Wiki page for TIER Packaging working group
Short URL for this document: https://goo.gl/yb6eQh

TIER Packaging Shib/Grouper/COmanage/midPoint Working Group

NOTE: All Internet2 Activities are governed by the Internet2 Intellectual Property Policy.
https://www.internet2.edu/policies/internet2-intellectual-property-policy/

Regular call schedule: Every Monday, 4pm ET

Coordinates that began April 16, 2018 still the same

https://internet2.zoom.us/j/7343238623

Or iPhone one-tap :
US: +16699006833,,7343238623# or +16465588656,,7343238623#
Or Telephone:
Dial(for higher quality, dial a number based on your current location):
US: +1 669 900 6833 or +1 646 558 8656
Meeting ID: 734 323 8623
International numbers available: https://zoom.us/u/D1M2z

Or Skype for Business (Lync):
https://internet2.zoom.us/skype/7343238623

March 4, 2019
No packaging call scheduled. Today is a travel day to Global Summit.

February 25, 2019
Attendees (please add yourself):

- Jim Jokl - Virginia
- Bill Kaufman - Internet2
- Colin Thompson - UC Merced
- Ethan Kromhout - UNC
- Paul Caskey - Internet2

Regrets:

- Scott Koranda - SCG
-

Call Agenda and Notes

1. Quick topics
    a. Please add any quick topics here
    b. …
2. (Final) look at TIER Reference Implementations
    a. ..
    b. ..
    c.

February 18, 2019
Attendees (please add yourself):

- Jim Jokl - Virginia
- Bill Kaufman - Internet2
- Ethan Kromhout - UNC
- Keith Hazelton - Internet2
- David Walker - Internet2
- Sara Jeanes - Internet2
- Scott Koranda - SCG
- Paul Caskey - Internet2

Regrets:

-

Call Agenda and Notes

1. Quick topics
    a. Please add any quick topics here
    b. …

2. Reference Implementations
    a. Verify the TIER Reference Implementations updates
    b. Any final changes
    c. Current draft version is here
    d. Changes needed
        i. Slide

3.

January 28, 2019
Attendees (please add yourself):

- Jim Jokl - Virginia
- Keith Hazelton - Internet2
- Colin Thompson - UC Merced
- John Gasper - Unicon
- Ethan Kromhout - UNC
- Chris Hubing - Internet2
- Sara Jeanes - Internet2

Regrets:

- Chris P / CACTI
- Scott Koranda - SCG
- Bill on Vacation

Call Agenda and Notes

1. Quick topics
    a. Please add any quick topics here
    b. …

2. Base Reference Implementations
    a. Any discussion?

3. Finalize TIER as a Suite Reference Implementations
    a. **[AI]** Jim to update drawings and language

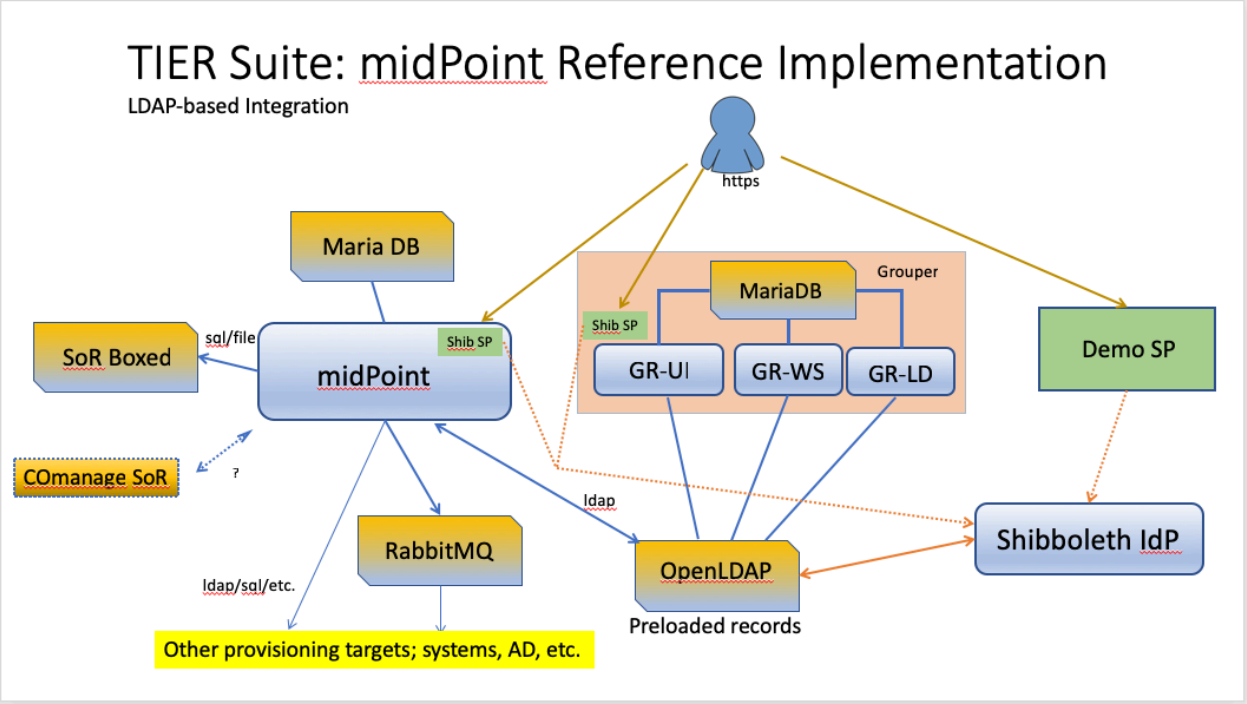

4. midPoint-Grouper Suite - integration via LDAP
    a. Backup plan - preferred solution next
    b. midPoint parses openldap change logs for transactions
    c. Picture Changes
        i. …
    d. Flow definitions
    e.

