

Implementing the Shibboleth - EZproxy Hybrid

Is this document for me?

This document is for you if:

1. you are a staff member at an academic institution that uses EZProxy
2. you wish to begin implementing Shibboleth access to library resources
3. you are unsure about what is required and what is involved

By laying out the necessary pre-requisites and the steps involved, this document will help you decide *if* moving to Shibboleth access to library resources is feasible at your institution. The document also provides a "cookbook" for the process.

What are the pre-requisites for integrating Shibboleth and EZproxy?

1. An institution-wide (enterprise) directory service with the appropriate information.
2. An Identity management environment (policies and business practices) that governs the management of identity information in the enterprise directory. This is necessary to build and maintain the trust necessary to participate in a federation such as InCommon.
3. A Shibboleth Identity Provider (IdP) from which Service Providers (EZproxy itself, JSTOR, OCLC, Elsevier, etc) can obtain sufficient identity information about each user to authorize access.
4. An EZproxy installation that provides authenticated remote access to library resources.
5. Institutional membership in a federation such as InCommon

What are the steps?

Step 1: Configure the Identity Provider (IdP) to release standard entitlement attributes (eduPersonEntitlement)

The Best Practices document,

<https://spaces.internet2.edu/display/inclibrary/Best+Practices>, produced by the InCommon Library Collaboration Group, details the standard entitlement attribute to enable authorization for common library licensed resource contracts. In brief, the standard entitlement attribute is eduPersonEntitlement, with a standard value of 'urn:mace:dir:entitlement:common-lib-terms'. More details as to the reason for this particular attribute are available in the Best Practices document.

In order for resource providers to make authorization decisions based on this entitlement attribute the Identity Provider must be configured to release this attribute to that Service Provider (SP). This step must be performed by the administrator of the institution's Identity Provider, which is typically the central IT office, and can differ from SP to SP. These policies are referred to as attribute release policies (ARPs). The attribute release policy is the mechanism by which the IdP administrator will release the eduPersonEntitlement attribute with the common-lib-terms value. In order to release this attribute properly, the IdP administrator must know two pieces of information: 1) the users for which this attribute should be released, and 2) the Service Providers to which the attribute should be released.

1. Appropriate users: The common-lib-terms entitlement value was established to represent the members of an institution that are included in the terms of typical

library contract with a library resource provider. The interpretation may vary slightly at each institution, but this typically includes students, faculty, staff and people physically present in the library. The common-lib-terms value is registered with MACE and documented here: <http://middleware.internet2.edu/urn-mace/urn-mace-dir-entitlement.html>.

To determine the users covered under the common-lib-terms entitlement, the administrator of the Identity Provider will need to use some logic based on available user attributes from the enterprise directory. The library and the Identity Provider administrator will need to work together to interpret common-lib-terms entitlement and the logic needed for implementation.

2. Appropriate Service Providers (Resource Providers): Given that this attribute is agreed to as a Best Practice for the InCommon community, this entitlement attribute *can* be released to all Service Providers in the InCommon Federation. IdP administrators may choose, though, to release this selectively to Service Providers. If the latter is the case, the entitlement attribute will need to be released to each resource provider with which the institution enables Shibboleth access, as well as the library's EZproxy installation (to be discussed in next step).

Step 2: Shibboleth-enable the EZproxy installation

The next step in integrating Shibboleth and EZproxy is to configure EZproxy to be a Shibboleth Service Provider. This is also referred to as Shibboleth-enabling EZproxy.

OCLC, in the EZproxy support documentation, provides instruction on how to configure EZproxy as a Shibboleth Service Provider (<http://www.oclc.org/support/documentation/ezproxy/usr/shibboleth.htm>). This document does a good job of introducing the concepts and configuration options, and provides a "quick configuration" section. The shibuser.txt section can be ignored for initial setup, but will be used in step 3.

Once EZproxy is configured as a Shibboleth SP, the IdP's attribute release policy should be configured to release the eduPersonEntitlement attribute to EZproxy.

Step 3: EZproxy - authorization based on user attributes

Once EZproxy is configured as a Shibboleth SP, EZproxy can be configured to provide authorization based on user attributes. This means that, once configured, EZproxy can make decisions as to which databases a user can gain access, based on user attributes made available to EZproxy by the authenticating IdP.

This allows for EZproxy and Shibboleth to enforce policies, such as "deny all alumni users to all databases" or "allow only nursing students access to journal x." The most basic and widely used policy, though, may be to allow access to users that have the common-lib-terms entitlement.

The shibuser.txt configuration file is where EZproxy can be configured to make authorization decisions based on a user's attributes. There are two types of configurations that are useful to understand. The first is the ability to deny access to all databases based on a user's attributes. The second is to selectively allow users to access certain databases based on specific user attributes (e.g., nursing student's access to journal x).

