InCommon TAC Meeting 2014-07-17

InCommon TAC Meeting Minutes - July 17, 2014

Minutes

Attending: Steve Carmody, Tom Barton, David Walker, Scott Cantor, Jim Jokl, Ian Young, Steve Olshansky
With: Nate Klingenstein, Tom Scavo, John Krienke

Action Items

AI: TomS will provide an outline of a multifactor deployment plan and a timeline.
AI: TomS will provide sequence diagrams and other documentation for the multifactor migration.
AI: Scott, David, and TomB will reformulate the text in section 7.2 of the FOPP. The group will also look at the text in section 9 in light of the Google Gateway and eduGAIN but that may be deemed out of scope initially.
AI: Steve will rev the Metadata Annotation WG document.

Phase 1 Implementation Plan

- TomS reports the fallback aggregate was synced with the production aggregate on July 1, 2014. This implies that **all** metadata aggregates are now signed using the SHA-256 digest algorithm. Thus the Phase 1 Implementation Plan of the Metadata Distribution WG is now complete.
- No user feedback was received, neither positive nor negative. (One user reported a metadata outage since June 9 but that seems to be unrelated.)
- Note: The redirect from the legacy metadata aggregate (which no longer exists) to the fallback aggregate will remain in place indefinitely.

Multifactor update

- John Krienke reports that persistent discussions with Comodo have resulted in a cost effective business agreement and a corresponding technical plan to federate the CM login interface. Overall, Comodo appears to be comfortable with mobile-based multifactor authentication.
- A staging instance of the CM with a federated login interface was tested last year. The next phase of the project will allow MRAOs to log into the production CM with two factors, a federated password and a mobile device. Comodo claims they can complete this work in less than two months.
- Recall that the InCommon RAs have been logging into the FM with two factors since March 26, 2014 (which is when the Multifactor IdP Proxy was moved to production). The MRAOs will leverage the same infrastructure that the RAs are using today.
- A Statement of Work with Cirrus Identity for the next version of the MF IdP Proxy has been in the hands of the lawyers for over two months. Once the SOW has been blessed by legal, work can begin, which will take approximately two months.
- The next version of the MF IdP Proxy will allow us to begin migrating Site Admins to multifactor. RAOs will follow almost immediately since the enrollment process for Site Admins and RAOs is exactly the same.
- Q: Will the deployment leverage multifactor at the campus IdP if it exists? A: Initially, the second mobile-based factor will be situated at the MF IdP Proxy, but yes, eventually multifactor at the campus will be supported.
- Q: Is identity proofing (ala Silver) required? A: No, identity proofing at the campus is not required. InCommon Operations vets all Site Admins and RAOs but not in the sense of Silver. In particular, the real name of a Site Admin or RAO is of no consequence.
- AI: TomS will provide an outline of a multifactor deployment plan and a timeline.
- AI: TomS will provide sequence diagrams and other documentation for the multifactor migration.

Changes to the FOPP

- Steve Carmody started a discussion on the mailing list regarding proposed changes to the FOPP. Scott, David, and TomB responded to the email. The focus of the discussion is on section 7 (especially 7.2) and section 9 of the FOPP.
- There is consensus that the text in section 7 needs to change to reflect current practice in the Federation.
- There is some concern about making statements only about the IdP while avoiding statements about the underlying IdM system. OTOH, the IdP is all we control; we don't control the underlying IdM system.
- Q: Is there a difference between technical control and policy control?
- Q: Is it okay for a participant to authenticate guests with Google? A: Yes, that's why the text needs to change.
- Even though the text suggests otherwise, we shouldn't dictate to participants how they manage their IdM infrastructure. Today the POP is used to communicate how a participant's IdM system functions.
- AI: Scott, David, and TomB will reformulate the text in section 7.2 of the FOPP. The group will also look at the text in section 9 in light of the Google Gateway and eduGAIN but that may be deemed out of scope initially.

Operationalizing eduGAIN

- TomS, IJ, and Ian have been meeting regularly.
- IJ created a mini-aggregate containing three LIGO SP entity descriptors, which has been ready to go for a couple of weeks now. John and Ann are working on the policy implications of exporting the mini-aggregate to eduGAIN and interfederation in general.
- At the same, Ian was working on a long-term technical solution.
- The UKf metadata aggregation tooling has been published at GitHub for some time now. Ian forked the GitHub repository and customized it for this particular use case. The forked repository is now mature enough to be deployed to production if and when the policy issues are resolved.
- There are links in the agenda to two versions of the mini-aggregate: 1) a snapshot of the three LIGO SP entity descriptors manually produced by IJ, and 2) the output of Ian's new process, which runs against the InCommon production aggregate. The former is signed using the trusted InCommon metadata signing key while the latter is signed using a test key. Once Ian's new infrastructure has been integrated with InCommon production infrastructure, the output of the latter will be signed with the trusted InCommon metadata signing key.
- Ian's process operates on the basis of a whitelist of entities. It translates InCommon R&S entity attribute values to REFEDS R&S entity attribute values. That is a temporary measure. There is considerable policy work to do before we can export the mini-aggregate to eduGAIN.
- John reported on the status of our policy efforts. Full steering met recently to discuss the policy implications of interfederation. Ann and John laid some groundwork at that meeting. Some changes to the Participation Agreement are required, which will take some time.

- The fact that we have a concrete, production mini-aggregate ready to go is helping the policy discussion. The question is how to export the mini-aggregate quickly without getting bogged down in a lengthy policy discussion. Essentially the question is how to become more agile.
- Q: Do the FOPP changes required for eduGAIN have anything to do with the changes to the FOPP proposed above?
- Ian has also developed infrastructure to produce per-entity metadata, again built on top of the Shibboleth Metadata Aggregator. As a demo, click the link in the agenda to receive the OSU IdP entity descriptor signed with a new key generated specifically for this purpose. This works for all InCommon IdPs (just include the IdP entityID in the URL path as in the example). The infrastructure still has some bugs that need to be addressed but it is nearly operational.
- Note that the implementation has caching built in and does lazy signing. The metadata server (that supports the md-query protocol) is new.
- Next steps: Deploy the per-entity metadata infrastructure on AWS.
- Q: What about per-entity SP metadata, is that on the roadmap? A: No, the focus is on IdP metadata primarily because the Shibboleth SP is the only known software that can take advantage of per-entity metadata.

REFEDS R&S Migration

This topic has been deferred until a later time.

Metadata Annotation Working Group

- The suggestion to annotate metadata enters the conversation quite often, so maybe this is the time for a WG to try to address this more generally?
- Steven has floated a WG proposal.
- The consensus is that folks still don't have a crisp idea of what the goals of this WG are. Could the deliverables be more clearly defined?
- As an example, InCommon participation should be called out in metadata, either directly or indirectly, since some IdP operators base their attribute release policy on the privacy considerations in section 9 of the Participation Agreement.
- Suggestion: focus on the relationships and facts that need to be represented in metadata (not the technical mechanisms).
- AI: Steve will rev the Metadata Annotation WG document.

File	Modified

No files shared here yet.