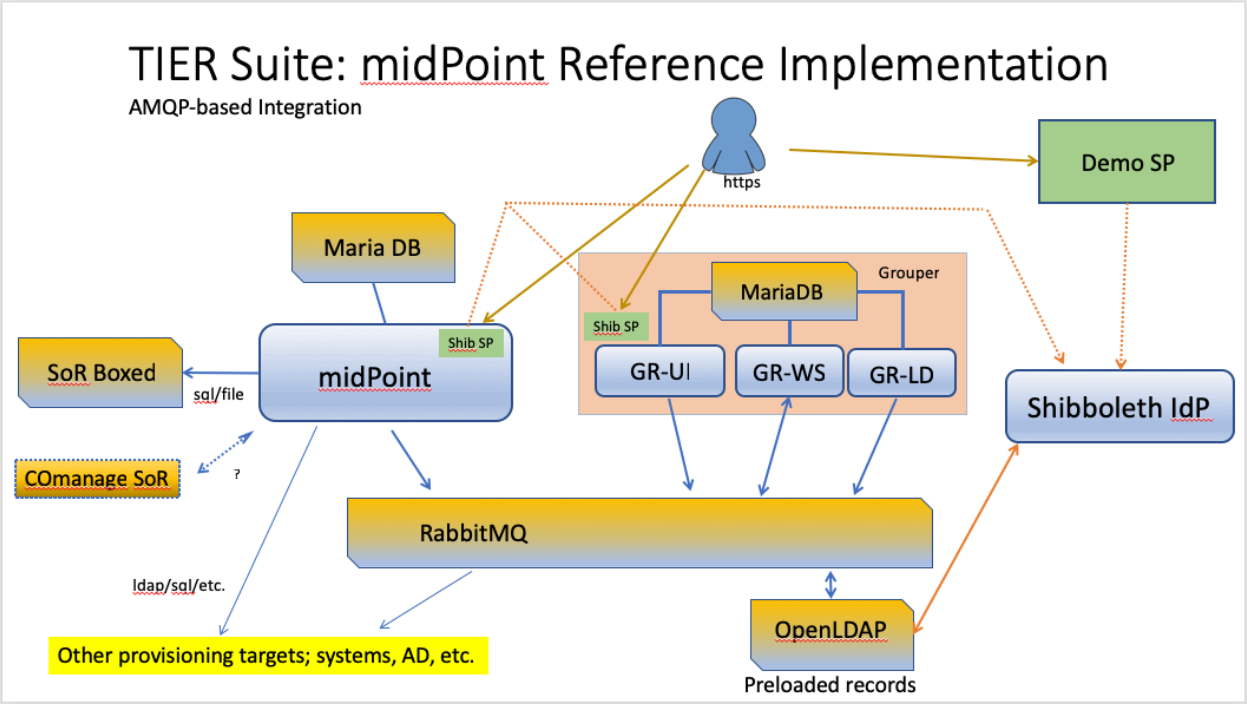

5. midPoint-Grouper Suite -- integration via AMQP
    a. Preferred solution
    b. midPoint messages changes to Grouper - due in March; not for Global Summit
    c. Picture Changes
        i. midPoint writes person records to LDAP, midPoint owns ou=People
        ii. All Grouper links go directly to OpenLDAP
        iii. Grouper Loader places messages on the bus
        iv. No link between LDAP and Messaging
    d. COmanage integration (one of our SoR use cases)
    e. Flow definitions
        i. midPoint writes person records to LDAP; midPoint owns ou=People, includes all affiliation, etc., information.
        ii. Grouper provisions new users via a loader job from data in LDAP; typically schedule based.
        iii. Grouper maintains the ismemberof portion of the user record
        iv. midPoint will listen for Group membership changes (now); group adds/deletes (future), etc.
        v. Midpoint handles provisioning
        vi. Our COmanage SoR will include LDAP; midPoint will be configured to poll this LDAP on some regular basis, looking for new users.
        vii.

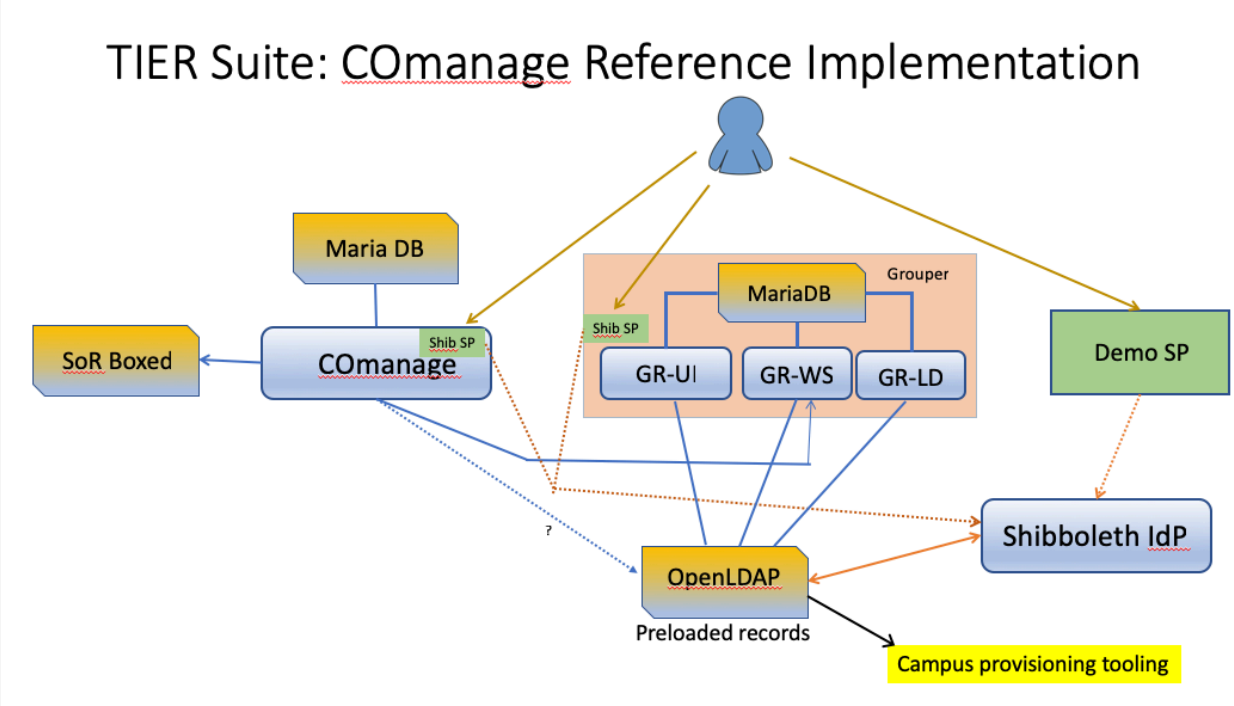

6. COmanage Grouper Suite
    a. Picture Changes
        i. Make dotted line from COmanage to LDAP solid; COmanage owns ou=People except for ismemberof (these attributes owned by Grouper).
        ii.
    b. Flow Definitions
        i. COmanage owns ou=People except for ismemberof (these attributes owned by Grouper).
        ii. Assumption: Campus provisioning processes LDAP change log and/or does a nightly sync.
    c.

January 21, 2019
Holiday - No call scheduled

January 7, 2019
Attendees (please add yourself):

- Jim Jokl - Virginia
- Bill Kaufman - Internet2
- Scott Koranda - SCG
- Colin Thompson - UC Merced
- John Gasper - Unicon

- Chris Hubing - Internet2
- Paul Caskey - Internet2
- Sara Jeanes - Internet2
- Keith Hazelton - Internet2

Regrets:

-

Call Agenda and Notes

7. Quick topics
    a. Add here
    b. ...
8. Base Reference Implementations
    a. Any discussion
9. TIER as a Suite Reference Implementations
    a. Continue the December 17 Section 3 discussions
    b. We will keep notes and check on action items directly in the December 17 notes below.
10. Check-in questions -- semi-TIER-as-a-Whole Reference Implementations
    a. Do we need a midPoint/Grouper and/or a COmanage/Grouper implementation?
        i. Research Universities and VO: COmanage / Grouper
        ii. Research and Smaller schools: midPoint / Grouper implementation
            1. midPoint as main IAM suite, lifecycle, etc; a few schools are interested in this possibility
    b. Yes, we will add the two scenarios above to the list of Reference Implementations
11. Action Items
    a. [AI] Jim to update / create / final drawings of Reference Implementations
    b. [AI] Keith to check in with midPoint re: Dec 17 Section 3.b.3.2 (ldap vs. messaging).
12. Other Components
    a. ShibUI
    b. COmanage Match
        i. Sits between systems of record and in front of the main ID solution
            1. See https://spaces.at.internet2.edu/display/COmanage/About+Identity+Matching

        ii. Generates a single unique identifier for each person affiliated with the institution regardless the number of “Systems of Record” submit data on the individual. Each person receives a single “unique ID” for the institution.
        iii. Queues records from Systems of Record matched or until resolved by a human
        iv. **[AI]** Scott - Architecturally this is the system that creates and authoritatively maintains unique Institutional Identifiers The generated identifier is opaque.

