IdM Cloud Services Cookbooks

Table of Contents

IdM Cloud Services Cookbook

Executive Summary (EXECSUM)

The Big Ten Academic Alliance Identity Management (IdM) Working Group has identified federated integration with "cloud" services as an area which needs special attention by the group in order to lower barriers to adoption by Big Ten Academic Alliance institutions. These integrations are often complex, requiring large amounts of staff time in the policy decision-making process, the architecture of technical integrations, and result in duplicate effort between Big Ten Academic Alliance institutions. It is important for the Big Ten Academic Alliance institutions to present consist procedures and practices (when possible) to vendors in order to simplify the integration process and ensure that best practices are followed on the part of the vendor and the institution.

What follows is a DRAFT set of recommendations in a "do's and don't's" format that we hope will result in a recipe for success for the institutions and vendors. The guidance will be most relevant to those who sponsor, administer or support Identity Provider and Service Provider services on both the campus and the vendor sides. The content is being contributed by staff from across the Big Ten Academic Alliance. Where broad consensus was not achieved (or is not anticipated) the word "consider" is used in place of "do" or "don't".

Problem Statement (PROB)

Partnering with a new cloud service provider can be complicated and time-consuming. Nearly all of the hurdles can be overcome with technical or policy solutions, but forming those solutions can be an expensive process, which diminishes the value proposition of adopting cloud services.

Overview of Higher Education IDM Landscape (LAND)

The identity management systems of colleges and research universities are by rule more varied and complex than in most similarly-sized corporations. Many universities source identity data from multiple authoritative systems, including HR, student records, and other systems used to identify more loosely affiliated individuals. Data from these competing sources is often reconciled into a single, master directory intended to offer the "best" identity data for the organization.

Higher education user role definitions also differ greatly from private industry. A particular user will often have multiple organizational roles (for example, a graduate student may also be considered faculty if assigned certain teaching responsibilities), and in some cases the user's *primary* role may be difficult or impossible to discern. Furthermore, user roles change very frequently (imagine a staff member who enrolls for classes one semester, but not the next), and most universities track a large number of "edge cases" (visiting faculty, adjunct faculty, special student types, emeritus, allied staff, and so-forth) each of which come with their own business rules and related service provisioning challenges. Faculty and staff may have a number of affiliations with different operating units, which in some cases may involve multiple email addresses and other identity data complexities. Depending on their responsibilities, some users may require service resources that fall outside of their primary role (for example, staff who support student systems), and this access may be achieved through one-off methods or the use of ad hoc provisioning processes. The breadth and complexity of this landscape is often surprising to vendors and commercial service providers who are new to the university environment.

Student and staff turnover is a challenge for higher ed identity management, and requires markedly different processes for user creation and service provisioning and de-provisioning. Each semester or quarter a university may process thousands of new student records, resulting in predictable but significant traffic spikes to service enrollment systems and help resources. (A cloud vendor that provides direct support to higher ed users may see a 100x spike in new user support requests during the opening days of a new semester.) During the same periods many existing users will undergo role changes (for example, undergraduate students matriculating to graduate student status) that also contribute to start-of-the-semester provisioning and support loads.

At the other end of the pipeline, student and staff de-provisioning presents a particularly difficult challenge for universities. It is sometimes said that "no one ever really leaves a university," and this viewpoint is reflected in higher ed identity management practice. Unlike the corporate world where when an employee leaves, the company retains their data and the user is expunged from the directory, de-provisioning in the higher education setting is usually carried out in large, end-of-the-semester batches (sometimes referred to as "de-provisioning runs") that can hinge on varied business rules with many exceptions and special cases. Privacy rules prevent universities from asserting control over an existing user's data, and graduating students and retired faculty may retain access to many IT resources. Students or faculty/staff may leave the university, then return later in the same or different role with the understanding that they will reclaim their previous credentials.

