Identity and Access Management – a primer

Students, faculty and staff require a digital identity to gain access to UBC resources. Frequently, controlling access to each new resource or system has meant another username and password. For any one person, this collection of usernames and passwords can represent a daunting list to remember. Controlling access to each new resource also meant the resource administrators must know which students, faculty and staff members were allowed to access the resource. Although stored in computer systems, **access lists are often manually maintained and can quickly go out of date**.

The need for better security

When does access stop?

Granting access to authorized users can be easy to do – a person requesting access to a system will have various credentials or the approval of authorized individuals. A more difficult challenge is determining when to **revoke** access to systems because the person's status has changed – this is the **deprovisioning challenge.**

Security of passwords

Periodic changing of passwords increases overall security. Password changes ensure that should the password in one system be compromised, this same password cannot be used to compromise other systems. Recently Sony revealed that millions of their customers' passwords were stolen. For users not regularly changing passwords on other systems, these passwords can now be used compromise other systems.

Changing passwords

Passwords that are longer, case-sensitive, mixed alpha, numeric and special characters are inherently more difficult to guess and thus more secure. Regular changes of these passwords increases security, however for those individuals using many systems with many different password rules and change frequency, this can be very unwieldy.

The need for better usability

Unfriendly password changes

UBC's policy is that passwords must be changed periodically. Users are often reluctant to change a password since it means they must also change passwords in every other UBC system to avoid having a different password in every system to remember. This not inconsequential inconvenience makes password changes difficult.

Unfriendly personal identity management

Many systems store personal identity information such as names, phone numbers, personal email addresses, and so on. Every person therefore has multiple identity profiles but it is not clear what the authoritative identity source is. Even if system owners can designate an authoritative source, it is unlikely that the authoritative source is used in all other systems. A simple example for our staff is his/her phone number which may be stored in the on-line directory, the email system, and in the HRMS system.

UBC's iDentity and Access Management Program (iDAMP)

Simple goals

The challenges describe above can be addressed through 3 components:

- Improving our security through automated deprovisoining
- Reducing the number of usernames and passwords
- Single authoritative identity repository, with the UBC master identity first

What needs to be done to meet these goals

Deprovisioning

Automated deprovisioning of user in a system requires:

- There is a universal and consistent method of identifying a user in multiple systems
- The entitlement status of a user is known
- The system is configured to take action based on this uniquely identified individual, using his/her current entitlement status

Usernames/passwords

Reducing the number of usernames and passwords for an individual requires:

- For each system, using a universal identity (or shielding the user from the system's underlying identifier)
- For each system, using the universal identity password (or synchronizing the system's password underlying password change mechanism)

Single authoritative identity repository

Creating a *single authoritative identity repository* that is consumed by all systems requires:

- There is single identity for a person that is owned at the UBC level. This identity is independent of the person's past, current, or future roles.
- This single identity is consumed by any system that currently has, or will require, identity data.
- The single identity is created before any other electronic record.

iDAMP work streams and milestones

There are 6 general work streams. These are:

- Security risk awareness (for system owners)
- IT communications and change management
- Facilitating identity system changes for system owners and IT system owners:
    - Using the UBC master identifier as the primary user ID
    - Defining deprovisioning business criteria
    - Implementation of deprovisioning business criteria using dynamic groups within iDAMP
    - Using these dynamic group memberships to determine system access rights (entitlements)
    - Administrative simplification business rules and automation process
    - Username and password unification (e.g., password synchronization)
- Facilitating access data system changes for system owners and IT system owners:
    - Defining access data required, and its authoritative (external) source
    - Defining access data integrity requirements, and potential automations

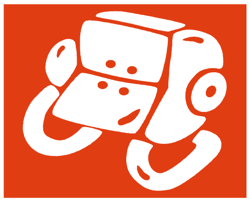

- Single, authoritative identity repository
    - System owner and IT system owner awareness of concept
    - Identification of key systems containing identity data
    - Multi-year identity data management approach, moving from multiple, fragmented identity sources to single authoritative identity source
    - Creation of a the single authoritative identity source
    - Use of the single authoritative identity source by system owners and IT system owners
- iDAMP technology deployments to support the above
    - Identity- and access-enabled Active Directory (EAD), Open LDAP (eLDAP), Crowd, Shibboleth, custom connectors
    - Enterprise group management tools

Additional information iDAMP work streams and milestones can be found in Identity and Access Management – a primer for system owners.