            Preferred text: Architecturally this is the system that uniquely identifies

individuals across multiple authoritative systems of record. The match engine assigns a unique "Reference Identifier" to canonically identify the individual across systems. The Reference Identifier is opaque, and typically not known to the individual.

v. This is not the “user” part of user@example.edu, i.e., not the NetID.. In our architecture, the NetID would be generated by midPoint or COmanage.
vi. Note that the above description doesn’t match the flow in 6.b.i - ID Match there does not sit between SoR and Registry

December 17, 2018 (includes notes below (in Section 3) for January 7, 2019)
Attendees (please add yourself):

- Jim Jokl - Virginia
- Bill Kaufman - Internet2
- Scott Koranda - SCG
- Keith Hazelton - Internet2
- John Gasper - Unicon
- Colin Thompson - UC Merced
- Paul Caskey - Internet2

Regrets:

- Chris Hubing - Internet2 (vacation)
- James Babb - Internet2 (vacation)
-

Call Agenda and Notes

1. Quick topics
    a. Add here
    b. ...
2. Base Reference Implementations
    a. Minor updates based on our past discussion
    b. https://spaces.at.internet2.edu/download/attachments/93653771/ReferenceImplementations2.pdf?api=v2
        i. Alternative link on Gdrive with no need to download the pdf https://drive.google.com/file/d/1cpGcWplgd6y_vrwocsRjhsITV7pZUQ7e/view?usp=sharing
3. TIER as a suite Reference Implementation(s)
    a. Focus for today’s discussion: Solutions that include COmanage and midPoint
    b. Which scenarios should be supported in a Reference Implementation
    c. https://spaces.at.internet2.edu/display/COmanage/COmanage+midPoint+Integration+Approaches

i. See also Slide 11

d. How does Grouper fit into these scenarios

e. AARC Blueprint Architecture see graphic below - Include link to this from the current "COmanage-only" Reference Implementation

f. Discussion

i. Include #2 and #3 as Reference Implementations

1. #2: COmanage Primary, midPoint Downstream

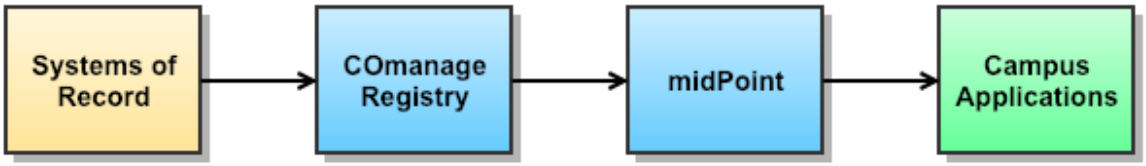

2. #3: midPoint Primary, COmanage Upstream

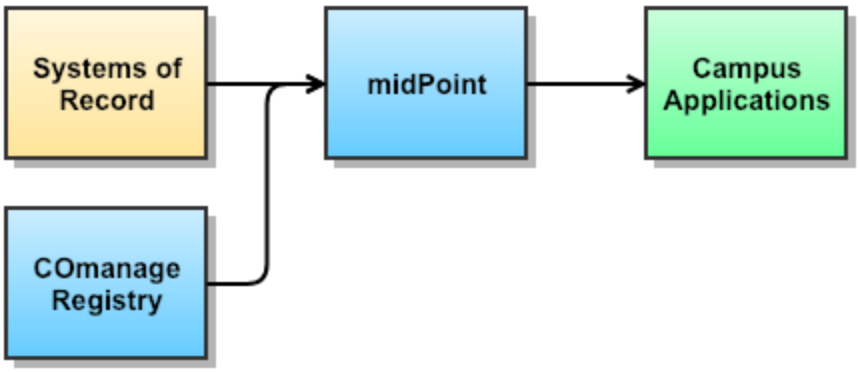

ii. COmanage can provision to Grouper, then Grouper to LDAP

iii. midPoint and Grouper integration - can be LDAP; would prefer messaging

iv. The openldap log reader to sync groups updated by Grouper back into midPoint is not ready yet --

v. Reminder to jaj - update slide 11 to point ldap updates from both Grouper and midPoint

g. Solutions

i. COmanage for VOs (the current "reference implementation")

ii. #2 with addition of Grouper and LDAP

1. COmanage to Grouper - web services to put CO Groups in Grouper
    a. Likely as a flat set of basis or reference groups
2. COmanage to LDAP - Identity information for Grouper to pick up
3. Grouper to LDAP - PSPNG - push complex Group info to LDAP
4. midPoint is doing provisioning of Apps based on data in LDAP
5. Satosa can get group/id info for provisioning Apps via SAML assertions

iii. #3 with addition of Grouper and LDAP

1. COmanage is one more System of Record

    a. midPoint pulls person-data from COmanage like a normal SoR
    b. COmanage provisions CO groups to its own stem within Grouper
2. midPoint provisions to
    a. People data to LDAP
    b. People and Groups to Applications
    c. midPoint to Grouper (preferably via messaging; could be via LDAP branch; could be via ConnId)
        i. **AI** [evolveum discussions / Keith] how best to do this now
            1. Draft messaging proposal https://wiki.evolveum.com/display/midPoint/Messaging+Resources
            2. This is the planned/desired long-term solution -- will take more than time available before Global Summit for delivery.
        ii. **AI** [evolveum discussions / Keith] can evolveum read openLdap change logs and keep groups in sync if there are multiple LDAP writers (if so, what about active directory, 389, etc.)
            1. Change log processing now available for AD, 389, and OpenDJ. They (midPoint) believe that they can add OpenLDAP easily (and may switch to it as their default directory).
        iii. **AI** [evolveum / Keith] answer the question: can we commit to a working test/beta version of the messaging solution by Global Summit. If so, we'll go down that path, otherwise we'll work on the LDAP piece.
3. Grouper
    a. Provisions groups to LDAP
    b. Draws on data from LDAP

