Instrumentation Breakout Notes
Reference Architecture Breakout Notes

Notes from TIER Meeting May 19, 2016

9:00am- noon in Chicago, following the 2016 Global Summit

List of Attendees (35)

Scribes for main session

- Emily Eisbruch, Internet2
- Mike Zawacki, Internet2

DISCUSSION:

Steve Zoppi and Ann West expressed appreciation to all present for their work on TIER. Tremendous community work and outstanding collaboration went into the TIER first release. This Face-to-Face meeting provides the TIER developers and working group members a chance to meet and plan for the work ahead.

Report Outs from TIER Working Groups

Security and Audit Working Group Report
-Helen Patton, Ohio State University, Chair

- The Security and Audit Working Group kicked off in January 2016
- See Charter and Work Priorities
- From the charter: "The TIER Security and Audit working group is charged with providing ongoing recommendations, oversight, and support of the TIER project through identification and review of security and audit standards and best practices for the TIER application suite
- TIER products must be a strong link in the campus armor against security threats"
- Phase 1 of charter goes through August
- Looking at how we develop tools
- The group is gathering info
    - What kinds of testing is being done?
    - What testing tools being used?
    - What is the software delivery process?
    - What threats to the product sets face?
    - Is Security built into how you deliver the products?
- The Security and Audit Working Group will be asking the component teams for proof of what is being done in security and how it is being done.

- The Security and Audit Working Group plans to make recommendations by August 2016.
- After August, the Security WG will be looking at how campuses can leverage TIER to help with their security operations
- Need the working group members help to assess security threats
    - Comment: There are a lot of 3rd party dependencies, lots of open source which creates a huge attack surface. Rare and/or unheard of to apply rigorous security Bulk of security issues with Shib have been attributed to 3rd party dependencies
    - Helen: our goal is to mitigate as much as possible
    - Comment: Also important to include incident response plans in our security posture
    - There is a precedent for applying security risk assessments to open source projects
    - Also consider what level of testing. E.G. Duke applied Shib to Verisign assessment/testing. Need to consider at what level does it make sense to apply something similar here.

Data Structure and APIs Working Group and Entity Registry Working Group Report
-Keith Hazelton, U. Wisconsin - Madison, Chair of Data Structure and APIs Working Group
-Warren Curry, U Florida, Chair of Entity Registry Working Group

- See Data Structures and APIs Working Gtroup Wiki
- The Data Structures and APIs Working Group has developed and shared guidelines and recommendations, including:
    - TIER Standards and Guidelines
    - TIER API: Basic Group Management Operations
    - TIER API: Basic Person Management Operations
    - Instrumenting and Monitoring TIER Components
- Plan to do similar work on the Registry side.
- Need to get messaging based event management incorporated
- Will be producing roadmaps. Project portal may play a role in that
- At this point we're planning on going forward with API and Entity Registry working in parallel.
- Work with other TIER Working groups:
    - Need to sync up with the TIER Packaging working group on assessing implementation options, etc.
    - Security and Audit WG:: Looking for recommendations on securing API calls. Instrumentation and schemas for sys logs
    - From TIER leadership: Governance, documentation for same to refer to during development

- We've achieved consensus on base functionalities for **entity registries**. Now dividing up/defining those functions. We will engage contractors to work out interoperability between campuses' native systems and the API
    - Provisioning
    - Loading items onto Grouper, other software components
    - More info on TIER Entity Registry Working Group wiki
    - Comment: our campus finds the need to build data cubes about people - recommendations for that?
        - Dependant on what tools campuses are using. As much as we want to standardize there's a huge variety of options
    - Concerned about attributes like departments, interactions with other groups, etc. Tools don't currently create searchable structures.