Identity and Access Management – additional information for service owners

Security risks and usability

The need for better security is described in Identity and Access Management – a primer. To deal with these issues, there are 3 goals to be met:

- Improving our security through automated deprovisoining
- Reducing the number of usernames and passwords for any one user
- Single authoritative identity repository, with the UBC master identity first

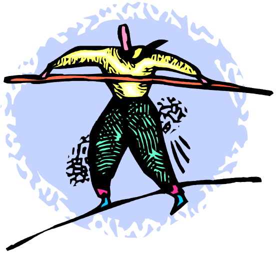

Automated deprovisioning

Automated deprovisioning is a simple concept where a change in a user's entitlements or access rights triggers a deprovisioning action in a system. For example, if employees are entitled to an Exchange account, an employee's resignation would trigger automated deprovisioning of the associated Exchange mailbox. To accomplish this, iDAMP supplies user entitlement information to handshake with the system doing the automated deprovisioning. System/service owners need to consider the following:

- *Entitlements*. To automate entitlements and thus deprovisioning, there must be a **precise functional definition of the candidate group**. For example, employees may not be a sufficient definition if there are others such as contractors, vendors, volunteers, individuals accessing before an appointment, persons requiring access after the appointment end date, non-payroll employees, retirees, etc. This is the entitlement definition phase.
- *Matching the identity authentication source*. Identities in the system must **precisely match those within the enterprise identity database.** This is the data reconciliation phase.

After entitlements definitions and data reconciliations have been completed by the system owner, the necessary technologies are implemented and the supporting processes are designed and rolled out.

Service onboarding consists of:

- Entitlement definitions
- Data reconciliation
- Technology implementation
- Process roll out

Common misconceptions

- *Choosing EAD enables deprovisioning*. Use of EAD alone does **not** enable deprovisoining. Applications wishing to use EAD must also make use of the appropriate enterprise-enabled group within Active Directory to enable deprovisioning.
- *Choosing Shibboleth enables deprovisioning*. As with EAD, above, Shibboleth-enabled applications **must** also appropriate use of the correct enterprise-enabled Shibboleth attribute.

- *Departments or system owners do not need to define entitlements*. A precise functional definition by the system owner of what constitutes an entitlement, and what are the business conditions to revoke access is required.

What about automated provisioning?

Automated provisioning helps both improve data integrity and reduce administration. It is not, however, strictly speaking, required to address security risks. The good news is that iDAMP fully supports automated provisioning, which simply stated, is the automated instantiation of a new user in the target system as a result of an authenticated and authorized enterprise user requesting access to the target system.

Many systems have closed architectures which do not allow “external” systems like iDAMP to create users. If your system has this capability, however, automated provisioning should be considered.

Usernames and passwords

Simplification of usernames and passwords for a user involves using enterprise authentication rather than system-specific username and password management. Any system that is Shibboleth, EAD or CWL-enabled allows the user to manage his/her password in one place.

System owners must recognize that user authentication, even if through enterprise-identity enabled services such as Shibboleth or EAD, does **not** in itself enable automated deprovisioning.

Single authoritative identity repository

Although there are many usable sources of identity information, there is no one authoritative source. Teaching Assistants, for example, are in both the SIS and HRMS systems. Identity data such as phone numbers, email addresses, physical address, nicknames, departmental affiliations and so forth is scattered about in these and many other systems. Quality varies widely, and is dependent upon both the process to capture the information, and the time period it was captured. For our TA example, we could decide that one system trumps the other, however this does not tell us which system has the most accurate information. In addition, it is not easy for either system owner or the TA herself to determine identity data quality.

The solution to this situation is to have a single authoritative source of identity information. **iDAMP is this authority**. Systems requiring identity data should source it, or plan to source it, from iDAMP. Currently iDAMP acts as a broker for many sources of identity data; gradually over the next 1-2 years, each data element will cease to be a *brokered* authoritative source and will transition to iDAMP as *the* authoritative source. System owners need to consider the following:

- *Source identity data from iDAMP*. Your system plans should incorporate accepting inbound identity data from iDAMP. This includes but is not limited to preferred first name, preferred last name, legal first name, legal last name, mailing address, residential address, personal emails, personal phone numbers, business phone numbers, emergency contact information, birthdate, SIN, etc.

