Archived User Provisioning High-Level Design

Archived User Provisioning High-Level Design

This document provides a high-level description of a UCTrust-based infrastructure to support user provisioning for inter-campus applications within the University of California. This infrastructure represents an extension to the existing Shibboleth-based UCTrust infrastructure to address use cases, such as those described in User Provisioning Use Cases.

For the purposes of this document, user provisioning is defined to be the processes, both human and automated, that authorize (and de-authorize) people to use application systems, when those processes occur at times other than the start of an online session. This is distinguished from application systems that use a "pure" single sign-on infrastructure (*e.g.*, Shibboleth), authorizing anyone with a defined set of attributes that are provided at the start of a session.

The infrastructure described in this document will support the exchange of identity information from campus Identity and Access Management (IAM) systems to application systems, not the entire set of provisioning processes. The Roles and Responsibilities section below describes where those other provisioning processes should be implemented.

While UCTrust is the first intercampus use of middleware in the University of California, this project is UC's first use of middleware as an application development paradigm. The infrastructure described is specific to the exchange of identity information for user provisioning. It does, however, embody many aspects of a more general-purpose infrastructure for data interchange among arbitrary systems that should be useful in the future.

Principles and Assumptions

- Campus identity and access management systems and the organizations that operate them are authoritative for information about the members of their respective communities. The same campus organization that currently operates Shibboleth will be the organization that operates the infrastructure described in this document.
    - (Note that much of the IAM's information will likely be aggregated from other systems of record on the campus.; Nevertheless, UCTrust designates the IAM as the authoritative contact for its campus.)
- This framework provides a common mechanism for application systems to obtain identity information from campus IAM systems. Merging the results from multiple IAM systems, however, is left to the application.
- The existing UCTrust agreements, policies, processes, and technology should be leveraged as much as possible. All participating campuses have implemented UCTrust and are operating a current version of Shibboleth.
- The design and implementation must make effective use of University resources. Where possible implementations should be shared and/or reused. Deployment plans should accommodate differing priorities and schedules at different campuses, allowing for inter-campus collaboration and partial implementations at each campus until the entire infrastructure is deployed.
    - This effective use of University resources extends beyond this project, in particular by being the first UC-wide deployment of common middleware that can be used by other projects in the future.

High-Level Design

The following diagram illustrates the high-level design of this infrastructure for two applications that retrieve identity information from four campuses.

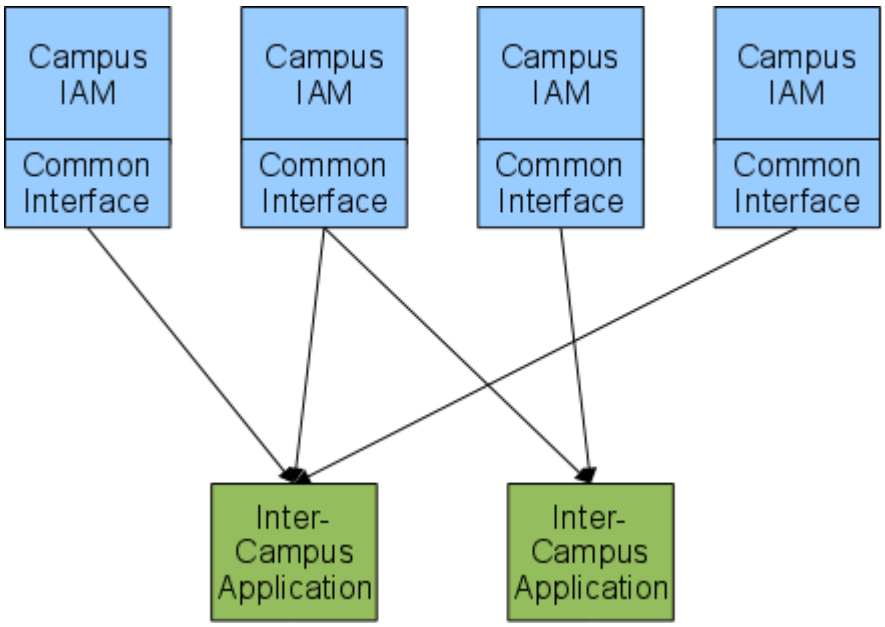

Just as with Shibboleth in UCTrust,

- Inter-campus applications obtain identity information about their users from IAMs through the use of standard network protocols and formats.
- All IAMs and inter-campus applications have unique names, called entityIDs, that are the same as those assigned for Shibboleth IdPs ("Identity Providers") and SPs ("Service Providers"), respectively.
- IAMs control the release of information to service providers through the use of *Attribute Release Policies*, which specify which identity attributes should be released to an application. In the case of user provisioning, however, the application's SP name will also determine the users for which the IAM will release those attributes.
- Software will be provided, written in Java, for integration into each campus IAM to implement the standard protocols and formats.

The following types of access will be supported. Other than SSO Event (Shibboleth), they will be supported by the Common Interface:

- **Snapshot**. All identity information allowed by the attribute release policy will be transmitted to the application.
- **Subscription.** Identity information will be transmitted to the application as add, delete, and update transactions on an event-driven basis. The transactions sent will be those that have occurred (or will occur) since the last Snapshot, Subscription, or Change Log access.
- **Change Log.** All add, delete, and update transactions that have been generated since the last Snapshot, Subscription, or Change Log access will be transmitted.
- **SSO Event.** Identity information about the current user is transmitted at the start of a session. This is the existing Shibboleth access type.

Roles and Responsibilities

- IAM Responsibilities
  - Accuracy and currency of identity information
  - Maintenance of identity attributes to enable selection of the users to transmit to each authorized application
  - Implementation of Grouper, the Internet2-sponsored open source group management system, to facilitate a common interface for specifying the users of intercampus applications throughout UC.
    - Individual campuses may propose alternatives to Grouper for implementation at their site.
  - Implementation of an unchanging and unique identifier for all identity records sent to a specific application.
    - eduPersonTargetedID should be considered for this during the detailed design phase of the project.
  - Deployment and operation of the Common Interface, as well as the Shibboleth interface
  - Deployment and operation of the middleware that will be utilized by the Common Interface
    - Kuali Rice should be considered for the middleware during the detailed design phase of the project.
  - The process for approving attribute release policies
- Application Administrator Responsibilities
  - Implementation of provisioning interfaces for the application
  - Implementation of appropriate protections for the identity information received
- UCTrust Responsibilities
  - Unique naming of all IdPs (IAMs) and SPs (inter-campus applications), as is already done for Shibboleth
  - Other UCTrust operational responsibilities, such as identification of support contacts, maintenance of logs, *etc.* These are described in UCTrust University of California Identity Management Federation Service Description and Policies.

Further Information

- Service Provider High-Level Design
- IAM Provider Notes and Observations
- Thoughts about Interchange Formats

Notes

- It should be noted that Internet2's COManage project is complementary to this project, as it focuses on authorizing and provisioning members of a Virtual Organization for LDAP-enabled applications. While it does include primitive exchange of user identity information via nightly LDAP queries, we believe COManage would benefit from our work on the exchange of identity information. Also, COManage provides an off-the-shelf solution for LDAP-enabled applications that can be leveraged within UC. Assuming implementation is approved for this project, potential collaboration with COManage should be pursued.