Universities also play by different policy and compliance rules than private industry. Expectations for data privacy is higher, most users (especially faculty and students) are considered the primary owners of their data, and higher ed security staff are often interested in keeping watch over service logs related to administrative activity and privilege escalation. FERPA and the proper handling of FERPA-related data is an ongoing concern for IT managers, and can have surprising consequences (for example, the need to offer group services with suppressed member listing) for new service deployments. Schools must also manage legal requirements surrounding service accessibility, and thus show strong preference toward vendors that offer accessible products or can demonstrate a clear accessibility roadmap.

University organizational culture tends to favor a consensus-driven decision process, and central IT units often control only a fraction of the infrastructure and services available to the campus. As a result, it is usually very difficult to mandate top-down change, and system-wide IT changes usually require buy-in from a diverse range of stakeholders in both technical and non-technical positions. This decentralized

whatever means they have established. It also allows the user to not only have a single password, but in the case of single sign-on, only need to authenticate once to log into several applications.

makes deployment of these technologies easy to extend to cloud services with little or no impact on vendors.

Take care however to think through how the various request, issue, activation, and reset processes are designed because it's relatively common to introduce dependencies on password authentication that can undermine the value of the additional factors.

Identifiers (IDENTS)

An *identifier* is a special kind of attribute that is specifically designed to distinguish each account/subject/user/thing from its peers in a particular set. While almost any attribute may contribute to differentiating a subject from similar subjects, identifiers are intentionally designed to do this by themselves. It is common for a subject to possess several different identifiers, used for different purposes or generated by different information sources.

With most applications, the primary function of identifiers is to identify its users. This process is more complicated when an application relies on a system other than itself to handle the authentication process, referred to above as *externalizing* it. Such an application no longer fully controls the way the identifiers look and behave, and must therefore make sure that its needs can be met by clearly defining the properties it needs from an identifier. (**Appendix A** discusses a number of these properties.)

The use of *federation* adds yet another layer of complexity because the application may not even be operated by the same organization that is providing the authentication function, and the identifiers. This makes the adoption of standards and common practices essential, because an IdP service cannot be tailored to meet the specific needs of every application; instead it should provide a sufficient set of identifiers of different properties so as to meet a wide range of application needs.

Further, an application supporting the use of multiple IdPs across many organizations at the same time has an even more compelling need for standards, because handling widely varying approaches to user identification at the same time greatly increases application complexity and creates opportunities for bugs and security exposures.

Both Campus and Vendor: DO support a varied set of identifiers.

Many standard attributes are available for obtaining
variety of standard and proposed eduPerson identifiers, using a core set of underlying identifiers managed by an IDM system. A vendor should base user identification on at least one of these attributes and, for improved flexibility, more than one. Particularly valuable are
the *eduPersonPrincipalName*, *eduPersonUniqueID*, and *eduPersonTargetedID* attributes. The latter two have quite well-defined properties, the former less so but it does have much wider adoption within the US higher education sector.

Both Campus and Vendor: CONSIDER the use of eduPersonTargetedID where appropriate.

The definition for *eduPersonTargetedID* can be intimidating, but once you take the time to understand it,

user's more traditional attributes, and should be as stable as possible, be of fixed (and reasonable) length, and ideally be alphanumeric to avoid confusion over numeric formats. Effort to minimize duplication of identities pays off in the prevention of remediation needed later to change an identifier that's already in use.

Campus: DO make eduPersonPrincipalName useful.

IdM

situation in which the cloud service provisioning model doesn't adequately support de-activation or deletion of its local user accounts.

Whether or not a service maintains local user accounts, information bearing on the allow/deny decision can be passed to the service at each instance of user access to the service. This is typically labeled attribute-based access control (ABAC). For example, an attribute that asserts a user's entitlement to access a service under a contract with the institution might be all a service needs to determine access. In that case, *eduPersonEntitlement* would be an appropriate attribute to convey the needed information.