- *User self serve*. Your system plans should consider that the iDAMP identity repository is based on the premise of self-serve. Users, for example, will be able to change their preferred name, make changes to their person email accounts, person phone numbers, and so forth. This will mean that your system must accept inbound changes as they occur. Some identity-specific capabilities existing with iDAMP, however many functions including self-managed user service requests are expected to be done with IT Service Management's Service-Now suite.
- *Identity data definitions*. iDAMP defines identity data precisely. For example, we are currently defining what constitutes a legal name and whether this is an appropriate label. Common usage at UBC is presently that it is a "short" legal name, i.e., not the full name that appears on a birth certificate but sufficient for contractual purposes.
- *LoA and single account type*. iDAMP will abandon the concept of [CWL] account types. Thus over time system owners should cease using the basic, guest, student or employee account tags. Accounts (or more precisely, the person) will be tagged with a level-of-assurance (LoA).
- *Level-of-assurance.* LoA is our certainty that the individual's identity data is what is claimed. Self-created accounts have a low LoA. However validation of government-issued picture identification by UBC staff will increase the LoA. System owners presently gathering identity data, or system owners that in some way define a level-of-assurance should plan on formalizing their role in defining a user's LoA. This applies to systems like SIS, the carding office, HRMS, etc.

Who is responsible for identity data?

iDAMP does not decide which department or function at UBC has responsibility for identity data. We do, however seek sufficient clarity on data stewardship to allow those authorized to manage data. For example, our current situation is that employee data is managed by Human Resources, and most student data by Enrolment Services. In our future state with a single identity repository, management of student employee identity data requires consistent business rules for both HR and ES. Continuing with this example, as these departments and staff are authorized to manage data, they are also responsible for the corresponding data quality.

Identity and Access Management – additional information for technology owners

IT professionals wishing to incorporate iDAMP principles into their planning should read Identity and Access Management – a primer and Identity and Access Management – additional information for service owners prior to this document.

Identity data

What does your system need (targets)?

Prior to discussing which iDAMP technology is appropriate, system owners must determine whether identity and access **data** is required for those operating the system, or for those who are using system services, or both. Our shorthand for the former is **operators** and for those consuming system services, **customers**.

What are your identity management goals?

- *Operators*. System owners should plan for Automated deprovisioning as described above. This is because the incremental technical effort to move from simple authentication to full automated deprovisoining is usually very low and the security risk reduction is very high. If you feel you are not yet ready for deprovisioning, you should likely defer your integration.
- *Customers*. System owners should define what authorization attributes they require to grant access to their system. In some cases there may be no authorization attributes, i.e., all authenticated users are allowed access, such as in the case of self-serve help desk. In other cases, authorization attributes such as course enrolment are needed to make access determinations such as after-hours building access.

In-band versus out-of-band

Each system requiring identity and access data has different capabilities. Identity and access management goals can completely or near-completely be accomplished with Microsoft-based technologies tightly coupled to Active Directory, or open-source system that have out-of-box SAML 2.0 support. This is **in-band** integration, and when combined with Integrations –technology selection and deprovisioning can result in a strong solution. Older systems may have only rudimentary awareness of Active Directory, or lack an ability to integrate out-of-box with Shibboleth. For these reasons, and because there are special cases based on the particular business needs, **out-of-band integration** may also be required.

Technology

What is available

iDAMP supports various authentication, authorization and Directory Services. These include Enterprise Active Directory (EAD), eLDAP (future), Crowd, Shibboleth and custom interfaces. Use of an iDAMP directory service alone does **not** enable deprovisioning and does **not** enable customer authorizations.

Directory service scope

Figure 1 shows the universe of UBC identities.

- Each directory service is a **subset** of the master identity store
- Each directory service receives identity data **after** the creation or update in the master identity store

An important consideration for Figure 1 is that the master identity store exists today. Other directory services are not yet fully populated today, and even once fully populated may not meet your identity requirements, or may have highly restricted access based on privacy legislation. For example, when evaluating whether to use EAD, there are a number of planning considerations:

- Users are in EAD solely based on their entitlement(s).
  - Do your users require systems access before/after the entitlements?
  - Have you considered situations where the person has multiple roles? (Student + employee)
- EAD is user-provisioned. Entitled users must explicitly provision this service, otherwise these users will not appear in the directory.
  - How does this affect your user community? For students? For employees?
  - How does this affect administration?
- Students may not be visible in the directory except to authorized UBC employees per privacy legislation.
  - Have you considered privacy legislation in your EAD and related systems planning?
  - Do hidden EAD entries affect your application?

One other consideration for is that there are presently just a handful of EAD users; many employees, faculty and staff are represented, and few students. The planned target of 62k users may take time to achieve.