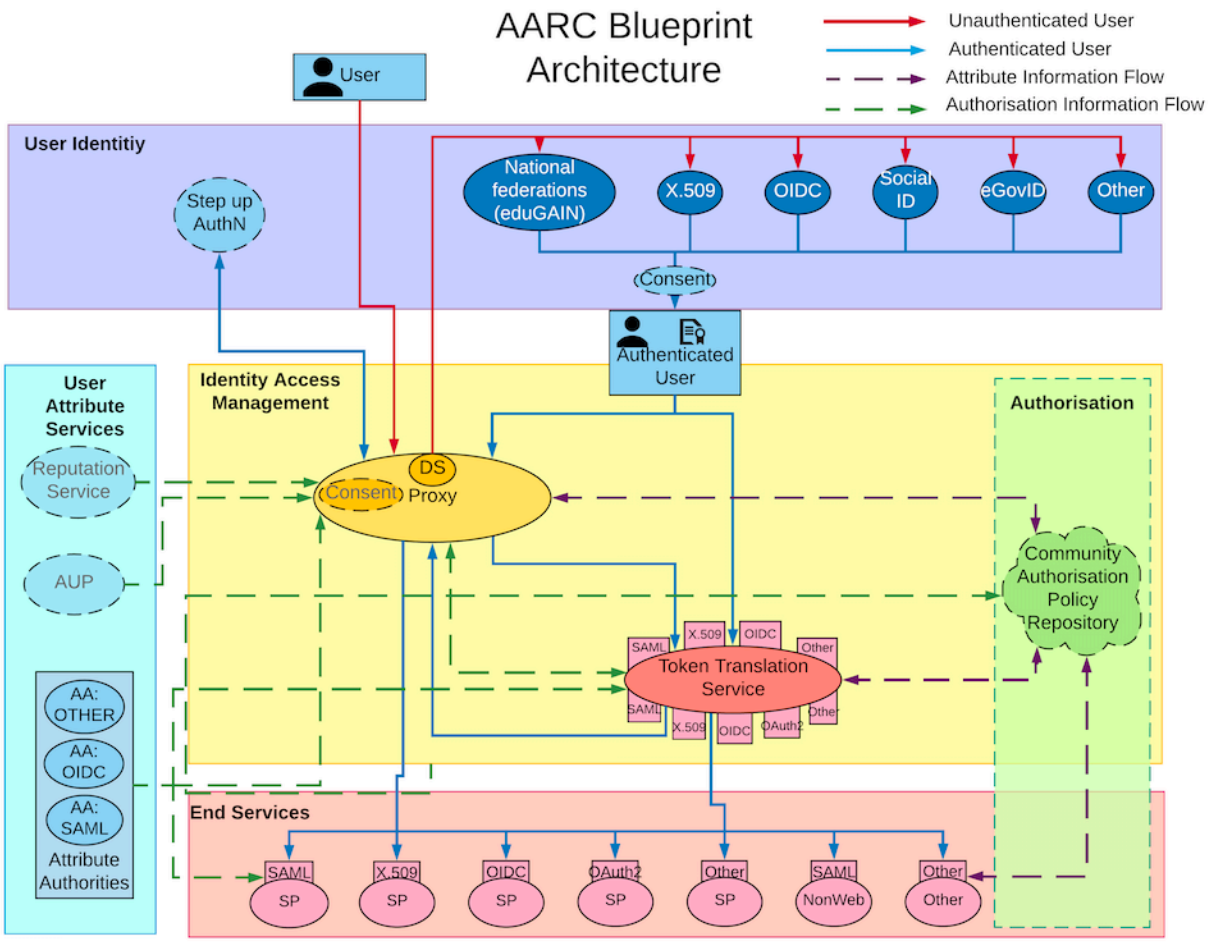

December 10, 2018

We will **not** hold a TIER Packaging call on Monday December 10. People who are critical to the Reference Implementation discussion are not able to make the call. The goal is to finalize the reference implementations during our Monday Dec 17 call.

December 3, 2018

Attendees (please add yourself):

- Jim Jokl - Virginia
- Scott Koranda - SCG
- Bill Kaufman - Internet2
- Colin Thompson - UC Merced
- Keith Hazelton - Internet2
- John Gasper - Unicon
- Blair Christensen - UChicago
- Paul Caskey - Internet2

Regrets:

- Michael Gettes
-

Call Agenda and Notes

4. Quick topics
5. Reference Implementations
    a. Our call will focus on specifying the set of combinations of the components that comprise common “whole TIER” deployment scenarios.
    b. This work is targeted at creating the final list of reference implementations.
    c. Reference Implementations for each component
    d. Reference Implementations for TIER as a whole
    e. Drawings to help start the discussion are here.
        i. COmanage Reference Implementation
            1. Add Pyff (metadata processing tool that also pairs nicely with MDQ) with RA21 discovery service
            2. SwitchWAYF instead of RA21
            3. For now, we’ll go with pyff and RA21
        ii. Grouper
            1. Openldap connectivity to each box
        iii. midPoint
            1. Can give people some notes and possible a minor exposed target application for provisioning
            2. Expose port 389 and password
            3. GTE has phpMyAdmin and phpLdapAdmin exposed that could be mimicked
            4. We can make a list / parking lot that we can draw from later given time and resource
        iv. Standard Reference Implementation
            1. Grouper is responsible for groups - not midPoint
            2. Basis groups come directly from SoR to Grouper;
            3. Both Grouper and midPoint deal with all SoRs
            4. midPoint will only see groups from LDAP, not create or manage any
            5. Assumption -- midPoint does all provisioning (delete orange Grouper to Other line)
            6. Fix grouper and midpoint to ShibIdp to different color; play with clearing up data flow vs. user AuthN
        v. Research Reference Implementation
            1. Change name to large sophisticated university -- we need the correct word for this - maybe “Advanced”
        vi. Potential to add a COmanage centric Reference Implementation
            1. Campus focused - COmanage as registry; midPoint for added Provisioning

2. Scott will check with Benn on this use case and if we should add this reference implementation

f.

g.

November 12, 2018

Attendees (please add yourself):

- Jim Jokl - Virginia
- Bill Kaufman - Internet2
- Keith Hazelton - Internet2
- Scott Koranda - SCG

Regrets:

- Paul Caskey
- Chris Phillips
- Chris Hubing - Internet2 (regrets)
-

Call Agenda and Notes

1. No call on Monday November 19 -- Thanksgiving
2. Quick topics and agenda bash
    a. .
3. Task Lists
    a. Discrete tasks awaiting completion
        i. https://docs.google.com/spreadsheets/d/1Cz3shX3FhfRR-IeUJtHiQbK8aWxBugzCuI9QD8WdSD4/edit#gid=0
        ii. What needs to be added?
        iii. TIER component Ref Implementations
            1. Would want to test drive the components so perhaps inject some test identities for each ref implementation
            2.
    b. Items/tasks being discussed and tracked
        i. Java support
            1. Likely several months until we know more
            2. Staying with Azul's Zulu Java for now
        ii. AWS Secrets and Shibboleth IdP
        iii. Container documentation requests
            1. What tests are applied during the build
            2. Making the tests available in some distribution
        iv. Curation of submitted / created Kubernetes reference implementations

v. Curation of campus documented implementations
vi.

