

University of Colorado at Boulder: Identity Management Governance

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.

The Case Studies are drawn from schools/consortiums with varying sizes, populations, and missions in an effort to provide examples of the diverse technology, policy, and project management approaches.

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 the authors Jon Giltner, Melinda Jones and Paula Vaughan. Ann West, editor.

Copyright © 2004, University of Colorado at Boulder. All rights reserved. Content may be used for non commercial purposes with attribution



Executive Summary

The University of Colorado at Boulder implemented an Enterprise Directory in the fall of 2001. However, the project began almost a year before the implementation with campus-wide information gathering, research and consultation across the higher education community and Internet2 landscape, and with significant efforts to bring together key players within the University's diverse culture of systems, data, functional experts, processes and policies.

Having built up some modicum of middleware expertise in-house and having achieved some unity among the campus players, the directory services infrastructure was designed and implemented. Its initial role was to provide a secure and authoritative source of data related to all

individuals with a CU-Boulder affiliation and to offer the campus a framework upon which additional middleware-dependent initiatives could be built. These goals were accomplished – and were successful to a level for which we were somewhat unprepared. The infrastructure itself is up to the task of supporting increasing demands for its many components. A bigger challenge has been adapting policies, processes, and technical staffing to keep up with our successes.

For more information about this University of Colorado at Boulder case study, contact Melinda Jones at Melinda.Jones@colorado.edu

NMI-EDIT Components Highlighted in this Case Study

A Recipe for Configuring and Operating LDAP Directories <http://middleware.internet2.edu/dir/>
This document outlines practices of common directory deployments within the Higher Education community. Also known as the "LDAP Recipe."

eduPerson Directory Schema <http://www.educause.edu/eduperson>
eduPerson contains identity-related attributes for higher-education institutions to deploy for enabling 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.

LDAP Analyzer <http://middleware.internet2.edu/dir/>
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 practice documents.



Metadirectory Practices for Enterprise Directories in Higher Education

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

This document outlines a set of metadirectory issues that are commonly considered in the deployment of enterprise directories and offers accompanying practices for higher education.

Practices in Directory Groups

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

This document offers recommendations to institutions embarking on the implementation of the use of groups to facilitate basic authorization.

Pubcookie

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

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

Shibboleth Software

<http://shibboleth.internet2.edu>

The Shibboleth support inter-institutional sharing of resources that are subject to access controls.



University of Colorado at Boulder: Identity Management Governance

The University of Colorado at Boulder is one campus of a multi-campus public university system. CU-Boulder serves 29,151 students (84% undergraduate, 16% graduate) and includes 2,174 faculty and 2,901 staff members as well as various affiliated individuals and agencies. The campus is dependent upon systems of record for student, employee and financial information that are managed for all of the campuses by a central system office. For many departmental-specific computing applications, CU-Boulder relies upon its campus-specific Information Technology Services (ITS) for systems development and support. ITS is also responsible for the IT infrastructure and security specific to the Boulder campus. As the population dependent upon computing continued to expand and diversify and as computing demands became more sophisticated, CU-Boulder recognized that a robust IT middleware infrastructure would be critical to the ability to deliver required services to campus.

Business Challenge

CU-Boulder witnessed an increasing number of applications springing up across the campus, each with an application-specific directory, and each with its own set of security concerns due to the proliferation

of data across multiple systems with varying levels of protection. Much of the application data was built through independent business rules and interpretations; little (if any) of the data was reconciled. Although many of these systems were departmentally-owned, ITS was called upon with increasing frequency to create the data extracts from various systems of record to feed these downstream systems. The result was a plethora of systems, diverse data streams and data repositories, multiple records per individual, no common agreement on data definition, no reconciliation between the sources or repositories of data, and no authoritative source of data for affiliation determination, provisioning or access control.

This environment was becoming unmanageable for developers (faced with inconsistent system and data requirements), provisioners of services (faced with inconsistent identities and undefined provisioning rules) as well as for end users who were continually faced with the complexities of designing and supporting their applications in a non-standard, disparate environment.



Solution

The first step in addressing this business challenge was to build an Enterprise Directory to provide the foundation for future identity and access management solutions. A commissioning statement and a list of Project Goals were created to establish the project's direction:

Commissioning Statement

Establish a framework for deploying and maintaining general-purpose directory services for the University of Colorado at Boulder, within the context of the University-wide environment. This "framework" should include a policy structure, data structure (including a unique identifier scheme), directory-enabling tools, and directory management process. Continue the project with a proof of concept by deploying the directory services framework to a defined set of systems.

