Big Ten Academic Alliance
Identity Management Task Force
September 29, 2016

Charter: Best Practices for Provisioning and De-provisioning

Executive Summary

Campus IDM systems are the central store for the institution's identity records. Out of these systems come all of the information that makes those identities useful: creating user accounts in on- and off-campus services, performing authorization decisions, and providing the attributes about users that any service needs to operate. The standards for central person registries, the core of identity systems, are fairly well-defined. The means for data to flow from the registry to services, however, is still a bit of a wild west. As important as it is for the right data to reach its destination at the right time, there are no widely-used standards to accomplish this essential task. The technology is there, at least in a basic form, but the standards aren't clearly defined. The result is a complex jumble of mechanisms for service and user account provisioning that is expensive to integrate and maintain.

The goal of this effort is to gather and document the standards that are available, and potentially propose additional development to take provisioning and de-provisioning to the next level. Some areas of interest are communicating group and role information, standardizing attribute release and mapping, and choosing the right method and protocol to communicate this information. The range of services out there is huge, and so is the functions and features of those services. This effort recognizes that there is no one-size-fits-all solution for provisioning and de-provisioning. A complete set of best practices will include multiple methods for distinct use cases along with documented benefits of each. The goal is a document that a campus IDM operator or service developer can use to leverage scalable, standards-based methods for provisioning and de-provisioning.

Problem Statement

The challenges of service provisioning are as plentiful as the services to which campus IDM systems communicate. Finding a scalable and standardized method to provision and de-provision on-campus services is challenge enough. When you add in cloud services that use methods out of your control, the challenge quickly balloons. Add to the complexity that many vendors will invent their own proprietary APIs to handle provisioning and de-provisioning tasks. Setting up mechanisms for each system or service separately is time-consuming and expensive, and once you do, there's even more cost associated with trying to get off of a certain vendor's model if you need to change service providers. The standards, though available, aren't widely used, and the time and energy spent as a result is considerable.

Some of the areas of challenge include:

- Just-in-time (JIT) versus just-in-case (JIC) provisioning: many services can create user accounts when the user first logs in, often with data from a SAML assertion – just in time. Other services provision users through back-channel processes, often with API calls – just in case. Many services, in fact, support both methods. JIT and JIC provisioning both have benefits, and it's not appropriate to say one method should always be used. It's not clear, however, from both the vendor's and campus's standpoint, when each method should be chosen.

- Account reconciliation and de-provisioning: the processes to query a vendor and learn what users have accounts and to de-provision accounts differ widely between service providers. The patterns of provisioning and de-provisioning from higher-ed are very different from corporate customers, and often the methods presented by a vendor don't hold up well under these patterns.

- Many and often changing APIs: There are few standards currently around the APIs used for JIC provisioning/de-provisioning. Those standards that are available are rarely used. Vendors often create their own APIs and, many times, change them with little notice. The end result is a lot of development to integrate each vendor's service.

- Attribute standards: JIT provisioning works well as long as the service provider only needs a basic set of attributes about the user and the identity provider can support those attributes. Basic attribute bundles only go so far, though, and as has been demonstrated with the InCommon Research and Scholarship category's slow adoption by identity providers, standard attribute release happens rarely. In addition, many vendors will invent their own attributes that they then require for users to be provisioned or granted access.

- Groups and roles: the wide range of group memberships and roles that a user can have and the wide range of use cases for those groups and roles in services is extremely challenging to express in a standard and scalable fashion. Coarse roles such as staff, student, or faculty are easy enough, but finer roles such as administrative privileges or enrollment in a class quickly makes this a very complex task.

- Security and privacy: Vendors often have on-boarding processes that include sending a large file of information about users to initially populate accounts or querying an LDAP directory on campus. Even vendors that support federated authentication often rely on some such back channel for populating user data. These methods, aside from lack of scalability, open up significant concerns about data security and privacy.

- Scalability: Many prov solutions require a central IdM group to implement and maintain all service-specific provisioning policies. A more scalable and distributed solution is to give the service owners the tools and information they need to put their own policies into action.

Proposal

We propose conducting this as a collaborative effort between Big Ten Academic Alliance and the Internet2 TIER Program. TIER has identified provisioning and de-provisioning as one of its key focus areas. Much of the work that we are proposing in this charter complements the initial phases of the TIER provisioning effort. Both groups plan to start by defining a comprehensive functional model (the what and how) of provisioning and de-provisioning. By meeting as a single group initially we can leverage the collective knowledge and experience of a broader cross-section of the higher education community.

This effort will consist of two phases:

Phase I: Create the functional model described above and map it to relevant standards and protocols. Compile these standards into a document that can be used by campuses and service developers.

Phase II: Where current standards can't solve the challenges, propose a roadmap to attain the next level of provisioning and de-provisioning. The roadmap could include such deliverables as software, protocols, connectors, methodology and/or a mix of all of these. Promulgation and Implementation of the items on this roadmap could be another collaborative effort of the Big Ten Academic Alliance and Internet2 TIER.