4. Other
    a. Reference Implementations
        i. Provide a container that includes a Shibboleth and LDAP for use instead of a test federation. A default set of users with passwords would be included. This container could be reused by all of the components.
    b. Reminder: https://spaces.at.internet2.edu/display/COmanage/COmanage+midPoint+Integration+Approaches
    c. Next Call:
        i. * Focus in on the small set of Full TIER implementations docker-compose files * Monday 26th
        ii.

November 5, 2018
Attendees (please add yourself):

- Jim Jokl - Virginia
- Carey Black - tOSU
- Scott Cantor - tOSU
- Colin Thompson - UC Merced
- Ethan Kromhout - UNC
- Bill Kaufman - Internet2
- Scott Koranda - SCG
- Keith Hazelton - Internet2
- Blair Christensen - uchicago
- John Gasper - Unicon
- Chris Hubing - Internet2
- Chad Redman - UNC

Regrets:

-

Call Agenda and Notes

1. Quick topics and agenda bash
    a.
2. Lessons & Topics from TechEx / ACAMP -- Recent traffic on Packaging List
    a. List Traffic - Grouper - UNC - Chad Redman's message
    b. Patching of Grouper - how do we keep it up to date
        i. Faction #1 - Updates annually - security continually
        ii. Faction #2 - Keep up to date with all patches - latest release
        iii. Containerized Grouper BoF
        iv. Need to discuss this again on a future call

c.

3. Consolidated Action Item list
    a. Grouper Shell return code fix - completed in 2.4
        i. Listed as fixed in 2.4 API patch 3
            1. https://spaces.at.internet2.edu/display/Grouper/v2.4+Release+Notes#v2.4ReleaseNotes-v2.4.0patches
        ii. Ready to add a few additional tests to the pipeline
        iii.
    b. Container Preview Release Program Implementation
        i. Language cleanup on page https://spaces.at.internet2.edu/x/PoAUCe [JimJ]
            1. Add that we remove older versions that are not supported by the Project as quickly as practical.
            2.
        ii. Shibboleth ✓
        iii. Grouper ✓
        iv. COmanage✓ (checkmark added by Scott K - thanks)
        v. midPoint
        vi. Newer products - Shibboleth UI, ID Match, etc
    c. Container Orchestration decision document [JimJ]
    d. OpenLDAP vs. 389 large group (> 30k) update performance [BertB]
        i. Add some simple documentation re: potential issues with OpenLDAP for large groups.
    e. AWS Secret Manager for Shibboleth Secrets
        i. Sealer Key - ScottK - the Shibboleth back end work needs a formal request to the Shibboleth Consortium; sufficient detail to scope the work.
        ii.
    f. TIER Container Specification
        i. Wording verification (Support vs. Ancillary) - [JimJ]
    g. …
    h. ...
4. Remaining Work

    What remains to be completed & how
    a. Packaging of COmanage Match
        i. A first preview release of COmanage Match is available and the Docker packaging effort can begin. Scott K expects to begin that work in early December.
    b. Shibboleth IdP/UI Reference Implementation(s)
        i. See October 8, Item #2
        ii. Release is available at https://spaces.at.internet2.edu/x/mgOMBw
    c. midPoint Reference Implementation
        i. mP Container Project Status
        ii. Internet2/Evolveum midPoint Container Home
    d. Removal of TIER VMs once Reference Implementations are ready

e. All TIER containers with Logging and TIER Beacon support
    i. Shibboleth IdP, Shibboleth SP, and COmanage - OK
    ii. Grouper is not currently sending the Beacon
    iii. midPoint - Keith will check on Tuesday
f. TIER Rabbit MQ Container Specification [EthanK]
    i. We are starting with the RabbitMQ container, adding our logging mods to it.
    ii. Unclear if its being used/tested yet.
    iii. It's available in the Internet2 Github and being built.
g. Container build documentation *** we need to come back to this topic
    i. Description of tests performed
    ii. See Aug 20, 2018 - Section 2.e
h. Build and document Reference Implementations *** NEXT time
    i. Individual Components as documented
        1. midPoint
        2. COmanage
        3. Grouper
        4. Shibboleth IdP
        5. ID Match
    ii. TIER as a whole solution

5. Open Items
    a. The future of Java
        i. Is Zulu the solution we thought it would be?
        ii. Zulu looks no better than plain old openjdk
        iii. We no of no one doing longer-term LTS support at no cost
        iv. No reason to change anything now - stay with Zulu until early next year and revisit

6. Other Topics
    a. Production Implementation Summary Curation
    b.

October 8, 2018
Attendees (please add yourself):

- Jim Jokl - Virginia
- Chris Hubing - Internet2
- Keith Hazelton - Internet2
- Colin Thompson - UC Merced
- Chris Phillips - CANARIE
- Jonathan (Jj!) Johnson - Unicon
- John Gasper - Unicon
- Bill Kaufman - Internet2

- Mike Grady, Unicon
- Paul Caskey - Internet2
- Scott Cantor - tOSU
- Sara Jeanes - Internet2

Regrets:

-

Call Agenda and Notes

1. Quick topics and agenda bash

2. Guests on call to discuss the Shibboleth UI project and its Packaging
    a. Project Wiki
    b. Expected TIER Initial Release -
        i.
    c. Several deployment options / possibilities
    d. Simplest way to get started with the application itself, is to just run it; ships with an embedded non-persistent database; use mysql, etc., in production; download the jar file and run it to test/demo/play/learn.
    e. Reference Implementation
        i. docker-compose with UI and Shib IdP Containers might be a start at the Reference Implementation?
            1. Shibboleth IdP with config changes to support the UI
                a. https://wiki.shibboleth.net/confluence/display/IDP30/MetadataDrivenConfiguration
            2. Persistent database - TIER mariadb
                a. Can not recover entire state via import of previous output files
            3. Container to run Shibboleth UI
        ii. non-Shibboleth authentication since we expect a small number of people at any one institution using the application
        iii. Scripting to move/copy generated files to IdP instances; shared Docker volume for demo or dev/test; etc.
        iv.

3. TechEx Preparation (likely quick)
    a. Ready to announce at TechEx that the TIER midPoint containers are ready for testing
        i. https://spaces.at.internet2.edu/display/MID/Dockerized+midPoint
    b.

