InCommon Per-Entity Metadata: Architecture, Status and Next Steps

IAM Online

January 23, 2019

Albert Wu, Federation Service Manager, InCommon
Nick Roy, Director of Technology and Strategy, InCommon
David Shafer, DevOps Manager, InCommon
Shannon Roddy, Security Lead, InCommon

Welcome!

Today's IAM Online explores InCommon's upcoming Per-Entity Metadata Service. In addition to a report of the service's origin and project status, we will spend most of the hour discussing the service's operating design and architecture.

- Origin and Introduction
- Project Timeline and Status
- Service Operational Plans
- Architecture
- Security and Business Continuity Considerations

InCommon has Grown Up

Since 2008, InCommon has grown from just over entities in its metadata aggregate to nearly 9,000 entities in the combined InCommon / global metadata aggregate (4,551 in InCommon)

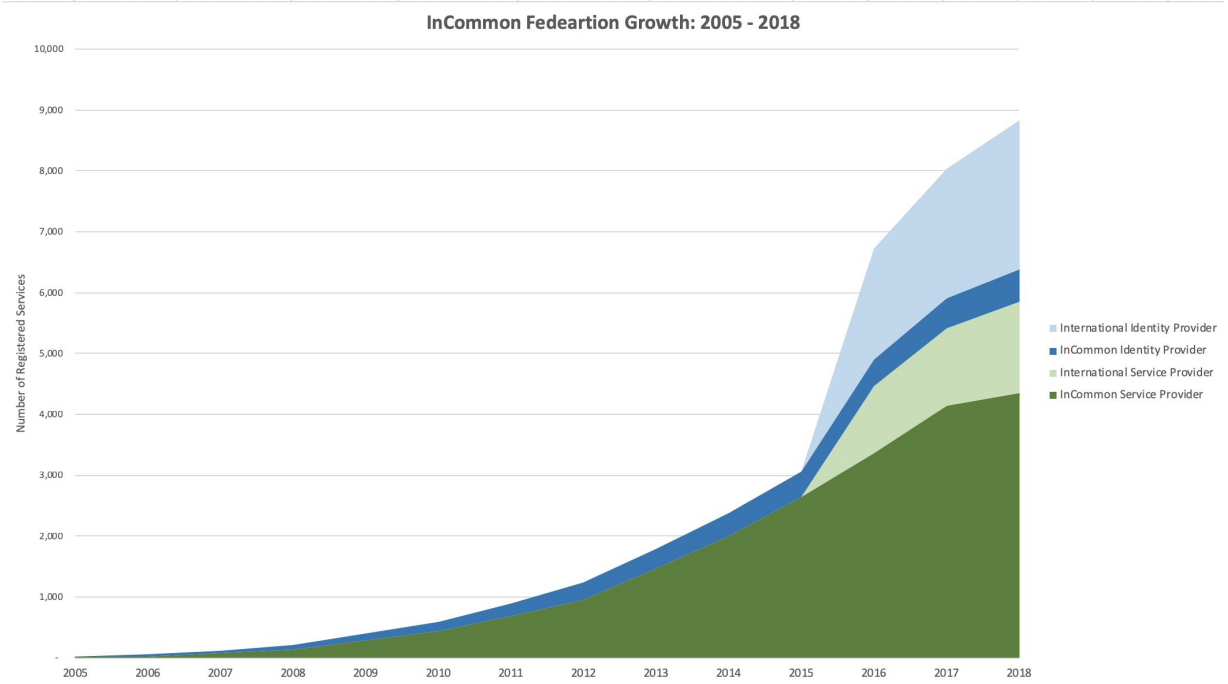

The Explosive Growth Exposed Scaling Issues

As we scale, major drawbacks surfaced in the current metata aggregate distribution strategy:

- An error in a single entity descriptor in the metadata aggregate can cause widespread outage for services depending on the metadata aggregate.
- Large metadata aggregate slows system startup time for IDP and SP’s
- Dramatically increased memory requirements wastes system resources and strains resource-constrained deployers’ ability to fully participate in federation.