Provisioning and De-provisioning (PROV)

When a cloud service design depends on a specific, dedicated source of user information, and an organization contracts for that service on behalf of its members, some process must be in place to keep the service-specific repository current and consistent with the organizational information on members. This process is referred to as provisioning. When organizational members lose their eligibility for a service, there should be a corresponding de-provisioning process to remove them from the service user store, or to mark them as inactive.

Provisioning is one of the most complex integration challenges for sites adopting a cloud service. Each service offering tends to come with its own proprietary provisioning approach. For an institution adopting multiple cloud services, this means devoting valuable staff resources to an ongoing series of one-off integration projects. Service providers and consumers would both benefit from standardizing the provisioning and de-provisioning process.

Just-in-time vs. Just-in-case approaches to user provisioning

There are two fundamental models for provisioning external services: Just-in-case and just-in-time. Just-in-
for the service, whether or not they may actually ever use it. This model can be simple, but will mean more traffic between institution and vendor, and more user information in the hands of the vendor. The primary alternative is just-in-time provisioning: Only when a user actually attempts to access the service does the provisioning step take place. This requires more sophisticated coding by the institution: Essentially a side-flow operation has to detect that a first-time service user is attempting access, and then

model of detecting first-time access at their end, but this is relatively rare.

Campus: DO expect the typical vendor to have a single, set model for creating user accounts on their systems.

Many cloud services depend on having a local record for each of its users. Often the design center of their original service offering was individual sign-up for the service. They may support institutional licenses for access to their services, but the model of a single user establishing an account on their side is often present just under the surface. At most they will have defined a single model by which a client institution can initiate provisioning of user accounts for their service. The burden is on the subscribing institution to understand the vendor's model and conform to it. This puts an increasing burden on the institution as the number of cloud services scales up. The more cloud services, the more integration models for user provisioning the institution will end up supporting. For example, a cloud service provider

If the institution supports two vendors, most likely two directories will have to be configured and populated.

Campus: DO practice "defensive programming" when setting up provisioning services.

Experience has taught us that in many cases, the documentation on provisioning models from the service provider is incomplete, or worse, simply incorrect. The institution-side integrators must then adopt a "defensive programming" style that features extensive testing for error conditions or simple uncaught failures, including setting up rich logging to support diagnostic work. Based on testing results under realistic conditions, campus staff will need to put in place counter-measures to detect and remediate failed operations. Seemingly random fluctuations in responsiveness under load are not uncommon and again require detection modes and counter-

Vendor: DO support just-in-

browsers has been a thorn in the side of federated identity since its inception, and there are a variety of imperfect solutions used today.

Campus: DO provide a discovery service if you operate multiple IdPs.

Some organizations operate administratively as separate divisions, or campuses in the case of some university systems, and operate multiple IdPs. One way of dealing with this is to operate a proxy IdP as a stand-in for the various real systems behind the scenes, hiding the multiplexing that's going on from federated partners. This can often be a good approach, but may not be practical in some organizations.

However, operating distinct IdPs will frequently cause problems if the parent organization negotiates with vendors as a unit, and particularly if all users share a common email domain. Vendors often use email domains as a proxy for connecting an access attempt to a specific IdP, and breaking that 1:1 relationship may be difficult.

It will be advantageous in such a case to be prepared to assist vendors by operating a discovery service as a standard shared service for different applications to use. This isn't a panacea; most vendors will not support the concept of a discovery service without customizing their services, and it can introduce extra steps for users to go through, but it's better to be prepared with a viable solution than to go in unprepared and expect to be accommodated.

Campus: DO socialize the use of organizational email addresses.

Some universities have a strong "email" culture in which use of campus email addresses is common and encouraged. Others tend toward a more laissez faire attitude that fosters use of personal email addresses, particularly by students. Those in the latter camp are going to encounter major difficulties with many cloud projects, as the use of organizational email addresses often double as a user identifier (frequently one that is displayed to other users), and in particular as a discovery mechanism.

