Identity and Access Management At Northwestern University

(*Abridged Version*)

Working Group Report

August 29, 2014

Working Group Membership

James Rich	Kellogg School of Management
Michael Satut	The Graduate School
Ken Woo	School of Continuing Studies
Kirsten Yehl	Feinberg School of Medicine
Stu Baker	Northwestern University Library
Serena Christian	Finance, Facilities, and Research Administration
Kristin McLean	Human Resources
Jody Reeme	Student Enterprise Systems
Tom Board	NUIT
David Keown	NUIT
Phil Tracy	NUIT

Executive Summary

The Context

The Northwestern University NetID management system was launched in 1993 to support electronic mail services. This marked the beginning of our Identity and Access Management (IAM) system as we have come to know it. (An IAM system is a set of applications, policies, and processes by which electronic identities and credentials are managed over their lifecycles, and the mechanisms by which business applications utilize that system to make decisions about permitting (or denying) access to their online services and resources.) Over the past twenty years, the majority of the electronic services within the University have adopted the NetID as their user identifier and authentication credential. This level of adoption has clearly benefited the University's ability to introduce new services in a relatively coordinated fashion.

During this time, the University's IAM infrastructure has grown organically without ever having benefitted from a systematic review of its functionality or how it aligns with the business needs of Northwestern. The decision to pause for a comprehensive review of this evolving and increasingly critical area was driven by multiple factors:

- the product "end of life" for the core Identity Management application (NUValidate)
- the growing importance of IAM functionality
- the frustration by the IT@NU community with the functional short-comings in this area
- the difficulty in maintaining the current, fragmented suite of systems

The University's IAM system is the primary hub of our ever-growing portfolio of online services that support a changing Northwestern community, and the context for this set of functionalities has changed qualitatively, particularly in the last 5-10 years. The community for which these identities are managed, and access decisions are made, is very different:

1. the University is entering into more partnership and affiliate agreements with external institutions;
2. the geographic scope of Northwestern is becoming increasingly distributed;
3. collaboration with people outside of the traditional boundaries of Northwestern is "the new normal";
4. there is an increasing interest in expanding the range of years during which the University maintains a relationship with "members of the Northwestern community", e.g., with talented and interested youth well in advance of the time they might apply to Northwestern, to people well past their young-adult student or even working days.

Concurrently, there has been a qualitative shift in technology. With the growth of online services and the rise of cloud computing, transactions and services need to happen online on a real-time basis, and the interaction of systems and the management of identities needs to happen "at scale" on a "hands off" basis.

The Path Forward

A cross-organizational working group (whose members are listed on the cover page) was formed in the fall of 2012 to compile a broad sampling of the IAM needs across the Northwestern community, and to recommend a path forward. One clear conclusion of this information gathering is that a reliance on episodic, just-in-time responses to changing circumstances has left the IAM system undervalued, and thus underinvested in, leaving it insufficient for the University's current and future needs. This insufficiency does not manifest itself in a big-bang, highly noticeable manner; the effects are felt repeatedly throughout the enterprise in user frustration, delays in getting new systems integrated and online, and staff time routinely wasted working around the system's deficiencies. Our IAM system is perhaps our most valuable enterprise system, enabling all of our online services, and it needs to be restructured and repositioned.

The degree of change that is needed to accommodate the trends listed above goes beyond isolated adjustments to one part of the system or another. The vision laid out in this document is designed to lead to an IAM system that will return much higher value for the University by being more integrated within itself, more integrated on a real-time basis with the applications that surround and depend on it, more secure where it needs to be, and more extensible and flexible via federated identities.

The report's recommendations are organized into three sections, each of which includes suggested changes based on key architectural cornerstones:

- The Identity Management System's (IdM) integration within itself needs to be improved via simplification and consolidation. Some of this work needs to be done as part of the process of replacing NUValidate. Areas connected to this replacement in which change is recommended include the processes for manual NetIDs and WildCARD procurement, the distributed Active Directory structure, and the practice of embedding access management logic within the IdM system.

- The Access Management system needs to be more integrated on a real-time basis with the Identity Management system, moving from a "heads down, internal to each system" process for authorizing access, to a process where online systems are more integrated with the IdM system on a real-time basis, "expose" information outside of their system about individuals' access status within their system so it can be used by other systems, and are "smarter" in the sense that they can make use of a new central registry (much of which will likely be virtual) when making their decisions about authorizing access. The University's new web services infrastructure (Service-Oriented Architecture (SOA)) and a commitment to enterprise web Single Sign-on (SSO) will be key to making these changes.

- The way our IAM system incorporates Identity assurance (the level of confidence that the credential is accurately associated with a real person, and the correct person) and trust (how do we know the person presenting the credential is really the person to whom it was issued) into its processes needs to be optimized vis a vis the resource/service being accessed. In some situations, this will mean the NetID is supplemented by multi-factor authentication during the login process, and in other situations it will mean there will be a reduced reliance on NetIDs via such techniques as Identity federation.

To achieve these goals, the entire IT@NU community will need to be involved. NUIT, distributed IT units, enterprise system development teams, and business application owners will need to be involved. The scope of the technological work should not be underestimated, but these technological changes cannot happen in a vacuum. New business rules and standard processes will have to be envisioned, refined, and adopted in order for the new technology to be selected, implemented, and work effectively.

NUIT's Identity Services team will be a pinch point in this initiative, and the Next Steps section (page 34) highlights work that will involve this team that is recommended for consideration for initial prioritization. Due to NUValidate's end of life status, preliminary envisioning of a new IAM model leads the list in order to know the functionality needed for its replacement. Also included for consideration are other sets of work that are more easily outsourced than the envisioning work is.

This is an abridged version of the report. Most of the appendices have been omitted, several of the sections in the original have been combined or omitted, and some of the detail in the text of the original report has been removed. Both versions attempt to cover a very complicated topic that has a lot of misunderstanding attached to it. A Quick Reference Guide to the report, which is a good starting point for both versions, is the only appendix carried over into this version

We hope we have articulated the need for change and have provided not only a beginning point for that change, but also a roadmap to be pursued over time in order to take advantage of different technological possibilities and keep pace with the University's changing environment and business aspirations.

I. Introduction

What is Identity and Access Management?

Two similar acronyms are used in this report: **IdM** and **IAM**. IdM stands for "Identity Management," which is a subset of IAM, or "Identity *and Access* Management." The two sets of functionality – the management of identities and the management of access – are tightly connected and therefore are often mistakenly conflated.

Identity Management (IdM) encompasses the maintenance tasks associated with the *lifecycle of electronic identities*: provisioning, de-provisioning, and handling changes in between. The IdM system also makes those identities, and a set of attributes for each identity, available via published directories, which can be used by surrounding applications to *authenticate* a person's credentials at the time of requested access and receive attributes about that person in return.

Access Management (the **"AM"** in **IAM**) encompasses the tasks associated with *providing access* to resources once a person's credentials have been authenticated. The identity management system makes no decisions about access to surrounding applications, only about the verification of credentials. The applications are, or should be, responsible for defining the business rules that *authorize* people's access to resources (e.g., read/create/update/delete data, gain access to a building) and implementing those rules based on personal attributes associated with an electronic identity. Together, these two sets of functionality – authentication and authorization - comprise **IAM – Identity and Access Management**.

(See Appendix A on page 60 of the unabridged report for a Glossary of Terms used in the report.)

Context for the IAM Working Group and the Following Report

Northwestern's Identity and Access Management infrastructure has grown organically over the last twenty years without ever having benefitted from a systematic review of its functionality or how it aligns with the business needs of Northwestern. The decision to pause for a comprehensive review of this evolving and increasingly critical area was driven by multiple factors:

1. the difficulty in maintaining the current, fragmented suite of systems;
2. the frustration expressed by the IT@NU community with the functional shortcomings in this area;
3. the growing importance of IAM functionality; and
4. the product "end of life" for the hub of the IAM system: NUValidate.

A special note is due regarding the status of NUValidate. In 2011, following Oracle's purchase of SUN, the identity management product was declared "end of life." It is no longer fully supported by the Oracle. We still retain a perpetual license to run the software, but there is some risk to this situation. The risk level is thought to be "low" to "medium low" because the software has been running for years without an incident, it is not widely accessible, and steps have been taken to reduce or eliminate storage of sensitive data in the system wherever possible. However, the status quo is not where we want to be, and the lack of ongoing vendor support is ultimately not tenable. Knowing that a system must be replaced makes it much less attractive to do further development and customization, which (a) will have to be re-done when the system is replaced and (b) reduces resources available for the replacement. This limits what can be done in the system to support other initiatives important to the University.

A cross-organizational working group (whose members are listed on the cover page) was formed in the fall of 2012 to compile a broad sampling of the IAM needs across the Northwestern community. That work was completed in the late spring of 2013, and work was begun on a summary report. The completion of this report has been delayed by two studies that now inform this report: working group reports on "Enterprise Content and Business Process Management," and on "A New Vision for Research Administrative Systems."

Because the topic and the research undertaken by the group is so broad, there are undoubtedly some omissions, oversights, and probably even some misstatements in this paper. For these we apologize in advance. Similarly, the focus groups were completed between November 2012 and April 2013. Time has passed since then, and some important work has taken place in the interim (e.g., iBuyNU is now enabled with the enterprise web Single Sign-on environment, and very important work is ongoing at the Feinberg School of Medicine and Northwestern Medicine). The authors of this report have not tried to cycle back with people to fully incorporate these developments.

Despite these caveats, we firmly believe we have captured the important details and the essence of IAM across the University. We do not intend to make this a living document to be updated as a result of the ensuing conversations or future developments. Instead, the intent is to grow the awareness of this critical area at this time, and to engage the community in a discussion about repositioning and improving this system in our IT portfolio.

Finally, it should be noted that once the focus groups were completed, the tasks of reviewing, congealing, and presenting the data necessarily had to narrow down to a smaller set of people, and these tasks became the responsibility of the NUIT representatives on the committee. Once a version of the report that was close to its publishable form was completed, the full working group was asked to provide feedback on the report. The members of the working group outside of NUIT dedicated many hours of effort and insight to this project, and the report has benefitted greatly for their commitment. Whatever shortcomings there are in the report, the primary authors within NUIT take full responsibility for them.

The Abridged Version of the Report

This is an abridged version of the complete 100+ page IAM at Northwestern University report. Most of the full report's appendices have been omitted, several of the sections in the original have been combined or omitted, and some of the detail in the text of the original report has been removed. A Quick Reference Guide to the report, which is a good starting point for both versions, is the only appendix carried over into this version. The unabridged version of IAM at Northwestern University is posted on the ASAC site.

II. The Changing Context for IAM

The Evolution of IAM at Northwestern

The first identity and access management (IAM) system at Northwestern was developed internally in 1993 to save labor in establishing email accounts and to support authentication through the University's modem pools. Known as SNAP (Simple Network Account Program), this initial IAM system received daily feeds from the HR and Student systems to create a combined census of who had rights to the services, and allowed administrators to "activate" the service for a requester. Rights to services were removed when the person was no longer present in either of the incoming data feeds.

Relative to the requirements today, the problem solved by the original deployment of SNAP was quite limited in scope, the community served was relatively small, and the application required only the most primitive expressions of the relationship between the person and the University. Since then, the use of NetIDs has expanded greatly to include the core "systems of record" (SES, FASIS, and NUFinancials), scores of other University applications (Blackboard, InfoEd, FAMIS, etc.), and many local applications in the schools.

As the growth of administrative systems increased, and the NetID spread to become the primary online identity credential, the variety of different situations requiring different ways of deciding whether or not to provide access to a person was growing as well. Rather than changing the role of the applications in this IAM relationship, the relation was left unchanged:

- the business applications still only expected a Y/N response to the authorization query about whether the NetID was active;
- the applications still only made a Y/N access decision on the result of that response; and
- business rules within the SNAP system, previously coded for the relatively simple original email application space, kept getting built out.

As a result, the SNAP system became crisscrossed with special cases, becoming increasingly difficult to modify to incorporate the next special case presented. Replacing the homegrown SNAP system with a commercial product in 2010, now called NUValidate, improved certain support efficiencies and addressed certain IAM system usability problems, but it did not address this tangled web of business logic.

Beyond just the increasingly tangled web of business logic within NUValidate, the IAM environment continued to become more complicated to accommodate growth in the number of online systems, their increasing complexity, and the need to manage access to their functionality.

The IAM system at Northwestern today is not a single system, as, for example, a room scheduling system is. Rather, it is a collection of applications:

1. **a core Identity Management (IdM) system** (NUValidate), which stores identities based on NetIDs that are in turn based on data fed primarily from authoritative identity sources such as the Faculty and Staff Information System (FASIS) and the Student Enterprise System (SES), allows people to manage those identities, and updates Northwestern's identity directories;
2. **identity directories** (e.g., LDAP, Active Directory, and Kerberos), which surrounding business applications use to authenticate users requesting access to their system;
3. **a physical identity system** (the WildCARD system), which provides proof of identity for access to buildings, events, etc.;
4. **a directory synchronization utility** (Radiant Logic), which keeps data in multiple Active Directory domains synchronized;
5. **a web Single Sign-on system** (SSO), which reduces the need to keep logging in with the same credentials for each Northwestern University application that is used;
6. **federation services** (e.g., Shibboleth), which allow people at trusted affiliate, partner, or peer institutions to use their home institution's credentials to gain access to Northwestern systems and services;
7. **a multi-factor authentication service**, which provides an extra layer of password protection using an application on a registered smart phone or answering a phone call to reduce the risk that personal information can be easily compromised should someone learn a NetID password; and
8. **an "Identity Provider" bridge service** (currently being run by the Alumni and Development Enterprise Applications team for the OurNorthwestern system), which enables alumni to log in with either an active Northwestern identity or with one of their own external social accounts (Gmail, Yahoo, Microsoft).