Paving a Path to Per-Entity Metadata Distribution Service

2014

Metadata Distribution WG: Recommendations:

- Expand use of multiple metadata aggregates
- Conduct pilot study to explore feasibility of per-entity metadata
- Conduct landscape study of needs and uses of hardware security modules
- Participate in the samlbits.org project
- https://spaces.at.internet2.edu/x/F4G8Ag

2016 - 2018

Per-Entity Metadata WG Report (2016):

- Defined requirements for a Per-Entity Metadata Distribution Service based on the MDQ protocol
- Outlined implementation and operational considerations.
- http://doi.org/10.26869/TI.5.1

In parallel, InCommon Operations:

- Conducted MDQ pilot study
- Developed strategy to scale to the cloud
- Designed solution to align with Per-Entity Metadata WG requirements

2019

Spring 2019:

Launch Technology Preview of Per-Entity Metadata Distribution Service (MDQ Service)

Summer 2019:

Per-Entity Metadata Distribution Service goes live.

Service Operational Plans

Nick Roy
Director of Technology and Strategy, InCommon

Service Operational Plans

Requirements from the Final Report of the 2016 InCommon Per-Entity Metadata Working Group (summarized):

Full report: http://doi.org/10.26869/TI.5.1

- Mitigate risk of man-in-the-middle attacks by maintaining document signature and using TLS
- Service uptime of 99.99% on a monthly basis
- Latency of no more than 200ms for 99% of queries ***from the Internet2 network***
- Monitor performance from geographically disparate locations, make publicly available

Security and Availability

- Metadata is the root of trust for a federation
  - Both legacy and MDQ services support XML signature verification and TLS
  - MDQ service will use a hardware security module for signing metadata

- Availability of the service is critical
  - Authentication could be blocked for clients if the service isn't 'up'
  - We have chosen a commercial CDN, Amazon CloudFront, to achieve at least 99.9% uptime
    - It was infeasible from a cost perspective to achieve a 'four nines' guarantee
  - We are doing geographically distributed monitoring using Pingdom
    - Monitoring URL will be included in our service documentation
    - Looking into other options as well
  - Testing indicates that metadata in the edge cache achieves <= 100ms latency, with > 99% of queries ***from a location on or near the Internet2 network***