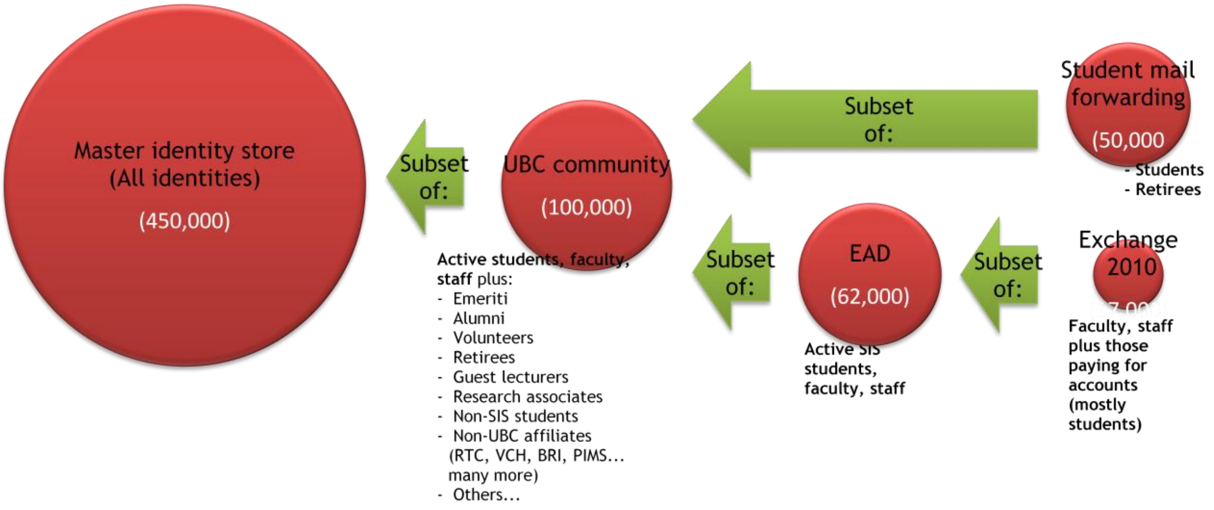

Figure 1

Integrations –technology selection and deprovisioning

System owners should confirm they are deprovisioning as part of their integration. System owners should ensure that Entitlements have been clearly defined as a lack of these *functional* definitions will greatly impede the *technology* implementation.

Be sure to plan your functional definitions **before** you plan your technology implementation. If you do not have your entitlements defined, there are effectively no deprovisioning business rules. You will then be forced to abandon your integration, or limit this integration to authentication only.

Entitlements

System owners must establish the scope and criteria for entitlements. Prior to the technical implementation this means:

- *Business rules* must be documented for any user that may be entitled. This includes all exception or edge cases, including emergencies identity additions, extensions past employee termination dates, access requirements prior to appointments, and so forth.
- *Access management processes* must be defined. If system accesses today are handled via a manual process with an administrative gate, these processes must be redefined to allow entitlements to be automated to enable deprovisioning.
- *Side effects* must be analyzed. System owners must analyze side effects to introducing identities particularly when they are visible in a directory service. There are different implications, say, to adding a member that is visible in a directory service (eLDAP, EAD) versus supplying attributes on an individual, as-authenticated basis (Shibboleth).
- *Anonymization.* Careful consideration should be given on how your integration complies with privacy legislation, and in particular FIPPA legislation. Identity information should be provided only on a need-to-know, as-requested basis. In many cases, identifiers must be anonymized (a unique identifier just for the requesting system) to ensure privacy is maintained.

Technology choices for operators and customers

Operators

Technology is selected for *operators* based on the *directory scope* and your *application constraints*. In many cases, applications may be tightly integrated with Microsoft's Active Directory, narrowing technology choices to Enterprise Active Directory (EAD) or a custom interface.

Customers

Technology is selected for *customers* based on *directory scope*, but rarely on application constraints. This is because most customer-side identities may be unsuitable for inclusion in institutional directories. Institutional directories generally are **not** suited to prospective students, retirees, past employees, community users, research associates, and so forth. Because of this, technologies which can draw on the complete master identity store such as Shibboleth are the best choice.

The technology implementation should include authentication and authorization.