See the section on "IAM in Action" in Appendix D of the unabridged report (page70) for a diagram and description of how these parts work together to provide IAM functionality when a person tries to log in to a Northwestern application. Appendix D also has an annotated diagram that shows how data flows within the IAM system.

The Increasing Importance of IAM in Today's World

These changes in the parts and complexity of the IAM system, and the surrounding web of applications it enables, reflect the changing nature of our world over the two decades since the most vexing problems were access to University email and the modem pool.

- Online services offered by the University are qualitatively greater in both number and complexity, and more parts of the daily activities on the University community are premised on easy access to these services.
- As services become available outside of proprietary systems and through web services, services and data are expected to be available for self-service in real time, and integrated with one another.
- More services are being offered from the cloud, and maintaining control over identity-related information is increasingly important in the face of increasing security threats and the increased regulations they engender.
- The Northwestern community is growing and becoming more complex as:
  - the University enters into more partnership and affiliate agreements with external institutions;
  - the scope of Northwestern becomes increasingly distributed geographically;
  - collaboration with people outside traditional boundaries of Northwestern becomes the new normal (e.g., fellow researchers and peer administrators at other institutions, consortiums of universities offering courses, practitioners outside the University, community engagement);
  - interest grows in expanding the range of years during which the University maintains a relationship with members of the Northwestern community (from earlier in life – youth and young adults participating in University programs such as the Center for Talent Development or the National High School Institute program – to later in life, via alumni programs and life-long learning).

As these trends continue, the IAM system will need to handle a wider variety of situations, offer more options, depend less on physical proximity, and be flexible enough to be deployed quickly and effectively at scale. Otherwise, it risks becoming the bottleneck for deployment of new services and an area of increased risk of compromise. In this sense, the IAM system has become one of the University's most important systems.

A Note on Northwestern Medicine

Before beginning this analysis, a special note on the situation at Northwestern Medicine is in order. The separate but inextricably intertwined relationship of the Feinberg School of Medicine, Northwestern Memorial Hospital, and the Northwestern Memorial Faculty Foundation has led to an IAM situation that is particularly complex, and becomes further complicated as partnerships are made with more medical institutions. The intertwining of these institutions creates an enhanced set of IAM challenges. For example, hospitals want more stringent standards around identity and access management policies and procedures, and budgets, policies, and resources are the responsibility of three mostly independent organizations, yet the doctors, researchers, and the administrators need to work fluidly across the organizational silos and resource redundancies. Undoubtedly, there are places in this report where those challenges could have been better highlighted.

In the time between the publication of this report and its beginning in 2012, Northwestern Medicine and the University as a whole have engaged in significant work on these fronts. Suffice it to say, that being a part of that work is a high priority of any plan for making progress within IAM at the University.

III. An Assessment of Northwestern's Current IAM Environment and a Recommended Path for Improvement

Northwestern's IAM system provides a wide range of functionality, and for the most part, it handles basic IAM functionality across the identity lifecycle for the traditional core constituencies of on-campus faculty, staff, and

students. However, there are many areas where the University's situation can and needs to be improved, given the changing nature of the Northwestern community and the changing landscape around us (technological, regulatory, security, and user expectations). Much work lies ahead if we are to leverage our Identity and Access Management infrastructure at scale, as we must to do to optimize the delivery and support of our growing portfolio of online services and resources.

Identity and Access Management (IAM) is a relationship, with a set of systems/applications in the center that manage identities and enable a surrounding portfolio of systems/applications to make intelligent, real-time decisions that authorize access to their resources/services. Based on the results of this study, we will posit that a highly functioning identity and access management system is typified by these nine characteristics:

1. Each person has a single electronic identity. There may be multiple credentials attached to that identity, but there is only one electronic identity.
2. The IdM infrastructure is integrated within itself, so that data about identities and personal attributes flows smoothly throughout the system.
3. Identities and access to resources are provisioned and de-provisioned rapidly in alignment with the need for their actual usage, with easily auditable trails.
4. Authorization is appropriately granular and based on robust identity information.
5. Surrounding business applications are integrated with the enterprise IdM system.
6. The level of rigor employed in identity proofing and authentication at the time of access is based on the risk and value of the transactions to be done.
7. Identities are protected and secure.
8. Each part of the IAM system is relatively easy to maintain and to replace.
9. Business applications and the IAM infrastructure are flexible and easily modified to take advantage of new IAM technologies as they emerge and become stable.

The following sections look at these areas, summarizing our current situation, laying out a vision for where we need to be going with our architecture, and describing the work that needs to be done. Two organizing frameworks are used for presenting this material.

First, the sets of work are organized into three categories, one category per section:

1. **Restructure Identity Management** – The identity management system needs to be restructured: reducing its fragmented complexity via consolidation where possible, improving the flow of data within its parts, and making more information available.

2. **Integrate Identity and Access Management** – The management of access that is done by business systems throughout the University needs to be more closely integrated with the identity management system, it needs to be capable of more sophisticated and more granular decision-making, and its processes need to be online rather than based in paper transactions.

3. **Optimize Levels of Assurance and Trust** – The overall IAM system needs to be optimized as to how it deals with identity assurance and trust, leveraging external identities or applying additional layers of security and awareness where feasible and appropriate.

Second, the vision for Identity and Access Management is identified within these sections by calling out five architectural cornerstones:

1. **A consolidated Identity Management System** – The identity management system needs to be consolidated at the center with delegated administrative functionality.

2. **Central Registry** – A central registry should be built to provide access to a more robust set of data (than is currently available via LDAP) about a broad spectrum of people with a relationship to the University (i.e., not just those with NetIDs). Each person's information should be tied to a unique identifier that is not an already existing University identity or identifier. Most of the data will be accessible virtually (rather than being replicated to a database).

3. **"Smarter," more Identity-aware Applications** – Authorization, the permission to access resources, needs to be handled by the surrounding business applications, not by the identity management system. Applications must become identity-aware pieces of an integrated portfolio, rather than heads-down, internally focused silos, and authorization should be flexible enough to open access to individual services as needed.

4. **Online Processes** – Identity and access management processes need to be done online utilizing web services that provide real-time and workflow functionality based on data for individuals.

5. **Northwestern NetID Supplements** – The Northwestern NetID will remain the core Northwestern electronic credential, but its role needs to be supplemented by external credentials and by other means of insuring identity trust and assurance.

Restructuring Identity Management

IAM Characteristics addressed in this section:

1. Each person has a single electronic identity. There may be multiple credentials attached to that identity, but there is only one electronic identity.
2. The IdM infrastructure is integrated within itself, and data about identities and personal attributes flows smoothly throughout the system.
4. Authorization is appropriately granular and based on robust identity information.
8. Each part of the IAM system is relatively easy to maintain and to replace.

As stated earlier, our current IAM system handles the IAM basics for the traditional core constituencies. However, a lot of hands-on attention is required to keep the system functioning, these traditional core constituencies are becoming less traditional and less the sole focus, and the expectations surrounding User Experience have grown. In order to meet these challenges, Northwestern's IdM system needs to be less complex (often via consolidation) and better able to provide a single unified identifier for each person and more robust information for surrounding business applications. The following discussion shows multiple ways in which the first architectural cornerstone should be incorporated:

IAM Architectural Cornerstone #1:
The identity management system needs to be consolidated at the center with delegated administrative functionality.

Specifically, six different forms of consolidation are discussed:

1. Use the NetID as the University authentication credential for system access when a University credential is called for.
2. Reduce the use of Manual NetIDs.
3. Merge provisioning of WildCARDS in with the core IdM processes they now mirror.
4. Consolidate Active Directory domains, but provide distributed ability to create and manage groups.

5. Centralize access to identity credentials and identifiers.
6. Centralize access to personal attributes that are now isolated in systems and local AD domains.

Reducing Complexity via Consolidation – the NetID as the primary University credential

One of the core tenets of Identity Management was a fundamental premise of multiple focus groups: individuals should have one electronic identity at Northwestern. The first aspect of being able to provide the correct information about a person, and being able to accurately provide access to resources, is being able to correctly match an identity with a person and their attribute data. To do this effectively, there has to be a unique identifier for each person. Duplicate IDs or separate IDs for access to different sets of services are barriers to many important goals: improved customer service, administrative efficiency, foundational information security, and integrated reporting. There is great value to be realized by applying effort on the front end of the identity management lifecycle to avoid issues in these areas later on.

With the widespread adoption of the NetID as the primary electronic identity at Northwestern, there is a basic adherence to this premise. However, the identity landscape is more complicated overall because there are other enterprise-wide identifiers: WildCARD barcodes, and EMPLIDs created in FASIS and SES and used in other systems. (An EMPLID is an employee ID, which is the unique key for a person in both FASIS and SES. If a person is both an employee and a student, their EMPLID in each system should match.)

In addition, there are system-specific identities at Northwestern. By definition, these systems are not connected to the University's identity management system and, therefore, cannot utilize functionality associated with it. Examples mentioned during the focus groups include: iBuyNU (since corrected), the I-9 Service Center, Galter Library, multiple systems for alumni that are not tied together, McCormick systems that grant students access to Microsoft products for their classwork, Quest login (UNIX ID), HR Benefit Systems (e.g., FSA), Vista, ProCard, FundDriver, and CBORD (cashless card system utilized by University Services for "Munch Money"). All of these identities have different username and password conventions, and different frequencies for password changes.

Northwestern's online alumni community system used to be by far the largest system using its own unique, non-NetID Northwestern identity. Now, however, the use of Google IDs for most students in the University has almost entirely eliminated this issue for recent graduates, and Our Northwestern's current ability to accept social identities (e.g., Facebook, or a non-Northwestern Google account) has further reduced the need for a system-specific Northwestern identity for alumni.

One of the exceptions to a standardization on the NetID is the management of identities with "elevated" privileges – i.e., they are responsible for other people's data, not just their own via self-service -- within FASIS and SES (but not for NUFinancials). Premised on the need for additional external controls, a separate set of IDs is maintained for these people outside the NetID process by SES and FASIS administrators. This practice is unable to take advantage of the benefits of integration with the normal IAM system (e.g., users only having to know one ID/PW, the management of identity lifecycles being less dependent on weekly reports and manual processing).

Wherever a Northwestern credential is required for authentication, we should work to make it the NetID, whether it is for third-party systems, home-grown systems, or people with elevated access to enterprise systems. The process of creating special IDs for people with elevated access to FASIS and SES should be discontinued, bringing it in line with NUFinancials, and eliminating one of the blockages to integrating these systems into the University's web Single Sign-on system. The current procedures should be replaced by a process that utilizes multi-factor authentication to enhance security on these identities, leverages online workflows to improve the speed and auditability of the process, and externalizes this status outside of the relevant system so that it can be used by other workflows. (See the discussion of secondary attributes on page 19.)

Improving the ability to have a single identity for each person

Due to limitations in the identity-creation processes within the two systems of record (FASIS and SES) and within the Manual NetID process, there are multiple ways duplicate identity credentials/identifiers can be created at Northwestern. (For a fuller discussion of duplicates, see page 12 of the unabridged report.)

Duplicate IDs are problematic. They cause negative experiences for members of the Northwestern community at the beginning of their relationships, giving them poor impressions of Northwestern, enterprise systems, and NUIT. Take, for instance, a student who is hired as an employee and ends up with two sets of NetIDs / EMPLIDs. Unless they are resolved, she will have to use one NetID for all of her coursework, and the other for entering her time in Kronos. That creates obvious confusion for her, and for other people who depend on her being represented by one identity (e.g., uncorrelated data in reports, two Online Directory listings with different info, official notices going to two separate email accounts, etc.). Administrators who need to remediate these duplicates face challenges because it's difficult to "move" privileges or resources from the old NetID to the new NetID. There are no tools to merge records in the enterprise systems or the Identity Management system, so most, if not all, things need to be moved manually, and there is no one-stop shopping: each system must be handled individually (e.g., Exchange mailbox, email address, NU Financials privileges, Blackboard courses, etc.).

Creation of duplicates within the systems of record is related to data issues, and access of each system to the most current data in the other system. To guard against duplicate EMPLIDs or NetIDs being created FASIS and SES have algorithms, run as part of the creation of a record for a new person, which check against other existing records in their own system. In order to verify a pre-existing record, the following fields are used in varying combinations that have varying levels of certainty attached to them: name, gender, DOB, SSN, citizenship, address, email. Missing information, typos, name changes, non-domestic names reversed by mistake, dates of birth entered incorrectly because of international differences in the order of m/d/y, multiple passport IDs, and gender changes can all cause a match to be overlooked, as can the "dummy SSNs" that are created for international students who do not yet have a SSN.

Clearly, one area where these mistakes can congregate is with international students. Another is people who have intermittent relationships with the University (e.g., CTD students, lifelong learners, adjunct faculty, etc.), where information gets entered many times and can also change over time. A key to reducing the number of duplicates here is improving the quality of data not only within systems, but provided across systems. For instance, SES gets a nightly feed of bio/demographical data from faculty and staff in FASIS. However, once a person's data is sent over, only the faculty data is subsequently updated with changes from FASIS. If a staff member comes over in the feed initially with a dummy SSN, and that person decides to take a course later after they have received a real SSN, their SSNs will not match. One focus group also talked about the relatively recent administrative access that had been given to the International Office so they could directly correct personal attribute errors in international student SES records, and how that had made such a difference in keeping this demographic data accurate.

