

Importing eduGAIN Metadata

How best to import eduGAIN metadata? There are at least two options:

Option 1. Offer two aggregates:

1. <http://md.incommon.org/InCommon/InCommon-metadata.xml> (current)
2. <http://md.incommon.org/InCommon/global-metadata.xml> (new)

where the latter contains eduGAIN metadata in addition to InCommon metadata.

Option 2. Tag every entity descriptor in the production aggregate with a new entity attribute value such as:

- <http://id.incommon.org/category/extended-validation-metadata>

and then import eduGAIN metadata directly into the production aggregate.

Option 2 is strongly preferred since this method of distinguishing between InCommon metadata and eduGAIN metadata persists even if the entity descriptors are exposed as signed, per-entity metadata. Option 2 has the following additional advantages:

1. InCommon SPs can continue exposing the same set of IdPs on their discovery interfaces by filtering all IdPs **not** having the new entity attribute.
2. InCommon IdPs can continue releasing attributes to the same set of SPs by referring to the new entity attribute in their attribute release policy.
3. Besides eduGAIN entities, other "foreign" entities can be safely introduced into InCommon metadata:
 - a. Participants can introduce arbitrary entity descriptors into InCommon metadata. Entity descriptors that are vetted by the InCommon RA get the above entity attribute while those that aren't vetted get another entity attribute (or no entity attribute at all). In other words, the entity attribute indicates the relevant [metadata registration practice statement](#) in effect.
 - b. Entities registered by regional federations are a special case of the above.



Recommendation

Import eduGAIN metadata directly into the production aggregate. Start by importing IdP metadata from eduGAIN since the impact on InCommon SPs is less than what it will be for InCommon IdPs.

As a general rule, a new entity attribute precludes the need for another aggregate. Consumers simply filter entities from the production aggregate on the basis of entity attributes.



Recommendation

Tag every entity descriptor in the production aggregate with a new entity attribute value (`extended-validation-metadata`). Resist the urge to create a new aggregate.

Since the production aggregate will grow without bound, InCommon should deploy a production-quality metadata query server (mdq.incommon.org) that serves all the metadata of the world. All deployments should be advised to configure the following chaining metadata provider:

```
<MetadataProvider type="Chaining" precedence="first">
  <MetadataProvider type="XML" ... />
  <MetadataProvider type="Dynamic" ... />
</MetadataProvider>
```

In other words, every deployment is provisioned with an aggregate, which it checks first. If the deployment finds the metadata it's looking for in the aggregate, it uses that, otherwise it calls out dynamically to the query server. This allows a deployment to initially consume some subset of the production aggregate based on entity attributes.



Recommendation

Bring the beta metadata query server (mdq-beta.incommon.org) to production (mdq.incommon.org) as a distinct server environment (i.e., distinct from md.incommon.org).