Project Goals:

- *To develop a scalable, flexible, robust directory service for the University of Colorado at Boulder.*
- *To build a trusted and authoritative data source for CU-Boulder resources (i.e., an enterprise-wide data model), and, in the course of the project, to identify these resources and their corresponding information requirements.*
- *To offer a source for distributing information, finding information and locating data objects.*

- *To provide authentication services, integrating with campus-wide and application-specific security services.*
- *To develop a general purpose service, usable by a variety of authorized, independent applications and services, thereby enabling relationships within and between communities, systems and services and resolving discrepancies in information between these communities, systems and services.*
- *To set a direction toward an enterprise-wide, general-purpose directory and away from special-purpose, system-specific directories.*
- *To develop a process for identity management, data management and relationship management resulting in a unified management model.*
- *To offer location-independent access to directory information.*

Our approach to meeting these goals was first to create a database of identity information (the registry) that would be fed by multiple systems of record. A key component of the registry building process was the reconciliation of data from independent systems (most notably, the Student Information System and the PeopleSoft Human Resources System). The success of the registry build and reconciliation process was dependent upon system owners and functional experts



agreeing upon the business rules that would be applied during the data extract, reconciliation and build processes. Once the data was reconciled and the registry was built, Boulder-specific information would then be extracted from the registry into the directory.

Project Structure

The success of the Enterprise Directory depended upon the commitment of data owners, system owners and functional experts to create a common directory infrastructure. To achieve this, it was critical to have a project champion with the ability to generate campus-wide philosophical, political *and* financial support. In our case, our project champion was Dennis Maloney, ITS Executive Director. He had the clout to gather the appropriate campus decision-makers around the table to affect change. He also, arguably, had the most vested interest in the project since so many ITS-related projects (such as portals, provisioning, identity management, etc.) were dependent upon the successful implementation of an Enterprise Directory. As champion, Dennis has fully supported this project since its inception and is instrumental in promoting the project to the rest of campus. His position also provides a source of exposure to and collaboration with all of CU's campus and system-level IT Directors. This has proven beneficial for providing a basic blueprint as the other campuses have begun their own directory or directory-related project efforts.

While the Champion generated support throughout the campus, the project was also dependent upon four other key project components for decisions, review, expertise and project support: the Steering Team (which evolved into the Directory Governance Board as the Directory moved into production), the Project Core Team, the "Big Team," and the Technical Team.

- **Steering Team.** The steering team consisted of various decision-makers from different areas of campus and was responsible for resolving policy issues and promoting the project. Members included:
 - Director of Housing
 - Dean of Libraries
 - Director of Enrollment Services
 - Registrar
 - Director of Human Resources
 - Vice President of University Systems

The Steering Team has since evolved into the Directory Governance Board, continuing as the decision-making board related to Directory Services production issues. The most critical element of this team has been their ability to make policy decisions including the Directory Policy (see Appendix A). Interestingly, the policy has proven to be a good idea as a reference point and ultimate goal but, in reality, is rather difficult to enforce. The Steering Team was also the primary decision-making body for establishing the criteria for affiliation



definitions based upon an individual's relationship with the university.

- **Core Team.** This hands-on team consisted of subject matter and functional experts from various campus departments including Libraries, Housing, Communications, Registrar's Office, Computer Science Department, IT Security, and the Systems Office. These representatives provided functional and technical expertise for the design considerations, recommendations and guidelines which were discussed at monthly meetings. In addition, their membership created a sense of ownership and support that was then passed back to their respective departments.
- **The "Big Team".** In addition to the Core Team, an expanded team of approximately 40 functional experts was consulted during the design and analysis stage of the project (and as needed throughout the project). This team, although it never met as a full team, had a significant influence on the shape of the final product.
- **Technical Team.** Recommendations from Steering and Core became specifications that were then turned over to the Technical Team who made it all a reality. The Technical Team included a database analyst, a programmer, UNIX and directory server support staff, web designers and, shortly after

implementation, our Directory Manager who was/is responsible for managing the directory, for bringing enhancements and services into the Directory Services environment and for dealing with all issues directory-related. With the exception of the Directory Manager, none of these technical team members were devoted full time to the project. In hindsight, the lean staffing was a significant drawback; a full-time analyst and additional development staff would have moved the project along both more quickly and further through subsequent initiatives.

Many others have had significant roles within the project:

- Our IT Architect at the time (now retired) had the vision to initiate and set the direction for the project.
- The Project Manager worked with all of the teams and with the campus, bringing silos together, translating business rules into specifications and communicating between all teams and partners. The Project Manager was the only project team member who was devoted full-time to the project during development and performed both Project Management and systems analysis duties during the development effort.
- The Vice Chancellor of Planning, Budget, & Analysis was our funding source.



- CU-Boulder's CIO teamed with our Project Champion in promoting the project across the University.
- Our current IT Architect has taken the directory ball and run with it, drafting an Identity and Access Management technical position that provides us with a straightforward identity and access management direction. (See Appendix B.)