4. Action Item Updates
    a. Grouper Shell return codes
    b. Implementation of Container Preview Release Program
        i. Ready for TechEx?

    ii. TIER Package Delivery Site changes?
c. Container Orchestration
    i. [jaj] Still working on document - will have by TechEx
d. Java distribution, again, (zulu may not solve the root problem)
e.
f. In progress - see notes below for now

September 24, 2018
Attendees (please add yourself):

- Jim Jokl - Virginia
- Colin Thompson - UC Merced
- Ethan Kromhout - UNC
- Kevin Ruderman - Boston University
- Chris Hubing - Internet2
- Scott Koranda - SCG
- John Gasper - Unicon
- Paul Caskey - Internet2
- Keith Hazelton - Internet2

Regrets:

- Michael Gettes
- James Babb

Call Agenda and Notes

1. Quick topics and agenda bash
    a. Grouper shell return code update
        i. Is logged as https://bugs.internet2.edu/jira/browse/GRP-1853
        ii. Sent failure scenarios to Shilen, and he is looking into it
    b. OpenLDAP vs. 389 and large group performance
    c. Update on September 10, 3.d.i, Shibboleth Sealer Key work - scheduled for 9/24
        i. [AI} Jim to ping Scott re: timing
    d. Question about ClickJacking vulnerability issue; What fixes are built in? If using Jetty (not TIER default, TIER uses Tomcat) a sample jetty-rewrite.xml is provided.
        i. In https://wiki.shibboleth.net/confluence/display/IDP30/Jetty93 see etc/jetty-rewrite.xml (*optional*)
        ii. From a Scott Koranda recommendation here: http://shibboleth.1660669.n2.nabble.com/Jetty-configuration-wiki-page-and-configuration-to-help-mitigate-clickjacking-td7638735.html#a7638754

iii. See https://issues.shibboleth.net/jira/browse/IDP-627 for a discussion and overview of the work that will probably make it into IdP release 3.4.

e. Progress on midPoint packaging under Evolveum SoW- Keith
   i. midPoint container work repo now at: https://github.internet2.edu/docker/midPoint_container
   ii. midPoint container is now in the TIER Jenkins workstream (still needs some additional testing logic during the post build stage); pushed to docker hub as tier/midpoint
   iii. Current goal is to announce TIER midPoint ready for testing at TechEx

f. … insert your item(s) here

2. Container Orchestration Decision
   a. Update on the Kubernetes test via docker-compose and Kompose? Scheduled for 9/24 call [AI] John Gasper
      i. Kubernetes is supposed to run a docker stack deploy docker-compose file natively (JG: This is `kompose`.)
      ii. A translation tool also exists. (JG: This is also `kompose`.)
      iii. An old-enough version of the docker-compose format is needed for these features to work properly.
   b. Decision
      i. Officially stay with docker swarm as our mechanism for TIER Reference Implementations
      ii. Curate donated Kubernetes versions
   c. Discussion on how to communicate decision
      i. TIER Packaging Site - quick page
      ii. Send email to lists that received the survey
      iii. Be prepared to answer questions at TechEx.

3. Container PRP program
   a. Initial Draft: https://spaces.at.internet2.edu/x/PoAUC
   b. Try to operationalize by TechEx?
      i. For Shib, Grouper, COmanage - the production versions
         1. Shibboleth - yes
         2. COmanage - yes
         3. Grouper - yes
      ii. midPoint
         1. Still under development
         2. Is not be part of PRP right now.
         3. **[AI]** Jim to clean up the language on the PRP page

4. TIER Package Delivery Site

    a. Agreed to delete the VMs and AMIs as soon as we have better docker-compose/scripts ready for the reference implementations
    b. On the next TPD site edit, we'll clearly mark that the VMs and AMIs will be removed after the reference implementations are ready.
5. COmanage Match
    a. Benn has delivered a preliminary version of the code (release candidate 1.0.0-RC1) and Scott K is preparing to package it, but does not expect to have it ready for TechX. If the need arises to have it packaged before TechX please let Scott K know.

6. September 10, Section 5

September 17, 2018
Attendees (please add yourself):

- Jim Jokl - Virginia
- Scott Koranda - SCG
- John Gasper - Unicon
- Carey Black - tOSU
- Keith Hazelton - Internet2
- Bill Kaufman - Internet2
- Sara Jeanes - Internet2
- Chris Hubing - Internet2
- Paul Caskey - Internet2

Regrets:

- Michael Gettes

Call Agenda and Notes

1. Quick topics and agenda bash
    a. Any update on the Kubernetes test via docker-compose and Kompose? Scheduled for 9/24 call [AI] John Gasper
    b. Any update on September 10, 3.d.i, Shibboleth Sealer Key work - hold for next call
    c. midPoint packaging update - see: https://spaces.at.internet2.edu/x/Xw-tBw
    d. TIER Docker Container Specification update for base OS and Java https://spaces.at.internet2.edu/display/TPWG/TIER+Docker+Container+Specification
    e. … insert your item(s) here
2. Grouper Shell Return Code Request
    a. Specific example that generated our request?
    b. Stop processing a command file when one command fails?

c. We need to document our specific technical need.
   During non-interactive operations:
   i. When we run a command file, if any command fails, it continues on and returns zero.
   ii. Grouper shell appears to always return zero; we need it to report back with a non-zero return code whenever it encounters any error anywhere in the process.
   iii. Grouper shell returns zero and accepts commands even when the start-up of pieces of it's execution environment were not successful.
   iv. All exceptions should return a non-zero return value.
   v. The issue may be in the gsh.sh wrapper shell script. Chris to ping Shilen

3. Container Preview Release Program
   a. Start with September 10, 2018, Sections 5.d and 5.e
   b. Proposal
      i. Discussion of new features on component-specific slack channel
      ii. Images are built from the pipeline and made available
      iii. Announcement of new build on slack channel
      iv. Discussion of testing on slack channel
         1. Minor / no known changes normally require a 3 day minimum discussion
         2. Releases with new features normally require a 7 day minimum discussion; perhaps longer for major version updates as needed.
         3. Critical security updates can be done immediately
      v. Merge to master, update TPD wiki (or start with new build)
         1. TPD process kicks off an email message
   c. Need to document in Red on TPD that you should just not use "latest".
      i.
4. Remaining items from Section 5 of 9/10

Future Reminders:

1. Next week - check in on September 10, 2.e on container orchestration, then discuss how to communicate decision
2. Continue on with other notes/action items in Section 5
3. (future) Shibboleth UI update and packaging
4. (future) Status of ID Match and packaging
5. Review Section 1.b above (Sealer Key)
6. Review Container Spec Section 1.a.ii
7. Wording on "Support vs. Ancillary vs. " in Container Spec
8. One last check in on 3.b above

September 10, 2018
Attendees (please add yourself):

- Jim Jokl - Virginia

- Nick Roy - InCommon/Internet2
- Keith Hazelton - UW-Madison/Internet2
- Scott Cantor - tOSU
- Paul Caskey - Internet2
- John Gasper - Unicon
- Chris Hubing - Internet2
- Colin Thompson - UC Merced
- Keith Wessel - Illinois
- Chris Phillips - Canarie
- Ethan Kromhout - UNC Chapel Hill
- Blair Christensen - U. Chicago
- Erik Coleman - Illinois
- Sara Jeanes - Internet2

Regrets:

- Scott Koranda
- Michael Gettes
- James Babb

Call Agenda and Notes

1. Quick topics and agenda bash
    a. …
    b. ...
2. Container Orchestration
    a. Question: should we change our default container orchestration framework for the reference implementations from Docker Swarm to Kubernetes?
        i. Seems like a good idea to support both
        ii.
    b. The survey received a total of 32 responses from 25 different schools , one consortium, and one commercial firm.
    c. The Survey Results are available on-line here.
        i. Original survey questions are here.
    d. Translation of a docker-compose file for use with Kubernetes does not appear to be hard. Is anyone aware of a tool that runs in the opposite direction?
    e. ** Pending a retest of the docker-compose to Kubernetes tool, we'll stick with Swarm and docker-compose for our base standard. We'll revisit this in two weeks but don't anticipate problems with the Kompose tool.
        i. Could we incorporate in the jenkins pipeline for generation from the docker-compose file?
3. AWS Secret Manager and IdP Sealer Key management
    a. Leverage AWS Secret manager for Sealer Key management
    b. How much development is needed?