Other scenarios that can result in duplicate EMPLIDs and NetIDs occur in the admissions process, with either former students applying to grad/professional school, a prospect submitting multiple applications, or a person applying for full-time undergrad status after being in Continuing Studies for a while. These situations are complicated by the fact that students do not need to provide a legal name or SSN unless they get financial aid. Duplicates can also occur within FASIS when hiring an affiliate, re-hiring former employees, or giving a current employee a second job. However, because FASIS must collect SSNs and verify I-9s, duplicate records in FASIS are less common than in SES.

These scenarios mostly focus on data issues within the two authoritative systems, where improvements are dependent on data entry, data quality assurance, and searching/matching. Manual NetIDs (also referred to in

some instances as “affiliate NetIDs”) also contribute to the problem of duplicate IDs. They are discussed below.

Manual NetIDs

Because the SNAP system was informed only by FASIS and SES information, a sizable set of people did not have a way to gain access to file shares, the internet, or online resources such as the Library, BlackBoard, or Alumni Relations applications. These people, whose relationship with the University was temporary or fell outside the normal HR employment or student matriculation processes (e. g., affiliates, contractors, Sodexo and Aramark employees, volunteers, summer enrichment program attendees, CTD students, NMFF, NMH, and other medical center staff who work closely with Northwestern, especially with the Feinberg School of Medicine), needed a different path to these resources.

This path was created by the Manual NetID (thus named because these IDs are “manually asserted” by distributed administrators rather than being created as part of the normal hiring/matriculation processes that have multiple identity assertions, such as I-9 data, passport numbers, SAT scores, transcripts, recommendation letters -- built into them). Unfortunately, the lack of controls placed around the Manual NetID process, and its lack of integration with the rest of the IdM process, heightens the likelihood that a person can end up with multiple NetIDs. For instance:

1. NUValidate’s Manual NetID creation functionality has limited search/match features. You can search for an existing NetID or EMPLID and add a manual assertion to keep it alive beyond termination of employment/enrollment, but when creating a new NetID, none of the standard search/matching is available (even by a person’s name), either within the IdM system itself or with SES or FASIS.
2. The batch processing of FASIS/SES data for new faculty, staff, or students matches only on EMPLIDs and SSNs, and manually asserted NetIDs have neither of these.
3. Because the Manual NetID functionality has such limited search/match capability, the ability to avoid creating a second manually-asserted NetID is dependent on the information that is available to the person creating the NetID. The creation process can match on EMPLID or NetID if the person’s current EMPLID/NetID is known and used by the NetID administrator. If they don’t have this information, for whatever reason, a duplicate NetID will be created for someone who already has one.

The Manual NetID is hard to manage it during its lifecycle because so little contextual information is captured when it is created (about the person, the reason for the identity creation, or the person responsible for the identity). This lack of contextual information makes these identities more likely to linger past the time the access is needed and appropriate, and it makes troubleshooting any identity and access problems difficult.

Despite the problems created for the identity-owner or the administrators having to sort them out later, the Manual NetID process has come to be seen as a tool to accelerate giving online access to people who still would go through the normal HR hiring or student matriculation processes. For example:

1. Faculty and graduate students sometimes need early access to systems in order to support grant applications or work in Blackboard. (New hires can now be entered up to 90 days in advance of their actual hire date, and adjunct faculty can be hired on an annual basis but only activated for the quarters in which they actually teach, but not everyone knows about these processes.)
2. Contract/temporary employees are issued affiliate (i.e., manual) NetIDs. If they are “converted” to a full-time NU position, they are then issued a new NetID.
3. Staff at NMFF/NMH who need access to NUFinancials and other systems are issued manual NetIDs.

In other words, the creation of the Manual NetID system addressed an acute problem, but in the process of doing so, it turned the acute pain into a less noticeable aggravation, and the diffusion of the pain took the urgency away from the need to remediate the root problems.

Reducing our Dependency on Manual NetIDs

While it is hard to imagine not needing a means for providing and managing identities for people who do not come to the University via the core hiring and matriculation processes, eliminating/reducing the use of Manual NetIDs would not only make a major contribution towards reducing the problem of duplicate NetIDs, it would also have positive implications for auditing, security, and reducing the overall complexity of the IdM system. With a goal to replace NUValidate by the end of 2016, it does not make sense to invest time enhancing this custom module now, and it seems unlikely that resources should be invested in recreating this functionality in whatever succeeds NUValidate. Instead, attention should be focused on two areas:

1. *Low-hanging fruit*: Improving the awareness of the problems surrounding Manual NetIDs, and the existence of alternatives to using manual NetIDs can help. Improvement in the business processes in the distributed units for creating NetIDs could also help reduce duplicate NetID issues. These could include: developing templates for keeping records, getting people to fill in all fields accurately and completely, improving communications with the person getting the NetID to see if they've ever had one before, or perhaps moving the responsibilities for maintaining these identities to the HR teams in the distributed units.

2. *Transformation*: Engaging in a conversation about the possibilities of reducing, and perhaps even eliminating, the need for Manual NetIDs will be beneficial in the long run. If there are resources where there is no longer a need to require the use of a Northwestern credential, such as was done with the guest wireless network, that transition should be done.

   Where the use of a credential is still desired, a bigger push on federation (see below, page 30) can help a lot. (It was stated in one focus group that 90% of the manual NetIDs that need NUFinancials access are from NMFF/NMH.) Where possible, federation should be with other institutions having solid identity management processes. For less sensitive situations, federation can be done with lower levels of identity vetting (e.g., with social identities). To increase the level of assurance associated with using such credentials, this federation can be coupled with active in-person vetting of people and identities. (See the section below on Assurance and Trust, page 27.)

   For people who don't fall into either of these categories, and we still want to give them a NetID, using an HR- or SES-like process for entry into the official systems of record (FASIS, SES), perhaps using the existing FASIS POI ("person of interest") category, could be an option. These might be people with a non-employment relationship (e.g., contractors, those needing access to NetID-authenticated resources such as legally/contractually restricted data, or those using the University's Virtual Private Network). However, to be clear: if the FASIS POI process came to have a role, it would not be simply to give them an identity. It would be because they need access to resources that are restricted to NetID access.

Reducing Complexity via Consolidation – revisiting "extensions" to the IdM system

As Northwestern's IAM needs have grown over the years, the parts that comprise the system have grown as well. Too often, however, these pieces have been added with short time lines, limited budgets, and/or no overall strategic discussions or agreement. One of these – Manual NetIDs – has just been discussed. The current implementation of the WildCARD system and the University's Active Directory infrastructure are two other extensions of the IdM system that have added great value, but have created User Experience frustrations and significant information disconnects, and have resulted in high support burdens related to day-to-day activities and ease of replacement. These two extensions are discussed next.

Integrate the WildCARD with the Rest of the IdM System

The WildCARD system represents a second separate identity management environment that issues a physical credential with attributes recorded upon it. Similar to NUValidate, the system receives its own data from both FASIS and SES to create most WildCARDs, and then that data is supplemented by the ability to create "manual" WildCARDs for people not within the primary systems of record (e.g., for alumni or spouses wishing library privileges, or spouses who would like to be able to get WildCARD discounts).

As part of NUValidate's retirement, the WildCARD provisioning process should be consolidated with the NetID provisioning process in order to reduce the number of moving parts that have to be kept in sync.

The community would like to see more integration between these credentials. For instance, WildCARDs do not store NetID information in their onboard data cache, and while WildCARDs are used to record attendance at an event or permit access to a building, the systems that store identity information have not contained the WildCARD barcode number. (NOTE: The WildCARD barcode information has recently been added to the attributes stored in LDAP, which makes it retrievable for these purposes along with the already stored NetID and EMPLID.) Because usage of the WildCARD will only increase (e.g., the addition of RFID functionality into the card will enable touchless access, it could be integrated into the parking system or Ventra), its unification with the rest of the IdM system should be made a priority.

Reduce the Number of Microsoft Active Directory Domain

With the rise of client-server technology in the mid-1990's, and the attendant growth in popularity of Microsoft applications, Microsoft environments (i.e., "domains" built around their own independent and proprietary directory, "Active Directory") began to spread through schools and business units across the University. The University chose to continue its focus on LDAP as its main online registry of identities for the University, but took two steps to address the existence of these two independent identity environments: (1) a provisioning process was set up to replicate NetIDs out to the Active Directory (AD) domains via a Radiant Logics server that acted as a synchronization utility; and (2) a bare-bones central University AD domain was created ("ads").

There are now twenty Active Directory domains at Northwestern synchronized via Radiant Logic, and although far better than having a University identities and separate identities in each domain, this solution is not without its flaws:

1. While NetIDs and their passwords are successfully synchronized across these domains on a daily basis, this infrastructure introduces glitches in the IdM system that have resisted attempts at remediation. For instance, users must change their NetID password as the only known way for a transferred employee to get access to the AD domain in her new department, or for a person to fix "lost" access permissions when their identity status changes. An upgrade to Radiant Logic is underway, which may address some of these issues, but a tighter integration directly with the IdM system is a better long-term solution.
2. People in schools or business units with their own AD domain, need to remember to preface their login credential with a different (and often strangely foreign) domain name ("ads") when trying to use a centralized shared service, or another unit's domain name when trying to use a federated service.
3. Group memberships and other data are core to providing access to online resources, and this data does not flow easily downstream from the LDAP registry to the AD domains, or upstream from the distributed domains to the central University domain.
4. The duplication of these environments adds security risks (data replicated to twenty different environments presents twenty opportunities for accidental or intentional security breaches),

redundant infrastructure and staff effort in the distributed IT units, and added complexity to be supported by the central identity management team.

In the early years of this architecture, the limitations associated with it were at acceptable. But as the University has deployed more services that assume use of Microsoft's AD environment, this arrangement has become more problematic. The fragmentation of the AD directory infrastructure, its limited integration with the rest of the IdM system, and the relatively undifferentiated group structure in the bare-bones, central University AD domain can limit and complicate deployment of global services that depend on the Microsoft AD infrastructure (e.g., the Microsoft collaboration suite, OnBase or ImageNow, or support tools utilized by IT organizations).

In recent years, the central AD domain has been extended in its functionality and utilization, a number of smaller schools/units have given up their own AD instance and moved to the central domain, and others are expressing similar interest. The consolidation of domains will become an increasingly viable option as more services begin to be utilized across domains, as the central domain becomes more robust, and as schools and units look to save or refocus local IT resources. This movement should be encouraged, with a goal of a single Active Directory domain, or at least at reduced number of domains that are linked via "trust."

However, local AD domains are a significant business resource for schools and business units, and a considerable number of design discussions will be needed in order to determine how the distributed units are actually using their domains, and what additional tools and overall AD structure must be in place to rationalize the twenty remaining AD domains. For instance, the current ability to manage local groups and resources will probably need to be replaced by a different software solution (e.g., Grouper), work needs to be done to insure that currently local AD schemas have not been extended in ways that will be harmed by collapsing into one central schema, and the local applications that rely upon local domains need to be evaluated for transitioning them to the central domain.

Reducing Complexity by Removing "Special Case" Authorization Logic from the IdM System

In the examples above, IdM complexity is reduced via consolidation. In another very important instance, complexity needs to be reduced by moving the "special-case" authorization code (See the section on the Evolution of IAM at Northwestern, which begins on page 6.) out of the NUValidate system so NUValidate can be more easily replaced, and more easily maintained thereafter.

Prior to actually replacing NUValidate, the tangled web of special-case logic needs to be placed outside of the IdM system. Software logic that is added into the identity management system itself should be restricted to functionality that is in direct service of actual identity life cycle and password management. Code related to applications' authorization needs belongs outside of the identity management system and is the responsibility of the unit owning the relevant application. Who actually writes and maintains this code for an application is, as always, open to discussion, but the business and financial responsibility for the application rules belongs with the application owners, and all work should be done in accordance with documentation provided by NUIT. Placing this code outside of the IdM system (and decentralizing it as much as possible) will reduce the difficulty in replacing NUValidate. It will reduce the day-to-day operational burdens (maintenance, incident troubleshooting and remediation) on the Identity Services team so it can focus efforts elsewhere. It places the knowledge of the logic closer to the application that depends on it, making that team more aware of their own business rules.

Improving IAM via the Creation of a Consolidated Central Registry of Identities and Personal Attributes

Discussions in the focus groups often returned to the need for better data in order to provide better access to online resources. Four areas of improvements in data were mentioned:

Persistent Information

It is not unusual for a person to come and go from the University, with varying amounts of time in between being connected and not. Examples include:

1. Intermittent instructors (SCS, adjuncts, CTD) who teach for a part, or parts, of a year, year after year, and for whom any delays in getting access to Blackboard can have a significant impact on their classes;
2. Retired learners (OLLI – Osher Lifelong Learning Institute) who need access to Blackboard on a recurring but often episodic basis, and often register with a very short turnaround time;
3. Students in part-time programs where continual enrollment is not required, or students who take a year off for a variety of reasons;
4. Consultants, contractors, or volunteers who may subsequently become students, staff, or faculty, and vice versa.
5. Students in the continuing professional education courses offered by multiple schools at Northwestern, whose identity is not retained from one session to the next.

