Per-Entity Metadata Pilot

Per-Entity Metadata: A Pilot Study

Project Update 2016-03-01

As of today, the Per-Entity Metadata Pilot has been extended for six additional months, until September 1, 2016. To find out how to participate in this effort, or to simply follow our progress in this emerging area, you are encouraged to subscribe to the metadata-support mailing list.

In June 2013, the InCommon Technical Advisory Committee convened the Metadata Distribution Working Group. The output of that Working Group included an initial set of Phase 1 Recommendations (the implementation of which is now complete) and another set of Phase 2 Recommendations (whose implementation is a work-in-progress).

One of the Working Group's Phase 2 Recommendations is to conduct a pilot study of per-entity metadata:

> *RECOMMENDATION: Conduct a pilot study that explores the utility of [signed, per-entity metadata] as an alternative to metadata aggregates, and evaluate current implementations of this model to discover problems or identify new requirements.*

The only SAML implementations known to support *dynamic metadata query* for per-entity metadata via the Metadata Query Protocol are:

1. Shibboleth SP 2.4 (or later)
2. Shibboleth IdP 3.0 (or later).
3. simpleSAMLphp 1.14.0 (or later)

If you know of other implementations that perform dynamic metadata query, please share your experiences to the metadata-support mailing list.

It is well known that Shibboleth, simpleSAMLphp, and other Metadata Client Software can automatically refresh the entire InCommon metadata aggregate. We expect that these software implementations can also refresh per-entity metadata but one of the goals of this pilot study is to demonstrate this capability. If you know of other metadata client software or SAML implementations with the ability to automatically refresh per-entity metadata, please document this fact on the metadata-support mailing list.

Please Join the Mailing List!

All questions, comments, and feedback regarding this Per-Entity Metadata Pilot should be directed to the metadata-support mailing list.

Contents

Configure your Shibboleth software deployment for dynamic metadata query!

Pilot Overview

This beta instance of the Metadata Query Server has been in operation since Sep 2014

This *Per-Entity Metadata Pilot Study* will last for two (2) years. Here is a rough timeline:

1. [**Sep 2014**] A beta instance of a Metadata Query Server will be deployed.
   - Initially, there will be support for per-IdP metadata only
2. [**Dec 2014**] The Metadata Query Server will be upgraded.
   - Support for per-SP metadata will be added (**note**: this feature was added on Jan 13, 2015)
3. [**Mar 2016**] The Metadata Query Server will be synchronized with the InCommon preview aggregate.
4. [**Sep 2016**] The beta instance of the Metadata Query Server will be decommissioned.

If the Pilot Study is successful, a production instance of a Metadata Query Server may be deployed. In any case, the beta instance of the Metadata Query Server is guaranteed to be decommissioned at the end of the Pilot Study period.

This Metadata Query Server instance is temporary!

The endpoint locations of metadata served by this beta instance of the Metadata Query Server are **not stable** and should not be relied upon in a production setting. At the end of the Pilot Study period, this beta instance of the Metadata Query Server will be decommissioned. A production server may or may not continue in its place.

Pilot Goals

Goals and possible deliverables of this *Per-Entity Metadata Pilot Study* include:

1. Test and document an instance of the Shibboleth SP software configured for dynamic metadata query
2. Test and document an instance of Shibboleth IdP V3 configured for dynamic metadata query
    1. Note: versions of the Shibboleth IdP software prior to V3 do not support dynamic metadata query
3. Determine if Microsoft AD FS (and other non-Shibboleth implementations) can leverage per-entity metadata
4. Determine the optimal configuration for this instance of the Metadata Query Server:
    1. refresh interval on metadata sources
    2. `validUntil`
    3. `cacheDuration`
5. Determine the need for a production instance of a Metadata Query Server:
    1. Is the mdq-server reference implementation production-ready?
    2. Is an alternative implementation necessary or desirable?
    3. How will a production deployment be managed and maintained in the long run?
6. Determine the impact of per-entity metadata on the IdP discovery process and suggest alternative approaches if necessary
    1. Note: the current implementation of embedded discovery at the Shibboleth SP requires a metadata aggregate
7. Discuss the advantages and disadvantages of asynchronous HTTP in relation to dynamic metadata query and whether this should be addressed in the Metadata Query Protocol
8. Investigate hybrid metadata models that incorporate both metadata aggregates (perhaps custom aggregates) and dynamic metadata query

If necessary, one or more conference calls will be scheduled to discuss concerns, recommendations, or topics of general interest.