c. Use native IDP code rather than AWS SDK for web service actions
d. Who is going to do the work?
    i. Shibboleth add-on - ScottC could tackle in 3 to 6 months with assistance from someone with the use case and infrastructure.
    ii. Chris(Hu) can help with the Amazon side ; assistance from KeithW and Illinois who have already made a start on this work.
        1. Infra template for code (lambda, cloudwatch cron alarm) to work, IAM role for container to be able to have authz for secrets call
    iii. AWS Secrets Manager Docs Link: https://aws.amazon.com/secrets-manager/resources/
        1. Yes, it returns JSON ("The JSON that AWS Secrets Manager expects as your request parameters and that the service returns as a response to HTTP query requests are single, long strings without line breaks or white space formatting.")
            a.
```
{
   "ARN": "string",
   "CreatedDate": number,
   "Name": "string",
   "SecretBinary": blob,
   "SecretString": "string",
   "VersionId": "string",
   "VersionStages": [ "string" ]
}
```
        2. Specifically: https://docs.aws.amazon.com/secretsmanager/latest/apireference/API_GetSecretValue.html
        3. Before querying the Secret Manager API, a machine (container, EC2 instance) with the proper IAM role will need to query the web API to get credentials: https://docs.aws.amazon.com/AmazonECS/latest/developerguide/task-iam-roles.html
4. Request for testers of the new Grouper 2.4 Container:
    a. Code/Docs: https://github.internet2.edu/docker/grouper/tree/2.4.0-a0-u0-w0-p0-test
    b. Dockerhub: From tier/grouper:2.4.0-a0-u0-w0-p0-test
5. Action Items and Follow-up
    a. [AI] (Chris Hyzer / Jim Jokl) - Grouper Shell return code request
    b. [AI] (Bert Bee-Lindgren) - OpenLDAP vs. 389 Directory Server performance for large groups
    c. [AI] (Ethan Kromhout) - Updates to TIER RabbitMQ container specification
    d. Publication of current production versions of the TIER distributions on the TIER Package Delivery site.
    e. [AI] Create the Preview Release Program for container build testing
        i. Build
        ii. Test in the internal training environments

iii. Make announcement on Component specific email list and/or Slack channel requesting testing?
    1. Wait for positive feedback from a tester before TPD
iv. If no negative feedback, move to TPD site (5.d) as current production version.

f. Component documentation requests
    i. Description of automated tests used in build process
    ii. Make actual tests available
g. (future) Shibboleth UI update and packaging
h. (future) Status of ID Match and packaging
i. (future) [AI] (Jim Jokl) Update container spec to match Centos decision (Aug 13 Section 2.b)

September 3, 2018
No Packaging Call is scheduled - Labor Day holiday

August 27, 2018
No Packaging call is scheduled. Please make a few minutes to look at the draft TIER container orchestration survey at the URL below and comment on our mail list re: needed changes, enhancements, etc., or if it just looks too one-sided in any particular direction.

https://virginia.az1.qualtrics.com/jfe/form/SV_6hf0DzRFV6SK2bj

August 20, 2018
Attendees (please add yourself):

- Jim Jokl - Virginia
- Bill Kaufman - Internet2
- Scott Koranda - SCG
- Sara Jeanes - Internet2
- Chris Hubing - Internet2
- Michael Gettes - Florida
- James Babb - UW Madison
- Blair Christensen - uchicago
- Colin Thompson - UC Merced
- Ethan Kromhout - UNC Chapel Hill
- Steve Zoppi - Internet2
- Paul Caskey - Internet2

Regrets:

-

Call Agenda and Notes

1. Agenda Bash - Additional Topics / Quick Items

a. Additional topics …
    i.
b. Quick Updates …
    i.

2. Agenda and Notes
    a. [AI] - JimJ - Grouper shell to process return codes request
    b. [AI] - BertB - 389 v. OpenLDAP large group update performance
    c. Decision on TIER Container Base OS - see proposal in August 13 notes, Section 2.b -
        i. If actual TIER component CentOS 7
        ii. If ancillary service (like RabbitMQ / MariaDB) can use whatever container is being maintained by a core group supporting that technology
            1. MG: if what is maintained by another core group (ex:RabbitMQ) is found to be deficient then we would maintain under CentOS
            2. Minor additions, e.g., log format changes do not require a rebuild to Centos.
    d. RabbitMQ Container Specification
        i. Has been pending 2.c
        ii. Will use the standard container released by the the RabbitMQ team
        iii. Changes needed [AI] Ethan
            1. TIER logging standard modifications
            2. No other known changes at this point in time
            3.
    e. Docker Hackathon Feedback
        i. Automated testing, sharing of existing tests, etc.
            1. Component vs. Container tests
        ii. Generically - completing the reference implementations
        iii. Shibboleth: verifies test page (automatically checks several possibilities)
        iv. Grouper suggestions
            1. ChrisHubing: Jenkins will compose and try to bring up the entire environment successfully. Currently have a zero return code that ChrisHyzer is working to fix
            2.
        v. COmanage: Jenkins pipeline is for build only. Startup scripts checks for database access and LDAP access on startup. Post-build checks are manual.
        vi. How to respond to request:
            1. Current: mixed
            2. Goal:
                a. Document level of automated testing for each component
                b. Add additional build tests over time
        vii. How to publish latest known-good versions of containers:
            1. Document on TPD site

2. Also document test versions; preview release vs. production.
3. And a naming convention including timestamp and qualifiers expressive of what the “stability level” is (see 2-e-vii-2)
4. Get [AI] people to join the Preview Release Program (PRP)
5. See our discussion below from July 16 Section 1.b.iii where we came to many of the same high-level decisions.

f. Logging Discussion
g. Reference implementations including evaluation environment
    i. Adding full suite integration - midPoint
h. Container Orchestration Framework
    i. Need to get a poll out
    ii. There may be those who just want to run the containers in non-K environments. How to support those who want to get up and running in simple ways vs. Kube which may involve more effort? Need to confirm Kiube configs vs. non-Kube. [AI] - Jim will draft a poll, focusing on when and motivation and keeping in mind to be careful what you ask for.
i. Shibboleth UI container specification
    i. SHIBUI Packaging meeting on Thursday, August 23rd - 2pm Eastern, 11am PT
    ii. https://unicon.zoom.us/j/270290441

j. ID Match container specification - pending work completion
k. MARIAdb - container - next week

August 13, 2018
No call today - if you have action items, please work on them for next week.
Pending Agenda