It's increasingly common to find applications that combine federation with traditional consumer access by using an email address entered by the user as a hook to recognize a SSO-enabled account and assume the IdP to use. Organizations that want, or require, users to use their institutional identity may need to increase user education and socialization around their own email domains or be faced with a de facto migration of user access to personal identities.

Vendor: DO provide a discovery mechanism for federated login.

Cloud services often tend not to address discovery because of a focus on the enterprise outsourcing use case of a single set of users accessing a compartmentalized service. In such cases, discovery is often "optimized out" by tying an instance of a service at a particular URL to the IdP to use. Even when this makes sense in the context of a particular service, a single organization may have multiple IdPs representing different campuses or user populations. Thus, avoiding discovery is ultimately a losing battle, particularly (perhaps uniquely) in the university sector.

In addition, many cloud services supporting collaborative use cases (file and media sharing, scheduling, surveys, and so forth) intrinsically need to consider the fact that a resource's location should **not** dictate the IdP to use. Doing so, as is quite common today, has a very negative effect on the viability of federation because it makes sharing resource links among colleagues difficult or impossible.

IdM Cloud Services Cookbook

While it is quite common for commercial SPs, especially larger ones, to self-publish metadata documents, the practical fact is that this approach does not scale unless the information is accepted blindly, which in turn acts to limit the level of trust or the number of SPs supported (or both). Federations can elevate trust without sacrificing scale.

Both Campus and Vendor: DON'T expose untrusted URLs to users.

An IdP or SP may have occasion to leverage information supplied by a partner in its user interface, to describe a service, link to information about the service, render a logo, etc. Care should be taken to ensure that the source of this information is trustworthy because of the ease with which an attacker could take advantage.

In particular this use case demonstrates the risk of remotely acquiring metadata directly from a partner unless it's signed, and we discuss above the scalability limitations of directly establishing trust in a signing key for every individual system. Federations offload the cost of establishing that trust and of vetting the safety of data that might be incorporated into system operations, and of responding to cases in which unsafe data might appear.

Technical Interoperability (INTOP)

SAML is a wide-ranging specification covering many potential usage scenarios. By itself the SAML specification does not guarantee functional interoperability.

When we speak of operational agility, we refer to an understanding of what needs to be stable, what can be more easily changed, and what techniques can be used to facilitate changes when they do occur.

As an example of how not to do this (or maybe more to the point what you have to do when software is

IdM Cloud Services Cookbook

management procedures and whether they create barriers to following your own. Create a game plan for how changes have to happen, and a clear timeline for how long they would take.

For campuses in particular, when possible, push heavily to rely on metadata as your interface to their processes. Otherwise their manual work and system deficiencies become your burden.

Sharing the Burden (BURDEN)

Federated Incident Response (INCRESP)

The following material is based on guidelines from the InCommon Federation and are presented here for completeness. For further information see: https://spaces.internet2.edu/display/InCFederation/Federated+Security+Incident+Response

More information about the InCommon Participant Organization Practices (POP) can be found at: http://wwww.incommon.org/docs/policies/incommonpop_20080208.html

Both Campus and Vendor: DO publish federated incident response contact information in the InCommon metadata.

Federated incident reporting and response is difficult without a standard place to find security contacts. The InCommon federation provides a place to include Security Contact information in the Identity Provider's metadata (alongside administrative and technical contact details). Other federations also include this functionality. Make sure to publish a security contact for your organization, and make sure services that federate with your identity provider know about it.

Campus: DO update your Participant Operational Practices document annually.

participant operational practices is important as it lets others understand your processes when federating. An annual audit of this documentation is the best way to remember to keep it up-to-date.

Campus: DO include the URL of your institution's privacy statement in the InCommon metadata.

The InCommon federation manager allows organization administrators to supply their privacy URL to be included in InCommon metadata. Other federations allow for similar inclusion in metadata. Make sure to publish an up-to-
not in metadata.