Cached vs. Non-Cached Metadata Response Times*

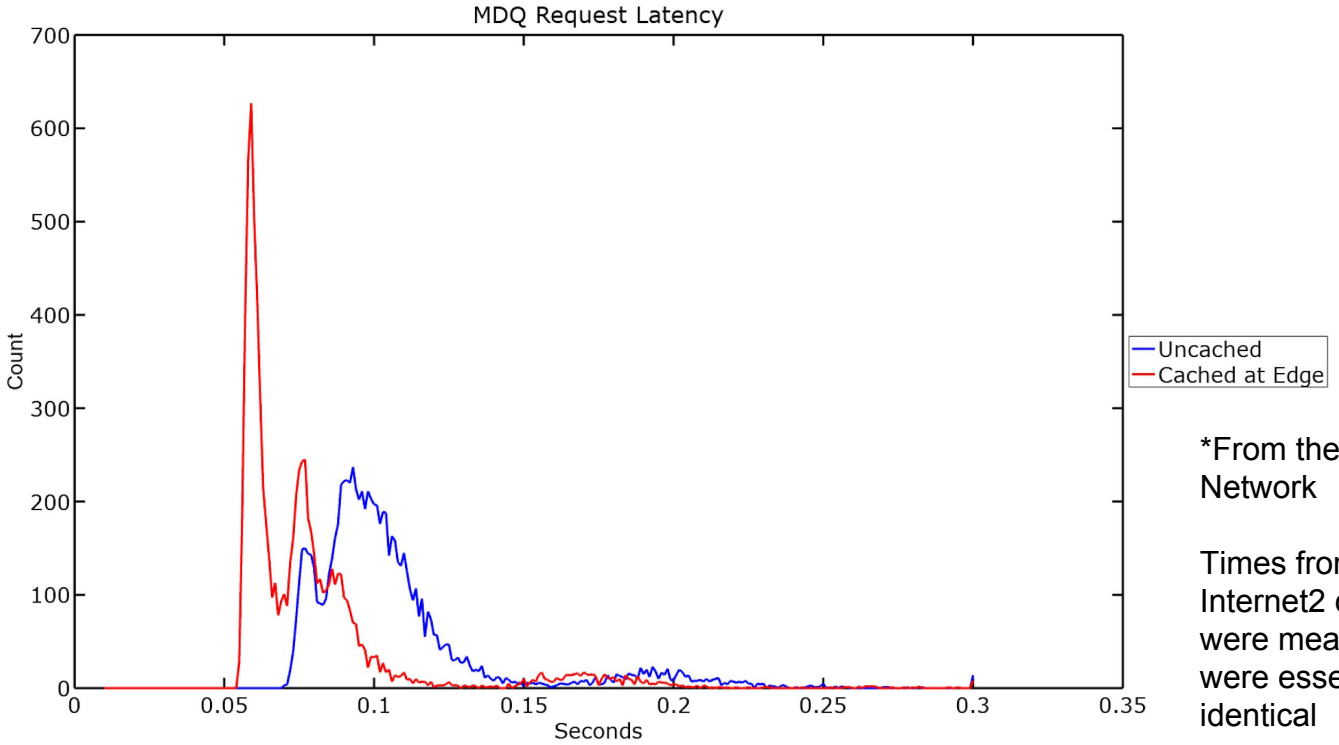

*From the Internet2 Network

Times from an Internet2 connector were measured and were essentially identical

Availability

A Note About Performance Measurement

- Pingdom shows significantly higher response times than our measurement from on the Internet2 network, and from an I2 network connector
- Likely due to a combination of timing (cache expired), network latency, spreading out their checks across different locations over a long period of time (each check hitting from a different location so edge cache takes a cache hit)
- We are looking at different monitoring options to try to show a more realistic view of actual latency

Security

Retiring Old Signing Key

- The original InCommon metadata signing key was created in a secure ceremony on March 30, 2004
- It has been in use since then
- It is time for a new, larger key
- We'll take the opportunity of changing where all InCommon SAML deployments get their metadata from to switch to a new signing key, 3072 bits in length, to provide some additional breathing room

Cost Analysis

- We want to make sure we are spending Participant fees wisely
- In order to do this, we first settled on an architectural pattern that would meet the requirements of the TAC Per-Entity Metadata Working Group as much as possible
- We then did cost analysis on the resources we would use to build this platform

It turns out that:

- CloudFront, Lambda, S3, Elastic Container Service are dirt cheap
- Hardware Security Modules are expensive, ***but!***
  - There is huge value in not sinking capital costs into on-premises HSMs and having to manage them
  - Amazon's solution is about $1,000/month, which means on a monthly basis it is still cheaper than buying our own

Bottom line: It will cost about **$2,300/month** not including staff time to support up to **100M** requests

Testing and Feedback

Technology Preview

- We plan to target current users of mdq-beta.incommon.org (the original MDQ beta instance built on a significantly different technology stack) to ask them to start using a new Technology Preview MDQ service, which has been built using the design for the production service
- We'll use this to gain operational knowledge and develop our logging/monitoring/metrics/key performance indicators for the production service
- The Technology Preview will become publicly available before production go-live (currently targeting February 4), and replace the old "preview" InCommon metadata aggregate. New features in InCommon's metadata will be introduced here first going forward

Migration Plans

Migration from Existing Metadata Distribution Service to New MDS

- Aggregates are not going away - we will still offer the production aggregate and the IdP-only aggregate on the new service
- Target go-live for the production MDQ service and new signing key is June, 2019
- At that time, we will communicate with all InCommon Site Administrators to ask them to start switching from the old InCommon aggregate and keys to the new service
- We will publish documentation for this service, the new key, etc. ahead of the launch and will point at that documentation in our communications
- The old MDS and key will be retired between six to twelve months later
- We'll monitor use of the old metadata distribution endpoints and publish the domain names that are still consuming metadata there on a wiki page, and send targeted communications to sites still using the old MDS.