Implementation summary

Integrations should use the table below to plan their integrations. Application *customers* should minimally be NZ; Application *operators* should be, at a minimum, NZD. However for both situations, your goal should be NZPD.

	eLDAP	Crowd	EAD	Shibboleth	Custom
Customers	*NZPD*	*NZPD*	*NZPD*	*NZPD*	*NZPD*
Operators	*NZPD*	*NZPD*	*NZPD*	*NZPD*	*NZPD*

N = Authe**N**tication
Z = Authori**Z**ation
P = **P**rovisioning
D = **D**eprovisioning

Where does CWL fit into this picture?

CWL is actually many system components. It includes:

- A legacy identity database schema
- A user management framework (UMF)
- A user, management and administrative interface
- An authentication engine (Auth2)
- The data itself
- Service entitlement management

Here is where we are heading with each component:

Legacy identity database schema

iDAMP uses the legacy identity database as the most complete, but flawed, identity store. iDAMP has created a parallel IAM schema that will gradually supersede the legacy CWL schema. This change **does not impact system owners and integrations**, but is tightly bound to the design and development of the identity repository and corresponding UIs.

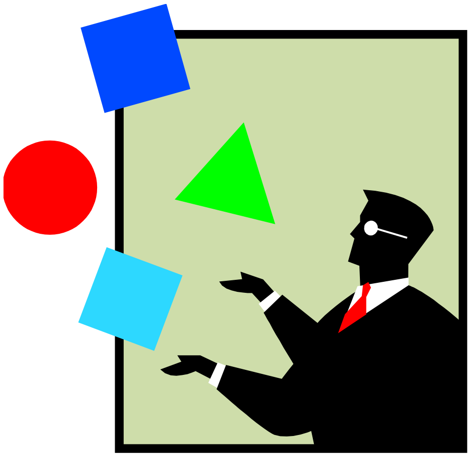

CWL user management framework

The user management framework (UMF) consists of the CWL account creation, account type changes, username recovery, password recovery, secret question management, password complexity logic and password policy change enforcement. The goal of iDAMP is to minimize the number of changes to the UMF. Any changes to the UMF do not affect system owners, but may have an impact on end-users and related support functions.

CWL user interface, management interface, administrative interface

Again with the goal of minimizing changes to CWL, changes to the various UIs will be largely aimed at ensuring higher quality user login ID data, and for urgent changes to support other iDAMP initiatives.

Auth2

The authentication engine of CWL is Auth2. Auth2 is a sunset technology in that system owners should plan for using one of the supported iDAMP authentication technologies, and that no new Auth2 integrations will be done. System owners should consider the following:

- *Authentication conversion and deprovisioning*. Your planning for Auth2 retirement should consider:
  - Replacing Auth2 with a supported iDAMP authentication technology, likely Shibboleth
  - Defining what authorization data elements that your system uses currently or requires in the future
  - Enabling deprovisioning through the use of SAML 2.0 attribute(s)
- *Write-back.* Your application may return data to CWL. If so, system owners should specifically call these out when planning their Auth2 retirement.

CWL data

Users may have more than one CWL login ID, thus there are often many CWL login IDs related to a single individual. System owners need to plan for the following:

- *Designation of a primary CWL log ID*. All identity-enabled services must link to one and only one CWL login ID.
- *Tying your system IDs to each primary CWL login ID*. Particularly for systems that are not CWL-authenticated, system owners should work towards mapping local system identifiers to the user's primary CWL login ID. For example, if you plan to migrate from your local Active Directory to EAD, you should prepare a mapping table of your local AD sAMAccount IDs to each user's primary CWL login ID.

CWL service entitlement management

CWL myAccounts currently has a limited capability to do user self management of provisioning. For example, a user can provision his/her EAD or Exchange 2010 account with the myAccounts self-provisioning function once an entitlement has been granted in within the iDAMP group management tool, Grouper.

iDAMP plans to retire this capability (other than identity-specific User-managed entitlements) and pass data to the IT Service Management tool.

Single authoritative identity repository (SAIDr)

The functional capabilities of the Single authoritative identity repository are described above. Because there is now a single repository of identity attributes, there must be methods to:

- Manage identity data
- Identity data exposure

Identity data management – "person hub"

The Single authoritative identity repository conceptually has three classes of data:

- Self serve, or user managed identity data
- User-managed entitlements
- Restricted, or administratively controlled data

User-managed data

This is data where the user can make changes at will. An example here is the user's preferred first name. If the legal first name is Robert, then the user can change the preferred name, say, to Bob. iDAMP can optionally add user policy compliance assertions to any user-managed data field.