Campus: DO implement a log retention policy.

can become a problem when investigating an issue such as a security incident. Keeping too many logs or never rotating and deleting old logs can become a burden for storage and log parsing. Make a policy and implement it. Ensure your local policy aligns with any institutional record retention policies.

IdM

ources in the cloud, the impact of the compromised account can be much more far

that SPs can rely on it with confidence that it's a legitimate resource. For InCommon federation members, there's even an error handling service used by some SPs that can generate links to an IdP's error page.

Vendor: DO display user-friendly error messages.

It is good practice to always make your error messages user friendly: What happened, why it happened, and what the user might do to remedy the situation. This is especially true for federated logins. The most common causes of these errors are the user not beign authorized to access a service and the identity provider not releasing required attributes. In eit
authorized, explain why. If the error is the fault of the identity provider, direct the user to a help contact from their home organization.

Vendor: DO make use of IDP error URLs in the metadata.

If your federation provides a place in metadata for an error URL, use it. It can be a very appropriate place
t be

externally maintained script to forward the proprietary logout into the IdP because this maintains local flexibility without losing any function. Note that often applications will allow a custom logout URL to contain arbitrary parameters, so it may be possible to tailor feedback based on the redirecting application.

Most IdP software, even if SAML logout is supported, also support a redirect-based logout that, if nothing else, will remove a session such that any subsequent requests for authentication will re-authenticate the user. This limits the damage of a session left unattended such that only applications already accessed will be at risk.

This kind of logout approach is simple to integrate with the proprietary logout capabilities supported by most services, and can usually be implemented easily even without explicit support in an IdP when the underlying session is based on cookies

as part of authentication is not the same as hosting it in a directory, it is also not materially different, and arguably constitutes a more controlled release of data at the discretion of the data subject.

InCommon members should consider the adoption of policies to release Directory Information attributes to federation services on at least some ad hoc basis. The Research and Scholarship program is an alternative to federation-wide policies and may also be of interest. Work with data owners such as the HR department and Registrar (or others as appropriate) to establish consensus about the appropriate policy to use for default release.

Recognize that a default-deny policy will act to inherently limit the value provided by an IdP, though it will likely decrease ongoing support costs. Make these decisions deliberately.

Campus: CONSIDER a user consent mechanism for attribute release.

Consider a consent-based approach to provide user control of data release. Many federation products include functionality, or support add-ons, to support user consent to the release of data to services. A consent process may be a good way of expanding the reach of an IdP by enabling users to make the decision about whether to release data, in lieu of requiring contracts for every individual service a user might access. Consent systems also provide an auditable record of approvals, and can help address a requirement for acceptance of usage terms.

Note that consent approaches have limits. They should never be used if a student or employees failure to consent would bar them from a service they are required to use, nor do they work well for attributes beyond a user's capability to understand.

A typical consent UI should be able to identify the relevant service, and attributes, concisely. InCommon membership, and the use of SAML metadata, provides a good source of reliable information to use in driving such a UI.

Campus: DO establish procedures for managing the release of restricted data.

necessary to provide healthcare services and to protect the nations overall public health. Like the previous section of FERPA this is not legal advice, but provides general guidance to be considered by campuses and cloud service providers that offer services that may be used to store HIPAA protected data.

Campus: DO require that cloud service providers sign a Business Associates Agreement (BAA) if it is intended to be used with HIPAA protected data.

HIPAA's provisions for Business Associates Agreement is particularly important in regards to the use of cloud services that a campus intends to use to store protected data. A "Business Associate" is generally defined as an entity "... that performs certain functions or activities on behalf of, or provides certain services to, a covered entity that involve the use or disclosure of individually identifiable health information."• (http://www.hhs.gov/ocr/privacy/hipaa/understanding/summary/)

Campus: DO articulate a strategy for dealing with the difference between institutional and personal user data.