These types of community members will only grow, as will the University's desire to reliably retain a positive connection to potential prospects, particularly those already in one of the growing number of programs the University has for bright and talented youth.

In most of these situations, the IdM system does not facilitate the re-integration of these people when they return to the University. At best, the result is that the same amount of effort is required to provision their ID each time they come; at worst, duplicate NetIDs can be provisioned.

Information about Multiple Relationships

Another area in which personal attribute data is missing is when a person has multiple relationships with the University (e.g., joint appointments for faculty, students in dual degree programs, staff who become students, and grad students becoming faculty members – sometimes for only one quarter). In this common scenario, each of these relationships has a different set of services, or different levels of the same services, tied to it. The IAM process should be able to handle these situations as part of its core competency. Multiple focus groups asked for the ability to know about the multiple relationships of a person, to be able to manage access to services for all roles (not just a primary role), and to not have a change in one relationship affect services provided for the other role(s).

Information that is Globally Contextual

Growth in the international nature of our community also puts strains on basic data elements in our systems that are identity attributes. For instance, international students/employees often have multiple visas, and renewed passports sometimes get new numbers, but there is only one data element for a visa number (and the federal government only allows one within the U.S.). Similarly, personal names are different in number and order around the world, which can also create confusion and visa verification issues. These situations are at root "systems of record" data issues that then become IAM issues. They can cause confusion in international campuses such as Qatar, can cause confusion for international citizens at domestic campuses, and can lead to the mistaken issuance of multiple credentials or conjoined records, which combine two distinct students under one record.

Information about Group Membership

Access to most services, including basic communications, is based on memberships in different groups. The need for a finer granularity in available group information for provisioning and de-provisioning is important not only at the enterprise system level, but also at the local school level where many services are delivered and controlled. Basic role and organizational information (e.g., staff, tenure-track faculty, chemistry major, works within WCAS, member of Central HR, member of Medill IT, etc.) needs to be

generally available (i.e., not just in LDAP but in Active Directory too). The ability to create and manage many of these groups needs to be available locally, and in many cases, with the emergence of centrally-managed shared collaboration services, the data also needs to be stored or replicated in central directories.

To cite only a couple of examples from the focus groups, the Library needs access to this type of information all the time to create custom groups (e.g., all Music faculty, a professor and her graduate students, thesis and dissertation committees) for providing access to a custom set of library resources, and if WCAS had better access to this information, they could maintain email lists dynamically rather than manually, as they do now.

In order to provide this type of information, a new entity needs to be created. Today, the information that is available centrally is tied to the NetID and is a relatively narrow set of information contained in LDAP and (to a lesser extent) in the various Active Directory domains. The constraints in the IAM dynamic supported by this infrastructure exist on both sides of the IAM relationship. For instance, some LDAP attributes (e.g., 'mail', 'displayName', 'postalAddress', and 'eduPersonPrimaryAffiliation') are defined by the LDAP protocol specs as "single-value" attributes that force limiting the choice of available values to those attached to the primary role. There are other LDAP attributes that store values of other sources, but many applications can't be configured to look anywhere besides the default LDAP attribute, or cannot parse a list of multiple values and pick the one they want.

If we are to address the desires expressed in the focus groups for managing access to a complex set of services more appropriately, we are going to have to find a way to (1) provide access to this data centrally, and (2) refashion our applications so they can access it and utilize it to make authorization decisions.

To address these needs, a second architectural cornerstone is proposed: a common census of persons, accessible through a central registry:

IAM Architectural Cornerstone #2:
A central registry should be built to provide access to a more robust set of data (than is currently available via LDAP) about a broad spectrum of people with a relationship to the University (i.e., not just those with NetIDs). Each person's information should be tied to a unique identifier that is not an already existing University identity or identifier. Most of the data will be accessible virtually (rather than being replicated to a database).

This registry would be managed centrally as a service to all business functions. It would include all people who have relationships that the University wants to track (e.g., not just those with a NetID, not just alumni, not just students in degree programs, not just faculty and staff). The NetID – while still being the primary identity of the University – would become one of multiple identities and identifiers that could be attached to a person, all of which would be accessible via this registry. The registry would also make available a more robust set of data (described below) than is currently available via LDAP.

The following sections more fully describe the content of this registry and its technological basis.

Better Data

Generally, identity attributes are assigned to, or removed from, particular entities through three types of assertions:

1. authoritative system assertions;
2. distributed processes, which approve access to resources not administered centrally; or
3. formal administrative processes which approve specific access requests.

Currently, only the first type of assertions are available centrally, and there are relatively few of them (e.g., basic demographic information such as faculty, staff, Kellogg, Medill, etc.). The set of data that is collected in this registry needs to be expanded, including, for instance, the profile information detailed above (e.g., information on a person's multiple relationships with the University, both present and past).

Assertions of the second type are currently confined mostly to local Active Directory instances where they instantiate school or departmental group definitions and other attributes for organizing assets or controlling access to file systems. If access is to be provided more granularly by more applications, an expanded set of attributes about the person and his/her status needs to be readily available. For instance, having access to a more detailed set of organizational data about people becomes increasingly important as more shared services are deployed centrally. This is particularly true for collaboration services, which are deployed to smaller groups of people (e.g., Box.net and SharePoint), some of which are formed on an unpredictable, as-needed basis.

The third type of assertions are currently held separately within each system's security tables and are not visible outside those systems. If these assertions are instead promoted to what we'll call a "***secondary attribute***" within a central registry service, many business rules could be written for manipulating the access permissions and identity lifecycles of entities with access to these core systems. For instance, approved access for SES, FASIS, and NUFinancials could become attributes in this registry (e.g., with secondary attributes names of OK-SES, OK-FASIS, and OK-NUFIN). Today, such a status is known only to the owning system, lying buried in its internal security tables. Externalizing it (outside the owning system) as a flag within the identity management system allows other systems to incorporate it into their own business rules. For example, the identity management system could build logic into its system that says if either OK-SES or OK-FASIS is flagged, password lifetime could be shrunk from 365 days to 180 days, and if OK-NUFIN is checked, it would shrink to 90 days. Similarly, the identity management system could use this information to adjust rules on unsuccessful login's, varying the number of permitted unsuccessful attempts and the action taken once the threshold is reached (e.g., extending the wait period versus disabling the ID) depending on the levels of access associated with an ID. In these cases, the identity management system is changing its own behavior based upon permissions granted to the entity and reflected in the assigned attributes.

It should be noted that this central registry is not intended to archive permissions of people when they leave. While some participants in the focus groups wanted to be able to check a box that said "give this person the same permissions that the previous person had," the risk of misapplying old permissions has a potentially dangerous downside to it. The goal recommended here is to not do this. Instead, the recommended path is to set aside the NetID to reduce the number of mistakenly duplicated identities, and to help people pick up where they left off (e.g., a person who left comes back in the same role, or a new person is hired to replace them). The ability to provision permissions quickly (via online forms and workflows, preferably with options for permissions based on definitions of basic roles) should be emphasized, instead of trying to retain a person's permissions. Work of this type is beginning to happen around the NUFinancials access permission processes, and to the extent these roles are able to be defined and utilized, they should exist in the central registry, outside of the application that develops them, so they can be used by other purposes when they are relevant.

New Data Structures

To make the widest possible use of identity attributes – regardless of source –they should all be visible or obtainable through standard means. For traditional directory services, this means all attributes would have to be present within central directory services. The University's new web services infrastructure should make rethinking how data is made accessible and maintained in this new context a high priority. While some attribute data may continue to be replicated into directories from authoritative sources, this is not required. In fact, this repository could be an entirely virtual database containing links to multiple different sources of data (e.g., the central identity database managed by the identity system, the centrally stored organization and personal attributes, and even links to distributed local attributes or attributes collected specifically for

functional areas such as the researcher-specific information discussed in the working group paper on "A New Vision for Research Administrative Systems"). It could also be a portfolio of web services that can be used as building blocks as needed. The application making an access decision should not have to know where the authoritative data resides, and the data need not be replicated in this centralized registry (though the registry could be used to house additional sets of data (e.g., multiple visa numbers, if the data schemas in the authoritative systems can't incorporate these extensions).

The ability of schools and business units to create and maintain local identities on their own will need to be preserved, possibly through utilities such as Grouper, which could synchronize group memberships across different directories, databases and other repositories. (See also the discussion of SOA as an architectural cornerstone in the section on Integrating Identity and Access Management, page 24.)

A Different Focus for Identity Management

The central registry should include all people who have relationships with the University that should be tracked, not just those with NetIDs. By definition, then, the unique key that identifies people in this database will not be the NetID, though there will still be the premise that there will be no more than one NetID per person, just as there should be no more than one WildCARD barcode and one EMPLID per person.

There will also be other identities stored in this database (e.g., federated IDs from peer institutions or from consumer-oriented businesses such as Facebook, LinkedIn, Microsoft, Google, etc.). How these identities might be used within Northwestern systems is described below, but the point here is simply that one person will have multiple identities, each of which will be stored in this central registry. (See the discussion below on Optimizing Assurance and Trust, page 27.)

As part of creating this centralized common census, the tasks of avoiding duplicate identities will grow. More people will be included in this census, they will come from more sources, and it is possible that each person will have multiple identities attached to them. This means that the task of avoiding duplicate identities will also become larger and more complex, and the value of doing it well initially, and over the lifecycle of an identity, will become even greater than it is today.

NUValidate currently does this work when requests for NetIDs are forwarded on to it from authoritative systems. In the new model, the NetID will explicitly be just another credential attached to an identity, rather than implicitly assumed to be *the* identity. A NetID will still need to be managed throughout its lifecycle, but the identity system will need to be a repository where additional credentials are stored and managed as well (e.g., the WildCARD, multi-factor systems, biometrics, perhaps a digital signature, and registered third-party credentials). Some people in the registry will not have a NetID at all. Theoretically, a person could use different identities and credentials in instances where different levels of trustworthiness are required. Systems could theoretically allow access with a variety of credentials – some may restrict different types of access to different types of credentials, and some systems may choose to accept only certain types of credentials (e.g., only NetIDs).

Because of its increased complexity and importance, this census of persons should be managed more closely to avoid wasting labor resolving errors and questionable practices, and where possible, tools should be available to correct mistakes.

While all of it does not have to be built simultaneously (for instance, the management of identities could be built before the personal attributes parts), moving to this common census of people is obviously not a trivial task. Not only does the registry need to be conceptualized and built, the surrounding applications need to be retooled to make use of it. However, given the changes in the environment in which we live, we see this as a fundamental piece of the new architecture that will qualitatively change how the University's is able to do its business.

Integrating Identity and Access Management

IAM Characteristics addressed in this section:

3. Identities and access to resources are provisioned and de-provisioned rapidly in alignment with the need for their actual usage, with easily auditable trails.
4. Authorization is appropriately granular and based on robust identity information.
5. Surrounding business applications are integrated with the enterprise IdM system.
9. Business applications and the IAM infrastructure are flexible and easily modified to take advantage of new IAM technologies as they emerge and become stable.

The IAM "system" is a set of relationships between the IdM system at the core and the surrounding business systems that depend on it for authorization functionality and personal attribute data for access decisions. Just as the IdM system needs transforming within itself, the Access Management side of the relationship also needs transformation so that it has a tighter integration with the Identity Management system. This set of work is all about improving User Experience and maintaining security.

Provisioning and De-provisioning

Whether it is provisioning or de-provisioning access, the themes in the focus groups were the same: it needs to be faster, it needs to be more tied to the timing of the business needs for access to the online resources, and it needs to be more granular. (For overviews of the provisioning and de-provisioning processes, see Appendix E on page 73 of the unabridged version of the report.)

One of the issues related to the speed of provisioning and de-provisioning is that each piece of the IdM system, and each of the surrounding business applications that depend on the IdM system, are like independent islands, unaware of the Identity Management system other than a binary Y/N response to an authentication attempt, and these islands are connected only by periodic batch shipments of data to update the identity and personal attribute data they store independently. This builds delays into the process, resulting in lags in gaining access to resources such as library privileges, online course materials, or getting a WildCARD. These delays become more problematic when the normal processing cycles are compressed (e.g., for late admits right before school starts who need to review financial aid packages, make payments, and get access to course materials very quickly).

Sometimes, the lag between identity and access provisioning is due to more than the delays associated with batch data transfers. If people need non-automatic access permissions and the completion of training before they can do their jobs, the less automated the process, the longer it will take them to do their job. This could be due to not knowing what permissions they need to have, not knowing how to request the permissions or the training, having the processes based in paper without good windows into the status of the requests, etc. One group said it could take weeks for new employees to get properly positioned to do their jobs, and it could take weeks for a director-level person to get all permissions and training.

Speed is also a concern when identities and/or access need to be de-provisioned, particularly when access to financial systems or sensitive research is involved. Identities can linger today for what some would consider too long in order to accommodate the unknowns within the off-boarding business process. For example, even the current system could completely and automatically deactivate/de-provision NetIDs for terminated employees within a day or two. In fact, the SES system has a deactivation routine that runs nightly, removing SES access from any employee who has been terminated. But the process is intentionally not set up to move that quickly because employees who are actually leaving may have a lot of back-and-forth in their actual departure date, and annual or academic year appointments often have delays in getting renewed. To avoid the error of deactivating a NetID prematurely or mistakenly, the system is intentionally slow with email notifications built into it to accommodate these situations.