In addition, "sponsored" individuals can be entered by hand. This data is reconciled (based upon well-established identity matching business rules), logic is applied to determine affiliation(s), and the Registry (an Oracle database) is populated. Each day, a metadirectory (IBM Directory Integrator v4.7) routine runs, moving data from the Registry to CU-Boulder's Enterprise Directory (an ldapv3 directory) for those individuals who are currently affiliated with the Boulder or Central System campuses. The directory

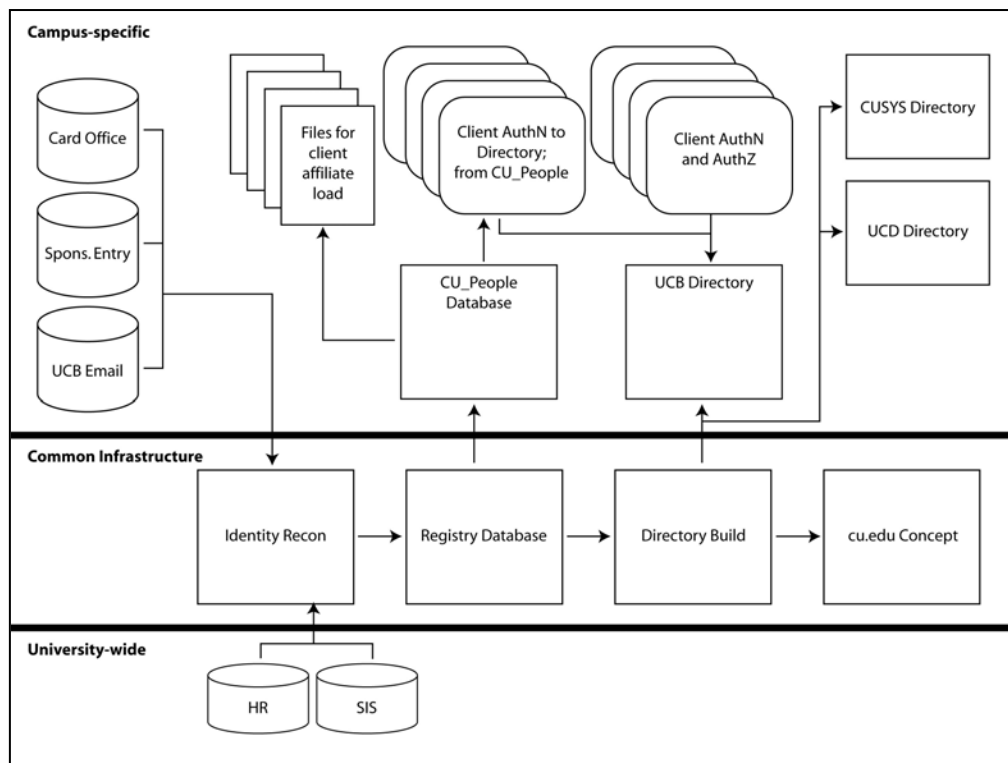


Figure 1: Current Directory Services Architecture at the University of Colorado at Boulder

Architecture

In general, data is extracted nightly from our systems of record for student information, human resources and for computing accounts. (See Figure 1 for a diagram of the current directory services architecture.)

has been defined using the eduPerson objectclass with the addition of a university-wide objectclass (cuEduPerson) as well as a campus-specific objectclass (coloradoPerson).



As noted above, one of the project goals was to provide authentication services, integrating with campus-wide and application-specific security services. This was accomplished by weaving together our Directory Services with various applications and with our already-established Kerberos-based authentication. This has resulted in single sign-on capability from the recently implemented student portal to WebCT, WebCal (SunOne calendaring), on-line software resource downloads, and student web-based personal look up services. Additional portal services planned for implementation this fall and spring that will take advantage of this authentication mechanism include computer-based training, on-line bill pay and Student WebFiles (Xythos). CU-Boulder has gone even further with authentication capabilities, however, with the planned rollout of CUAccess, an identity and access management service based on HP's SelectAccess product. This product proved to be a good fit in CU-Boulder's environment and offers robust access management to web services while taking advantage of our existing campus architecture including the directory and the Kerberos server. (See Figure 2 for future architecture plans.)

Authorization is still handled, for the most part, through individual applications. Our Directory Services does offer groups and roles and authorization based upon directory attribute checking; these can be managed using CUAccess.