Infrastructure Design

David Shafer
DevOps Manager, Internet2 Trust and Identity Services

Cloud Metadata Pipeline Prototype: Metadata Signing & Publishing Process

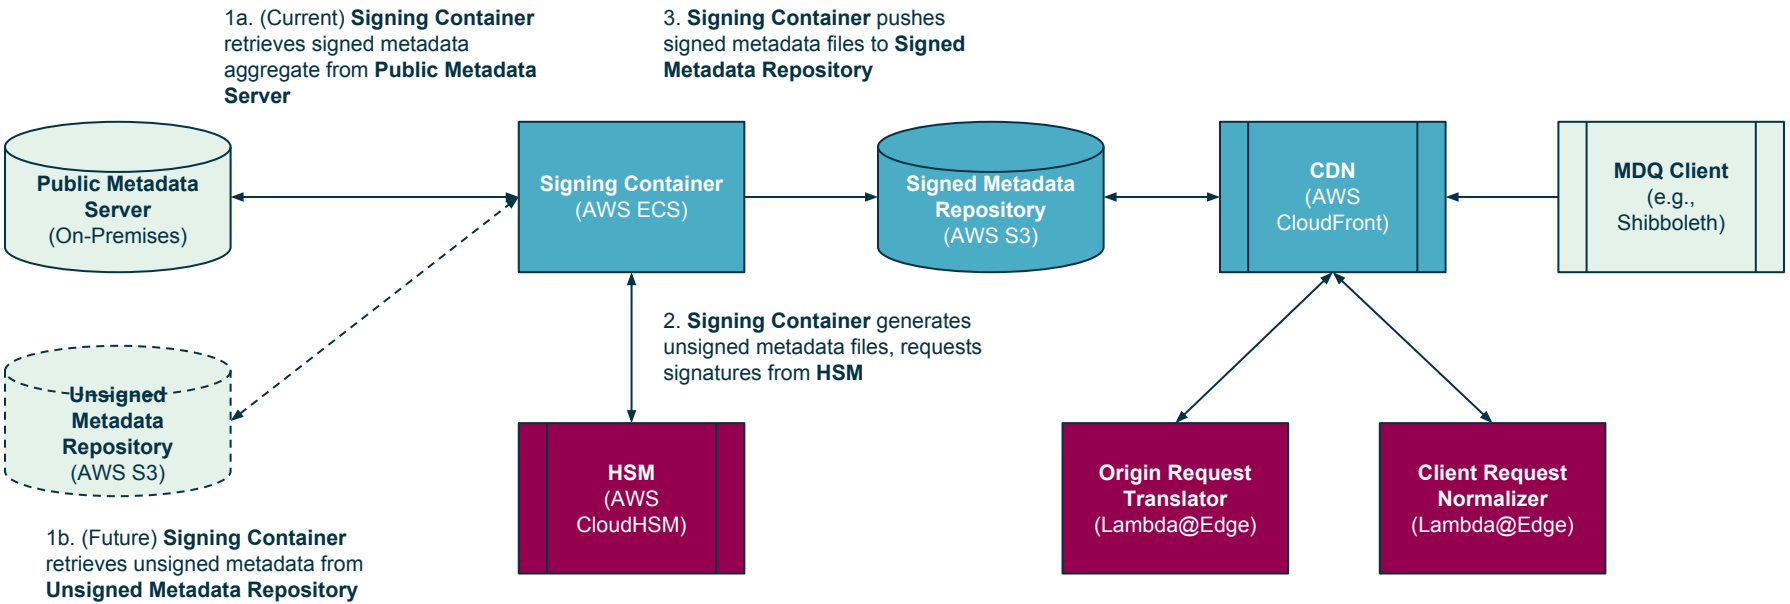

