

migrate-to-mdq

Transitioning away from the legacy metadata service and toward MDQ

InCommon participants are advised to begin to plan for a transition away from the legacy metadata service and to the MDQ Service. New participants are strongly encouraged to initiate their metadata consumption exclusively with the MDQ Service. **While there is currently not an official date set**, in the future, there will be a point where all metadata clients that currently consume metadata from the legacy service **must** migrate to use the new metadata endpoints described in our [Metadata Distribution Service](#) documentation. At that time, please note that will mean:



1. You **must** change all your InCommon-facing SAML software to consume metadata from <https://mdq.incommon.org> (either MDQ or one of the aggregates)
2. You **must** change all your InCommon-facing SAML software to use the new metadata signing key to verify the integrity of metadata

Instructions on these two critical changes are documented in the wiki linked above. We recommend that all InCommon participants engage with their Site Administrators and Delegated Administrators to understand the implications of this change, and to develop a plan to update **all federation-facing deployments** to use the new service before the old service is retired. Failure to do so will guarantee that your federation software fails to operate correctly starting with the lapse in metadata validity for the last aggregates published on the old service.

Frequently Asked Questions

Why would InCommon retire the legacy metadata service?

With InCommon metadata approaching the 100MB mark when aggregated, it is becoming no longer sustainable for the vast majority of deployments to continue to consume metadata from an aggregate. This would require rapidly increasing memory and CPU time to parse the metadata and verify the signature. The existing metadata signing key has been in use for over 15 years, and it's time to retire it, even though it has been carefully protected over the years. The new service is also automated, which promises faster and more frequent metadata updates in the future, and reduces staff time required to sign metadata.

Do I have to use metadata query (per-entity metadata) or can I still consume an aggregate if my metadata client isn't capable of dealing with per-entity metadata?

A) InCommon will continue to provide aggregates at the new metadata location, signed with the new key, for the foreseeable future (at least the next five years), barring an unforeseen event that forces us to retire aggregates.

How do I know if my software can use the new service?

A) If your software is currently using InCommon metadata, it is virtually guaranteed that it can use either the per-entity metadata service (Shibboleth, SimpleSAMLphp and PySAML/SATOSA are all known to work with per-entity metadata. We believe software such as CAS and ADFS will also work with per-entity metadata, to a greater or lesser extent). For clients that can't work with per-entity metadata or MDQ, you can still download aggregates.

How do I know that this new service will be stable?

The new service has been designed using modern practices and infrastructure, to offer high availability and performance. Status over the lifetime of the service can be seen at: <https://status.incommon.org/>

Related content

- [Migrating to the MDQ Service](#)
- [Metadata Distribution Service Documentation](#)
- [Introducing per-entity metadata service](#)
- [Configure Shibboleth service provider](#)
- [Configure other software](#)
- [Prefetch an entity with Shibboleth](#)
- [Locating the production metadata](#)
- [Metadata signing key for the Production environment](#)
- [Metadata Service](#)
- [Configure Shibboleth identity provider](#)

Get help

Can't find what you are looking for?

[help](#) [Ask the community](#)