Implementation

The Enterprise Directory was introduced gradually to the campus, beginning with a White Pages application and the support for email client address book lookups. After settling in a bit and being touted (and proven) as a strong component of the campus middleware infrastructure, requests for directory-based services began rolling in. Today, the following services are dependent upon our Directory Services:

- Mac OS authentication
- WebCT (our course management tool) authentication
- Our "cupeople" database built from the Registry and used by many applications that are dependent upon a complete and authoritative data source for affiliates of the university – where the application needs the relational features that a database provides (and a directory doesn't)
- WebCal (SunOne calendar directory built in parallel with the Enterprise Directory)
- Student Portal and, soon, Faculty/Staff Portal (relies upon the directory entry for authentication and for profile information as well as directory attributes for building portal announcement groups)
- Libraries for pre-population of their patron database and for on-line reservations authentication and



authentication and authorization for a web-based tutorial

- Judicial Affairs for pre-population of incident response forms

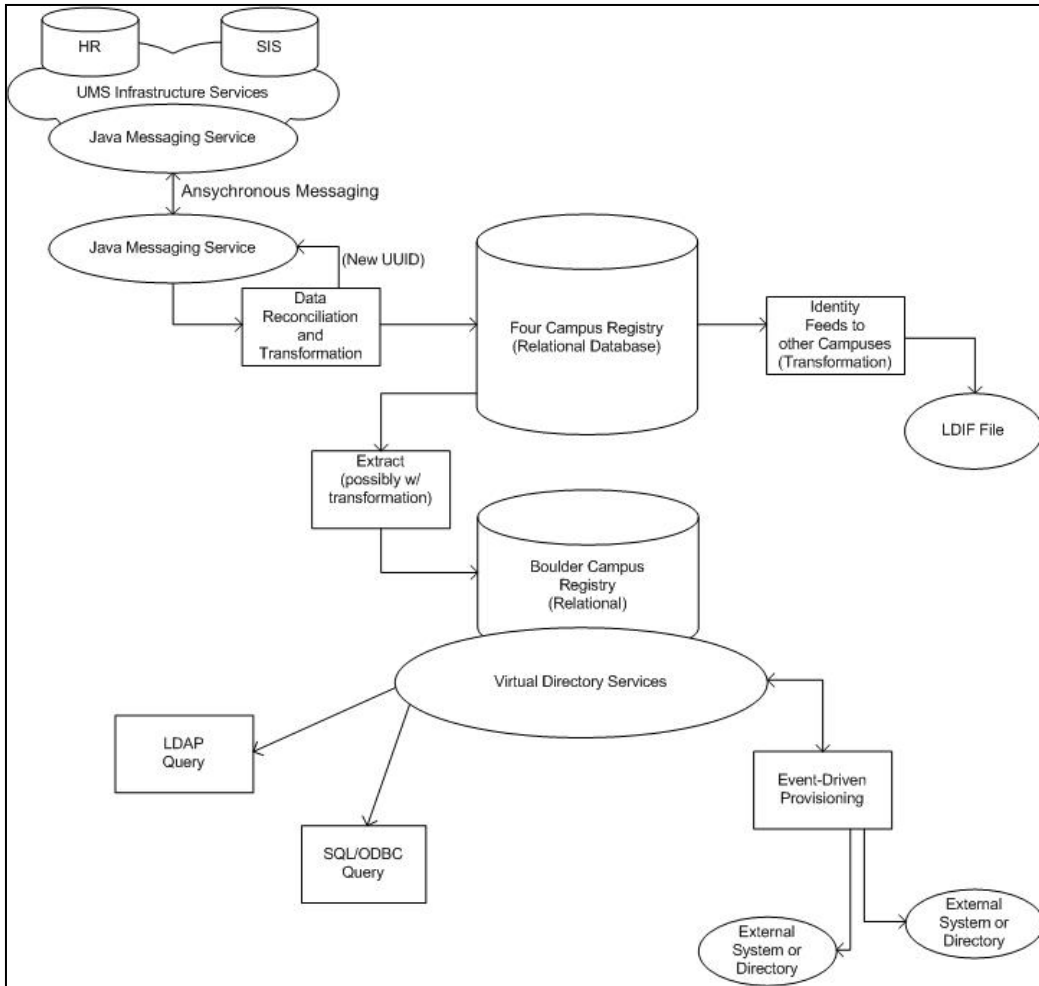


Figure 2: Potential Future Architecture

- Business School web services infrastructure (where our current identity and access management infrastructure is used to protect the Business School's internal web site)
- Web-based photo course rosters
- ALTEC language lab materials check-out authentication
- Parking Services for affiliation verification
- Recreation Center authorization to use their facilities
- Off Campus housing for authorization to their on-line services
- Secure look-up for affiliation (affiliation checking before provisioning)

- Secure look-up by our campus police for response follow-up
- Athletics for affiliation data for ticket sales prioritization applications
- Campus Idif files built specific to the Denver and Health Sciences Center campuses for their directory services
- Account creation and deletion activity based upon directory affiliation.

Of course, the list continues to grow. Email/messaging, federated authentication (Shibboleth), automated provisioning, the Admitted Student Portal and self-update capabilities are all waiting in the wings.

Resources

During the course of the project, ITS contributed a full-time project manager who doubled as lead systems analyst for the project. Other staffing resources contributed by ITS included a DBA, programmer, and server support staff (each roughly 30% time). Our now-retired IT Architect worked approximately 50% time on the project during its inception phase. Funding for the project came through funds set aside for infrastructure initiatives. Today, these 'pots of money' no longer exist and all projects must be justified and funded individually.

