

IdP-only Aggregate

This wiki topic describes the IdP-only Aggregate, one of several alternative [metadata aggregates](#) available to SP deployers.

<http://md.incommon.org/InCommon/InCommon-metadata-idp-only.xml> (IdP-only)

In terms of content, *the IdP-only Aggregate is a proper subset of the Main Aggregate*. Any entity descriptor that includes an `<md:IDPSSODescriptor>` element is included in the IdP-only Aggregate. Note that some entities include both an `<md:IDPSSODescriptor>` element and an `<md:SPSSODescriptor>` element—such entities are included in the IdP-only Aggregate as well.



Why is there an IdP-only aggregate for SP deployments?

The IdP-only Aggregate will help relieve issues some SPs have experienced as the size of the InCommon metadata aggregate continues to grow. It also becomes an essential part of our overall strategy for the foreseeable future since we know there are SP deployments that won't be able to leverage per-entity metadata. See the FAQ at the end of this article for more information.

The IdP-only Aggregate is signed with the same metadata signing key used to sign other InCommon aggregates. To verify the signature on the metadata, a consumer must obtain an authentic copy of the InCommon [Metadata Signing Certificate](#).



Who should consume the IdP-only aggregate?

All new SP deployments should consume the IdP-only Aggregate. Existing SP deployments are advised to migrate to the IdP-only Aggregate as necessary. An SP that implements a dynamic discovery interface is especially encouraged to migrate to the IdP-only Aggregate. See the FAQ for details.

To consume the IdP-only Aggregate, a Shibboleth SP deployment changes its metadata configuration from this:

Consuming the main production aggregate

```
<MetadataProvider type="XML"
  url="http://md.incommon.org/InCommon/InCommon-metadata.xml"
  backingFilePath="InCommon-metadata.xml"
  maxRefreshDelay="3600">
```

to this:

Consuming the IdP-only aggregate

```
<MetadataProvider type="XML"
  url="http://md.incommon.org/InCommon/InCommon-metadata-idp-only.xml"
  backingFilePath="InCommon-metadata-idp-only.xml"
  maxRefreshDelay="3600">
```

See the [Shibboleth Metadata Config](#) page for complete configuration instructions.



What are the risks of using the IdP-only aggregate?

Be advised there is no fallback aggregate of IdP-only metadata. In that sense, there is some risk associated with the use of the IdP-only Aggregate. If you must fall back, you will have no choice but to fall back to the full Fallback Aggregate described on the [Metadata Aggregates](#) parent page.

Frequently Asked Questions

I thought the plan was to distribute per-entity metadata?

Per-entity metadata is definitely the path forward but we're not quite there yet. The good news is the [Per-Entity Metadata Working Group](#) has published its final report. Please stay tuned!

Why did Ops publish an IdP-only aggregate for SP deployments?

An IdP-only aggregate helps relieve issues some SPs have experienced as the size of the InCommon metadata aggregate continues to grow. It also becomes an essential part of our overall strategy for the foreseeable future since we know there are SP deployments that won't be able to leverage per-entity metadata.

Why won't some SP deployments be able to leverage per-entity metadata?

The vast majority of SPs do not have a dynamic discovery interface (i.e., a discovery interface that depends on published metadata) and so these SPs will be able to leverage per-entity metadata without delay. In fact, many of these SPs depend on a small number of fixed IdPs so the migration to per-entity metadata will be straightforward for them.

On the other hand, for the relatively few SPs that implement a dynamic discovery interface, the benefit of per-entity metadata is less clear since these SPs require an aggregate for IdP discovery. We expect these SPs to consume the IdP-only aggregate until we address the IdP discovery problem.

Who should consume the IdP-only aggregate?

All new SP deployments should consume the IdP-only Aggregate. Existing SP deployments are advised to migrate to the IdP-only Aggregate as necessary. An SP that implements a dynamic discovery interface is especially encouraged to migrate to the IdP-only Aggregate.

What are the risks of using the IdP-only aggregate?

Be advised there is no fallback aggregate of IdP-only metadata. In that sense, there is some risk associated with the use of the IdP-only Aggregate. If you must fall back, you will have no choice but to fall back to the full Fallback Aggregate described on the [Metadata Aggregates](#) parent page.

Why didn't Ops publish an SP-only aggregate as well?

The decision not to publish an SP-only aggregate (for IdP deployments) was difficult but carefully considered:

1. Most of the entities in metadata are SPs so the benefit of an SP-only aggregate would be considerably less than that of an IdP-only aggregate.
2. An SP-only aggregate would disrupt the eventual migration of IdP deployments to per-entity metadata.

We believe that IdPs are poised to receive the most benefit from per-entity metadata. We will encourage all IdP deployments to migrate to per-entity metadata as soon as that option becomes available.