Cloud Metadata Pipeline Prototype: MDQ Responder Process

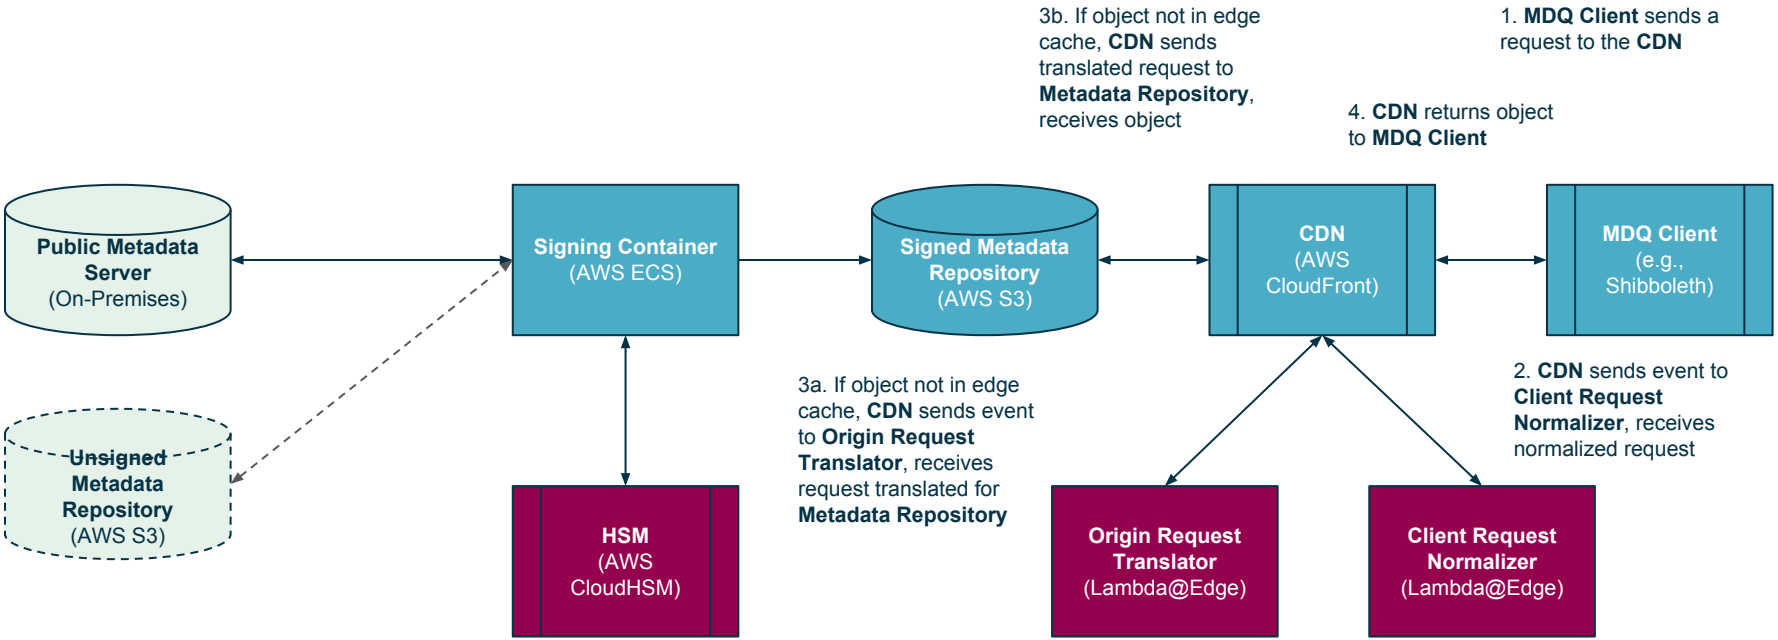

Unit and Component Tests

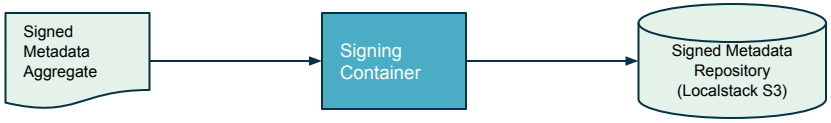

Metadata Aggregator Container
ShellCheck
docker-compose

Node.js
Mocha
(ESLint)

