

# Interfederation Lessons Learned

- [Methods of interoperability across federation boundaries](#)
- [Challenges](#)
  - [Of joining multiple federations](#)
  - [Of Gateways](#)
  - [Of Non-legal entities](#)
- [Metadata aggregation / exchange](#)
- [Interfederation opt-out / opt-in](#)
- [Trust issues](#)
- [Requirements for entity registration](#)
- [Entity names](#)
- [Privacy](#)
- [Interest in interfederation among InCommon participants](#)
- [Potential future work](#)

On this page, the members of the [Interfederation TAC Subgroup](#) are documenting lessons learned, recommendations, and potential future work areas /items for InCommon to consider on the topic of interfederation. This page is a companion to our [Interfederation Use Cases](#) page.

**This is a work in progress. Feel free to edit.**

## Methods of interoperability across federation boundaries

We observe five practical approaches to interoperability of federated identities across federation boundaries:

**Entity-to-entity metadata exchange:** In this method, entities (IdPs and SPs) exchange metadata directly with each other to enable interoperability that is not constrained by federation boundaries. This method avoids federation boundaries by avoiding federations, thereby sacrificing the many benefits of federations (scalability, community practice, etc.). Therefore this is not "true interfederation". Examples include 1) Google Apps and 2) LIGO's exchange of metadata with the KAGRA IdP in Japan.

**Entity joining multiple federations:** In this method, an entity (IdP or SP) joins multiple federations to enable interoperability. This also falls short of "true interfederation" because interoperability is achieved only at the entity level and not the federation level. As one example, LIGO (an InCommon member) is also joining the Italian Identity Federation (IDEM) to enable authentication from Italian IdPs to LIGO SPs.

**Bilateral interfederation:** In this method, two federations form a bilateral arrangement that enables interoperability across their combined membership. Each entity need only join its "home" federation. Scalability is improved over per-entity methods, though there is still a combinatorial cost if each federation negotiates separate agreements with every other federation. Examples include 1) [UK Access Management Interfederation trials with Edugate](#) and 2) [InCommon and UK Interfederation](#).

**Multilateral interfederation:** In this method, multiple federations join together under a common technical infrastructure and/or policy framework to enable interoperability across their combined membership. Examples include [eduGAIN](#) and [Kalmar Union](#).

**Hierarchical federation:** In this method, "smaller" federations join "larger" federations to enable interoperability across the participants in the "larger" federation. Examples include 1) the [UT System federation](#) joining InCommon and 2) regional network providers joining InCommon as envisioned in the [Quilt / InCommon Federation](#) activity.

## Challenges

### Of joining multiple federations

When an entity joins multiple federations, legal and policy issues (including signed contracts and fees) often dwarf technical issues. That said, a global metadata registration service (such as [REEP](#)) would ease the pain of joining multiple federations. Then federations would have to devise processes that consume metadata managed at the central service.

### Of Gateways

Like ADFS which masks entities from the federation

Like IdP Proxies which potentially shift problems/responsibilities around to different owners and may save, or may not save costs. (SP Proxies have similar issues).

Interfederating gateways lean toward taking all entities behind it which may not always be desirable or have entities behind the gateway comply with policies it should.

### Of Non-legal entities

Virtual organizations or research projects may be unable to sign contracts/agreements with federations which pose a challenge.

## Metadata aggregation / exchange

The group has found that the [Shibboleth Metadata Aggregator](#) works well for producing metadata aggregates for interfederation. We used it to produce aggregates containing InCommon and UK entities for an interfederation pilot project with a LIGO wiki SP.

A [SAML metadata aggregator written in python](#) by Leif Johansson is also available.

Given the potential for InCommon to participate in multiple bilateral and multilateral interfederation agreements, the group recommends that InCommon provide a single "interfederation" metadata aggregate to InCommon members that includes all external entities, rather than requiring InCommon members to keep track of different metadata aggregates for different agreements.

## Interfederation opt-out / opt-in

As InCommon adopts interfederation agreements, which InCommon entities (IdPs, SPs) should be included in the "export" aggregate for consumption by members of other federations? Requiring entities to opt-in to interfederation is a conservative approach that will limit adoption. Enabling interfederation for all entities by default and supporting some form of opt-out is the recommended long-term approach. For comparison, the UK federation is doing opt-in in the short-term and moving to opt-out in the long-term.

The issues around opt-in and opt-out are vastly different for IdPs and SPs. We see no support issues around exposing non-federated SPs to IdPs in other federations, because each SP controls (via the SP's discovery interface) what IdPs are offered to the SP's users. Exposing IdPs to SPs in other federations can introduce new support costs and raise attribute release issues (i.e., SPs outside InCommon are not bound by InCommon's Participation Agreement). Rather than defining new opt-in or opt-out tags for IdPs, the existing R&S tag could serve as a good interfederation opt-in indicator in the short term, and in the long-term, a "non-discoverable" tag could serve to indicate that IdPs should not be used by SPs both internal and external to InCommon.

Metadata scaling issues should not drive adoption of opt-out or opt-in mechanisms. Metadata scaling issues should be addressed directly by the federation and its participants and should not be used as an excuse to hinder adoption of interfederation.

## Trust issues

Trust is a core federation concern that applies also to interfederation. Both technical trust (ex. domain ownership) and behavioral trust (ex. privacy policy) are important. The federation's signature over metadata enables participants to believe in the veracity of the information contained in the metadata when making connections with business partners. Metadata tags (such as [R&S](#)) can provide trust qualifiers in metadata, indicating for example levels of assurance and categories of entities.