TIER Packaging Working Group Update
- Jim Jokl, University of Virginia, Chair

- For TIER Release one, the TIER Packaging Working Group focused on these two main areas
    - Analyze the TIER components and define needs of each
    - How to best package the TIER components
- See Packaging Working Group Wiki
- Still plan to analyze data around Shibboleth, and develop default configurations
- The Packaging Working Group chose the Docker container based on survey results - many wanted to use Docker but had no expertise. We went with DOcker based on that and decided to deliver the Docker containers as a VM
    - First look version of those containers is available
    - Long term plan is to have Docker containers be standalone modules but able to be dropped into a VM
- Looking ahead to possibility of component obsolescence - e.g. how to replace a component if it stops being supported or falls out of favor
- Request to group: We want to do more automated testing of current and possible components. Automation will save human cycles. Getting people to test has been a challenge.
- Looking at log analysis tools, messaging, and security
- Goal on Shib side is to make deployment as easy as possible. Need to determine default configs. Will ask group for help with that
- Q: as groups test, Is there a common test bed?
    - A: The next step for the Packaging Working Group is to move from handbuilt containers to a production pipeline. Automated testing will figure heavily into that. We will do manual testing at the end of the pipeline so automation doesn't have to be perfect from day one
    - Help with testing recommendations, default configurations for components,

- Q: What's the timeline for the automation/pipeline?
    - A: Looking at six-month timeframe. Publicly we've been saying we'll ramp up to automation but haven't publicly committed to timeline.
- Comment: At our school we have not yet tested the containers. How many have tested?
    - As of early May, it was a couple on Shib, one on Grouper, none on COmanage. It's really important to get that first look tested as by as many of the WG teams as possible. But only internal testing. Not intended for public testing.
- Q: Is there a Docker container inside the VMs? If so how much Docker do you need to understand?
    - A: Currently it's intended to function as a black box. There's very little exposed from a configuration standpoint. No need to know Docker to use the VM version of TIER.

Instrumentation Working Group (new WG being explored)

- Steve Zoppi, Internet2

- A goal is to put in place tools and processes in place for rigorous monitoring and testing of TIER components and releases.
- There is tendency to underestimate time and effort needed for TIER development
- We have contracts with professional partners (Unicon, Spherical Cow)
- Instrumentation WG will help address the measurement issue
- By TechEx we want rigorous, repeatable numbers. Also want to consider how log output/system exhaust can be fed into Splunk.
- Need to consider how heavily components are used and look at opportunities to optimize based on that analysis
- Need to craft a pipeline that allows for intersection of testing, optimizing, etc. Instrumentation will be key for that
- Need to develop and awareness in the community of what all the components do, how they can be adapted, deployed. (Consent work especially needs education to the community)
- We're moving away from big bang TIER releases - too disruptive to development,
- moving towards smaller, point in time, regular TIER releases/updates
- Comment, In the Scalable Consent work, we need to do instrumentation on user behavior, what screens they look at for how long, what they chose to release, etc. to improve user experience and inform development.
- Q: What is the level of enthusiasm in the community around instrumentation?
    - A: There may be some concerns, but will allow the community to see results, statistics, etc. The plan is to do this transparently and demonstrate the value of the data
    - Comment: At our campus we've found huge development wins in instrumenting usage. Also valuable for security and in demonstrating issues

    - The Danish federation presents their measures on user experience, provided data driven evidence of issues and opportunities for improvement
    - Q: Why is this not as easy as using Google Analytics on everything?
        - A: There are direct internal statistics that Google Analytics can't capture. Not useful for the types of system calls, events we need. Google is good for watching user flow..

Informed Consent
- Ken Klingenstein, Internet2

- Scalable Consent wiki for more info
- Scalable consent: Scalability of users, resources accessed. Needs allow the user to set consent options once and then forget about them.
- Funded in part by federal grant but will be migrated into TIER over time
- Multi-protocol, of which Shib is just one.
- Many campuses run a consent structure internally (like Duke). Make it challenging to develop/deploy
- Privacy Lens: Carnegie Mellon project/proof of concept
- Next phase - development with Duke of changing APIs to working code (?).
- Informed Content - MDUIs that will allow users to see linkage between IdPs to SPs
- Reputation tracking to drive user decisions around attribute release
- Geek to English translator - presenting system info in a user readable form
    - Alternate title - "What Is This?"
- Ken: Will be presenting to other segments of the privacy community, informing and educating
- Q: What are some use cases for compartmentalizing items that are on a single list currently?
    - A: Need to write up characterization on our findings to date.
    - Comment: Example: releasing only relevant courses to users with appropriate, matching attributes.

