User Provisioning - ITLC Project Proposal

User Provisioning - ITLC Project Proposal

ITLC Project Proposal - Initial Overview
1. Project Title
User Provisioning for Multi-campus Applications

2. ITLC Sponsors

Richard Kogut (UCM)
Peter Siegel (UCD)

3. ITLC Project Criteria

- Avoids duplicative efforts to provide mandatory facilities.
- System-wide.
- Commodity service (not strategic campus differentiator).
- Can succeed; builds on established success.
- Relatively short delivery time.

4. Project Overview and Goals

UCTrust has provided a common infrastructure for identifying users to multi-campus applications, including information that can be used to make authorization decisions. The Shibboleth technology used by UCTrust provides that information at the start of an online session. However, we have learned that many applications need to incorporate contextual information about their users *before* those users' first sessions.

For example, the LMS employee training system must be able to send email to various categories of employees to inform them that they are due for legally-mandated training (such as sexual harassment prevention training for supervisors). While UCOP has built software to upload basic employee information to the LMS, each campus is required to create a process to upload basic information about non-employees who must (or can) take training, as well as complementary information not maintained by UCOP. A similar need exists for Connexxus and UC Ready, and would have been needed for At Your Service Online if the identity information had not already been in place within the payroll system.

The goal of the proposed project is to leverage the policies and agreements already established for UCTrust to develop a second software infrastructure that supports the exchange of identity information before, and independently of, the establishment of a user session. While it is unlikely that all campus-level implementation in support of multi-campus applications can be eliminated, not addressing this issue will mean that future applications will continue to require 10-12 separate development efforts to provision users.

The development of detailed specifications should be a collaborative effort of the ITAG and the UCTrust Work Group so as to maximally leverage the perspectives of both architecture and implementation experience perspectives (although there is considerable overlap in the membership in any case).

5. Key Benefits and Deliverables

The key benefit of a common framework is more rapid deployment and greater agility in launching multi-campus applications, with lower costs in time, effort, and resources both centrally and at the campus level.

The key deliverable is the common framework/infrastructure, probably based on a service bus, that can facilitate these exchanges of identity information. This includes a standard interface specification to be used by applications that require such identity information.

6. Timeline

The target for completion of this work is the end of calendar year 2011. This will, however, require work by each of the campuses, so the first milestone should be a checkpoint to giving the ITLC a go / no-go decision on whether to proceed to subsequent milestones. The first milestone, to be complete by the end of 2010, should deliver:

- A technical and business architecture to address these issues
- An approach / plan with the subsequent milestones to implement this architecture throughout UC
- An estimate of resources required for those subsequent milestones, both UC-wide and at each campus.

7. Key Stakeholders

- Primarily campus IT organizations who have to create the various processes and run them on an ongoing basis.
- Business officers seeking to implement new applications in a timely manner.

8. Advisory Groups and Their Role

- ITAG
- UCTrust Work Group

9. Role of the ITLC in Implementation

At the first milestone, the ITLC will provide a go / no-go decision for the final phases of the project. In the event of a "go" decision, the ITLC, and its members, will also allocate resources to complete those final phases.

10. Related UC Initiatives

List other projects (whether underway, planned, or under consideration) that impact or depend on this project.

11. Required Resources

Completion of this project will require effort at each campus, as well as "central" effort. Details of this will be an outcome of the first milestone.

12. Marketability

Given the probable low cost, the constrained financial situation probably wouldn't weigh negatively. However, the facility would not add functionality at the user level or create short term cost savings, so it probably wouldn't excite leadership.

13. ITAG Discussion

Of the various potential projects that were suggested for a middleware leveraging demonstration, we have selected the initial identity provisioning problem. This is primarily due to the nature of the problem requiring a bus routed workflow that can progressively build a body of information from a variable number of sources. The majority of our UC shared software functions primarily as SaaS, with customer usage taking the form of a visual interface. Once the customer is provisioned in such systems, any new information generated is through the visual interface and saved internal to the application system. The event wherein the customer is first placed within the application system is where all the interesting details occur.

It is in the nature of business oriented application delivery systems that the population group eligible to use such systems can not usually be specified in a strict, logical rules based fashion. The users of such systems usually result from an unrelated collection of decentralized delegations. Additionally, different levels of authority within the same application system often require level specific attribute assertions. This contributes to the complexity of creating a common mechanism for marshaling the necessary data in any initial provisioning transaction. Oftentimes, the prerequisite data is not all found in an identity purveyor's (IdP) repository, and multiple data repositories must be interrogated.

Another factor outside the provisioning transaction itself, is the application of privacy constraints. In some designs, we see a baseline approach taken, wherein the owner of the attributes being shared specifies in advance which are sharable and which are restricted. One downside of such an approach is that all requestors of such attribute sharing get treated identically. The attribute owner can not, for example, choose to share their home address in one application system, but not share that same address in another. Such decision making is often unique and lends itself to a workflow.

14. Technical Environment

Once a decision is made to request that a specific individual be made a member of an application system's community of users, an automated process may be used for bringing together all necessary information and approvals. After the conclusion of such a process, the individual in question should then be able to sign on to the application system for the first time. Attributes may come from three potential sources - An Identity Provider who carries some global information such as employee id and the like; A data repository not associated with an IdP; Information that must be hand entered for this particular event.

The consumer of this data kicks off the process by enumerating the specific attributes desired, and where it is believed that they come from. The person who is the object of this provisioning operation must both ascertain that they are in favor of sharing the attributes specified with the consuming system, and must also supply any unique information not found in any data repository. Subsequent queries of both IdP and non-IdP data repositories may complete the transaction, or they may result in additional manual authorization of some steps. The resulting successful aggregation of information becomes a formatted input to the application system requiring this data. Such activity requires a bus, a workflow engine, and a dynamically specified document type.

Large scale application systems (ERP) have internal processes for create, update, and delete of authorized users. They often also have an internal exit point for such functionality, and it is to such an exit point that a generic provisioning process might be connected. Potential problem areas might be that such processes are thought of as synchronous. This might result in sessioning issues for the person initiating the process. A less automatic workaround would be to independently kick off the provisioning information gathering process, and then manually feed the returned information into the application system. Such a generic approach would still be more efficient than the status quo.