In addition to staffing and dollars, many other resources were tapped into during the course of the project. The Burton Group, Internet2 (including I2 institution business

cases), EDUCAUSE and peer institutions were referenced often for strategic insights. Tactical planning was done with contributions from Internet2, peer institutions, current literature, information gathered from conferences (EDUCAUSE, CUMREC, NMI-EDIT CAMP middleware workshops and Internet2 meetings) and listservs (Internet2's mace-dir and mace-dir-groups, and EDUCAUSE), and in collaboration with departments within the university. Valuable assistance from outside CU came from discussions with individuals from other Internet2 campuses, and from many of the initiatives produced through the NMI-EDIT/Internet2 Middleware initiatives. Most notable of these include the [eduPerson Directory Schema](#), [A Recipe for Configuring and Operating LDAP Directories](#), [Enterprise Directory Implementation Roadmap](#) (to which we contributed), [Metadirectory Practices for Enterprise Directories in Higher Education](#), [Practices in Directory Groups](#), [LDAP Analyzer](#), and [Shibboleth](#) and [Pubcookie](#) software, and video conferences and discussions related to multi-campus issues.

Lessons Learned, Recommendations and Conclusions

Go into a project like this with eyes wide open, examine the entire landscape, bring *all* parties together, realize technology is only one piece of the puzzle (the other 4,999 pieces are people, processes, policies, culture), and prepare for success – where



preparations should consider staffing, funding, and planning.

In addition:

- At the beginning of this project, we thought we would end up with a thin registry and a fat directory. Both have gotten fatter since then but overall our registry has been more developed than our directory. (Our rule of thumb is that, for an attribute to be added to the directory or registry, it must be needed by at least three applications.)
- One of our goals was to make a framework upon which other campuses could build their directory services. Does it make sense, however, for the Boulder campus to be supporting a multi-campus service?
- The project structure was instrumental to this project's success. Campus-wide, comprehensive and enthusiastic representation is a necessary component. The project must have insights from experts as well as a solid commitment from the institution's key decision-makers.
- Our registry/directory became the most consolidated and accurate source of identity on campus; consequently campus users are interested in using its information. However, the builders of the registry (ITS staff) are not the data owners. As a result we've established and nurtured an ongoing relationship with the data owners to work out privacy, data, reconciliation and access issues. The directory and registry would not thrive without this.
- We added two new object classes, `cuEduPerson` and `coloradoPerson`, to define university-wide and campus-specific attributes. The basic `eduPerson` definition laid good groundwork for this; the subsequent object classes have fit well within our multi-campus environment.
- The set of directory services is a continuum of development that must evolve to meet the needs of an increasingly complex array of applications. The work never stops.
- Directory Services must be thought of as a core service on the same scale as email, telephony, DNS, the campus network, and other traditional IT services that campus users rely on just to be there. Your Directory Services must be treated as a core infrastructure service from the beginning, and it must be built, managed and supported accordingly throughout its life, or its life will be very short indeed.

For More Information

For more information on The University of Colorado at Boulder case study, contact Melinda Jones at Melinda.Jones@colorado.edu.



Appendix A: Directory Policy

A. Rationale and Purpose of Policy:

The University of Colorado at Boulder (CU-Boulder) Enterprise Directory is a trusted and authoritative data source for CU-Boulder resources. As an Enterprise Directory, it is used by a variety of authorized, independent system applications and services; enables relationships within and between communities, system applications and services; and seeks to resolve discrepancies in information between these communities, systems, and services. This policy sets guidelines for the consistent use of the Directory and addresses the issues of inclusion definitions, source definitions, and uses in order to insure an accurate, secure, and functional enterprise Directory.

B. Policy and Scope

The Directory policy provides guidelines about the following aspects of CU-Boulder Directory:

- Directory governance
- Directory inclusion (categories of people who will be included in the CU-Boulder Directory);
- Official data sources (the information systems from which the Directory will

extract its data, create entries, and update entries, and upon which it will base its reconciliation);

- Directory uses (privacy requirements; who may have authenticated access to the Directory; who may pull data from the directory and for what purposes; and who must use the Directory).

C. Definitions

1. Directory Governance

The Directory and its policies, operation and evolution will be overseen by the Directory Governance Board (DGB). This Board will be an outgrowth of the Directory Services Project Steering Team and will include representatives of the Boulder Campus and the CU System and of each major constituency represented by the Directory (i.e., students, faculty, staff, and Information Technology Services). As the Directory expands its constituency base, so, too, will the Board expand its representation.

2. Directory Inclusion

The following categories of CU-Boulder and CU System Office people will be included in the CU-Boulder Directory. Status determination and/or affiliation duration are included as definitional items of each.



