

Splitting the Aggregate

Splitting the InCommon Metadata Aggregate

This document gives an InCommon Operations perspective on the question of *splitting the metadata aggregate* into one or more smaller files.

Executive Summary

Ops would prefer not to invest resources to create additional aggregates but would rather devote our limited resources to realize the future of metadata distribution. That said, since a significant fraction of SP deployments will be unable to leverage per-entity metadata until we solve the discovery problem, *Ops recommends that a production-quality IdP-only aggregate be published* at a permanent location.

Background

InCommon deploys a pipeline of three [metadata aggregates](#) that help mitigate the inherent brittleness of the aggregate distribution mechanism. If we were to split the aggregate into separate IdP and SP components, six (6) new aggregates would be needed to maintain the advantages of the pipeline construct.

Costs and risks associated with splitting the InCommon metadata aggregate:

1. The production of aggregate metadata is an expensive manual process. That said, adding six (6) new aggregates to the mix would not substantially increase the time required to produce and sign metadata. Indeed, Ops believes we can produce and sign 1000s of per-entity metadata files without significantly impacting our current process.
2. If we choose to deploy any new aggregates, a relatively insignificant amount of development time would be required to modify the scripts that orchestrate the metadata production process.
3. A new aggregate (or a new aggregate pipeline) is a **permanent** solution with both short-term and long-term migration issues. In the short term, we will tend to confuse deployers even more than they already are. In the long term, we will have created an additional set of metadata locations that must persist indefinitely.
4. If we create a new aggregate (or a new aggregate pipeline), we put the eventual migration to per-entity metadata at risk. Claim: we only get one chance to migrate deployers to a new metadata configuration.

Remember, the following benefits are coupled with per-entity metadata: a (possibly new) highly secure metadata signing key (in hardware); a distributed, more resilient metadata production infrastructure; and an augmented security model that embraces TLS as well as XML signature. We should invest what resources we have to promote per-entity metadata and the MDQ protocol.

Strawman Proposal

1. In general, resist the urge to publish new aggregates in production.
2. In particular, *do not publish an SP-only aggregate*. Push all IdP deployments towards per-entity metadata.
 - a. Most of the entities in metadata are SPs so the benefit of an SP-only aggregate is marginal.
 - b. IdPs are poised to benefit most from per-entity metadata. An SP-only aggregate will disrupt and confuse the migration of IdPs to per-entity metadata.
3. Since the vast majority of SPs do not have a dynamic discovery interface (i.e., a discovery interface that depends on published metadata) push these SPs towards per-entity metadata (which is an easy sell since most of these SPs depend on a small number of fixed IdPs).
4. For the relatively few SPs that implement a dynamic discovery interface, consider publishing a centralized JSON metadata feed that conforms to the [published JSON schema](#) associated with the Shibboleth Embedded Discovery Service.
5. Alternatively, if a JSON metadata feed turns out to be infeasible at this time, *publish a standalone aggregate of IdP-only metadata* (but no pipeline).

Although a centralized JSON metadata feed would be fairly easy to create, there are issues (most importantly, security issues) that need to be addressed and it is doubtful that these issues can be resolved in the short term. Alternatively, a production-quality IdP-only aggregate could be published in a matter of weeks, and moreover, clients could leverage this new aggregate immediately, with no changes to the client software.



Ops Recommendation

Deploy an IdP-only aggregate in production. This not only helps alleviate the pain felt by SP owners in the short term but it becomes an essential part of our overall strategy in the foreseeable future since we know there are SP deployments that won't be able to leverage per-entity metadata until we solve the discovery problem.

Once an IdP-only aggregate is deployed, Ops will focus its efforts on the production and distribution of per-entity metadata.