TI.142.1 MDQ Service Project Close-Out Report

Document Title: MDQ Service Project Close-Out Report
Version: 1.0
Repository ID: TI.142.1
DOI: 10.26869/TI.142.1
Persistent URL: http://doi.org/10.26869/TI.142.1
Authors: Romy Bolton, Nicholas Roy, Albert Wu
Publication Date: January 30, 2020
Sponsor: InCommon TAC
Superseded documents: None
Proposed future review date: NA
Subject tags: federation, metadata, security, operations, incommon



---

Project Closeout Review Approval: 1/30/2020 by vote of InCommon Technical Advisory Committee

## KEY PROJECT INFORMATION:

| | | | |
|---|---|---|---|
| Sponsor Name | Ann West, Steve Zoppi | Program Manager | Albert Wu |
| Project Manager | Nick Roy | Target Group | Internal Ops/External Use |
| Program | Federation | Service | Metadata Query (MDQ) |
| Document Author: | Romy Bolton, Nick Roy, Albert Wu | Document Date: | 12/19/2019 |

## Background:

In December 2016 the final report of InCommon's Per-Entity Metadata Working Group was approved by the InCommon Technical Advisory Committee. The working group charter and final report addressed a number of items:

---

1. Develop a roadmap for addressing the immediate needs for reduced aggregate size, as well as intermediate milestones along a trajectory to a sustainable future state, based on the MDQ protocol for per-entity distribution of federation metadata.

2. Address issues related to reliance on this new model, including but not limited to:
    a. Availability
    b. Performance
    c. Status measurement/reporting

3. Develop requirements, risks, and recommended risk mitigation strategies for a production per-entity metadata service delivered by InCommon, including a firm definition of the scope of the service, aligned with the immediate needs addressed in the roadmap.

4. Advise InCommon staff on implementation of a solution, based on the requirements of the service.

5. Compile the outcomes of these investigations into a report to the TAC.

## Accomplishment Summary:

The Per-Entity Metadata Service Implementation Project, aka “MDQ”, was launched by InCommon Operations in August, 2018, with the bulk of the work taking place between the fall of 2018 and summer of 2019. Per-entity metadata allows an IdP or SP to consume only metadata for specific entities as needed instead of having to load the entire aggregate. Metadata is delivered through a protocol called MDQ ("Metadata Query," details at: https://datatracker.ietf.org/doc/draft-young-md-query/

**Requirements and Deliverables**

| **TAC Requirement**[1] | **As Implemented** |
|---|---|
| Equivalent security to aggregates | Better security than aggregates |
| 99.99% uptime | Architected for >= 99.9% uptime; actual uptime to date (190 days) has been 100% |

[1] http://doi.org/10.26869/TI.5.1

| | |
|---|---|
| 200ms response time for >=99% of queries on the Internet2 network | 200ms response time for >=99% of queries on the Internet2 network |
| Multi-site monitoring and notification | Multi-site monitoring and notification |

**To date, there has been zero service downtime since go-live of production in July, 2019 (as of this writing, 190 days)** While we cannot offer a specific SLA, we aim to achieve very high levels of availability and performance.
The service was designed and implemented in an entirely “cloud native” yet portable way, and served as the basis for Trust and Identity operations to undertake and prove the merit of dev-ops methodologies in the delivery of core infrastructure. In turn, these form the basis for our use of scalable infrastructure-as-as-service and iterative approach to design and implementation.

## Detailed Accomplishments:

- A number of rounds of updates and changes to production, which were thoroughly tested using our continuous integration (CI) process and then seamlessly rolled out
- Developed completely serverless infrastructure to support the service, including cloud-native approaches to lightweight web services fronted by a commercial content distribution network
- Securely created and transported a new production signing key to the cloud-hosted HSM, documented and witnessed by multiple people
- Developed key recovery strategy and HSM modification strategy backed by sharding and quorum operations (multiple people need to agree, in order for sensitive operations to take place)
- Created new service documentation to allow participants to easily switch to the new system
- Communicated design, implementation and availability with the community through a series of blogs, webinars and open office hours
- Design and implement a 24x7x365 on-call rotation and supporting infrastructure
- Learning and adaptation from pre-production experience:
    - Increased frequency of signing from once a day to hourly to accommodate emergency signing requests
    - Determined an initial path forward for fully automated signing
    - Identified and worked-around AWS CloudHSM issues
    - Identified and implemented multi-account structure to support separation of concerns/privileges
    - Identified and implemented techniques to allow multiple team members to work on script-based infrastructure updates at a time

- Identified and implemented features to accommodate monitoring, alerting and transparent reporting on status to the community (https://status.incommon.org)
- Determined a strategy for and implemented a solution for streaming AWS logs to on-premises log storage, analysis and reporting facilities
- Edge-cases such as use of mdrpi:RegistrationInfo/@RegistrationAuthority to filter metadata or configure attribute release were worked around and added to the documentation, thanks to community contributions.