User-managed entitlements

Most entitlements are managed within the IT Service Management function, however a limited number of entitlements directly related to identity are managed within the identity repository. For example, email aliases are managed within SAIDr to allow users to select valid aliases, and to allow them to tie user-defined aliases to target mailboxes.

Restricted or administratively controlled data

Some identity data elements such as legal last name are currently not changeable by an end-user. In some situations, the user must present proof of name change for an authorized UBC individual to effect the change. Within SAIDr, this is controlled by designating fields as restricted or administratively controlled. (A future enhancement may be to allow the user to change restricted fields subject to approvals embedded in workflow.)

Identity data exposure

Current state

Identity and access data is exposed through a number of technologies including database views, Perl-based ETL scripts, Forefront Identity Manager and other tools. Near-term transport of identity data will continued to be done using these technologies.

Future-state architecture

A number of different technologies are being considered for publish/subscribe architecture and an enterprise service bus. See Figure 2.

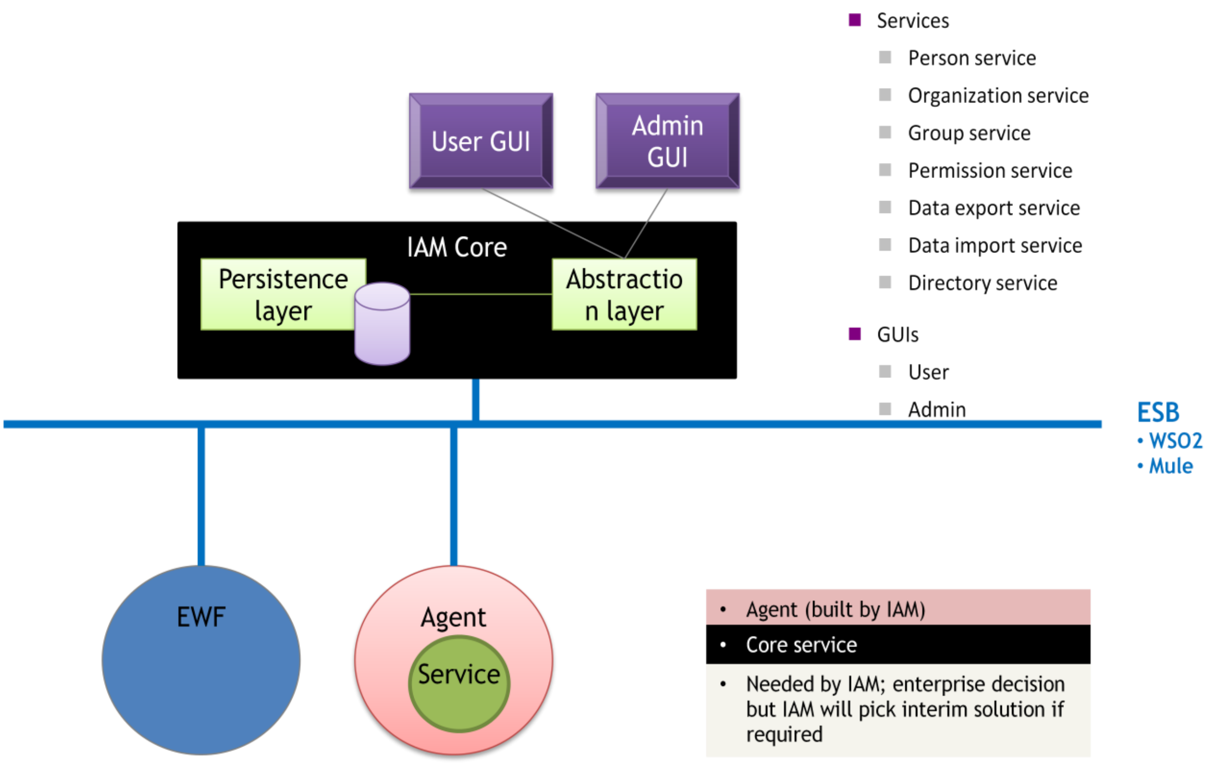

Figure 2

IT Service Management

An underlying assumption of iDAMP is that service entitlements will be managed through the IT service management tool, Service-Now. This will replace the CWL service entitlement management capability described above. This requires:

- Service information pushed to Service-Now
- Person-based entitlements pushed to Service-Now
- Revocations of entitlements pushed to Service-Now
- Service requests pushed to iDAMP
- Service enablement acknowledgements pushed to Service-Now.

Additionally, Service-Now will require both Operators and Customers to be integrated to this as an identity-enabled service.