided.
4. Build UF Directory Services from existing components. UF has the components and is developing the policy infrastructure to implement directory services. Resources will be needed as well as cross-unit coordination.

The Directory Planning team recommends that UF build its directory services. This approach will enable UF to take advantage of its existing components, own its solution and manage its policies independent of any particular vendor system. ERP and LAN systems can be integrated with UF Directory Services as needed.

Proposed Architecture

The figure below shows the proposed architecture for UF directory services.



Proposed UF Directory Architecture

The core of the system is a collection of Enterprise Directory Management Application Programming Interfaces (APIs). These standardized software components ensure that directory functions are performed in a consistent manner. All systems providing data to the UF directory will use these APIs.

Many systems supply data to the enterprise directory. Miscellaneous UF screens, new Self Service Screens, Personnel and Registrar Screens, Shands, Phone Book Coordinator and LDAP services all provide data. Some of these systems are interactive while others provide bulk data (batch processing). Many of these systems currently exist, while others are new. The Self Service Screens are a critical new component. These screens will permit individuals to update fields that they control. Giving our faculty, staff and students the ability to personally update their directory information will improve the quality of the information in the directory both directly (by their input) and indirectly (increased demand for improved institutional data). Additional systems may supply data to the enterprise directory in the future.

The APIs manage the consistent flow of information from the originating systems that supply data, to and from the enterprise directory database, and the enterprise directory message queues. In the language of the Internet2 middleware initiative, the APIs play the role of a “meta-directory” by ensuring that all directory operations are performed in a consistent manner and all systems relying on enterprise directory services have consistent data.

The enterprise directory database is a collection of database tables stored on the mainframe using DB2. These tables currently exist and will be upgraded to support the additional data and services described in this proposal. All access to and from the enterprise directory database is governed by the enterprise APIs. This ensures data and process integrity across all the services providing and using enterprise directory information. The correspondence of one identifier to another is supported by the database and controlled by the APIs. The database stores the multiple identifiers. The APIs regulate access to the mapping to ensure appropriate use of the data. This improves security and improves our ability to manage the multiple identifiers.

The message queues are holders for directory updates that can be timed to provide improved information to systems relying on enterprise directory information. The queues allow for different systems to be updated with varying frequencies, which is important for the different business practices these systems support. For example, Shands replication may need to occur only once per week in keeping with their business needs and the operation of their enterprise systems. LDAP replication may need to occur in near real-time as self-service users who have made an update to their information in the enterprise directory will expect to see that update in place immediately. Additional systems can be fed using message queues. In particular, Netware Directory Services (NDS) and Microsoft Active Directory (AD) are local area network services that would benefit from integration into the enterprise directory services architecture.

Implementation Plan

Required Tasks

The planning team identified the following list of tasks needed to provide effective enterprise directory services:

1. Assess user needs in relation to Directory Services.
2. Identify desired data elements needed in new directory & LDAP.
3. Develop needed data models.
4. Develop new schema for LDAP.
5. Identify current inventory of API's and interfaces to WCGOLD & LDAP.
6. Identify desired data flows and related APIs.
7. Develop or refine related applications and interfaces.
8. Define and create web interface to allow self-service.
9. Determine and implement needed policies, for both service provider and users, to ensure and manage Directory Services that are flexible, effective, responsive, and secure.
10. Develop procedures for implementation of new services.
11. Develop maintenance procedures.
12. Develop training.
13. Develop communications/information.
14. Develop tools to use Directory.

Gantt Chart

See attached Gantt chart for best estimate of the time and resources required to create UF Directory Services as described in this proposal.

Planning Assumptions

The following planning assumptions were made in the construction of the Gantt chart:

1. Project management expertise is available within UF for this project.
2. Effort from related business units will be provided as indicated.
3. Additional support for the process is available.
4. On-going support of the directory is available.

Staffing - Implementation and Ongoing:

Personnel costs in the Gantt chart are based on an average \$50,000 annual salary plus fringe, and 5 hours effort per workday. Estimate 3-4 FTE's on average over 14-month development period, and 5-6 FTE's during peak effort. Estimate need for Software Development Manager and equivalent of 2.0 FTE for ongoing support.

Functions/skills needed:

- 1.0 FTE Software Development Manager
- 0.5 FTE Communications, Writing (Documentation-Procedures, etc)
- 0.25 FTE Database Administration
- 2.0 FTE Application Developers
- 1.0 FTE Systems Programmers (LDAP & Code)
- 0.25 FTE Trainer

Project Manager

Given the range and type of work needed, including organizational, communications, and training, and the number of units involved in managing various types of directory information, a senior project manager with experience in developing enterprise IT services will be needed. The project manager should also be a member of the Steering Committee (see below) and should directly supervise the software development manager (see below).

Software Development Manager

Recommend dedicating a senior IS staff member full time to the implementation work. Recommend this staff member be designated as the 'Software Development Manager' to coordinate the work of programmers, developers, and database administrators.

Steering Committee

We recommend that the Vice Provost establish an Enterprise Directory Services Steering Committee including representatives from units that are data providers and IS units responsible for managing and maintaining directory services. The committee facilitates implementation activities; recommends policies and procedures; supports training and communications; and ensures ongoing revisions and updates as needed to directory structures, processes, and procedures. Suggested representative units include Personnel, Administrative Affairs, NERDC, Tigert IS, Academic Affairs, Shands IS, CIRCA, Registrar's Office, and Phone Book Coordinators. The Steering Committee will be responsible for assisting and supporting the project manager to ensure completion of work. The committee will meet monthly and report bi-monthly to the Vice Provost for Information Technology. As the University IT advisory structures are implemented and directory services are in place, it may be appropriate to transition the activities of the Enterprise Directory Services Steering Committee to the ITAC Data Infrastructure and Administration subcommittee.

Implementation leads to transition and operation. On-going operational responsibility will need to be determined by the Vice Provost for Information Technology.

Additional Resources

In addition to staff and oversight, implementation will require on-going hosting expenses for hardware, training materials, and computers, furniture, space for additional staff.

Parallel effort will be required by many offices to use University Directory Services.

Review and Approval Process

This recommendation will be shared with members of the CCI and UF IAIMS Steering Committee to obtain their input and suggestions for the future. The recommendation will be forwarded to the Vice Provost for Information Technology for consideration, revision, and approval and decisions regarding implementation activities.

Attachments

- Directory Planning Team Members
- Gantt Chart
- eduPerson schema
- UF Identifier Mapping
- Other Identifier Mappings – Johns Hopkins University
- Internet 2 Model Identifier Mapping
- Dr. Lombardi's "10 Commandments"
- Student ID Policy recommendation
- Memo from UF IAIMS Library Resources Planning Team
- Memo from UF IAIMS Email/Calendar Planning Team