Most cloud services house some kind of user data, and of course this is particularly true for applications that provide collaboration and file storage service. Many higher ed users will use cloud service accounts to store a mix of personal and institutional data. For example, a faculty member may have research work, family photographs, and departmental policy documents, all under a single account. If the user were to leave the university, the family photos and (probably) the research files would rightfully leave with the user, but the policy files should stay with the user's department. Untangling these document collections can be very time consuming, and lead to both institutional data loss and serious de-provisioning complexities.

Before launching a new cloud service, take some time to think through how users will store data on the service, and whether or not the service is prone to encourage the mixing of personal and institutional data. If so, do develop a strategy for dealing with this problem before the service goes live to your campus. In some cases providing a best practice guide may be the best solution. For other services you may want to consider leveraging a groups service or creating shared or non-personal service accounts that can be used to store data that should stay with the institution when an individual leaves their job.

Campus: DO publish a list of approved data types for your cloud services.

When a new cloud service goes live, the first question that many campus users will ask is, "What kind of data can I store on this service?" Consider publishing an approved data types chart that lists the institution's cloud services plus the approved data types for each service vendor. Data type categories might include FERPA protected data, HIPAA data, SSNs, IRB file, point-of-sale data and credit card numbers. You data classification policy should include enough categories to be useful, but not some many that the chart becomes cumbersome or confusing to your users.

Campus: DO manage data sensitivity and stewardship concerns from day one, and don't underestimate the potential impact to project scope.

This is another "look before you leap" recommendation. It's important to identify the data stewardship and compliance stakeholders on your campus at the beginning of your deployment project, bring them into your project team, and budget time for identifying approved data types as well as areas of concern that may need to be mitigated.

Summary Guidelines and Principles for Integration with Cloud Service Providers

Appendix: Identifier Properties (IDPROP)

Identifiers have a number of characteristics that help to determine appropriate usage and it's important to express application requirements in terms of properties that can be mapped back to those satisfied by particular identifiers to choose the best one for the job.

Persistence

Persistence is a measure of the length of time during which an identifier can be reliably associated with a particular subject. A very short-term identifier might be associated with an application session. A permanent identifier is associated with its entry for its lifetime (which is not necessarily "forever", so permanence is just a relative notion). Typically we refer to an identifier as having persistence if it is stable over a relatively lengthy period, usually measured in at least months.

Reassignment

Many identifiers do **not**

IdM

IdM

IdM Cloud Services Cookbooks
Vendor Guide

Partnering with a new cloud service customer can be complicated and time-consuming. Nearly all of the hurdles can be overcome with technical or policy solutions, but forming those solutions can be an expensive process that diminishes the value proposition of adopting cloud services in the first place. This can be particularly true of colleges and research universities. The identity management systems of colleges and research universities are, as a rule, more varied and complex than in most similarly-sized corporations. Some examples from the Overview of Higher Education IDM Landscape section of The Big Ten Academic Alliance Cloud Services Cookbook include:

Identity data is sourced from multiple authoritative systems, including HR, student records, and other systems used to identify more loosely affiliated individuals.Higher education user role definitions also differ greatly from private industry.
Student and staff turnover requires markedly different processes for user creation and service provisioning and de-provisioning.
Universities also play by different policy and compliance rules than private industry. Expectation for data privacy is higher, proper handling of FERPA-related data is an ongoing concern, and there are strong legal requirements surrounding service accessibility.
Inter-institutional collaboration is common in higher education, making multi-lateral identity federation a requirement.

The Cloud Services Cookbook suggests strategies for addressing these hurdles in a cost-effective manner for both the vendor and the customer. This document is a quick reference guide to the Cloud Services Cookbook for vendors, particularly those that have basic experience with SAML but are new to multi-lateral federated identity. We provide a brief introduction for each section and list the titles of the vendor-relevant criteria with links to explanations of the criteria in the Cookbook. Please use these links, as the explanations are crucial to understanding the issues.