- Current Staff and Faculty – current appointment in the Human Resources System
- Current Student – registered for current term or on time out status with future expected return date
- Continuing Education students – currently registered
- Retiree – as determined by classification in the Human Resources system
- Surviving Spouse – as determined by classification in the Human Resources system
- Formal Affiliates – as approved for inclusion by the Directory Governance Board (such as Regents and members of the CU Foundation)
- Libraries Public Patrons – as determined by a public patron entry in the Libraries database
- Sponsored Affiliates – i.e., individuals not affiliated with the CU-Boulder, but involved with activities directly associated with CU-Boulder functions. When requesting or renewing an affiliation, a current full-time faculty or staff member must identify him or herself as the sponsor or contact related to the individual's University activities. This "sponsor" will provide information describing his or her relationship to the individual and outlining the individual's

affiliation/benefit to CU-Boulder, including the amount of time the sponsorship will be in effect. Both the sponsor and the affiliate will affirm (through written consent) their understanding of their responsibilities related to the use of University resources. Examples of Sponsored Affiliates include visiting researchers, some vendors and contractors, and some conference attendees.

Additional groups (such as future students, former students, alumni and executive boards) may be added at the discretion of the Directory Governance Board. Inclusion parameters must include an affiliation definition and affiliation duration.

Affiliates listed above will be included in the directory for service authorization privileges as appropriate and/or for visible association with CU-Boulder. Affiliation-specific services and visibility will be determined by the Directory Governance Board in accordance with university policy. Regarding visibility in particular, students may elect to shield their public visibility according to the Family Educational Rights and Privacy Act. Faculty and Staff's association with the university is public information as defined by the Colorado Open Records Act.

3. Official Data Sources

Sources

Directory data is populated from the following sources:

- Human Resources System (HR) – for current faculty, staff, retirees, and surviving spouses.



- Student Information System (SIS) – for current or “save” students and current Continuing Education students
- Unix Unique Account System (Uniquid) – for current Unix account holders
- Identification Card System (BuffOne Card) – for the ID Cardholder’s ISO number, conference attendees and vendors/contractors
- Libraries System – for public patrons
- Telecommunications Management System (Telecomm) – for faculty/staff office numbers (building and room)
- Faculty Information System (FIS) – for faculty-specific information such as degrees and research
- Housing Resident Management System – housing resident identification
- Authenticated manual entry – for formal and sponsored affiliates and for a limited number of self-maintained attributes. All manually entered information must adhere to the University’s responsible conduct laws and policies.

It is the responsibility of source system owners to participate in the effort to successfully integrate the Directory’s data. Source system owners are also responsible for ensuring timely availability of source data to the Directory. It is the job of the Directory Services’ technical support team to gather

enterprise-wide requirements for the directory-dependent applications and for the campus’ Information Technology infrastructure and to develop a working strategy to meet all requirements.

Create rights

Directory entries may be created by any of the following sources and only the following sources: HR, SIS, Uniquid, BuffOne Card, Libraries, and authenticated manual entry sources.

Update rights

Appropriate attributes within Directory entries may be updated by the following sources: HR, SIS, BuffOne Card, Uniquid, Libraries, Telecomm, Faculty Information System and authenticated manual entry (including self update of specific attributes for the individual).

Reconciliation procedures

Directory update processes will flag conflicting data (such as mismatches of identifiers, name, and date of birth). The Directory Operations Manager will report these mismatches to source system owners for reconciliation. Corrected data must be posted through the source system for subsequent entry into the Directory.

4. Directory Uses

Privacy Statement

The Directory will reflect privacy standards as defined by federal, state, and university laws and regulations. The Directory Governance Board will review these laws and regulations on a yearly basis with input



from relevant campus units (e.g., Legal Counsel).

Access Privileges

Anonymous access (i.e., access which does not require user authentication) for all public Directory information will be available to any desktop client (for example, via white pages or address books).

Access required by services/systems that are dependent upon the Directory must be approved by the Directory Governance Board and formalized by Service Level Agreements specific to the service or system requesting authenticated access.

Mandatory Directory Usage

All CU-Boulder campus-specific systems implemented after the advent of the Directory must be directory-enabled if affiliation-check, authorization or enterprise data is required by the newly implemented campus system. "Directory enablement" means using the Directory for determining affiliation, authentication, authorization, or for data reference.

D. Procedures

These policies will be reviewed by the Directory Governance Board monthly during the first year of the Directory's existence and at least yearly thereafter. Changes will be authorized by the approval of the DGB, the Information Technology Council, and/or the Chancellor's Executive Committee.

Significant extensions to the Directory (for example, extending the Directory to include

additional campuses) will also initiate a review of the policy.

