Agenda and Notes - 2016-08-17

Per-Entity Metadata Working Group - 2016-08-17
Agenda and Notes

[EtherPad used to create these notes: Agenda_and_Notes_-_2016-08-17.etherpad]

Dial in from a Phone:
Dial one of the following numbers:
+1.408.740.7256
+1.888.240.2560
+1.408.317.9253
195646158 #
Meeting URL (for VOIP and video): https://bluejeans.com/195646158
Wiki space: https://spaces.at.internet2.edu/x/T4PmBQ

Attendees

- Scott Koranda (LIGO)
- Nick Roy (Internet2/InCommon)
- Ian Young
- Paul Engle (Rice U)
- Paul Caskey (Internet2)
- IJ (Internet2)
- Michael Domingues (University of Iowa)
- Scott Cantor (tOSU)
- Tom Scavo, InCommon/Internet2
- David Walker, Internet2
- John Kazmerzak, University of Iowa
- Phil Pishioneri, Penn State
- Chris Phillips / CANARIE (arrived late 10:30am EDT)

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 is being recorded.
3. Agenda bash
4. DRAFT slides for the 8/24/2016 InCommon TAC webinar
    a. https://docs.google.com/presentation/d/1YJiDpFUshWKpP77iBw1qvQeREHsRgVL8vTsvt3JEhfA/edit?usp=sharing
5. Tiered (no pun intended) architecture
    a. HA CDN-based solution operated by TSG
        i. SC: Sounds doable, assuming timeouts can be set short enough. We would want to add code to avoid servers that not behaving well for a while, then try them again.
        ii. This probably does not obviate the need for a local distribution server when there are very high availability needs on a campus (e.g., for local services, or critical off-campus services).
        iii. Could the IdP, for example, be configured to prefetch metadata for entities with a relying-party configuration? Or perhaps just from a list of "top five" (number arbitrary) critical SPs.
            1. Yes (though relying-party overrides don't always refer specifically to a single SP). It just requires code to be written... Could also be done just with a scheduled task / cron task to pull down per-entity files to the on-disk backing cache
        iv. DavidW offered to do some analysis of log files to determine the rate at which metadata is reused before it expires (i.e., how successful the client-side caching will be)
    b. Second tier operated by community? Perhaps also CDN based? Is this the role for samlbits.org?
        i. We can/should certainly recommend this. Final decision would be InCommon's.
        ii. Are we looking at primary/secondary CDNs, or two CDNs that are used relatively equally?
            1. samlbits.org is appropriate as secondary, but probably not primary.
    c. Clients can achieve higher availability goals configured with primary and then secondary as backup
        i. What's the practicality of achieving 5 9's by utilizing two independent CDNs?
        ii. Can it meet response time requirements, as well as availability?
    d. What requirements does that put on Shibboleth and SimpleSAMLphp?
    e. What are the current gaps in Shibboleth and SimpleSAMLphp?
6. Service Level Requirements
    a. Availability (How many 9's?)
        i. Achieving the best balance between what can reasonably be achieved with existing CDNs and what we can ask of Shib/SSP teams for caching
        ii. Consensus on whether (and what type of) persistent caching (between boots) is expected of IdPs and SPs?
            1. Perhaps different scenarios (*e.g.*, federation only for external services *vs.* federation for internal services, so availability of campus Internet connectivity is an issue)?
                a. We'll want to address these considerations in the report. There are more factors affecting availability to clients than just server reliability.
            2. What are the target platforms?
                a. Shibboleth, simpleSAMLphp
                b. Ping, AD?
                c. (DHW) Do we care about platforms that do not consume metadata automatically?
    b. Response time
        i. Retrieving metadata from the aggregate by an IdP or SP
        ii. Signing a new aggregate

iii. What if we move from daily to hourly signing? Separate question from how we actually deliver the service.
1. This affects cache timeouts and, therefore, the effectiveness of client-side caching - setting an overly long cache timeout could prevent upstream changes from being picked up, but this really depends on if the cache gets hit first or second

c. (DHW) Perhaps combine availability and response time? Without much thought, something like...
i. 99% of days in a year have 99.999% of response times less than 100 ms
1. (Response times during an outage are considered to be greater than 100 ms.)
ii. No day of the year has > 8.6 seconds outage/response (4 9's for a day)

d. Other service requirements?
i. All should be expressed as business requirements.

7. Distributing split aggregates
a. No time. We'll address this first next week.
b. Is this a good idea? How does it fit in our roadmap?
i. (SK input) Yes, good idea, should be in roadmap in near future
ii. (DHW) Does the end of our roadmap include aggregates?
1. (SK input) No

c. Should production of split aggregates have the same stages?