Agenda and Notes - 2016-07-20

Per-Entity Metadata Working Group - 2016-07-20
Agenda and Notes

[EtherPad used to create these notes: Agenda and Notes - 2016-07-20.etherpad]

Dial in from a Phone:
Dial one of the following numbers:
+1.408.740.7256
+1.888.240.2560
+1.408.317.9253
331718470 #
Meeting URL (for VOIP and video): https://bluejeans.com/331718470
Wiki space: https://spaces.at.internet2.edu/x/T4PmBQ

Attendees

- Nick Roy - InCommon/Internet2
- David Walker - InCommon / Internet2
- Michael Domingues - University of Iowa
- Scott Koranda
- John Kazmerzak - University of Iowa
- Scott Cantor - tOSU (will have to drop around 10:30)
- Ian Young
- Tom Scavo, InCommon/Internet2
- Paul Engle - Rice U
- Tommy Doan - Southern Methodist University
- Steve Carmody, Brown
- Chris Phillips / CANARIE

Agenda and Notes

1. NOTE WELL: All Internet2 Activities are governed by the Internet2 Intellectual Property Framework. - http://www.internet2.edu/policies/intellectual-property-framework/
2. NOTE WELL: The call may be recorded.
3. Agenda bash
4. Do we want recordings of these calls?
    a. We'll set the calls up for recording.
5. Ian Young on MDQ protocol and his server implementation
    a. Slides: 2016-07-20 Per-Entity.pdf
    b. Goal is to keep the protocol as simple as possible while satisfying the primary use case (retrieving a single entity's metadata)
    c. Protocol has provisions to support caching for better performance. It has not been extensively tested in practice.
    d. Can be implemented using any generic web server. The URLs for entities' metadata are fairly simple and can be mapped to individual files on the server.
        i. There are issues with IDs that can't easily be encoded into URLs (*e.g.*, those that include slashes).
    e. Web proxies and caches can be used.
    f. Shibboleth (SP and IdP) can use MDQ, as can the latest release of simpleSAMLphp.
    g. mdq-server
        i. Implementation of MDQ based on Java, Spring Framework, Spring Boot, Shibboleth MDA
        ii. Collects metadata on a timer from multiple sources and indexes it by ID(s)
        iii. Responds to queries by ID with rendered documents that are cached for future queries.
        iv. Can be deployed as an executable JAR, or a Docker container.
            1. Production system has 2GB in one container. It refreshes its sources once per hour. (The refresh is asynchronous; it doesn't interfere with queries.)
    h. Ian will post his slides on the wiki.
    i. The mdq-server software is released under the Apache 2 license.
6. Note the client-side tutorial from Tom Scavo at the bottom of https://spaces.at.internet2.edu/display/perentity/MDQ+Client+Software
7. Tom Scavo with recap of InCommon per-entity metadata activities to date
    a. Slides: https://docs.google.com/presentation/d/1KF2RSl33u_UjvKm7hA2cLV1r6c0_KsnkMMi6ZxAeLgg/edit?usp=sharing
    b. Our group was preceded by the Metadata Distribution Working Group, which produced two sets of recommendations.
    c. Per-entity metadata pilot was launched in September 2014 through September 2016
        i. Involved mdq-server, as well as the Shibboleth MDA for local infrastructure (to integrate with eduGAIN).
        ii. Deployed at shibboleth.net, various I2 SPs, NIH SPs. NIH deployed a local instance of mdq-server.
        iii. Lack of discovery with MDQ limited deployment for SPs. Deployment for IdPs became available only with the most recent Shibboleth IdP V3 release.
        iv. Open questions
            1. How does server perform under load?
            2. TLS is not yet supported; perhaps not needed?
            3. How can discovery be supported?
            4. How can the signing key be best protected?
                a. Perhaps sign metadata as it's collected, not when it's distributed via MDQ?
                    i. IJ has done an experiment to find that this doesn't significantly increase the time required for the (manual, daily) signing process.
                    ii. Chris observed that this is analogous to how SAMLbits works. He will send more information via email.
8. Begin in depth discussion: What are the risks for a per-entity metadata service and the possible mitigations?
    a. (We ran out of time for this.)

9. Next call is July 27, 2016 @ 10:00 AM (America/New York)
    a. Rhys Smith will present on his MDQ work for UK fed

File	Size	Creator	Created	Comment
2016-07-20 Per-Entity.pdf	33236	David Walker (internet2.edu)	Jul 20, 2016 18:04	
Agenda and Notes - 2016-07-20.etherpad	246224	David Walker (internet2.edu)	Jul 20, 2016 15:26	