

# Phase 2 Recommendations

## Phase 2 Report of the MD-Distro Subcommittee

- [Introduction](#)
- [Multiple Metadata Aggregates: Expanded Usage](#)
- [Per-Entity Metadata: A Pilot Study](#)
- [Key Management Practices: A Community-Driven Review](#)
- [Hardware Security Modules: A Landscape Study](#)
- [SAMLbits.org](#)

### Terminology

- *signing key*: a private key used for signing metadata; an RSA 2048-bit private key
- *signing certificate*: an X.509v3 certificate containing a public key used to verify the signature on a metadata file; a container for an RSA 2048-bit public key

### Introduction

The following Phase 2 deliverables were included in the [Phase 1 Implementation Plan](#):

1. Elicit and capture short to mid-term requirements for metadata aggregation
  - Multiple [metadata aggregates](#) will be deployed in conjunction with Phase 1
2. Devise a plan to transition the metadata signing algorithm to SHA-2
  - SHA-2 is an important driver for Phase 1
  - The [Phase 1 Implementation Plan](#) stipulates that all SAML deployments shall consume metadata signed with a SHA-2 digest algorithm by June 30, 2014
3. Determine the desirability, feasibility, and impact of changing the InCommon metadata distribution point
  - A new vhost for XML metadata distribution will be introduced in Phase 1

The following Phase 2 deliverables are included in this Phase 2 plan:

1. Expand the usage of multiple metadata aggregates.
2. Conduct a pilot study that explores the feasibility and utility of per-entity metadata
3. Conduct a community-driven review of InCommon key management practices
4. Conduct a landscape study of the potential needs and uses of hardware security modules
5. Participate in the [samlbits.org](#) project

The following issues identified in the [charter](#) were discussed but not addressed in this Phase 2 plan:

1. per-organization metadata
2. metadata aggregates based on self-asserted entity attributes
3. support for both XML and JSON formats (both signed)

### Multiple Metadata Aggregates: Expanded Usage

RECOMMENDATION: Investigate and subsequently expand the uses of multiple [metadata aggregates](#) to facilitate a broad range of metadata deployment scenarios.

- Multiple metadata aggregates were introduced in Phase 1 to facilitate the migration to SHA-2:
  1. *production metadata aggregate*
  2. *fallback metadata aggregate*
  3. *preview metadata aggregate*
- Initially, the *production metadata aggregate* will be signed using a SHA-2 digest algorithm while the *fallback metadata aggregate* will be signed using the SHA-1 digest algorithm. When the two aggregates are identical at the end of Phase 1, they become available for other uses as discussed on the [metadata aggregates](#) wiki page.
- The *preview metadata aggregate* has no actual purpose in Phase 1. It was introduced early for completeness. It is available for other uses at any time.
- The long-term intended uses of the [metadata aggregates](#) are documented in the wiki. (This wiki page is considered to be a Phase 2 deliverable.)

### Per-Entity Metadata: A Pilot Study

As the Federation grows, and campuses introduce larger numbers of SPs into InCommon metadata, the batch-oriented distribution model used today will become strained. While it's noteworthy that we are nowhere near real limits on that mechanism, there has already been progress on both specifications and software to supplement the use of more granular files for metadata distribution without losing the notion of third-party certification that underlies the current model.

While there are certainly many ways one might change this model, the concrete proposal that has so far been implemented in limited fashion is based on a REST-like API over HTTP for requesting and receiving signed, per-entity metadata in a negotiated format (typically SAML's XML format).

RECOMMENDATION: Conduct a pilot study that explores the utility of this approach as an alternative to metadata aggregates, and evaluate current implementations of this model to discover problems or identify new requirements.

1. The pilot will be focused on use of the Metadata Query Protocol ([Internet-Draft tracker](#), [working area](#)) and its SAML profile ([Internet-Draft tracker](#)).

2. A call for participation will be made to both deployers and software projects, the latter directed at projects that either have working code or are interested in developing it. The Shibboleth Project has already delivered production SP releases with per-entity metadata support, and is willing to work with the pilot study while developing the IdP capability, yet to be released.
3. The pilot may or may not make use of the existing InCommon metadata signing infrastructure and/or key.
4. Overlaps or synergies with the [samlbits.org](https://samlbits.org) conversation will be explored and leveraged if possible (but this subgoal shall not block completion of this pilot study).

## Key Management Practices: A Community-Driven Review

RECOMMENDATION: Publish an *InCommon Key Management Practice Statement* that describes current practices surrounding the InCommon metadata signing key.

1. Conduct a community-driven review of InCommon key management practices.
2. The scope of this review is limited to the current metadata signing key. Other keys managed by InCommon Operations are explicitly out of scope.
3. The InCommon Technical Advisory Committee (TAC) will appoint three (3) community members to review the key management practices employed by InCommon Operations. The review will result in a written report to be submitted to the InCommon TAC. The report will describe the current key management practices and recommend alternative practices if necessary.
4. The InCommon TAC will review the report and advise InCommon Operations accordingly.

## Hardware Security Modules: A Landscape Study

RECOMMENDATION: Conduct a study on the potential uses of Hardware Security Modules (HSMs) to secure XML signing keys and other high-value secrets.

Possible use cases for HSMs include:

1. The current metadata signing key
  - a. On-premise deployment
  - b. Impact: The current metadata production process that results in three (3) signed SAML metadata aggregates (production, preview, fallback)
2. A new metadata signing key
  - a. On-premise or cloud deployment
  - b. Impact: A new post-process that consumes the InCommon production metadata aggregate and produces a set of signed, per-entity metadata
  - c. Impact: A new post-process that consumes the InCommon production metadata aggregate and an alternate source of metadata (such as the eduGAIN metadata aggregate) to produce a combined metadata aggregate
3. A new IdP signing key
  - a. On-premise or cloud deployment
  - b. Impact: The production Multifactor IdP Proxy, an instance of simpleSAMLphp

## SAMLbits.org

RECOMMENDATION: Deploy a server node that participates in the experimental [samlbits.org](https://samlbits.org) metadata content delivery network.