The situation is further complicated by the fact that in some situations leaving is not an all or nothing situation. For example, sponsored research sometimes needs to have the NetID stay on when a researcher leaves so the researcher may continue working on a project after leaving NU and moving to another institution, or because resources may be tied to the NetID that are most expediently preserved by also preserving the NetID.

This need is one reflection of the qualifying phrase within IAM Characteristic #3: "*rapidly in alignment with the need for their actual usage*." It's indicative of another recurring theme in the focus groups: the current process is not fine-tuned enough to provide access to resources for all entering or departing members of the community. Examples cited included students needing to pay bills, request transcripts, or see credit balances or account history long after they graduate. (One comment was that it was like we disown the students after they graduate.) Other examples were faculty needing access to Blackboard, research file shares, or research proposals prior to getting their NetIDs.

On the provisioning side, identity and personal attribute data is available via periodic batch data export/import sequences instead of via real-time processes. Applications often look internally for the data that has been stored from these export/import routines, and the applications do not handle any special-case access logic on their own.

When applications depend almost solely on the IdM system for a binary On/Off authentication as part of their authorization needs, access requirements in one system can conflict with requirements in another, and providing more granular access becomes problematic.

On the de-provisioning side, times when access permissions are not taken away when they should be because of loose connections between the applications and the IdM system. Take, for example, a staff member with access to a variety of systems. Currently, when she leaves the University, her NetID is turned off, and access to these systems is effectively cut off. But if she is also a student, even though she is no longer a University employee, her NetID remains active, and her access is dependent on a tighter coupling of the business systems with the IdM system and personal attribute data. Similarly, should a staff member leave the University and subsequently return in a different role, any lingering permissions from the first period of employment would again be accessible. (Note: this is also a potential liability of the current system of handling people with elevated privileges in SES, which cannot be tightly coupled with the IdM system because it doesn't use the NetID and password.)

Many focus groups expressed bewilderment at how these decisions are made and whether they are coordinated (e.g., needing to see final paycheck information conflicts with disabling the NetID). Confusion about, and misunderstanding of, policies surrounding authorization to functions and information can only be resolved by service providers taking control of those decisions based upon identity and relationship information – and then documenting them and publishing them for their user community. Multiple focus groups said that simply knowing what services get turned on and off, for whom, and at what point would be very useful knowledge to be available, not only for the people affected by these changes in service, but also for those people who administer the services or assist the affected parties. But this information is not centrally stored or tracked. Adoption of SOA (Service-Oriented Architecture) will allow us to build the foundation for this sort of information, but it is only a necessary step, not a sufficient one. Creating a University-wide service catalog, complete with eligibility rules, would be a very large project. A better start might be identifying a handful (6-12) of critical systems, then documenting and publicizing the rules they use.

Addressing the Need for Improved Access Granularity and Responsiveness to Business Needs

In order to optimize agility, responsiveness, and the user experience within our portfolio of online services and resources, the responsibility for making access decisions ("authorization") needs to be more fully shifted to the

applications, and the applications have to be less inwardly focused and less passive about making these decisions.

This, then, is the third cornerstone of the new IAM architecture being recommended:

> **IAM Architectural Cornerstone #3:**
> Authorization, the permission to access resources, needs to be handled by the surrounding business applications, not by the identity management system. Applications must become identity-aware pieces of an integrated portfolio, rather than heads-down, internally focused silos, and authorization should be flexible enough to open access to individual services as needed.

"Smarter Applications"

The shorthand for saying what the applications need to do is to say that they need to become "smarter." They need to look outside themselves for identity and personal attribute data at the moment when the access request is made, or, next best case, when personal attribute information changes for people who use that system. Either way, the application needs to be more identity aware, on a real-time basis, of the status and attributes of the people requesting its services instead of trying to stockpile more attributes internally.

One example of the latter case is the Library's turnstile system, which might want to store information about who is eligible for library access locally to guard against network outages. However, rather than stockpiling that information via overnight feeds, the information should be updated on a per-individual, real-time basis. In these cases, the application must be capable of realizing a change has occurred in a person's relationship with the University, and be capable of acting upon that change. The declaration of a change may come from the IAM system, but could also easily come from an authoritative system (FASIS, SES, etc.).

Additionally, special-case business logic needs to be distributed, not concentrated in the center. Larger systems will need to be able to better grant access to selected sets of their services at different points in time. Systems should be able to fully control access to their services without relying on the Active/Inactive status of the NetID, so the access needs within one system don't step on the needs to block access by another. (There will still be a need to have a "single choke point" in some situations, which will be provided by the usage of the web Single Sign-of system.)

Ultimately, we should get to the place where NetIDs are no longer turned on or off. Instead, they will be an enterprise credential, which a person does or does not have, and each person's access to online resources will be handled by the system that provides the resource, based on attributes that are available to it on a real-time basis.

The Role of Roles

When people talk about an ideal identity management system, they usually speak glowingly of a system predicated on roles (groups of attributes that, when taken together, lead to the provisioning of access to sets of resources). These roles are usually projected as enabling the automatic provisioning and de-provisioning of access in real-time. While developing roles is a key factor in the IAM environment that is envisioned here, its scope is smaller than what is often implied in casual conversations about IAM.

It is our belief that once one needs to define access permissions for many real-world resources, the need to parse the role of someone often quickly becomes quite complex, taking one outside of what is typically defined as a role, and into more complex "If… then … else" sets of logic. A recent example of this is related to understanding who needs what access to various RCR (Responsible Conduct for Research) training. A simplistic role-based perspective is anyone who is a "Researcher." However, the reality is that it is much more complicated. Even the role of "Researcher" who is also "NIH" or "NSF" would fail to be granular enough for the stated business requirements of the campus. Adding a role of "School" would still not resolve this. Different departments have their own RCR training requirements and may apply them to people in different

fashions (even going beyond granting agency requirements). In other words, this logic cannot be resolved simply by adding roles to the IAM system; it can only be resolved by building the business logic in the business system. The detail that is needed starts with commonly used roles, but to achieve the specificity that is often assumed in talk about a proper system having the roles to make these determinations, using this type of complex logic to create roles would result in an unmanageable number of very granular roles.

Similarly, to build the overarching roles that are often posited in these discussions, one would need consolidated documentation on the relationships people may have with the University, and the permissions associated with each of them. While it would be good to begin to consolidate this information, the cost/likelihood-of-success ratio associated with spending time up front trying to define and agree on consolidated roles is usually not favorable. In most cases, it is more important to have the required personal attributes available, to have the application do the "If… then … else" logic with the attribute values, and to look for similarities as these sets of logic get defined so smaller roles can be defined when appropriate, and then expanded as more appropriate use cases get identified.

This is not to say that the use of roles will, or should be, non-existent. Some roles can be built on higher-level attributes, and this has been done in several functional areas already. However, each of these successes has also realized that the 80-20 rule definitely applies in this area. (For example, in NUFinancials the majority of permission requests are for users needing low-risk access, while the remainder of their users needs a variety of access levels.) In some cases it might even be possible to tie job categories to roles. For instance, the role of a department assistant 2 might contain OK-NUFIN because all department assistants will use that system in their work. Where it is possible to automate the association of roles to job categories, automating the work process for granting permissions can be envisioned, with necessary checkpoints and approvals built into the work flow. In other words, role-based systems could spawn access permission request workflows that start with a default set of permissions, but they would not automatically grant access. The automation would aid by having standard sets of permissions to start with, rather than relying upon someone to launch a set of requests separately and from scratch.

Integrating Applications Better – SOA and SSO

SOA – Service-Oriented Architecture

If more is expected from the applications, they need to be better connected to the IdM system, largely through the utilization of real-time service calls that will be available through the enterprise Service-Oriented Architecture (**SOA**). The growth of our SOA architecture will not have a great impact on the actual authentication step. This will probably continue to be via an LDAP bind, Active Directory, or Web SSO (using LDAP as a back-end data store). But SOA will play a very important role in providing personal attribute information to the applications on a real-time basis and helping with the optimization of assurance (discussed in the Optimizing Levels of Assurance and Trust section that follows this one).

Therefore, the web services that are the basis of SOA are the fourth cornerstone of the new IAM architecture being recommended:

> **IAM Architectural Cornerstone #4:**
> Identity and access management processes need to be done online utilizing web services that provide real-time and workflow functionality based on data for individuals.

As just noted, these might be service calls at the time of the access request to verify the status and attributes of the person making the request, or they might be a service that "publishes" a change in status (e.g., that a staff member has moved from the Dean's office in McCormick to the Bursar's Office). And then, because a system has "subscribed" to all such publications, it is coupled with a corresponding service within the subscribing system that harvests the relevant information and updates the personal attribute data and/or security tables it stores internally.

Moving towards a web services architecture will enable very different processes for identity and access management, but the technology is only an enabler. Any sort of qualitative change is also very much about how business processes are defined and executed.

Migrating to integration via services is not only driven by the desire to become more real time; it is driven by the movement of applications to the cloud. (To mention just a few of these systems at Northwestern: admissions, student career services, student email and collaboration, survey tools, payment of student fees, alumni community, athletic ticketing, the library's central administrative system, software test environments, file storage/sharing, and employee healthcare management functions.)

Third-party cloud applications increasingly presume the existence of web services and API's (application programming interfaces) for connecting to the University's identity management system or updating data in the enterprise systems. When a vendor is equipped to use the identity services we have in place, deployment can be a matter of days. When we don't have what they use, or when they require a data feed just like our on-campus applications currently do, integration will be much more protracted.

Not only do we need to be emphasizing our ability to integrate via web services and APIs, we also need to be pushing our vendors to incorporate these standards and approaches. Just because an application is in the cloud, it doesn't mean that it uses the optimal solutions for data transfer or identity management. But these building blocks are what all optimized applications use for integration and data accessibility.

Single Sign-on (SSO)

As stated earlier, the NetID is the accepted standard for online identities (though there are notable exceptions to this rule, see page 11). Appreciation of the productivity benefits afforded by this standardization were repeated in multiple focus groups; however, it was always coupled with frustration because the expectation is not only that there will be a single electronic identity used for access to University services, but that this single identity should give access to an integrated suite of services via web Single Sign-on (SSO). This integration provides:

- An important convenience for users of Northwestern systems: no need to login more than once or twice per day in order to gain access to many systems.
- Additional authentication factors (smart phone, hardware token) can be integrated into the SSO system rather than each individual application, and they can be more finely tuned to specific services within applications when multi-factor authentication is done this way.
- A building block for expanding the functionality, and hence the value, of NUPortal.
- Two valuable security controls: 1) should there be a need to quickly suspend access across a wide set of systems, the SSO infrastructure provides a single "choke point" on an identity; and 2) applications using SSO don't need to hold, even for an instant, clear-text usernames (NetIDs) and passwords.

Not only should the applications change to be more integrated into the SSO infrastructure, the SSO environment needs to be enhanced to be more predictable and convenient. Entering or exiting from an application should not affect credentials already accepted by another application. If applications are not sensitive to the possibility that they might not be the only application in use during the session, then the local cookie holding the record of SSO authentication will be discarded when "logging out" of one of the applications, potentially causing unpredictable states with other applications (e.g., triggering a new authentication challenge). For instance, when a staff member logs into SES, it puts a cookie into her browser. When that staff member then logs into NUFinancials, the SSO system will see that cookie and will not require a new authentication. But if the staff member logs out of either application, how the other log-ins will be treated depends on how the cookie is treated. That is, do the other applications remain gracefully open, or does one or more of them get unceremoniously terminated? It needs to do the former rather than the latter.

Similarly, the installation process for applications using SSO should be simplified. At present, applications must install agent software directly into the web or application-server layer. This has proved difficult or impossible, sometimes for technical reasons and sometimes due to licensing restrictions. Setting up SSO on simple HTTP/S proxy servers is a likely solution, as is direct SAML integration with our federated authentication infrastructure. Other options should be investigated as well, and it would probably be good to increase in-house expertise on doing these integrations and/or look at budgeting for consulting assistance.

Eliminating Paper – Move Processing Online

Systems being able to get access to changes in status on a real-time basis is only one element of improving the speed of the IAM processes. On-boarding and off-boarding also involve people – people who make requests for access permissions, people who review and act upon those requests, and people who go through status changes (e.g., getting a new job, changing a job, ending a job) – and these people need a way to interact with these changes or requests for changes. The more these interactions are tied to paper, the longer they will take and the less transparent and interactive (i.e., the less predictable) they will be.

Moving these processes online immediately speeds them up and makes them available for everyone to see their status and raise a problem if one occurs (e.g., plans have changed and the separation date won't be until two weeks later). It also makes them reviewable to help improve processes, and it makes them easily auditable. Eliminating paper and moving all IAM processes online should be an overarching goal.

Similarly, multiple focus groups cited savings if there were electronic workflows where a person could review a document or a report of financial transactions and check a box as one does with University time sheets. ("If a Principal Investigator could electronically acknowledge their monthly budget statement has been reviewed and approved, it would save a staff person in every department lots of paper and storage.") The process of *paper -> pdf -> paper + signature -> pdf* is laborious and inefficient. Some 90% of the documents that go through this slow process are internally generated, and productivity could be increased by turning these into online forms with electronic approvals.

