Tues 2pm SalonsA-D

Scribing Template --Tues., Nov 12, 2013 at 2pm -- Salons A-D

TOPIC: Shibboleth V3 Update

(Notes are rough at this point, will be proofed/extended by Scott)

CONVENER: Scott Cantor

SCRIBE: Eric Goodman

# of ATTENDEES: 40+

MAIN ISSUES DISCUSSED: Status of the development process, what's done, what's available to play with or evaluate.

ACTIVITIES GOING FORWARD / NEXT STEPS:

A testbed project is available (usable via command line or Eclipse): https://svn.shibboleth.net/java-idp-testbed/trunk

See README.txt for instructions

Design Documentation: https://wiki.shibboleth.net/confluence/display/IDP30/Software+Design

Most of that's fairly mature/stable, particularly Authentication, Sessions

We are encouraging deployers with advanced authentication needs to start evaluating the new testbed for how their requirements can or can't be met. This is the real stuff in as much as we believe it will look at this point in time.

Scott Cantor Presentation

- Where we are right now
    - About 7-9 months out from delivery of V3
    - Plan has evolved ad hoc based on what becomes most critical to work on
        - SAML coding happening late in cycle
    - High Level Status
        - Libraries (OpenSAML) being polished but not really changed
        - Largely threw idp code out and rewriting
            - Good and bad
                - Probably will not be as stable as v2 on day one
                - Allowed looking at what's working and not working with current design
    - IdP Rewrite Areas of Focus
        - Much more flexible authentication system that's much more customizable
            - Step up authentication, cloud support, etc
            - Customization with less code/risk
        - Session layer and clustering support
            - Terrible under v2
            - Were planning to go to infinispan
                - Paul Hethmon did some testing with that and found that it probably wasn't the way to go
        - Also a factor: OpenID Connect
            - V2 was all about supporting SAML 1 and 2; so was very SAML and XML oriented
            - V3 is looking to decouple the SAML part, more about campus integration and less about protocol assumptions
            - Good proof was Marvin Addison developing CAS support fairly quickly, and working with dev team on design
        - OpenID support won't be here out of the gate unless a new contributor joins, but should be an easy addition (months of work)
    - Status of IdP rewrite
        - Based on top of Spring Web Flow
            - Typically SWF more used in web apps with complex UIs
            - IdP doesn't really use it that way; not really as much a UI component as a code orchestrator
            - Does support much better code decomposition
            - Possible to replace delivered elements, since not monolithic any more
                - Possible <> easy in any given case
        - Have working authn flows for existing Password/RemoteUser login handlers as well as some additional ones
            - One goal is to support ECP without a separate configuration step
            - Native support LDAP/Kerberos (not coded yet, but small work)
            - IP address based login
            - Built in username transforms (trim, case folding)
            - Kiosk configs based on ip addresses or user input (e.g. Don't cache session)
            - Proposes a new design for selection of authentication method
                - Mostly affects LoA support and those testing with the MCB
                - Makes having a form that can handle multiple methods easier to deploy
                - Not SAML specific, so authn capability can be portable across protocols
                - Full support for SAML feature allowing SPs to request "at least as strong as" or "better than" particular methods
                    - This is still a federation or local site issue to manage the relationships of the authentication methods
                    - Tracks within a flow how you used it (context wise), e.g., a flow that can be used with a weak or strong credential can be used and the system will track which method was used and only allow reuse as appropriate

        - Configuration and compatibility with v2
            - Conversion scripts or backwards compatibility
            - Backwards compatibility likely to be an issue for authentication configurations (handler.xml), since this has changed so much
                - Any customizations there will probably need to be redone.
                - Nice thing is that these flows are largely done, so clients with significant customizations can start working on this now
                - Brent Putman put together a testbed that allows testing the web flow engine, Scott added the authentication processing, which supports that testing
                    - Can be checked out of svn (see top of scribe notes)
            - Concerned about the configuration of Spring config files
                - Have split the Spring flows and bean files so that the to-be-editable parts are in separate, simpler files
                - Means customizations are done in smaller files that configure individual properties/behavior
                - Example: was asked yesterday can idp handle a password expiration event?
                    - During refeds sessions was able to generate a property switch mechanism that can call a no op flow that you can replace with your password reset/unlock process of choice.
        - Session management
            - In memory
            - Client side cookie storage
                - Common storage interface that reads/writes storage map with a cookie using same API used for server-side storage options
                    - Overall stores session info in less than 1Kb
                    - Cannot store logout-necessary info in this blob
                    - Limitation on not using Artifact binding with client-side storage may be removed for SAML 2.0
            - Likely will have memcache and JDBC back-ends also
            - The whole session layer can be bypassed globally with a simple property
            - Identity switch errors will no longer cause problems, the session layer detects any conflicts and will destroy the original user's session.
    - Attribute processing has been added to testbed since ACAMP session
    - Error handling
        - Currently just bubbles errors out as a SAML message so the testbed provides reasonable feedback

Question and Answer

Can attributes be used to drive some of the authentication processes?

- Desired, but not there yet. Attribute layer is now added to testbed, but haven't explored how to leverage that during authentication yet.

Issue with sessions is not just session but configuration. Any plan to share configuration through the clustering?

- Don't see this as an issue. Can use svn for management, among other ways.
- Did fix some issues with stateless client-side state management key rollover and added key versioning
- Question clarification: wants to be able to edit things in one place and have it provisioned
    - Scott uses rsync
    - Did add a programmatic way to force config syncs without restarts e.g. for updating metadata

Installation of the IdP assumes the install directory is off the root folder. Can this be easily fixed?

- Install directory should be modifiable
- Question clarification: there was an issue when running multiple virtual instances on one box out of one source tree
- Probably could tweak this in the installer but need an issue filed on this to understand requirements

You mentioned native Kerberos login handling. Does this mean we can expect better logging?

- Assume yes, but hasn't looked at yet. Seems unlikely it could be worse than Sun's JAAS behavior.

Mentioned a way to tickle the idp to reload changes. Is there a list of features, etc, listed on the wiki?

- Not yet
- Followup: Changing logging on the fly would be good
    - You can do that now

Any plans to move from svn to github?

- git yes, though probably not before release. Doesn't see doing it on github due to their terms of use.

Will the IdP support configuration check support (for the new spring-flow based configurations)?

- Good idea, please file a ticket.