Use Cases for External Identities

This page is being used to collect Use Cases. Per the usual policy on this wiki, any authenticated user should be able to edit this page, and submit their Use Cases for discussion. Please add your name and site within parentheses at the start of each of your Use Cases. Also, please include any risks or other concerns you may have related to your use case.

Please add new use cases above the "Use Cases from the Social Identities Work Group" section below. If you are modifying one of the Social Identities use cases, please do so where it is.

Use Cases from the Social Identities Work Group

- (Carnegie Mellon University, Russell J. Yount <rjy@cmu.edu>)

We are working towards using Incommon Federation, OpenID and other future social authentication systems to provide access for what we are terming "Low Level Of Assurance" LLOA access. Our first application to use this is for parental accounts. We would like to provide student billing information to parents based on permissions granted by their children.

The model we have decided on will send an email at the student's request to a parent that invites the parent to register with using one of a number of OpenID providers or Incommon Federation identity providers. We are using this model as we believe students are more likely to know their parents email address than know their parents InCommon or OpenID identity.

The implementation will be done with gateway service which will provide a CMU identities within local CMU shibboleth federation which are the translation of the OpenID or other provider's identity. InCommon federation identities will be used directly.

We hope to use outside identity providers for access to a number of other service in the future. These may include guest network access, preadmission students....

- (CILogon, Jim Basney <jbasney@ncsa.uiuc.edu>)

We are accepting both InCommon and OpenID authentication for the issuance of X.509 certificates for use with grids and cyberinfrastructure (for command-line, message-based, and batch workflows). For IGTF (high assurance) certificates, required by TeraGrid, EGI, and others, we can accept only InCommon Silver authentication. But representatives of other cyberinfrastructure projects (such as DataONE, FutureGrid, and OOI) have requested that we provide lower assurance certificates based on OpenID authentication for greater ease-of-access, especially in cases where the researcher's home campus isn't (yet) an InCommon member. The OIX Certified OpenID providers are of particular interest to us, so we have some minimum level of assurance.

A significant difference for us between InCommon and OpenID is that we can get the researcher's name and email address from InCommon identity providers to insert into the certificate, whereas we find many OpenID providers can not provide this information to us, so we just put the OpenID URL itself in the certificate (see http://ca.cilogon.org/names).

- (Internet2 COmanage, Benn Oshrin)

We anticipate that some members of some VOs will maintain relationships with different Real Organizations (ROs) over the course of their affiliation with the VO. We anticipate that different ROs will be able to support different identity protocols, and that in between affiliation with "traditional" education and research organizations, some members may only maintain identities at "social" organizations such as Facebook and Google. This should look pretty similar to Peter's original use case.

- (Shibboleth Project, Scott Cantor)

My interest is in ensuring that the use of OpenID in conjunction with Shibboleth is sensible and as much as possible seamless to applications. I prefer to think that Shibboleth integration models represent a coherent, consensus-based view of how federation and SSO ought to be integrated, and would just like to maintain that with OpenID in the picture. In particular, I think user identification, discovery, and provisioning should be addressed with a protocol-neutral view where possible, and I would like to see multiple deployment architectures possible without changing applications.

Some of the scenarios I can imagine:

A SAML IdP can leverage OpenID as an authentication source. This is addressed in SAML as "proxying". This might be done by hiding the details from the SAML SP, treating all OpenID sources as a single SAML IdP, and mapping OpenID identities either through links or transformations into familiar SAML-friendly attributes or subject identifiers. A challenge here is what to do about "ahead of time" provisioning, if knowing somebody's OpenID can't be used to predict the identifier the application will be getting. Another issue is discovery, since a naive approach might lead to a multiple hop selection of one IdP followed by an OP.

The OpenID information might also be passed through the gateway and carried in some new SAML form, so that the application can handle the OpenID identifier directly when needed.

Or the gateway might be replaced by a more capable SP stack that handles both SAML and OpenID (and possibly other protocols), in which case the software itself may need to know how to represent the different types of users. The issues there should be largely the same as in the gateway case, but it makes discovery less likely to cause "new" problems.

- (UCSF Library, Lucas Rockwell <lucas.rockwell@ucsf.edu>)

We would like to allow temporary and even long-term guests to authenticate to our wiki, which is used by faculty, students, researchers, and staff across campus. Most of the guests are temporary, as they might be a guest lecturer for one day, and then might need to post some content or comments on the wiki. These guests are most likely from a corporation or doctor's office, i.e., not from an InCommon institution.

As Scott points out, "ahead of time" provisioning is difficult, but once the user authenticates, the space admin (we use Confluence) could then add the guest to a particular space (perhaps by looking the user up by name or email address).