Note: This paper purposefully leaves aside any discussion of electronic signatures that are needed for legal documents and restricts itself to workflows and electronic "sign offs" that do not require this level of identity proofing. This restriction does not mean that pursuing this higher level of "signature" is not worthwhile. It simply means that this need was not a recurring topic in the focus groups, and it requires a significantly greater amount of effort and resources to deploy when there is already a large set of resource requirements to vet.

Auditability

The final part IAM Characteristic #3 -- "*easily auditable trails*" -- has not been a past requirement. However, with increasing security threats and regulations, the ability to audit both the management of a NetID and its use in accessing systems will become more important.

NUIT's Security and Compliance team receives requests regarding where and when a NetID was last used. The ability to log and review activities by a NetID is a compliance requirement of HIPAA/HITECH, and it is easy to see where this confidence may become necessary for certain security applications (e.g., laboratory entry, computer access to human-subject data, etc.).

Auditing the *use* of a NetID to access systems is most often implemented within the applications themselves. Auditing the *management* of a NetID ensures an ongoing level of confidence that the real person remains in control of the associated credentials and the credentials are not being used maliciously. For example, Financial Operations wants to minimize the chance that a single person can create and approve transactions. One of the ways creation and approval could be done within the current IdM structure is that a NetID administrator could create a manual NetID and use it to approve transactions entered by him with his real NetID. University auditors have shown increasing interest in the manual NetID process, and one result has been that every six

months, NUIT contacts a designated person in each school or department asking them to review all of the school's NetID administrators (i.e., anyone who can create NetIDs or reset passwords). The designated reviewer must contact NUIT and indicate that each administrator still requires those privileges. If this attestation is not done, the privileges are automatically removed.

Optimizing Levels of Assurance and Trust

IAM Characteristics addressed in this section:

5. Surrounding business applications are integrated with the enterprise IdM system.
6. The level of rigor employed in identity proofing and authentication at the time of access is based on the risk and value of the transactions to be done

In today's world, one size cannot fit all when it comes to security. The need for a credential such as the Manual NetID is testimony to that fact. However, as everything changes -- the nature of our world, our business, our Northwestern community, and the technology mediating all of it -- we need to provide greater security for our systems while also easing access to resources and services. This section is about the fifth cornerstone of the new IAM architecture being recommended:

> **IAM Architectural Cornerstone #5:**
> The Northwestern NetID will remain the core Northwestern electronic credential, but its role needs to be supplemented by external credentials and by other means of insuring identity trust and assurance.

Assurance and Trust Overview

Granting access to resources results from the successful completion of two steps: authentication and authorization.

- *Authentication* answers the question: "Who is this entity and how confident are we that this is the exact entity we believe it to be?"
- *Authorization* then answers the question: "Given the answer to the authentication question, and any other information available about the entity, what functions and data items should be made available for this entity's use?"

To adequately protect the access gateway, the process of answering the authentication question needs to optimize, in relation to the value of the resource to which access is being requested, the extent to which it can establish two types of certainty about the identity: the level of "assurance" that is attached to the identity, and the level of "trustworthiness" that is attached to the credential presented for authentication.

- "*Assurance*" refers to our level of confidence that our electronic identity of a person is accurately associated with a real person and the correct person.
- "*Trust*" refers to our level of confidence that the person offering the credentials for authentication is actually the person to whom the credentials were issued.

The more valuable the resource, the higher the levels of assurance and trust should be; and the higher these levels need to be, the more difficult it should be to get authenticated. Levels of assurance rise along with the number of independent attestations that are presented, and the basis for those attestations. For instance, having someone show up in person and present multiple government-issued photo IDs provides a much higher level of assurance than receiving a request for access based on the self-service online entry of a Facebook or Google credential. Levels of trust also rise along with the type and number of credentials being presented.

Let's take a hypothetical example from Northwestern. Professor Jones has the following credentials:

1. A WildCARD issued through standard processes
2. A NetID issued through standard processes
3. A second-factor password generator (e.g., a smart phone number for a phone with a multi-factor app)
4. A fingerprint scan, taken by a Northwestern office, and kept on record
5. A Gmail account

Professor Jones has the following access needs:

a. Buy NU athletic event tickets
b. Read email
c. Approve payroll timecards for staff
d. Purchase hazardous chemicals for research from a grant
e. Enter into laboratory space where hazardous materials are used

Each of the applications (associated with *a – e* above) must decide which of the credentials (*1 – 5* above) they will require to be confident Professor Jones is, in fact, making the request. The decision will rest upon a balance between convenience and security/compliance. For example, *b* and *c* above require a NetID (*2*). The level of assurance and trust provided by a NetIDs is balanced with the value of the resource being accessed. Because NU Athletics (*a* above) wants to serve the public, a NetID is out of the question for being their default credential. Instead they might allow Professor Jones to create an account within their system using Gmail proxy authentication (*5*) rather than issuing a new credential (NetID or otherwise) to her. On the other end of the continuum, a NetID or a WildCARD by itself may not be appropriate for either of the latter two cases (*d* and *e*). It may make sense for *d* to require a NetID and a second-factor authentication process, and *e* might require a combination of swiping a physical WildCARD and having one's fingerprint matched with it.

Each of these credentials has a different level of trust, and combining them in different ways can increase the trust in a particular authentication event. It is vital that assurance be high when issuing what will be considered high-trust credentials. For example, it makes no sense to use biometric security if you are unsure if the person involved is really the individual who is supposed to be granted access at so high a level of trust.

Identity assurance ranges from very low (self-attestation), to moderate (one or more third-party attestations), to high (photo IDs, biometric cross-check with government databases, background checks). There may be applications that will not want to provide services to an entity if the assurance is only at a minimum level.

Increasing Trust

In these times of phishing, social engineering to steal credentials, and increased concern about compliance and the security of PII and PHI, it is important to improve our ability to increase the level of trust in certain authentication situations. As noted above, there are multiple ways this can be done utilizing bio-metric methods of authentication and multi-factor methods of authentication (MFA). Both of these authentication methods have improved and have become easier to deploy, and with the near ubiquity of cell phones these days, purchasing separate key fobs to provide a physical authentication factor is no longer necessary. Interest in MFA was a repeated topic in the focus groups, and this interest has coalesced into a pilot project on multi-factor authentication using a product called DUO.

Just as a one-size-fits-all approach does not work with NetIDs, adding in these levels of trust needs to be done selectively, adhering to the principle of needing to optimize levels of assurance and trust in relation to the value of the resource. Interestingly, utilizing some of these higher trust authentication methods can also decrease the difficulty of authenticating. For instance, some universities are now using bio-metric scanning devices for residence halls, dining halls, laundry rooms, etc., eliminating the need to carry a physical credential such as the WildCARD, and also speeding up the authentication process.

Multi-factor authentication may also help in the transition away from an authentication dependence on physical location. Many of our most sensitive resources require a physical presence on campus (as attested to by the use of a network IP address that is a Northwestern IP address). If one is not physically on campus, then these resources can only be accessed via a more secure virtual private network (VPN), which was frequently mentioned as an irritant to members of the community and a burden to support teams. While auditors love the requirement that access to these systems is restricted by physical proximity, we should seek to reduce our dependence on physical location alone and move towards methods of authenticating that are person- and risk-based (e.g., logins from unusual locations or at unusual times may require additional means of authentication).

The interest in MFA and biometrics reflects an interest in insuring confidence in the identity at the time of request for access. There is also growing interest in insuring confidence in identities themselves during their lifetimes. Identity assurance (the confidence that a person is who he or she claims to be) has become increasingly important to research and academic functions. This has been driven by: the growth in regulations surrounding student records, personal information, and health information; the heightened sensitivity of protecting valuable research data; and the growing role of online education. Granting agencies are beginning to require improvements in this area, and external auditors are paying more attention.

At Northwestern, employees who go through the basic hiring process have fairly robust identity vetting as part of the federal I-9 employment eligibility process, as do TGS students in order to receive stipend funds. Students may never present a photo ID between the time they take the ACT/SAT and the time they pick up an official transcript at graduation. Manually asserted NetIDs are created with little or no identity vetting.

Reducing our Dependence on the NetID

We need to strengthen our authentication processes for some resources, supplementing the levels of trust that accompany a NetID. In other situations we need to reduce our reliance on the NetID.

For people who are less directly connected into the daily fabric of the core Northwestern community, being forced to rely on a NetID for access to Northwestern resources – getting one, remembering it and its associated password, and then dealing with on-campus support if problems arise with it – often reduces their productivity and positive attitude about their Northwestern experience. It also increases support loads for Northwestern administrators. There are two ways the reliance on NetIDs can be reduced in appropriate situations: moving the control of access management closer to the resources being controlled, and using identities offered by other entities. Each is described below.

Moving the Control of Access Management Closer to the Resources being Controlled

In some instances, requiring owners of resources to use a centrally-vetted and controlled identity to permit access to their own resources is counterproductive to the shared goal of protecting University resources. File sharing is a classic example of this with the rise of Dropbox and the availability of free gigabytes of storage from whichever major cloud vendor (Apple, Google, Microsoft, Amazon) one prefers. Without a well-developed sense of risk, and a commitment to support University requirements that institutional information be managed from an institutional perspective (institutionally accessible, secure, backed up, etc.), gaps in functionality between these consumer-oriented tools and tools offered by the University will often lead Northwestern faculty and staff to choose personal productivity over the risk of compromised confidentiality.

Unfortunately, in an age of increasing collaboration beyond the traditional on-campus boundaries of the University community – in research, community involvement/projects, experiential learning, and work with consultants – limiting University file sharing tools to NetIDs often drives people to these less secure, but highly available alternatives. In some situations (e.g., research projects with regulated data) it may well be appropriate to continue to require this set up, and to continue to educate people about why there are these expectations. However, in many other relatively similar situations, it may be more appropriate to delegate the

control of access management to the person who controls the resources, and rely on their direct connections to their collaborators to supply the needed levels of assurance and trust.

Suppose, for instance, Professor Jones wants to share her research files with a co-researcher at the University of Arkansas and a former student now in the private sector, and they want to use either their work email or a Gmail account as a credential. Unlike an anonymous relation handled by a centralized service, Professor Jones has a close working relationship with them, and the odds that the email address is not used by that person, or that the person is not really the person they say they are, are quite low. And when the basis of authentication needs to be changed (say, the colleague gets a new job at the University of Wisconsin) or the need to collaborate ends, who is going to know quicker, or have more of an incentive to update the authorized IDs, than Professor Jones? In effect this ongoing, personal relationship gives a higher level of assurance to this external identity.

These type of situations are similar to the idea behind Manual NetIDs – delegating control over identities beyond the processes attached to the core systems of record. However, the above scenario is fundamentally different in several ways:

- The people in the assertion process are closely linked instead of being anonymous to one another.
- The process is handled directly by the person who has a vested interest in, and control over, the resources.
- The identity to which access is given is already very familiar to the external person, and often it is one that is also known to the person controlling the access.

The same concept, but taken to a further extreme, underlies the deployment of a guest wireless network at the University. Prior to this implementation, anyone using the University's wireless network had to have a NetID. This always caused problems for people temporarily on campus (e.g., contractors, consultants, guest lecturers, recruiters, parents, prospective students, or visitors). With the growth in mobile devices (laptops, smartphones, tablets) and the spread of wireless networks everywhere else, the problems only became worse. Now, a separate guest network (separate from the regular University wireless network, which is still NetID authenticated) is available, with only a self-reported name and email address required for access. This approach has worked great with parents of students and prospective students who are visiting campus, with guest lecturers, recruiters, and contractors, and for on-campus conference and meeting attendees. In the past, all used to need a NetID to get wireless reception on campus.

In each of these cases, the NetID, with its centralized control, is no longer being utilized as the credential, and the levels of assurance and trust attached to the credentials have been adjusted, though not eliminated, to balance risk with trust and ease of access.

Identity Federation

Another way to reduce our dependence on the NetID is also built around using a credential with which the external person is already familiar: extending Northwestern University's identity management system via federation. When institutions federate their IdM systems, members of one institution can use their own institutional credentials for accessing services in another institution. This enables people who have a more removed connection to our community to use a credential they are already familiar with to access Northwestern resources. And vice versa: when Northwestern federates with other institutions, members of Northwestern can use their NetID to use services in the federated institution. It also reduces labor required to grant and support appropriate access to University assets, enabling quick integration of people and systems.

The wave of the future is being able to connect to and integrate third-party systems hosted in the cloud by someone else. We need to be able to connect to these seamlessly and have our community use their NetIDs to authenticate. In the last year, Shibboleth, the University's primary federation application, was brought to the most current release, and about 40 cloud-hosted applications are now being accessed with NetID and

password. Examples include: TeraGrid (research computing), Student Conduct (Student Affairs), CareerCat (Career Services), Qualtrics (surveys for Feinberg, Weinberg), Orbitz (Travel Services), Primo, Illiad and Ares (University Library), and the Canvas pilot (Provost, NUIT). Planned deployments include the Box.net file sharing system, several University Library systems and purchasing from SciQuest via NU Financials.

More remains to be done:

- There are no well documented best practices to follow, and we do not prioritize this capability when new applications are being vetted. This needs to become a standard operational process on all fronts.
- Our federation infrastructure handles our current load, but should be made fault-tolerant with ample capacity to increase its load.
- As emerging standard protocols gain traction in the federation world, we need to be adopting them to improve our ability to federate. (See Section IV on page 26 of the unabridged report for a discussion of standards within Identity Management.)
- Medical campus partners must be incorporated into the overall access and authorization plan. Integrated MS Active Directory services appear to be the best approach for direct collaborations between Feinberg, the hospitals and medical practices.
- Research collaborators at other institutions should be supported through InCommon or equivalent services. While the University is a member of the Internet2 InCommon federation and has external partners with whom this service allows convenient NetID authentication, the InCommon technology is a niche higher-education solution that the University cannot assume will be implemented by potential partners.