Agenda creation for second half of the meeting

- Suggestions for topics to cover
    - Meaning of high availability. What components need to be HA.
    - Configuration steps for our components.
        - Might scope that down to how the components can be made to mesh. Options for 3rd party integration.
        - Not looking for "Tomcat vs Jetty" type conversation.
        - From Packaging standpoint we needed to consider how to make GUIs, backends, etc work together.

        - AD(active directory) is the one of the only things we've found that impacted configuration and deployment.
            - Could be a topic for a breakout group.
    - Ensuring security of underlying building blocks of all components.
        - procedures for test/response and a comprehensive look at security tools.
        - comment: All of those things are included in the charter of the Security & Audit working group.
    - Instrumentation. Maybe a short list of goals, formation of a project plan?
    - How we paint a picture for others/ourselves on how components fit together. Diagram? Looking for difference in ideas among group on what a fully installed instance.
        - We've found that you can do that best in a series of diagrams which lay out each component.
        - Agreed - critical for this group
    - What are the mandatory components, what are optional, and how would we handle, say, a school with legacy architecture that only wants to deploy portions of the components?
    - How to ensure information flows cleanly between TIER and Working Groups? Maybe not for this meeting but should be discussed.

Group decided to hold two breakout sessions:

- **Instrumentation**, facilitated by Steve Zoppi
- **Reference Architecture for TIER**, facilitated by Tom Jordan, U-Wisc Madison

=========

REPORT OUT FROM THE BREAKOUTS ON Instrumentation and Ref Architecture

Instrumentation - Breakout Notes

- Next Step from the Instrumentation Breakout
    - [AI] Charter and launch a new TIER Working Group to address Instrumentation
        - The goals is by Internet2 Technology Exchange (Sept 25-28, 2016) to have rigorous, repeatable numbers around TIER Instrumentation. Also want to consider how log output/system exhaust can be fed into Splunk.

Reference Architecture for TIER - Breakout Notes

- Next Step from the Reference Architecture Breakout
    - [AI] Ask the component architects group (meets weekly) to continue this conversation around TIER Reference Architecture

Discussion with the entire group about the Reference Architecture Breakout

- Started around the idea that we had a common understanding of what “Reference Architecture” is
- Identified where we are, where we need to be, looked at that gap
- How do we communicate to institutions and how TIER would fit into their architecture
    - Which function is served by which TIER component
        - Helps institutions understand what they’re will not have if they opt out of installing a given component
    - Another need is a clear list of what components are mandatory. Again, helps campuses wanting to selectively install components.

- Q: How would the list of functions be identified and communicated?
    - Warren has an example of how that might accomplished
    - Warren: Model is to lay out all the components and their relationships grouped by function. Analogous to building plans with different sheets for wiring, plumbing, etc.

[AI] Ask the component architects to continue this conversation around TIER Reference Architecture

- comment - could start there, but will likely migrate up to governance/advisory minded bodies like Ad Hoc Advisory Group (morphing soon to CACTI)

Feedback on the TIER Developers and WG Members Meeting

- Request for more F2F meetings for TIER Developers and WG members
- Could be a couple of days of intense focus
- Airport or campus hotel in Chicago? Big Ten Center near O’Hare Airport
- Comment: would like a chance to work intensely on components (such as Grouper)
- Opportunities for ReFactoring we should consider
- Internet2 Tech Ex in Miami Sept 25-28, 2016 is next obvious time to meet, but can we schedule prior to that?
- Suggestion for next meeting to give homework and make decisions when we are F2F
    - Tuesday afternoon is open at 2016 Tech Ex

Closing Comments

-Kevin Morooney, Internet2 Vice President of Trust and Identity Programs

- We are a community that gets things done and does things that matter.

- On Monday May 16, 2016 InCommon Steering and TCIC (TIER Community Investor Council) met together for 4 hours and discussed priorities and

resourcing, among other topics. There are further meetings being planned that will help maintain the communication around TIER progress, activities, and resource needs.