## Benefit/Impact of the Solution Implemented:

The benefit to InCommon participants includes:

Instead of pre-loading and verifying thousands of entity descriptors in the InCommon aggregate, an SP or IdP will only load those entity descriptors it needs, on-demand. An IdP, for example, may only regularly interact with a fraction of the total SPs in InCommon and will not have to load metadata for every single SP into memory. This results in a significantly lower memory footprint in addition to quicker start up times, since deployments no longer need to load, canonicalize and verify a signature on the entirety of the aggregate at startup. Using MDQ, the UK Federation reports that it is possible to run the Shibboleth IdP with as little as 500MB heap size compared to the current recommendation of a minimum heap size of 1.5GB.

Additionally, because metadata for a specific entity is available at a specific URL, it is possible to pre-fetch metadata for entities you may consider high-value, such as commonly-used services or identity providers. For example, an SP in InCommon that primarily interacts with a single IdP could pre-fetch and cache that one IdP on startup and refresh it on a regular basis (just as you would with the aggregate). This setup can help minimize risks specified below.

Per-entity metadata also alleviates the brittleness of metadata aggregates, since introducing a change into a single entity descriptor no longer runs the risk of breaking *all* other federation metadata. This means, in the future, InCommon could conceivably publish “test federation” metadata in a production per-entity metadata service alongside production federation metadata, and use that to introduce new metadata features, test new ideas, and allow participants a “playground” for experimentation, in a low-risk but production-like environment.

Participants have provided several examples of their benefits from using MDQ:

- *Lafayette College was able to reduce IdP memory footprint from multiple gigabytes to less than a gigabyte.*

- *Multiple service providers report radically decreased memory usage and load times. One reported load times going from 15 minutes to 5 seconds.*

- *University of Washington commented "The UW has been using the InCommon MDQ service for our production IdP cluster for a few weeks now. Not much to report, other than it works as expected and we've not encountered any issues with it so far. I was looking forward to this for a long time and now it's here. Thanks!"*

## Project Contributors:

| |
|---|
| David Shafer - Technical Lead |
| Shannon Roddy - Security Lead, Technical Team Member |
| Chris Hubing - IAM Architect, Cloud Services Architect, Technical Team Member |
| Paul Caskey - IAM Architect, Cloud Services Architect, Technical Team Member |
| Ian Young - SME, Proof-of-Concept Lead, Shibboleth Metadata Aggregator Lead |
| IJ Kim - SME |
| James Babb - Documentation Lead, Technical Team Member (departed) |
| Rhys Smith - IAM Architect @ UK Access Management Federation, SME |
| Albert Wu - Service Owner, Documentation Lead |
| Bill Kaufman - Project Manager |
| Erin Murtha - Project Manager |
| Romy Bolton - Project Manager |
| John Krienke - SME |
| Mike LaHaye - SME |
| InCommon Technical Advisory Committee and Per-Entity Metadata Working Groups (standards and deployment WGs) - Advisory |

TI.142.1 MDQ Service Project Close-Out Report

| |
|---|
| Nick Roy - Project Lead, Operations / Strategy Lead |
| Ann West - Project Sponsor |
| Steve Zoppi - Project Sponsor |

## Lessons Learned:

Projects of this magnitude always take longer than expected or hoped. A large chunk of time was invested in making sure the new key material was created and handled securely, and that disaster recovery, chain of custody, and related processes and procedures were carefully documented and executed. In addition, we forged new ground in this project in terms of TI operational maturity, use of cloud infrastructure, adherence to and leverage of devops techniques and technologies.

When we started designing this service, we were thinking about VMs, on-premises HSMs, shell scripts, etc. When we finished, we had designed and implemented a completely serverless solution that is inherently robust, secure, testable, autoscaling, self-healing, idempotent, runtime-configurable, software-defined, reproducible, and adheres to best practices. This puts us on a solid foundation for future work to underpin the Federation and other InCommon services.

Our budget estimates for AWS are within a 5-10% margin of error of what we forecast a year and a half ago. This team has proven to be careful stewards of participant resources.

We need to get better at time estimation, and specifically accounting for the inevitable and necessary interruptions that happen in the course of a project that distract from the work of the project. We were also affected by the loss of a core technical team member as the delivery deadline approached.

We continue to face challenges with insight into security systems, processes, information and other factors across Internet2 divisions which, if resolved in cooperation with all Internet2 divisions, could act as a force-multiplying factor for improved security posture (proactive versus reactive, identification of blended threats, shared security infrastructure and analytics) across the organization.

## Remaining Deliverables:

| |
|---|
| **Transition community to the new service** - This was planned as a future scope of work although we continue to encourage participants to move to the new service when the opportunity arises. |
| **Retire legacy metadata signing process** - Future scope of work, including a move of the Federation Manager to AWS. |
| **24x7x365 channel for community to report issues** - Future scope of work, requires organizational cooperation and resourcing. Note that we already have a self-alerting capability and a 24x7x365 on-call rotation which is automatically notified of outages. |
| **Process and channels for handling non-outage/non-security incidents** - Future scope of work in collaboration with other Service Owners. |

## Future Releases

| Future Plans | Key enhancements | Rough Date |
|---|---|---|
| FM migration and retire legacy pipeline | Automated metadata signing and publication, future near-real-time signing and publication. | 2/1/2020 |
| Transition community | TBD | 8/31/2020 |
| Non-security incident handling | TBD | TBD |