Items marked with ✅ or ⛔ (depending on whether the item says DO or DON'T) are considered fundamental to a successful deployment. Every effort should be made to satisfy those specifications. (Note: For people using screen readers, fundamental items are also flagged with " ", a white plus-sign that is not visible but can be spoken by the screen reader just prior to the item's DO or DON'T.)

We start with requirements and guidance related to authentication, authorization, and provisioning ("Determining Who Our Users Are and What They're Allowed to Do") and then move on to a number of issues related to multi-lateral federation ("Working within the Federation" and "Sharing the Burden"). Finally, we close with help for assuring a good user experience.

Determining Who Our Users Are and What They're Allowed to Do (USERS)

Authentication, Identifiers, and Attributes (AUTHN, IDENTS)

Authentication services are used by a cloud service to obtain information about its current user. Typically, this information includes a unique identifier for that user, plus other attributes that are requested by the service. This information may be used in various ways in the provision of the service, as well as determining the eligibility of the user to access the service or specific functions within the service. For security reasons, this information should not include passwords or other secret credentials that may be used by the campus or its community members.

It is much preferred for cloud services to leverage existing campus single sign-on (SSO) infrastructure for authentication.

Working within the Federation (FED)

Working within a multi-lateral federation provides you with a highly scalable infrastructure for delivering your service to many campuses. In particular, federation metadata is a trusted common repository of all service endpoints, public keys, and other information needed to interoperate with your service. It also provides you with the information you need to interoperate with your customers' identity providers. Once you have registered your service into the federation, adding new customers requires very little work, and period security maintenance tasks, such as key rollover can be accomplished without careful coordination with each of your customers.

Federation does, however, require you to address some issues beyond authentication and authorization.

Identity Provider Discovery (IDPDISC)

In order for your service to accept identifiers and attributes from multiple campuses, users must have a

DO use self-signed certificates on non-user-facing endpoints.
DO

Logout (LOGOUT)

DO more than just destroy your local session, as appropriate.
CONSIDER supporting logout requests from IDPs.

IdM Cloud Services Cookbooks
Campus Guide

IdM Cloud Services Cookbooks
Campus Guide

DO expect the typical vendor to have a single, set model for creating user accounts on their systems.
DO practice "defensive programming" when setting up provisioning services.
DON'T require out-of-band acceptance of Terms of Use.
DON'T expect robust de-provisioning support
DO handle username changes

Working within the Federation (FED)

Identity Provider Discovery (IDPDISC)

In a federated environment, a service must accept logins from multiple organizations. A "discovery service" is one term for a standard way of letting a user select their home organization or chosen authentication source and be guided to the right IdP. The lack of built-in support for this concept in web browsers has been a thorn in the side of federated identity since its inception, and there are a variety of imperfect solutions used today. For example, vendors will often use electronic mail addresses for discovery. Also, it is often the case that a vendor can support only one IdP per enterprise customer.

DO provide a discovery service if you operate multiple IdPs.
DO socialize the use of organizational email addresses.

Technical Trust Framework (TECHTR)

The multi-lateral trust afforded by a federation requires a strong technical infrastructure to ensure that trust is not subverted by untrusted actors.

DO be prepared for the case in which a campus or vendor drops their membership in a formal identity federation
DO register SAML metadata with the InCommon federation.
DO define a process for maintaining up-to-date SAML Identity Provider and Service Provider metadata.
DON'T expose untrusted URLs to users.

Technical Interoperability (INTOP)

It is imperative that all federation participants utilize the same protocols and formats.

DO conform to the SAML2 Web Browser SSO Interop Profile
DO

resources appropriately. Also, session management, particularly handling logout, can be challenging and exhibit surprising behavior.x

Error Handling (ERRHAN)

DO provide an error web page where an SP can send users.

Logout (LOGOUT)

DO focus on desktop and mobile security. q0.000009 2 0 6 2 2 reW*nQ EMC /P #MCID 10mlnETQ84 509 2 0 6 2 2 reW*nl