

# Entity Naming

Shibboleth identity and service providers are used in SAML deployments, and as such, they are assigned a unique name known as an "entityID". (Older Shibboleth versions referred to this name as a "providerId".) Any entity implementing SAML profiles is required to identify itself using an entityID. These names/identifiers are used throughout various configuration files, in [Metadata](#), as cookie values, and on the wire inside SAML messages and assertions to tie the entire system together.

## Background

Many XML-based technologies rely on the concept of a URI as a name. A URI is a general term that includes both URLs and URNs.

URLs are well known on the web as the means of identifying and locating resources or documents. In the XML world (for example XML namespaces), URLs used as names don't necessarily have to refer to a specific, resolvable document on a web server. They may do so, often as a means of documenting whatever the name represents, but their usage as a name simply relies on the URL being a unique string of characters that is "self-administrated" through the reliance on a DNS-based hostname in the URL. In other words, URLs allow for unique naming without relying on an additional registry to assign control over the names. For example, it's assumed that a name like "https://www.osu.edu/stuff" is something that only Ohio State University can use, because the DNS domain "osu.edu" is owned by that organization.

URNs are a special kind of URI that are less commonly encountered, but look similar to URLs (and begin with the characters "urn:"). URNs are different from URLs primarily because they tend to rely on some other mechanism to exert control over them.

## Persistence

The most important attribute an entityID needs to have is persistence. The entityID is the public identifier for a deployment and is not only used throughout the deployment's own configuration, but more importantly will be used throughout all of the **other** deployments that it interoperates with. As such, changing it will have ripple effects throughout many systems you don't control, and will often take time to propagate (possibly a lot of time).

For this reason, it is strongly recommended **NOT** to use the physical hostname of a server running Shibboleth as the entityID. As time passes, things get moved and that deployment may not always live on the same box. Additionally there may be multiple logical deployments of Shibboleth on a single physical server, each requiring their own unique entityID, so using the server's name doesn't scale beyond a single one.

Instead, use an entityID that describes the service itself. For example, if the School of Engineering at Example University is protecting their Blackboard installation, a reasonable entityID for the SP might be `https://engineering.example.edu/blackboard/shibboleth`.

## Resolution

As mentioned earlier, whether an entityID can actually be resolved into something is generally a secondary issue. However, SAML V2.0 defines a fairly obvious way of obtaining metadata about a given entity by resolving an entityID URL (see section 4.1 of the [SAML Metadata Specification](#)).

For this reason, select a URI that you control directly and **could** resolve at some future date. This is generally not difficult to do because a well-chosen name that has good persistence will usually correspond to a service's public/logical DNS name. When you offer a service to a significant numbers of users, getting them to switch to a different name after they're used to one is effectively impossible.

## Selection/Assignment

Some Shibboleth federations have strict policies governing the selection of an entityID, though this is more common with IdPs than SPs. In other federations, selection is up to the federation participant, but operators may enforce basic conventions or react negatively to obviously poor choices. In general, you should check with the federation(s) you plan to join, and follow the advice above.