- (UC Berkeley, Dedra Chamberlin <dedra@berkeley.edu>)

At UC Berkeley, we have been getting an increasing number of queries from campus departments who would like to provide services to short-term guests. Most of them aren't concerned about LoA and just want to establish a lightweight profile for guests and keep a record of what they do with the service (how often they visit, what types of features
they use, etc). Examples include our Matterhorn/OpenCast developers, our campus Sakai team, a group implementing the Alfresco content management system, and managers of our campus SharePoint service. I'm sure there are others.

They would very much like to avoid the overhead of forcing these users to create an account before they log in, whether it be setting up an account in our identity management systems or creating an account with something like ProtectNetwork or openidp.org to accept Shibboleth or SAML2 assertions.

Many of the campus departments I deal with would like to allow guests to authenticate using any number of existing identifiers, from OpenID, Facebook account, Google account, etc. They would especially like to set up their apps to use a single authentication protocol (Shib or CAS), and still provide users the choice to use some existing external credential.

- (Umeå University/SWAMI, Roland Hedberg (roland.hedberg@adm.umu.se)

I was commissioned by SWAMI to build a IdPproxy. Something that to the SAML2 world looked like an ordinary IdP but when the user wanted to authenticate provided her with a choice between Facebook, Google, OpenID, Twitter and Windows Live ID.
When I demo'ed (https://sp-test.swamid.se/ 'Logga in via SocialProxy') this at a SWAMI conference the other day we had a discussion about possible use cases. The following patterns emerged:

1. Allowing presumptive students access to the universities services in a such a way that we could keep the history of what that person had done and later if needed bind that history to the student identifier she got when enrolled as student.

Here we didn't feel we really needed to know who was at the other end, the knowledge that it was probably the same person returning was enough.

2. Allowing short term acquaintances specific access

This fell into two cases: one was short term work items (you needed input from someone outside the community and you wanted this in a specific system) and the other loose affiliations (one example given was people supervising student on trainee assignments) that for a specific time needed access to a specific system. This would need an out-of-band method for assigning, in a more secure way, an specific identifier to a specific person.

3. Keeping in contact with old acquaintance.

In Sweden we have up to now not treated alumni as in for instance the USA. Here after a grace period students and former employees loose the account that the organization has given them and thereby access to the university services. The possibility of letting persons leaving the university specifying a Facebook, Google, .. account that could be used by the person afterwards as an identity when accessing services appealed to the group.

- (Project Bamboo, Keith Hazelton <hazelton@wisc.edu>)

UW-Madison is one of several institutions participating in Mellon-funded research into collaborative platforms (VOs) for humanities disciplines. There is a requirement for an extensible "Bamboo Person" Registry but this is explicitly decoupled from a need to issue credentials to people in the registry. SAML IdPs will be essential to support several use cases in which institutionally vetted attributes, groups and entitlements are required, but for other use cases, very low assurance credentials (such as OpenID and OAuth ids) are perfectly adequate. The latter use cases involve "serial identity" by which user preferences, bookmarks, user notations and other information can be persisted from session to session.

Some users may begin by using low-assurance applications and services but later transition to researcher status with institutionally-maintained credentials and attributes. In this case, the ability to link OpenID accounts with SAML IdP-asserted identities would be a great convenience for Bamboo users.

- (The University of Chicago, David Langenberg <davel@uchicago.edu>)

We're interested in augmenting services like mailing lists, wiki, etc to allow an individual who is not affiliated with Chicago to authenticate and use those services without needing to go through the processes of creating a local account. In an ideal process the service would allow you to use any credential you could come up with (OID, Facebook, SAML, others) to register and login.

- (The Pennsylvania State University, Chris Hubing <cjh@psu.edu>)

We are enabling OpenID authentication for our wiki service.

I have a prototype up at the following URL: https://wikispaces.psu.edu/. Click login, select OpenID providers. If the email attribute is passed from the OpenID IDP (e.g. Google and Yahoo) and the domain matches, it is used as EPPN.

The Earth Science Women's Network (ESWN) is a VO that is distributed around the country. Not all institutions are in InCommon, but they still need to be able to collaborate. We will be enabling OpenID auth for their web site (https://eswn.et-test.psu.edu/).

- what architectural model do you use (a gateway mapping multiple input protocols to SAML, SP providing native support for multiple protocols, other ?)

Shib IDP behind a OpenID relying party running modified mod_auth_openid

- do you store info about the person in some sort of database or Registry (ie remember information about them), or is everyone "new" every time they arrive ?

OpenID identifier or email address

- do you tell the application which protocol was used for authentication, or convey some information about LoA ?

not as of yet.

