User Provisioning Project Objectives and Principles

User Provisioning Project Objectives and Principles

[The following was originally drafted by Arlene Allen in her electronic mail of 6/28/2010. It has since been modified by the project team.]

1. The same codebase can run in an arbitrary location (campus) as long as prerequisite resources are available to it. Define those resources.
2. The codebase can run successfully with zero customization OR it can be tailored to successfully front end a particular SP it is being customized to represent. Such a customized front end to a particular SP would also be capable of running unaltered in arbitrary locations (campuses). This excludes customizations associated with branding.
3. Since this project is contextual to UC ITLC, it is both acceptable and prudent to build inherent knowledge of UCTrust technology into it. Stated differently, it does not need to be agnostic to Identity authentication processes currently in use, nor does it need to be agnostic to the general organizational architecture of the UC.
4. There are three mandatory parties involved in this workflow and an arbitrary additional number of parties to the approval process. Mandatory parties are:
    a. Functional service initiating the process of adding a previously unknown person to an SP's internal repository with the associated attributes and permissions. This is the incep point of the workflow.
    b. The specified IdP with pre-existing knowledge of the personal identity being added to the SP's repository, i.e. this is a pre-existing identity, somewhere. Might be in multiple IdPs.
    c. The person in question being added to the SP's repository. I'm thinking we assume some permissioning for attributes the SP requests. This finesses the IdPs not doing it.
    d. 4th through n'th parties are any secondary approvals beyond the three basic parties, and also the possibility of an attribute repository that is *not* IdP in nature.
        i. Hmmm... I think it'd be easier if we assume all sources are the IdPs (even though they would generally broker on behalf of other sources on their campuses). Maybe start with the 12 IdPs and allow for other sources of identity information at some later time? We'd need to understand the use cases where it didn't make sense to have the existing IdPs represent all sources on their campuses. - DHW
        ii. Agreed. In the interest of a restricted scope for POC, lets keep it strictly IdP in nature. - AA
5. The basic flow assumes that all such provisioning starts with a business process that involves the delegation of authority. For example, no one is ever part of a financial system without being specified by another person that is already a member of that system or a delegated admin thereof. There is always a chain.
    a. I think the issue of human intervention depends on the specific role (and authorizations that are derived from that role). From our (user provisioining) point of view, maybe we should just assume that *some* process has been established for deciding whether a person should be authorized? That process doesn't even have to be the same across campuses, for that matter. For example, Some campuses might want human intervention to identify travelers to Connexxus, while others might want to do an automatic process based on job classification. - DHW
    b. That was my point, so I see this as: Arbitrary business process begets an authorization/provisioning IT process being designed herein.
6. Decoupled presentation technology is pretty widely accepted, but I'm stating it here anyway.
7. What ever technology is chosen to implement this should be somewhat ubiquitous. It is acceptable for that technology to have one or more general prerequisites that are reasonably easy to meet. This is a rather vague statement, and I assume it will be arrived at via consensus rather than anything more formulaic in nature.
    a. I don't recall if we agreed on a technology in our meeting. Lets agree on a basic technical framework herein. - AA
8. The process will, by definition, be handling some amount of PII type data. For adequate security in a peer to peer type environment, we need some form of code signing that assures this process has not been tampered with.
    a. If this is handled within Shib 2 our data exchange is de facto protected. The out-of-band ftp type transfer we discussed would use its own encrypted communications. - AA
9. All communications pathways must be appropriately encrypted.
10. No application layer process ever sees or has access to a password. Passwords must always be handled by trusted third party authentication technologies. We might see an implied violation of this in the vendor software itself that is the heart of the SP functionality. We will need to talk that scenario out.
    a. I propose we not worry about SPs that demand an internal authn. By definition, they would not be natively Shib enabled. I was too focused on generality when I wrote that. - AA
11. There needs to be a designated code repository for sourcing and security purposes.
12. The minimum acceptable scenario for proof of concept is that there be at least one SP and two IdPs.
13. The only middleware technology specifically implied here is the workflow engine. All of the above could be done without an ESB. Question: Do we want to force more than the workflow engine?
    a. My recollection from our meeting is that we didn't answer this question. Are we going to design around interoperable campus ESBs connected to each other for our subscription service? - AA
14. As much as possible, user provisioning processes should leverage existing UCTrust policies, processes, and technology.