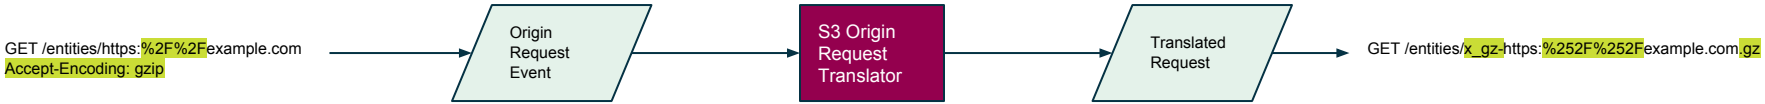

Node.js
Mocha
(ESLint)

Acceptance Tests

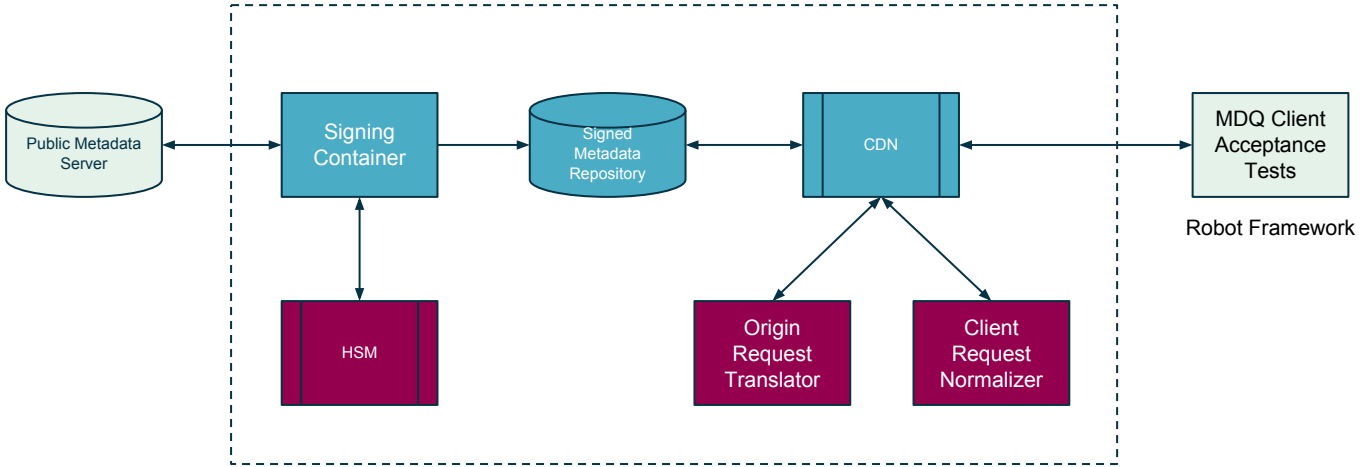

Infrastructure Code

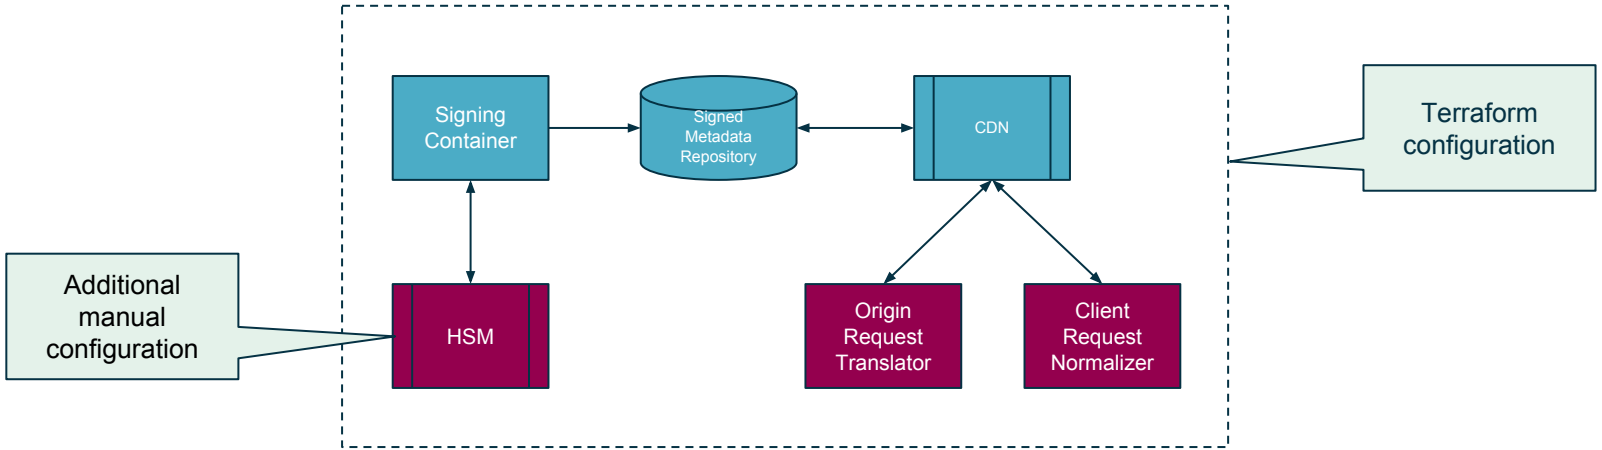

Service Security

Shannon Roddy
Security Lead, Internet2 Trust and Identity Services

CloudHSM Security

- FIPS 140-2 Level 3 validated
- Networking to the HSM is strictly controlled
  - Communication with HSM encrypted end-to-end
  - Privilege separations for operations related to HSM & VPC environment
- PKCS#11, Java Cryptography Extensions (JCE)
  - Integrates with OpenSSL through PKCS#11 libraries
  - MDA uses JCE for signing (Ian Young developed needed extensions to MDA, Thank You!)
  - JCE Library cryptographically verifies HSM
- Signing performance with a 4096 bit key ~15ms per-signing operation
  - ~8500 entities can be processed in ~130 seconds
- Security controls such that quorum is required for key and user operations
  - Export, import, user modifications, etc.

Disaster Recovery

- Key will be AES encrypted for long term storage
  - Safe deposit box, geographical redundancy
- Decryption key for encrypted key will be sharded and will require quorum to decrypt
  - Access to the stored encrypted key does not compromise the key
  - Access to stored decryption shards will not compromise the key
- Key could have been generated on an HSM and marked non-exportable in hardware, however this locks the key to that particular vendor and HSM solution and also trusts the vendor RNG

- **This scheme will be resilient to localized natural disaster, vendor lock-in, and will require quorum for all sensitive operations on the key**