Policy compliance will be enforced throughout the University of Colorado at Boulder campus by Information Technology Services (ITS) in collaboration with the Office of the Associate Vice Chancellor for Academic and Campus Technology (AVCACT). Requests for exceptions will be reviewed by the DGB, which will communicate exceptions and/or policy changes regularly to ITS and the Office of the AVCACT.

E. References

This policy complies with the guidelines as found in:

- Family Educational Rights and Privacy Act:
<http://registrar.colorado.edu/FacStaff/privacy.htm>
- Colorado Open Records Act (C.R.S. 24-72-201)
- University of Colorado at Boulder Information Technology Services, Access and Authorization Policy
- University of Colorado Laws and Policies:
http://www.cu.edu/Pres_Ofc/Policies/
<http://www.cu.edu/regents/LawsPolicies/>



- University Computing Use Responsibilities:
<http://www.Colorado.EDU/its/docs/responsibilities.html>

F. Responsible Organization

Information Technology Services in collaboration with the Office of the Associate Vice Chancellor for Academic and Campus Technology will be responsible for the maintenance and enforcement of this policy.

Appendix – Best Practices

1. **Directory Services** – Registry Data Requiring Mutual Oversight
(<http://www.Colorado.EDU/committees/DirectoryServices/>)
2. **Internet 2 Middleware Initiative**
(middleware.internet2.edu)
3. **The Burton Group, Network Strategy Services** (www.tbgroup.com)
4. **NMI-EDIT** (www.nmi-edit.org)



Appendix B: ITS Technical Position – Identity and Access Management

University of Colorado

Information Technology Services

ITS Technical Position

Identity and Access Management

Table of Contents

1	STATEMENT OF PROBLEM	1
2	BACKGROUND.....	1
3	ALTERNATIVES.....	1
4	STATEMENT AND BASIS FOR POSITION	4

[This page intentionally left blank.]

1 Statement of Problem

How should any campus group deploying an on-line service handle authentication, authorization, and personalization? For existing services, what, if any, modifications should be made to authentication and/or authorization schemes and when?

2 Background

Many on-line services that students, faculty, and staff use on a regular basis require the user to authenticate. The underlying applications then use available identity information to further authorize or personalize access to the application features.

Historically, providers of on-line services to students have either requested copies of student IDs and PINs that the application could reference for authentication, or they have made the services available through PLUS, assuring pre-authenticated users. This isn't always the case, however, as evidenced by a significant number of ad-hoc services such as course web sites that inflict a new and different username/password combination on the student.

On-line services for faculty and staff typically use some combination of HR data for authentication. Examples in use are combinations of SSN, Employee ID, and birth date.

The advent of the enterprise directory came with a policy that stipulates that any application requiring identity data would get that data via LDAP from the directory. Enforcement of this policy has been difficult and local copies of identity data prevail.

3 Alternatives

The overall strategic direction is for all on-line campus services to eventually be delivered through the campus portal framework. This framework can provide the authentication mechanism, an authorization scheme, the ability to personalize the delivery of information in a secure manner, and access to data from a number of sources.

In the short term, however, applications need to authenticate users and access identity information before they can be incorporated into a portal. Even in the long term, there are many services that will not be provided using logic that exists in a portal channel, but rather will be stand alone applications and simply have a link from within the portal or be incorporated into portal channels using web proxy or similar mechanism.

What follows are the alternatives for application authentication, authorization, and identity based personalization:

1. Store authentication and identity information local to the application:

Typically this identity data would be stored in a local relational database, but not always. This alternative is often combined with another. For example, an external authentication service could be used with some specific user attributes that are managed locally. This alternative has numerous severe drawbacks:

- Issues with copies of data from systems of record include timely updates, data inconsistencies, stability of extract mechanisms and data used out of context.
- If the local data is not official, then it potentially creates an inconsistent version of data that exists elsewhere. For example, the local application may be using a username and password for the user that is different (or at least not officially based upon) others that the user has.
- The data being managed locally is sometimes sensitive and, regardless of sensitivity, must be properly secured and managed according to federal guidelines and university policy. For example, local data must adhere to all FERPA regulations and management and display of this information must be handled accordingly.
- Often, local data obtained for one purpose gets used for other purposes in a violation of policy and possibly in violation of federal regulation. This behavior is extremely hard to protect against once the data is made available to a local independent system.

2. The application authenticates to the enterprise directory and retrieves identity information from the directory for authorization and personalization:

Because of the way it is implemented, Enterprise Directory authentication uses IdentiKey, CU-Boulder's the well-established user authentication mechanism. If the application binds securely to the directory, then it can also access specific user identity information.

Most modern applications can do LDAP authentication and many language toolkits exist, making directory authentication and authorization simple for in-house developed applications.

There are a couple of drawbacks with this alternative:

- Some identity information that is commonly required by applications, such as course enrollment data, is not available in the ED. However, much of this data is available through the portal framework (see below).
- Using the directory for authorization often introduces the requirement for some directory object or attribute to be managed by the application administrator. An example may be an application-specific group, where membership in the group is required for access to some feature of the application. There is no native directory administration interface that allows for this type of delegated management.