While this wave is swelling, we also need to be able to authenticate using Northwestern credentials for all sizes of third-party solution providers. Some can do federation via Shibboleth, but many just want to authenticate straight to our AD / LDAP. We need to work to provide a straightforward set of tools that can be used securely across all of these situations.

Federation via Social Identities

Federation is also the concept behind using credentials from consumer-oriented vendors (e.g., Facebook, Google, LinkedIn, Twitter, Microsoft). These identities have much lower levels of assurance and trust: getting one only requires self-assertion of who you are, and these identities are recycled. Clearly they are not suitable for all IAM functions at the University. At the same time, it's a credential that is already very familiar to the person needing access, and for situations where there is less need for higher levels of assurance and trust, these can be useful (particularly when they are used in a context where the owner of the resources would know that the social identity is actually connected to the person being granted access). Higher levels of assurance could be attained by business processes (e.g., in-person or online identity proofing) to match an external credential to a real person. (Online identity proofing could be done. Banks already do this via links to public records databases: show 5 driver's license numbers, addresses, or car make/models, and then have the person pick the one that is actually theirs.)

Federation via social identities is already being leveraged for alumni using OurNorthwestern. Northwestern alumni used to be able to access the online alumni community application only via a Northwestern-issued alumni account. This was secure, but it was difficult for alumni to remember their ID and password, which created a barrier to participation and placed a support burden on the University alumni staff. By contrast, the new OurNorthwestern alumni application has an Identity Provider module that allows alumni to log in with their Facebook or Google account. This removes participation barriers and allows the University to get updated personal information stored in those external applications, all while lowering staff support burdens. In order to supplement assurance, OurNorthwestern requires the person to answer some basic questions about their relationship to Northwestern (questions to which we already know the answers) before allowing information to be accessed.

This approach could be used in multiple other situations around the University where the services are less sensitive and the users are more removed from the everyday life on campus. For example, other schools are using this approach to provide access for library patrons (e.g., the person who comes in off the street to read *The Chicago Tribune* for two hours, or the graduate student who needs to come for a week to research African art). They are also using it for parents to look at their children's grades and financial records, for people taking non-credit courses, for students sharing their portfolios with potential employers and friends outside the University, for recommendation letters, and for practitioners involved in experiential learning situations.

The idea being outlined here is not that core constituencies of the Northwestern campus – faculty, staff, and students – will begin to have the choice of using their NetID or a consumer credential to access University resources. Rather, it is that for specific combinations of constituencies and resources, using these external credentials can be a better fit for balancing risk with ease of use. In each situation where this alternative is considered, business owners will need to weigh the complexity of University-issued credentials (e.g., NetID) versus the need to maintain direct, institutionally-structured control of the identity process. Where the associated risk is acceptable, savings in labor, increased success of online service utilization, and increased constituency goodwill may be substantial.

Recording and Using Levels of Assurance

As we start diversifying our means of identity management, and improving our access management capabilities, we need to make the levels of assurance and trust that are attached to a credential available for use within the authorization process. For instance, we will be increasing our instances of federation, and the institutions with which we federate will undoubtedly have varying degrees of assurance and trust associated with their credentials.

We've already seen that, for instance, not all Google IDs are asserted with the same levels of trust or assurance. Some credentials are simply self-reported (the guest wireless network), some are submitted by a person who has answered multiple attribute questions about themselves for which we already know the answers (OurNorthwestern), and some might be submitted by a person who is applying in person to use a University resource (such as checking out a book) and could easily show a current driver's license.

Similarly, it may be that one will be able to have a NetID with varying levels of trust. For instance, it may be that NetIDs get issued initially with lower levels of assurance and trust as a default, which get raised after their credentials have better levels of attestation. Or, it may be that a person is traveling when their NetID password expires, and they are not able to renew it in person but they need to check their email. It's possible that the NetID could be reactivated with a lower level of trust that would allow them to check their email but not allow them to approve their direct reports' timesheets. However, for this to work, applications need to do real-time access provisioning that uses real-time identity attributes in this environment (assurance and trust levels) rather than relying on internal security tables with data lags built into their maintenance.

To be effective and efficient, the Northwestern IAM ecosystem must gracefully support this range of methods and situations. As access decision-making diversifies, the applications making these decisions will need to examine real-time assurance and trust levels, and wherever there are secure application requirements, those must trump convenient credentialing with lower levels of trust and/or assurance.

Maintaining a Secure Environment

IAM Characteristics addressed in this section:

7. Identities are protected and secure.

The University has an obligation to protect individual identities for legal and regulatory reasons (e.g., HIPAA, FERPA) as well as strategic reasons (e.g., institutional reputation, faculty/student/staff recruiting, avoiding

fines). In addition, the protection of individual identities also protects the institution: compromised individual credentials can lead to further breaches and unauthorized data disclosures.

Protecting identities, and the credentials associated with them, has to be pursued on many fronts.

- Northwestern's IAM systems must not only safeguard information they store internally, but must also provide data to other systems in a way that enables them to make appropriate access decisions.
- Sensitive information should be sparingly replicated. To this end, local and cloud-based applications that access data in directory-type repositories (AD, LDAP) or via federation protocols (Shibboleth/SAML) go through an approval process whereby data stewards explicitly approve each such data release. Additionally, IAM systems do not use credit card or other financial information at all, and the IAM ecosystem makes use of SSNs only in matching data between FASIS and SES, stores them in encrypted form in its internal database, and does not provision them into directories where they might be visible to applications or end users.

  However, most of Northwestern's enterprise and other systems still rely on bulk data replication. Even aggregation of data into AD or LDAP directories represents added risk compared with simply querying the source (authoritative) system for that data when it's needed. Moving to real-time access to information only when it is needed offers a much smaller target for intentional misuse or accidental disclosure versus making multiple copies of the data.
- Our IAM infrastructure and most other applications already use standard cryptographic techniques to protect data in transit (SSL/HTTPS, SSH/SFTP).
- Strong password management policies contribute to identity security by making passwords more difficult to guess or discover by other means. Our current policies regarding minimum password length, frequency of change and complexity are roughly in line with peers, but are weaker than desired for high-value or high-risk transactions conducted with HIPAA and other sensitive data. Security experts are also increasingly skeptical whether *any* password scheme can, by itself, provide adequate security.

  Use of the Web SSO system allows secure authentication of the user without the added risk of exposing credentials directly to applications during authentication operations.

  Multi-factor authentication (MFA) protects identities by requiring physical possession of a device in addition to knowing the NetID password. Broad adoption of this technology would reduce our vulnerability to phishing and other attacks involving the compromise of passwords. It will assist in improving the security around sensitive data and in meeting some compliance requirements (e.g., HIPAA). Ideally, MFA should be integrated with the SSO environment, and/or directly with applications, so the applications can decide whether to allow certain transactions based on their confidence (or lack thereof) in the strength and validity of the original authentication.
- The requirement that people have a NetID to access University resources is a "heavy" process, with requirements to be physically present to get identities and credentials changed, the need to go through a limited number of entry points to get into the system, etc. These requirements are complicated by the changing nature of the Northwestern community, where many of the people who comprise the core constituencies of the community – faculty, students, and staff – are increasingly dispersed geographically. As this trend continues, the mechanisms for maintaining security around our identities need to change as well.
- As more data becomes available to more partners, existing tools such as the Service Provider Security Assessment should continue to be used with all vendors hosting our data in the cloud. Processes and

procedures developed for data access via the LDAP Registry, and those being developed for the emerging SOA infrastructure, should be consolidated, reviewed, updated as needed, and widely publicized. Existing data categorizations should also be reviewed and updated as needed to aid in making decisions about what data may be stored in the cloud.

Similar policies and procedures should be developed for the use of identity assurance and trust. General guidelines are needed for assessing the risk level of various types of transactions. We also need to decide on a reasonable number of levels of assurance and trust (perhaps as simple as low/medium/high), then match each transaction type to the appropriate level of assurance and trust. Ultimately, those levels will be matched to technology, policies, and procedures in the IAM systems. Individual application and data owners will then have to determine their own specific policies within those guidelines, and ensure that they are enforced by their applications – in part by leveraging data made available via the IAM ecosystem.

- Some investigation of InCommon certification has already been done and needs to continue in light of discussions related to assurance possibly being required for access to granting agency systems. (For discussions of InCommon and its assurance levels, see the sections in the unabridged report that begin on pages 27, 31, and 50.)

IV. Next Steps

So often, the problems with "identity management" get flagged in conversations with the connotation that these shortcomings are purely technology issues, and sometimes there is the implication that the topic is the sole responsibility of NUIT. To the contrary, one of the major intents of this report is stating that improving IAM at Northwestern will take the coordinated effort of the entire, and very broadly defined, IT@NU community.

Certainly, much of this work is technical, and the effort not be underestimated, but it is important to recognize that a significant part of this work is outside the purview of NUIT. Whether it is consolidating Active Directory domains, or reworking the applications so they integrate with the identity management systems (e.g., processing identity data and information via services instead of via batch files, making their own more finely-tuned authorization decisions based on real-time data held elsewhere), much technical work will be required outside of NUIT.

This technical work also presupposes that existing business processes and needs have been documented and compared, and new processes, definitions, and policies have been articulated. These business analysis efforts, which must accompany the technological changes, are central to the implementation of the new architecture being envisioned in this report. Undoubtedly they are complicated and time-consuming and probably more challenging than the technological challenges that await.

The preceding pages have argued that the University's current identity and access management system is insufficient for its current and future needs. The coming changes in the number and range of persons to be served, the complexity of relationships in universities with very high research activity, the compliance requirements within academic medical centers, and the revolution in the means to deliver solutions are examples of the dimensions unanticipated by the current ecosystem that cannot be solved by changes in a single system. The entire complex of identity sources, identity and credential management, and applications themselves needs to be substantially redesigned and re-implemented to shore up the foundation of all of the University's online resources and services.

As just described, the new, restructured IAM system envisioned in this report rests on five architectural cornerstones:

5. **A consolidated Identity Management System** – The identity management system needs to be consolidated at the center with delegated administrative functionality.

6. **Central Registry** – A central registry should be built to provide access to a more robust set of data (than is currently available via LDAP) about a broad spectrum of people with a relationship to the University (i.e., not just those with NetIDs). Each person's information should be tied to a unique identifier that is not an already existing University identity or identifier. Most of the data will be accessible virtually (rather than being replicated to a database).

7. **"Smarter," more Identity-aware Applications** – Authorization, the permission to access resources, needs to be handled by the surrounding business applications, not by the identity management system. Applications must become identity-aware pieces of an integrated portfolio, rather than heads-down, internally focused silos, and authorization should be flexible enough to open access to individual services as needed.

8. **Online Processes** – Identity and access management processes need to be done online utilizing web services that provide real-time and workflow functionality based on data for individuals.

9. **Northwestern NetID Supplements** – The Northwestern NetID will remain the core Northwestern electronic credential, but its role needs to be supplemented by external credentials and by other means of insuring identity trust and assurance.

Included in the preceding pages is a very large set of work. More than could be accomplished in one year even if all other work was halted. We will need to discuss and prioritize what comes first and what is delayed.

Projects to be Considered Initially

The matrix that concludes this report on page 37hows two sets of work.

Work related to the Replacement of NUValidate

The first set of work consists of topics that are intertwined enough with NUValidate that they require a degree of envisioning before a replacement product is chosen. This set of work should be prioritized because of NUValidate's product end-of-life status. While the risk associated with its status is viewed as low to medium low (see page 5), the need to replace it is time insensitive.

The actual replacement of the system is expected to take one to two years. There are key discussions, envisioning, and planning to be done in addition to that time due to the many facets of the Identity Management System that are entwined with NUValidate. We cannot wait until everything is perfectly figured out before proceeding with choosing a replacement. However, we also cannot choose a replacement without having a sense of where we're going on some key issues.

The first three pieces of work in this set are relatively straightforward. The ability to do two of them (Wildcard consolidation and Online Directory restructuring) is, however, very much dependent on the ability of groups outside of NUIT to contribute to the effort. The last four pieces, bracketed by bold lines within the matrix, are less formed, and require broader discussions and input.

In some cases, the envisioning in this initial pass through these seven pieces of work may be short. For example, if the degree of difficulty in moving to a new and improved model is large, the resources required will be difficult to obtain, and the impact on the replacement project could be contained, then the envisioning could be limited to making a relatively quick decision that no work will be done until after the NUValidate replacement is in place. For most of these topics, however, more depth of analysis will be required. (The matrix gives orders of magnitude on their size (both envisioning and implementation) and their relationship to the project.) Nevertheless, the estimate is that with the right mix of people – combining business knowledge

with technological understanding – a six month window could provide the needed input for the NUValidate replacement product selection process to proceed.

Other IAM Work

Despite the small size of the NUIT Identity Services team, and their limited capacity for new undertakings, we do not think that all IAM work must stop while this envisioning takes place. The second set of projects are also connected to the Identity Services team, but they differ from the first set in that they are more amenable to being largely delegated to consultants/contractors with oversight by the Identity Services team. To the extent that this is possible, the projects related to the replacement of NUValidate do not take a higher priority than this second set of projects. However, we need to be diligent about moving forward with the NUValidate replacement, and not let it linger behind the scenes.