- **Compromise of the stored DR key would require BOTH compromise of the sharded decryption key AND the encrypted key.**

Key Generation

- New unused & dedicated shrink wrapped hardware will be used
- Duplicate hardware is undergoing testing for HWRNG “quality”
  - FIPS 140-2 tests (rngtest)
  - Open source test suite (dieharder)
- Key generation in secure area, offline
  - Unannounced location & time
- Key will not leave generation ceremony in unencrypted form
- Chain of custody, two-man-rule, etc. will be used during critical phases of key generation & transport
- New 3072 & 4096 bit key generated simultaneously for future proofing
  - 3072 bit key will be used for production plans
  - 4096 bit key available for future rollover if needed
- Will retire old signing key & process as part of migration to MDQ

Thank You and Questions

Contact Info:

Albert: awu@internet2.edu
Nick: nroy@internet2.edu
David: dshafer@internet2.edu
Shannon: sroddy@internet2.edu

Evaluation

Please take a minute and evaluate today's IAM Online

https://www.surveymonkey.com/r/IAMOnline-Jan2019

Upcoming Events

InCommon Fee Change Open Office Hour - Thurs., Jan 24, 2 pm ET
Details: https://spaces.at.internet2.edu/x/DQJ0C

March IAM Online - CILogon and eduTEAMS – Collaboration and Virtual Organizations Made Easy
March 20, 2019 - 2 pm ET

April IAM Online - Update from the Shibboleth Consortium
April 10, 2019 - 2 pm ET

BaseCAMP - Training and Learning Opportunity - InCommon Federation
August 13-15 - Milwaukee, Wisconsin