The ability to deny access to all databases based on a user's attributes will be explained by way of example. EZproxy can be configured to deny access to all users that do not have the common-lib-terms entitlement. And these users can be directed to a particular error page, if desired. Below is the configuration command that can be added to shibuser.txt in order to enforce this policy:

```
If !(auth:urn:mace:dir:attribute-def:eduPersonEntitlement eq
"urn:mace:dir:entitlement:common-lib-terms"); Deny unaffiliated.html
```

The ability to allow selective access to databases based on user's attributes is done using EZproxy's Group directives. This works by mapping users to groups based on their attributes, and defining databases as assigned to groups as well. When users attempt to access a database through EZproxy, they must belong to the group to which the database is assigned.

Again, an explanation by way of example. Let's say that there exists a journal, Nursing Weekly, and the institution has contracted for access to this journal for its nursing department only. And let's also say that departments are represented at this institution by the user attribute, eduPersonOrgUnitDN, and the Nursing Department value for that attribute is 'Nursing.' In order to enforce the policy, EZproxy is configured such that users from the Nursing Department are assigned to group, Nursing, and the Nursing Weekly database is assigned to the Nursing group as well. Here are snippets from the shibuser.txt and config.txt.

shibuser.txt:

```
If auth:urn:mace:dir:attribute-def:eduPersonOrgUnitDN eq
"Nursing";
  Group +Nursing
```

config.txt:

```
Group Nursing

Title Nursing Weekly
URL ...
```

The concepts here can be extended to create more fine-grained policies if desired, but this example provides a sense for how this might be accomplished.

Step 4: Configure EZproxy to seamlessly enable Shibboleth access to SPs

The last step in the integration of Shibboleth and EZproxy is to configure EZproxy to be aware of resources that are Shib-enabled. For these Shib-enabled resources, EZproxy can be configured to hand off to Shibboleth for authentication and authorization, rather than proxy the user's entire session. One of the main benefits to this approach is to lower the amount of traffic that is proxied through EZproxy, which usually results in better performance and end-user experience. Another main benefit is that it allows, through Shibboleth, for resource providers to create personalized services for users in their interfaces, while maintaining the user's privacy and seamless experience. [Note, however, that to integrate that level of personalization, the IdP will need to be configured to release a personally identifiable attribute instead of the generic eduPersonEntitlement.

For more information, refer to Best Practice #1 at <https://spaces.internet2.edu/display/inclibrary/Best+Practices> .]

This step must be performed for each resource for which you wish to enable Shibboleth access. The InCommon Library Registry of Resources (<https://spaces.internet2.edu/display/inclibrary/RegistryOfResources>) has been created to help facilitate this step (however, it is always useful to contact the resource provider to confirm or if you need additional specifics). In order to set up Shibboleth access to a resource through EZproxy, the resource provider should adhere to the Best Practices. The Registry of Resources provides information as to which resources adhere to the Best Practices, where to go for help, and sample configurations. The text below explains how to use the Registry to enable Shibboleth access through EZproxy.

Resource providers generally have to manage a handful of mechanisms for providing access control to their content, including username/password, IP address validation, Referring URL, Athens and Shibboleth. The resource providers have to configure, for each of their customers, the mechanism(s) being used. Some resource providers have full-blown administrative interfaces for their customers to use to configure access; others have support via chat and still others require the help of an account rep. The Registry provides contact information for some resources, and also relevant documentation for setting up your account for Shibboleth.

If the administrator for your Identity Provider has chosen (in Step 1) to selectively release the eduPersonEntitlement attribute to resource providers, the IdP will need to be configured to release the eduPersonEntitlement attribute for each resource provider that is Shib-enabled.

Once Shibboleth is enabled for your institution's account, the resource provider needs to be configured in EZproxy. The configuration for Shibboleth-enabled resources is done with SPU (Starting Point URL) commands in EZproxy. These are SPUEdit and SPUEditVar. These commands are documented in the EZproxy documentation, <http://www.oclc.org/support/documentation/ezproxy/cfg/spuedit/>. The Registry also provides sample EZproxy configurations for each resource provider. The intent of these sample configurations is for you to replace the appropriate institution-specific variables with local variables.

Note that SPUEdit commands for a particular resource need to REPLACE any existing EZproxy configurations for a resource. Once EZproxy is configured with SPUEdit commands for a Shib-enabled resource, existing software that routes URLs through EZproxy should not need to be modified. This includes OpenURL resolvers, federated search engines, course home pages, etc.

Conclusion

The Shibboleth/EZproxy hybrid (or, if you prefer, Shibbolized EZproxy) takes advantage of the wide deployment of EZproxy among libraries. It combines the rewrite proxy's benefits with those of Shibboleth, including leveraging the university's central identity management system and eliminating the need for IP-based authN and authZ decisions.