


eduGAIN

Deprecated

Note that this page has been deprecated. The information it contains is no longer current. See [Interfederation and eduGAIN](#) for current information.

What is eduGAIN?

eduGAIN is a metadata aggregation service run by [GÉANT](#) for the higher ed community worldwide. The [eduGAIN service](#) routinely imports metadata from dozens of participating federations, aggregates that metadata into a single file, and then serves the signed aggregate file from a well-known HTTP location. Participating federations consume and distribute metadata in the eduGAIN aggregate at their discretion.


 InCommon signed the eduGAIN Declaration in April 2014. A small group of InCommon participants began exporting their metadata in pilot mode in the summer of 2014. Full eduGAIN participation began in February 2016.


Importing eduGAIN Metadata

InCommon Operations downloads a fresh copy of eduGAIN metadata once a day, in conjunction with our daily metadata signing process. Briefly, the import process includes the following steps:

1. Fetch the eduGAIN metadata aggregate
2. Verify the XML signature
3. Validate the expiration date
4. Enforce InCommon import policy rules
5. Ensure schema validity and metadata correctness of imported metadata

Finally, eduGAIN metadata is merged into the InCommon metadata aggregate. Thus the metadata file consumed by InCommon SAML deployments contains both metadata registered by InCommon as well as metadata registered by other federations.

 The eduGAIN metadata aggregate is not intended to be downloaded directly by IdP operators and SP owners, so please refrain from doing so out of respect for our agreement with GÉANT.

 The [Per-Entity Metadata Pilot](#) has been aggregating and serving combined InCommon and eduGAIN metadata for over a year. Browse this comprehensive [list of entities currently served from mdq-beta.incommon.org](#). (The previous page may take a long time to load.)


Entity descriptors in the InCommon metadata aggregate are distinguished by various entity attributes associated with so-called [entity categories](#), especially the [Registered By InCommon Category](#). In addition, all entity descriptors contain a registrar ID, a globally unique identifier for the registrar that initially registered the entity descriptor.

Entities in Multiple Federations

Before eduGAIN, the process of interfederation was quite difficult. As a workaround, some organizations published their metadata in multiple federations. This is no longer strictly necessary. Over time, we expect most of these organizations to consolidate their metadata in a single home federation (as long as that federation is willing and able to export their metadata to eduGAIN).

There are entities in InCommon metadata that are likewise published and distributed by other federations. Some of those entities are now exported to eduGAIN. To prevent interoperability issues, duplicate entities are wholly filtered by the InCommon import process. In other words, entity metadata published by the InCommon Federation always takes precedence.

An organization chooses a home federation so that its metadata can be maintained as a single source. To do this, the organization must synchronize its metadata sources across federations. Once all of its metadata sources are in sync, the organization's migration to eduGAIN may begin.

 Some organizations will continue to publish their metadata in multiple federations in spite of eduGAIN.

Exporting eduGAIN Metadata

Recommended Deployment Practices

The introduction of metadata registered by other federations affords InCommon IdP operators and SP owners the opportunity to review their deployment practices in preparation for federation at scale. As is often the case in federated scenarios, there's an asymmetric set of recommended deployment practices, one for IdPs and one for SPs.

A interoperable IdP typically consumes all SP metadata distributed by InCommon. Such an IdP does not filter metadata. Instead an interoperable IdP implements a rational set of attribute release rules, subject to local policy.



An IdP that does not [refresh and verify metadata at least daily](#) as recommended by InCommon should declare itself a member of the [Hide From Discovery Category](#).

A globally interoperable IdP exports its metadata to eduGAIN. Such an IdP may extend its default attribute release policy to include global SPs or it may implement individual policies for SPs based on registrar ID, entityID, or more generally, entity attributes.

The situation is quite different at the SP. An SP with a federated login interface may incur significant risk, and therefore an SP is inclined to consume exactly the IdP metadata it needs, subject to local policy. Consequently an SP may filter metadata to limit the IdP partners it is willing to interoperate with and/or expose on its discovery interface.



WARNING: Filtering metadata vs. filtering IdPs from a discovery interface are not equivalent operations since the discovery interface is easily bypassed. For instance, IdP-initiated SSO bypasses the SP's discovery interface.

The Shibboleth SP software has a remarkable set of features for filtering metadata and/or filtering IdPs from its discovery interface. A combination of whitelists and blacklists may be used, based on entityIDs or entity attributes (or both).

Both sets of deployment practices—attribute release at the IdP and metadata filtering at the SP—traditionally have been based on the entityID, but today's deployments increasingly depend on entity attributes, which are generally more flexible than entity IDs.

With respect to global metadata, the primary entity attribute of interest is the `registered-by-incommon` attribute, which signals membership in the Registered By InCommon Category. Visit our wiki for recommended [uses of the Registered By InCommon Category](#), for both IdPs and SPs.