- when granting permissions, when and how do service managers identify users and do the granting ?

out of band

- do you have some sort of invitation mechanism? An existing user can cause an email to be sent to someone, inviting them to use the service ?

no

- what's the syntax of the "unique name identifier" that you present to the application ?

email or openid identifier

- (Clinical Informatics Research Group (CIRG), University of Washington, Kim Goldov <kgoldov@uw.edu>)

CIRG has created an Apache module, "Apache2::AuthAny" as an extensible authentication/authorization system. The module currently protects the Distribute project. Distribute is used to collect influenza data from local health departments across the USA, and disseminate it to authorized 3rd parties.

Prior to the use of Apache2::AuthAny, the site was duplicated at two URLs; one protected by Shibboleth, and the other by basic authentication. Basic authentication was needed for programmatic access to web services, and as backup protection against outages to Protect Network. Even though accounts can be created in minutes though Protect Network, Protect Network account creation was seen as a deterrent for some potential users. The initial Apache2::AuthAny requirements developed from this predecessor system.

– *Allow Google authentication* - Many of our users already used Gmail and other Google services.
– *Maintain existing accounts* - It would not have been feasible to ask our existing users to switch from using their "UW NetID" or other InCommon account.
– *Protect multiple applications* - Protect portions of the Distribute application written in Python (Django), fast cgi (Perl), PHP, as well as static images.

The system now in production supports these requirements along with the following features:

– Authentication through Google, basic authentication, Shibboleth (UW and Protect Network), and LDAP.
– Log out without closing the browser window after logging in with any of the authentication providers/mechanisms.
– Linking of accounts to a single identity.
– Role based authorization subsystem (based on the single identity or the IdP if the user cannot be identified)
– Apache "Require" directives for user or role authorization. Environment variables are passed allowing the protected application to make authorization choices and control behavior by user/role.
– Command-line tool for adding and updating users.

- (University of Washington, saml gateway, Jim Fox <fox@washington.edu>)

We provide a social-saml gateway (Roland's) that allows our SPs to send users to either Google or Facebook for login. Registration and configuration is self-serve, so I don't know what other people are doing with this capability. Thus far about a half-dozen non-test SPs have signed up.

The social gateway returns identity provider entity ids of https://idp.u.washington.edu/google or https://idp.u.washington.edu/facebook. I think most people use the email address as the identifier.

Of the applications I do know about:

- Our SP registry itself allows non-university people to register and manage their SPs (e.g., request attributes) with a Google login. In this case authorization comes from their email being a member of an associated group.
- We allow non-university people to participate in Canvas courses by using their Google identity. In this case the course list is pre-populated with the user's gmail address.
- I know at least one group also allows Facebook login through our gateway.

- (Virginia Tech, Account Recovery, Mary Dunker <dunker@vt.edu>

We allow our users to set up account recovery options that are used for self-service password reset. Google and Yahoo OpenID providers are allowable options. If the user forgets his/her Virginia Tech personal ID password, they can login to the password reset application with the OpenID account. The user must also know their account name and Virginia Tech ID number in order to reset the password.

- (University of California, Various Use Cases, Eric Goodman <eric.goodman@ucop.edu)

We have a multiple use cases that appear to be good fits to be addressed with the "external identities" paradigm:

1: Cross campus enrollment

We have an increasing number of students with accounts at one UC campus who are taking courses at a second UC campus. The goal is to allow the student to use their "home" campus account to access their "host" campus resources. Federated authentication here is not in itself sufficient; the "host" campus needs to recognize the student not just as an "external" student but also as an "internal" student. That is, by logging in with a "campusA" account,

the student needs to be identified as a "campusB" student. How would one campus assert "student-ness" at another campus (and why would SPs trust such an assertion)? (While in theory this could be done with entitlements, entitlements imply there would need to be a provisioning process between the host campus' source student systems and all the potential home campuses students may come from, so it seems a poor solution).

2: Incoming students, especially long-term tracking of prospects

This would include applicants, but the UC system is also working to develop support for CCC and CSU (community and state college) students who want to transfer into UC. The goal is to provide access to these students, but then also to track them through the transfer process when they become UC students (and get local UC identities).

3: Separated employees

We have an online system that allows employees to access their paystubs and W-2s. Currently this system uses a local account, but it is moving to use federated authentication against campus credentials. This works fine until an employee separates; several campuses close accounts once the account holder is no longer affiliated with the UC system. In these cases we are looking at how to continue to allow the separated employees access to their data. While creation of local accounts would work, we are interested in considering use of external IDs, and how much trust is appropriate to place on such accounts.