3. The application is an IIS service and authenticates to the Active Directory and uses AD group membership for authorization:

Applications deployed using Microsoft technologies on servers participating in the campus Active Directory can authenticate against the AD. Most users in the ED are in the AD, though the AD contains a minimal amount of identity information (significantly less than what's available in the ED).

OU administrators do have the ability to easily create and manage groups within the AD.

The AD also does not, by default, have a meaningful password for users. Any user needing to authenticate against the AD will need to explicitly set their AD password.

4. The service is written as a portal channel and authentication comes through the portal framework:

Any functionally created as a portal channel will inherently use the portal authentication mechanism and have access to the authorization and personalization capabilities of the portal framework.

5. The service is accessed through the portal framework and “trusts” the portal to authenticate users and pass some identity information:

External applications may or may not be able to accept “pass through” authentication from the portal, either for technical or logistical reasons. However, authorization remains the responsibility of the external application. The application may use this “pass through” mechanism for authentication and another alternative, such as a local database or an ED lookup, for authorization and personalization.

6. The application uses the CUAccess service for authentication, authorization, and possibly retrieval of identity data for personalization:

CUAccess is an identity and access management service that uses HP's SelectAccess software to allow both authentication and authorization to be externalized from applications. CUAccess is tightly integrated with the ED and can make identity information from the ED available to applications that have no ability to perform LDAP operations.

CUAccess provides a delegated administrative interface to the ED for managing user groups and roles. These groups and roles can be used directly by an LDAP enabled application even if the application does not use CUAccess for authentication and authorization.

Applications that use the CUAccess authentication service can enjoy single sign-on functionality between other CUAccess enabled services.

4 Statement and Basis for Position

IF the web application is new development or can be subject to re-design
AND the set of users is a subset of the users having portal access (or portal access can easily be expanded to include the desired set of users)
AND the service does not need to be provided external to a portal for any reason
THEN

⊕ **Write new web applications as portal channels; re-implement existing services as portal channels.**

By fully leveraging the portal framework, web applications can be readily available to portal-authenticated users and will have access to not only the most complete set of identity data, but also pre-established mechanisms for performing transactions and interfacing with a number of UMS and campus systems.

IF the web application exists and cannot be re-implemented at this time
OR the service must be provided external to a portal or is a packaged software product
AND the set of users is a subset of users in the enterprise directory
 OR can be sponsored users reasonably managed in the ED
AND a SelectAccess enforcer exists or can easily be created for the service platform
THEN

⊕ **Use CUAccess for authentication and access to directory identity attributes**

The goal is to incorporate CUAccess into the portal framework (by using CUAccess as the authentication mechanism for the portal). From a portal perspective, any service using CUAccess for authentication and authorization can be incorporated into the portal. An authenticated portal user will not need to re-authenticate and identity data specific to the service can be made available without the service having to perform an LDAP lookup.

Web services that have no relationship with a portal still benefit from CUAccess by externalizing authentication, authorization, and retrieval of identity attributes, providing a consistent authentication interface and mechanism that's well known to users, and being able to offer single sign-on between multiple services.

NOTE: The CUAccess service will not be available until fall 2004 and will not be incorporated into the portal framework until spring 2005. Services that must be deployed before CUAccess is available will need be postponed (preferred) or use another approach. Services that do use another approach should have a goal of being converted to CUAccess as soon as possible.

IF no SelectAccess enforcer is available for the service platform

AND the set of users is a subset of users in the enterprise directory
AND the application can perform LDAP lookups and/or LDAP authentication
THEN

⊕ **Use LDAP authentication against the ED and perform LDAP queries for identity data. Use SelectAccess to manage user groups and roles in the ED.**

When the service is not a Web-based service or can't use CUAccess, but can leverage the ED for authentication and authorization, do so. For authorization, LDAP enabled services should rely upon user group and role membership information from the enterprise directory. User group and role information can be accessed for lookup using LDAP, but delegated administration of this information is provided by CUAccess.

IF the set of users for a service do not have entries in the ED and cannot reasonably be managed as sponsored affiliates (due to volume or degree of management)
THEN

⊕ **Store authentication and identity data local to the application**

Additional security measures may be required *and audited* for systems containing local data so as to ensure protection of personal or sensitive data.

NOTE: The set of users having identities in the ED is continually expanding. Before taking storing authentication and identity data locally in an application, the needs of the service should carefully be considered along with known plans for the ED to determine if some schedule adjustment would eliminate the need for local data.

Authenticating against the AD or using a “pass through” authentication from the portal are NOT recommended approaches going forward, though in the near-term, before SelectAccess is available and incorporated into the portal framework it may be expedient to use one of these approaches with the understanding that they will need to be re-visited at a later date.