Return to Advance CAMP wiki

Advance CAMP Wednesday, Sept. 28, 2016

2:10pm-3:00pm

Hibiscus Room

SIRTFI - Security Incident Response Trust Framework for Federated Identity

CONVENER: Jim Basney

MAIN SCRIBE: Heather Flanagan

ADDITIONAL CONTRIBUTORS:

# of ATTENDEES: 24
Rich Nagle - Ohio State nagle.51@osu.edu
Brett Bieber - Nebraska
Mark Scheible - MCNC
Derek D Owens - University of Notre Dame
Eric Goodman - University of California

DISCUSSION:

https://refeds.org/sirtfi

SIRTFI is a Very Popular Topic, and we hope to share the session between the various proposers.

See the refeds website for more of an explanation and justification around SIRTFI. Also, note that SIRTFI is now a strict requirement for CERN.

We do have part of the SIRTFI framework out for consultation now (see REFEDS wiki, consultation pages). This tells you what you need to do to really implement SIRTFI, and how to be compliant. The consultation is open until the end of October. We really need to iron out any issues before we say this is ‘done’.

If you declare having SIRTFI as an entity attribute, this Means Something. See the specs linked to in the consultation.

If there is time today, we'll discuss specific requirements and look to the roadmap going forward. Will we need a management tool for SIRTFI? If we have a case for that, there is money to make it happen.

Have actual incidents run through the SIRTFI process? U. Chicago did have some experience with this, where someone found them to handle a security incident. There is heavy interest from the SPs to have a regular exercise to determine that people are paying attention. The SPs were hoping this would evolve out of the federations, but it's not clear that the fed ops have the resources to do this. It may end up being a subset that tries this.

Are there guidelines/plans for use cases of major incidents? Trigger points for when you'd expect to invoke it, when an IdP would initiate communication to an SP? Hannah has been leading some role playing exercises to explore this. We do have experience as to when this would have been useful (e.g., Heartbleed) and a managed process helpful. The current version of SIRTFI does NOT have any requirement that IdPs contact SPs when they think there has been a security incident. The SPs can contact IdPs and say they suspect a credential is compromised; the only requirement on the IdP is that they respond to requests (not necessarily do anything with it). The bar was deliberately set low. The response thing is important.

Do we want to test security contacts? Three times a year, C-SIRT does a response test, and that has improved the quality of the data.

There is a further phase (3) whose objective is to enable proactive notification, but there is a lot of detail still to be worked out.

What did we accomplish in phase 1? Created the spec. The consultation underway takes us into phase 2. The bulk of what phase 2 will be is beyond the WG; it's uptake in the community.

https://wiki.refeds.org/display/CON/Sirtfi+Consultation%3A+Sirtfi+Identity+Assurance+Certification+Description

2.1 Operational Security - comments
This is very loosely specified/weak. Alternatively, it's encouraging broad adoption. It would be nice if there were eventually a list of acceptable mechanisms.
Note that things like US federal guidelines have similar vague language, and the auditors then develop their own specific interpretations that are usually poorly informed.
Everyone in the room is saying they want and can do more.
The Dutch federation has already implemented this for all their IdPs across the board; they knew they could do this. That was what we were hoping for at this level.

One of the things that came up in an slightly different exercise was an FAQ that offered more guidance (but not official spec) -> See existing FAQ; it's there, but could be fleshed out. Remember this is global, and so we have to keep this broadly applicable.

Is there a script for self-attestation? No, it's just having an entity attribute in your metadata. It's between you and your fed op to get it there.

These are statements that the organization is saying is true about themselves. The org knows how to contact its own users. New FAQ entry?

2.2 Incident Response - comments
This has the requirement for contact information, and the security contact in metadata should fulfil that requirement. Also, that you have an incident response procedure and follow it.

2.3 Traceability - comments
This just covers logs.
Some sites have had challenges in the format of the logs that could in turn be used by a security response team. The traceability requirements were the most time consuming requirement. There should be a RFE to the default Shib logging levels (though that may in and of itself a problem). New FAQ entry on how to configure Shib to support these requirements.
If you aren't doing this to this level of detail (or perhaps cannot), are you stuck about asserting SIRTFI? Yes, though you can still put your security contact details in there and can try to get to SIRTFI compliance later.

2.4 Participant responsibilities - comments
No comments.

Note that it is a one-time compliance statement; participants are expected to remain at this status. That said, there is strong support for a test exercise. Maybe something the CISO's can encourage further? While security incidents happen all the time, the contact information for SPs and the federated nature of this isn't tested.

What would happen if one federation participant thinks another fed participant isn't in compliance? Who do they appeal to?

SIRTFI v2 needs?
See Phase 3 on the wiki
We've talked about implementing a better query service so the SP knows which IdP to contact in case of an incident, and vice versa. Also making these notifications in a privacy preserving way.

What should we pay attention to in phase 2?

Is registering with SIRTFI the equivalent of being registered as being at-risk? Is this a risk profile separate from just participating? No, it’s identifying at-risk accounts. And there are ways that accounts can be “at-risk” that don’t involve compromise.

How far back does one go in terms of account compromises? Sometime they don’t know when it happened. This comes back to traceability.

SIRTFI doesn’t require that participants require SIRTFI of others.

ACTIVITIES GOING FORWARD / NEXT STEPS:

FAQ updates:

- Clarify users and organizations in 2.2
- Add an entry on how to configure Shib for appropriate logging to satisfy 2.3

Next step for others in the room:

- Assert SIRTFI to your registrar
- Feedback in the consultation thread
- Coordination with Splunk to get something that can parse a Shib log format for the necessary information