The second set of projects contains work that is either already underway, addresses core applications within the IAM experience at Northwestern, or is in high visibility business areas that have key IAM needs. There are many other projects on the list of IAM work that are not included in either of these lists. Sometimes that is because they are second-order tasks. In other situations it is because the work required is only tangentially related to the NUIT pinch point, and prioritizing them would suppose a broader discussion within the business units or IT Government advisory committees. The projects listed in this second set are being recommended as the leading candidates for the additional bandwidth of the Identity Services team that might be available for IAM projects.

A More Comprehensive Listing of Work

These projects discussed above are a small subset of the more comprehensive list of work included in Appendix C of the unabridged report that starts on page 61. This listing of work summarizes the work that is either explicitly recommended, or implied in the pages of this report. It is divided into three organizing concepts: Integrating Identity Management, Integrating Identity and Access Management, Optimizing Levels of Assurance and Trust. No attempt has been made to attach beginning or ending dates to these tasks at this point. Those will come out of the discussions that follow the release of this report.

Initial Recommendation for Governance

This initiative will have ties to both the Administrative Systems Advisory Committee (ASAC) and the Infrastructure Advisory Committee (IAC). Much of the business-related discussions will happen within ASAC, but the IAC will be the best body to advise on the technology decisions. Because the IAC has all of the school IT leaders in it, it will be a good body to supplement the perspective and input of the ASAC school representatives.

The initiative will be overseen by one person, who will be supported by NUIT personnel to assist with managing and keeping the work on track.

To begin, we are recommending a steering group for the initiative, similar to the group that is guiding the SOA Initiative, comprised of 6-10 representatives from schools, business units, and NUIT. It will need to have business perspectives and technological familiarity represented within its members.

Envisioning the Prioritization of Initial IAM Work

	Outsource-able?	Topic	Involved Parties (in addition to NUIT)	The Business Value attached to changing the way this work is currently done	Impact on Selection Process if the decision is "Just do what you're doing now."	Must Do? (yes, no, maybe)	Envisioning Effort (s,m,l)	Impact on NUV Replacement Selection Process (s,m,l)	Impact on NUV Replacement work (s,m,l)
Tied to Replacing NUValidate	N	Email Provisioning	NUIT only	Makes IdM more flexible - easier and safer to modify.	Moderate, will need to make sure new system is flexible enough to allow us to translate the business logic from current IdM.	Y	S	M	M
	N	Online Directory	UR, HR, Registrar	Makes IdM more flexible - easier and safer to modify.	Minimal, most vendors can handle this level of back-end complexity and front-end UI just fine.	Y	S	S	S
	N	Wildcard/IdM Consolidation	Univ Svcs, FASIS, SES, downstream WildCARD users	Reduces need to maintain two parallel systems. Reduces chances for errors to surface. More consistent data across systems that rely on IdM and Wildcard.	Zero, if consolidation is not planned before the predicted end of life of the new IdM system.	N	M/L	M	L
	N	Manual NetIDs	FASIS, FSM, SCS, ...	Eliminates one of the main sources of duplicate NetIDs, and reduces complexity of overall system. Addresses (or at least shifts) many audit-related concerns over separation of duties and adequate operational contols in financial and other systems.	Moderate, need to ensure that new system has flexible way to delegate appropriate admin privileges and has a good UI for doing this.	Y	M/L	M	S / L
	N	Group Management	Schools	The shortcomings in being able to manage groups both ad hoc and "automatically" has become even more of an issue as more collaboration services have been offered centrally.	Moderate to Large - the current system has many shortcomings, both behind the scenes (what it can do) and in front (bad UI, clumsy & slow tools). Replacing IdM without at least marginal improvements here will seem like a step backwards to many people.	Y	M/L	L	L
	N	Central Registry	Schools, Enterprise apps	Key request from multiple groups: make more data (e.g., historical relationships, multiple roles, organizational attributes, group membership) available centrally.	Negligible. If we never plan to do this work, we don't need to prioritize features to support it when selecting a new system.	Y	M/L	M	L
	N	Data Privacy Considerations with LDAP vs. AD	ASAC, IMC, OGC?	Active Directory has fewer privacy options on attributes than LDAP does. Need to know how to balance functionality and privacy.	Minimal, most vendors can handle the type of AD & LDAP provisioning we do now, and anything we might choose to do in the future, with no problem.	Y	M/L	S	S/M
Selected Other Priorities	Y	Create SOA services	SES, FASIS, NUL, ...	Fundamental building block of new architecture.	NA				
	Y	SSO	SES, FASIS, schools, ...	this is a key building block for user experience improvements (fewer logins and ability to integrate an app into the NUPortal) and for deployment of MFA on an enterprise basis.	NA				
	Y	MFA/SSO Integration	FASIS,SES	Allows more granular application of Multi-Factor Authentication, and allows it to scale to an enterprise level.	NA				
	Y	Social/SAML IdP	ARD, NUL, TGS, IAC, ...	Allows for the use of social identities for more "peripheral" constituencies. Key functionality for the Library's Alma project and desired by multiple other places in the University.	NA				
	Y	Federation enhancement	IAC	Changes the setting up federation between systems to more of a routine operations process from an exception. Makes the infrastructure more robust.	NA				
	?	Pressing School-based IAM needs	Relevant schools	For instance, improving the state of IAM is a key priority for Northwestern Medicine.	NA				

Appendix A - Quick Reference Guide to the IAM at Northwestern Report

There are eight main "applications" that work together to comprise the IAM "system at Northwestern:

1. **a core Identity Management (IdM) system** (NUValidate), which stores identities based on NetIDs that are in turn based on data fed primarily from authoritative identity sources such as the Faculty and Staff Information System (FASIS) and the Student Enterprise System (SES), allows people to manage those identities, and updates Northwestern's identity directories;
2. **identity directories** (e.g., LDAP, Active Directory, and Kerberos), which surrounding business applications use to authenticate users requesting access to their system;
3. **a physical identity system** (the WildCARD system), which provides proof of identity for access to buildings, events, etc.;
4. **a directory synchronization utility** (Radiant Logic), which keeps data in multiple active directory domains synchronized;
5. **a web Single Sign-on system** (SSO), which reduces the need to keep logging in with the same credentials for each Northwestern University application that is used;
6. **federation services** (e.g., Shibboleth), which allow people at trusted affiliate, partner, or peer institutions to use their home institution's credentials to gain access to Northwestern systems and services;
7. **a multi-factor authentication service**, which provides an extra layer of password protection using an application on a registered smart phone or answering a phone call to reduce the risk that personal information can be easily compromised should someone learn a NetID password;
8. **an "Identity Provider" bridge service** (currently being run by the Alumni and Development Enterprise Applications team for the OurNorthwestern system), which enables alumni to log in with either an active Northwestern identity or with one of their own external social accounts (Gmail, Yahoo, Microsoft).

See the section on "IAM in Action" in Appendix D of the unabridged report (page70) for a diagram and description of how these parts work together to provide IAM functionality when a person tries to log in to a Northwestern application. Appendix D also has an annotated diagram that shows how data flows within the IAM system.

If these pieces comprised a highly-functioning IAM system, they would be typified by the following nine characteristics:

1. Each person has a single electronic identity. There may be multiple credentials attached to that identity, but there is only one electronic identity.
2. The IdM infrastructure is integrated within itself, so that data about identities and personal attributes flows smoothly throughout the system.
3. Identities and access to resources are provisioned and de-provisioned rapidly in alignment with the need for their actual usage, with easily auditable trails.
4. Authorization is appropriately granular and based on robust identity information.
5. Surrounding business applications are integrated with the enterprise IdM system.

6. The level of rigor employed in identity proofing and authentication at the time of access is based on the risk and value of the transactions to be done.

7. Identities are protected and secure.

8. Each part of the IAM system is relatively easy to maintain and to replace.

9. Business applications and the IAM infrastructure are flexible and easily modified to take advantage of new IAM technologies as they emerge and become stable.

Northwestern's IAM system has grown organically over the past twenty years without benefit of an overarching architectural strategy. Work needs to be in three main areas, based around five architectural cornerstones:

1. Identity Management (IdM) needs to be restructured, reducing its complexity, better integrating its data flow, and making more identity and personal attribute information available.

IAM Architectural Cornerstone #1: The identity management system needs to be consolidated at the center with delegated administrative functionality.

IAM Architectural Cornerstone #2: A central registry should be built to provide access to a more robust set of data (than is currently available via LDAP) about a broad spectrum of people with a relationship to the University (i.e., not just those with NetIDs). Each person's information should be tied to a unique identifier that is not an already existing University identity or identifier. Most of the data will be accessible virtually (rather than being replicated to a database).

Six areas of consolidation are discussed as ways to improve data flow, User Experience, and resource efficiency:

7. Instead of system-specific credentials, use the NetID as the University authentication credential for system access when a University credential is called for
8. Reduce the use of Manual NetIDs
9. Merge the provisioning of WildCARDS in with the core IdM processes they now mirror
10. Consolidate AD domains, but provide distributed ability to create and manage groups
11. Centralize access to personal attributes that are now isolated in systems and local AD domains
12. Centralize access to identity credentials and identifiers that are now

Notable associated sets of work:

- Envision a new IdM model so that a replacement for NUValidate can be chosen, including:
    - reviews of current and desired functionality associated with manual NetIDs, group management, central registry, and data privacy in access directories
    - conceptualizing the relationship between the central registry and the NUValidate replacement
- Build web services to handle Identity Management
- Replace NUValidate

- Exploring the consolidation of Active Directory domains
- Explore the consolidation of the provisioning process for WildiCARD into the IdM system's

2. Identity Management and Access Management need to be better integrated.

IAM Architectural Cornerstone #3: Authorization, the permission to access resources, needs to be handled by the surrounding business applications, not by the identity management system. Applications must become identity-aware pieces of an integrated portfolio, rather than heads-down, internally focused silos, and authorization should be flexible enough to open access to individual services as needed.

IAM Architectural Cornerstone #4: Identity and access management processes need to be done online utilizing web services that provide real-time and workflow functionality based on data for individuals..

Notable associated sets of work:

- Utilization of web services to make access-related tasks occur in real-time
- Enterprise commitment to web Single Sign-on
- Move identity and access processes online and eliminate paper
- Promote the movement to online processes (e.g., make status attributes available centrally, outside of systems; investigate the usage of roles at the University already; investigate the connection between access permission requests and the provision of training)
- Conceptualize the personal attribute side of the central registry
- Conceptualize a new methodology for applications to make use of more robust data in order to make more informed and, where needed, more granular authorization decisions ("smarter, more identity-aware" applications)

3. The IAM system needs to optimize its ability to leverage Assurance and Trust levels attached to an identity and its credential.

IAM Architectural Cornerstone #5: The Northwestern NetID will remain the core Northwestern electronic credential, but its role needs to be supplemented by external credentials and by other means of insuring identity trust and assurance.

Notable associated sets of work:

- Identify areas where levels of Assurance and Trust need to be optimized
- Deploy Multi-factor Authentication
- Improve our capacity and ease of using Identity Federation
- Expand the usage of Social Identities where appropriate for constituencies less tightly connected to the University
- Associate Levels of Assurance and Trust with electronic identities, and have applications utilize them when making authorization decisions about access to resources and services
- Reduce the need to have a person be physically present in order to change their credentials with appropriate levels of Assurance and Trust

Authentication Overview

Area	Today	Proposed
Main Authentication Mechanism	– LDAP or AD	– LDAP or AD
Web Single Sign-On	– Limited Deployment	– University Standard

Federation	– Limited usage – Exception rather than standard operational process – Shibboleth only	– Standard operational process – Robust infrastructure – Other protocols, e.g., OAuth?
Federation via Social Identities	– OurNorthwestern only	– Multiple cases across the University for constituencies with more limited/removed relationships with the University
Role of NetID	– The Northwestern credential and primary electronic identity – Gateway to all Northwestern online resources for all constituencies (except for OurNorthwestern)	– The Northwestern credential. Gateway to Northwestern resources for core constituencies – Supplemented by credentials from other institutions and by social identities.
Single Northwestern Identity	– NetID comes closest and is widely adopted by applications – Fragmented multiple credentials and identifiers	– NetID becomes "just" a credential – Single identifier gets created, to which all credentials (e.g. NetID, WildCARD barcode, EMPLID, LinkedIn ID, Facebook ID) get tied
Attribute Data	– Replicated in LDAP – Some isolated in distributed Active Directories	– Available centrally – Mostly available via services or virtually
Assurance and Trust Levels	– Not utilized (e.g. no difference between Manual NetIDs and regular NetIDs) – NetIDs are either active or not	– Tied to each credential – Can be changed as a result of a change in status or a verification process (e.g. physically showing up and presenting proof of identity) – Applications should use in determining access to their resources

Authorization Overview

Area	Today	Proposed
Most common Application Relationship to IdM for Authorization	– Authentication = Authorization + Internal Security Table	– Applications should be "Identity aware"
Timing of Personal Status and Attribute data Updates	– Usually overnight	– Real-time
Available attribute data	– Relatively limited	– Wide-ranging (e.g. multiple roles, history, externalized role/permission variables, functional area specific)
Authorization Granularity	– All or nothing based on coarsely-defined roles stored within applications	– Basic lower-level roles defined and stored external to applications – Applications should be capable of refined granularity.