Regarding behavioral trust, the InCommon participation agreement places constraints on use of attributes by SPs.

Certificates in metadata are a troublesome technical issue for federations that insist on the PKIX trust model rather than self-signed certificates.

## Requirements for entity registration

InCommon and UK federations have similar entity registration procedures that include validation of DNS ownership of entity IDs and entity [scope](#) as well as validation of a "canonical name" of each entity for use when locating entities in metadata (i.e., finding your business partner's entities) and displaying IdPs in discovery interfaces.

## Entity names

Registrar validation of IdP OrganizationDisplayName and mdui:DisplayName is important because SPs show these strings in discovery interfaces. When importing names from multiple sources, name conflicts may not always be solvable at the source (for example, Tsinghua University in different national federations), and metadata aggregators and/or SPs need to be prepared to address this through name mapping.

## Privacy

As is well known, sharing of metadata is not sufficient to enable multilateral federation. Privacy concerns rooted in FERPA or EU privacy laws inhibit the release of attributes and make it difficult to realize "federation that just works." The same is true of interfederation.

The InCommon [Research and Scholarship Category](#) helps to scale attribute release within a federation. Working through REFEDS to agree on categories across federations could help with attribute release for interfederation.

The Code of Conduct approach relies on a self-asserted metadata tag that SPs use to indicate conformance to the policy. IDPs could conclude that their liability is sufficiently low if they release attributes that meet the **"necessary for the legitimate interests legal grounds"** to an SP that is asserting conformance to the CoC. This is a good approach in general to scaling attribute release across federations, related to InCommon's R&S tag.

The [Code of Conduct](#) currently does not apply outside EU, but this is a future work item for REFEDS. EU privacy laws are a challenge for US higher ed entities, which are not eligible for Safe Harbor. There has been some conversation with the Brussels-based lawyer who helped to develop the CoC, and he has made some recommendations on how to proceed with extending the CoC.

## Interest in interfederation among InCommon participants

The [Interfederation Use Cases](#) page documents interfederation use cases of interest to InCommon participants. Currently it seems interfederation is not a high priority for most InCommon participants. Of the 60 attendees of the November 2012 InCommon TAC Priorities Webinar, only 14 voted for interfederation as a priority "most helpful for you/your institution", making it the second lowest ranked, ahead only of improved metadata administration. The group's interfed email list has 32 subscribers as of April 2013.

## Potential future work

**InCommon support for <mdrpi:PublicationInfo> and <mdrpi:RegistrationInfo> in metadata:** Addition of <mdrpi:PublicationInfo> to InCommon metadata is now planned. Assuming that goes well, adding <mdrpi:RegistrationInfo> to each entity in InCommon metadata will happen later. This will help with metadata aggregation by clearly identifying the registrationAuthority and publisher for each entity. When an aggregator publishes metadata, the registrationAuthority won't change but the publisher will identify the aggregator.

**Documentation of InCommon registration practices:** Building on [InCommon FOPP](#), document InCommon registration practices to a level similar to UK Federation Technical Specifications. This documentation will be useful as input to eduGAIN. REFEDS may develop a template for registration practice statements, and if/when that happens, InCommon should conform to the template. In terms of priority, this work item is good to do but is not blocking interfederation work.

**InCommon support for hierarchical federation:** For example, automated publishing of UT metadata aggregate as input to InCommon metadata. Starting step could be: Paul logs in to InCommon and gives metadata URL. Register certificate that signed it. Related to XML submission. Or dynamic referral - InCommon delegates lookups to UT? Also support REEP? Doing this helps prepare InCommon for working with regional federations (via Quilt). NCTrust is another federation that could pilot hierarchical interfederation capabilities.

**InCommon support for bilateral interfederation:** InCommon could publish metadata aggregate containing InCommon and UK entities for [InCommon and UK Interfederation](#). This work could serve as a technology prototype for work towards eduGAIN, see below.

**InCommon joining eduGAIN:** First step is to review current eduGAIN policy framework documents at <http://edugain.org/policy>.

**InCommon providing a production interfederation metadata aggregate for its members:** As discussed above, a single InCommon "interfederation" metadata aggregate would provide a stable source of entity metadata for consumption by InCommon members wishing to interoperate with external entities based on current bilateral, hierarchical, and multilateral interfederation agreements.

**InCommon enabling the export of a subset of the InCommon metadata to be consumed by eduGAIN:** To enable federation and interoperability InCommon entities need to be consumed by eduGAIN after such time that InCommon joins eduGAIN. Initially InCommon entities could "opt-in" but for the future an "opt-out" policy would promote more interoperability and support for international virtual organizations or VOs. A possible first approach is to export InCommon entities in the Research and Scholarship category.

**InCommon support for additional entity tags:** As REFEDS and other groups develop standard entity tags, indicating (for example) whether an IdP should be included in discovery interfaces or indicating an SP's privacy policy, InCommon should provide the ability for InCommon entities to self-assert these tags.

**InCommon support for a version of the Code of Conduct that extends beyond the EU:** There is a DRAFT of an extension to the CoC that would allow EU-based IDPs to release attributes to SPs that are InCommon members if those SPs were to assert compliance. This draft should be forwarded to the InCommon lawyers for review.

See also: [June 2013 Recommendations to TAC](#)