1. [AI] Jim - follow-up on Grouper gsh return codes request
2. TIER Base OS Decision
    a. No strong feedback, positive or negative, from the BTAA Docker Hackathon
    b. We appear to be converging on requiring Centos 7 for TIER Core Components and any containers that we build but allowing other base container operating systems when the container is for an external component that is maintained by that component’s project team.
3. RabbitMQ Container Specification
    a. On hold, pending #2 above
4. OpenLDAP vs. 389 server for Grouper
    a. On hold for test data
5. Container specification for the Shibboleth IdP UI
6. Reparse BTAA Docker Hackathon notes for AIs
7. Progress on reference implementations

August 6, 2018
Attendees (please add yourself):

- Jim Jokl - Virginia
- Ethan Kromhout - UNC Chapel Hill
- James Babb - UW Madison
- Bill Kaufman - Internet2 (might need to leave to take wife to clinic for bruised foot)
- Colin Thompson - UC Merced
- Jon Miner - UW Madison
- Chris Hubing - Internet2
- Blair Christensen - uchicago
- John Gasper - Unicon
- Scott Cantor - tOSU

Regrets:

-

Call Agenda and Notes

3. Agenda Bash - Additional Topics / Quick Items
    a. Additional topics …
        i.
    b. Quick Updates …
        i.
4. Agenda
    a. Grouper Action Items
        i. See 2.a and 2.b from July 30 below
        ii. See notes in July 30 minutes.
    b. Other Action Items?
    c. Feedback from TIER Docker Hackathon
        i. See July 23, Section 2.c.i
        ii. Summary page / Uncut Notes:
            1. https://docs.google.com/document/d/1j26AVzOfUgPGYKjKQ3vGgFTDVgEXXKDBne16Y2bKPII/edit?usp=sharing
            2. Container testing (specifically SSO, raw Geb could be used for Grouper, etc): https://github.com/Unicon/avus-testing-framework (documentation is coming, but examples are fully working)
    d. TIER Container Base OS requirement
    e. RabbitMQ container spec - See July 23, Section 2.c.iii
    f. TIER Containers: Configuration vs. Build Customization
        i. See July 2, Section 2.c.ii
    g. Shibboleth UI Packaging (depending on who makes call)
        i. Initial Discussion
        ii.

h.

July 30, 2018

1. We will **not** hold our regular TIER Packaging call today and will pick up with our agenda next week.
2. Interim Action Item Updates
    a. Grouper Requests
        i. Can we receive a non-zero return code if a patch fails when running the grouper installer?
        - Response: Yes, they can do this.

        ii. Can we receive a non-zero return code from gsh whenever the process that it kicks off fails?
        - Response: A request to check on this will be made
        - **[AI]** jaj to check on status

        iii. Does the Grouper team have any unit or other testing procedures that we might be able to incorporate into the automated docker builds for us to validate functionality where possible before pushing containers?
        - Response: There are Grouper junit tests ... they might need some care and feeding if you want to run them every time ... takes a while (8 hours?)
        ---> We would add very little value by running these tests again. We'll skip this, at least for now.

        iv. The final question is the one that we were unable to reconstruct on our July 23 call. The best guess is that it was related to an older version of some library impacting logging.
        - Response: Grouper 2.4 updates all libraries
        ----> Unsure of original need; hopefully addressed in updated libraries.

    b. Grouper and OpenLDAP vs. 389 Server

July 23, 2018

Attendees (please add yourself):

- Jim Jokl - Virginia
- Scott Koranda - SCG
- Bill Kaufman - Internet2
- Michael Gettes - UFL
- John Gasper - Unicon
- Ethan Kromhout - UNC Chapel Hill
- David Bickel - Indiana
- Scott Cantor - tOSU

- Keith Hazelton - Internet2
- Paul Caskey - Internet2
- Blair Christensen - uchicago

Regrets:

- Chris Phillips

Call Agenda and Notes

1. Agenda Bash - Additional Topics / Quick Items
    a. Additional topics …
        i.
    b. Quick Updates …
        i.
2. Agenda
    a. Action Item Review
        i. Jim - Grouper Change Request Section 1.b.2.5 of July 16 notes
        ii. Use of TIER Package Delivery Confluence site to highlight/document the current stable releases. [AI] - Jim still needs to deal with 1.b.2.5.d
        iii. **[AI]** OpenLDAP vs. 389 performance testing - Jim to reconnect with Bert during the week of July 23 when he is back in the office.
        iv. Jim - Update the TIER container specification to document that default time in logs is UTC. -- Done. Leave this statement alone using “should” instead of must” Documentation should exist on how users can change this behavior.
    b. TIER Operating System Container Specification Discussion
        i. See Section 2.c from July 16. We started this discussion last week.
        ii. Question: should we retain our explicit requirement for TIER components to be based on Centos 7.
        iii. TIER Core Components
            1. e.g., Shib, Grouper, COmanage, midPoint, etc.
            2. Retain requirement for Centos
        iv. Ancillary Components
            1. e.g., OpenLDAP for COmanage deployments, GNU Mailman 3 for COmanage deployments, MariaDB, RabbitMQ, etc.
            2. Hold for right now --
    c. Carryforward Topics
        i. Update on potential BTAA TIER Docker Hackathon on Aug 1
            1. Add 2.b.iv.2 - [AI] Jim to draft paragraph - see if interest
            2. Docker Hackathon ideas
                a. Packaging enhancements and/or refinements
                b. Automated Testing
                c. Container Orchestration and TIER Containers
                d. Spinning up TIER Containers in an integrated way

        e. Any feedback on current documentation would be useful
    ii. Discussion topic: can all/most TIER components be configurable enough such that they used without build-level customization
        1. [AI] - Scott will ask what they believe issues are
        2. **Ended meeting here** - discussion remains on TIER container goal - perhaps 80/20 rule on available via configuration vs. available via customization.
    iii. TIER RabbitMQ Container Specification
        1. TIER from source vs. Pre-built docker image
        2. Plugins
            a. Tracing
            b. Management
        3. Erlang/OTP package version
        4. Tuning and configuration for Ethan's current build
            a. Local additions: Firehose Tracer java app
            b. Supervisord
            c. A couple of logging items are still needed
        5. https://www.rabbitmq.com/install-rpm.html
        6. https://github.com/docker-library/rabbitmq - appears to be based on (both Alpine and Debian are available). Bitnami also publishes a server - appears to be based on their mini-Debian distribution.
        7. Added to Ancillary Components list - Topic on hold until after Aug 1 BTAA Docker Hackathon.
    iv. Initial discussion: Shibboleth UI Packaging
    v.

July 16, 2018
Attendees (please add yourself):

- Jim Jokl - Virginia
- Keith Hazelton - Internet2
- Michael Gettes - University of Florida
- Bill Kaufman - Internet2
- James Babb - UW Madison
- Kevin Ruderman - Boston University
- Scott Koranda - SCG
- Jon Miner - UW-Madison
- Dusty Edenfield - Georgia Tech
- Chris Phillips - CANARIE / CACTI Chair
- Sara Jeanes - Internet2
- Paul Caskey - Internet2
- John Gasper - Unicon
- Colin Thompson - UC Merced
- Chris Hubing - Internet2

Call Agenda and Notes

1. Agenda Bash - Additional Topics / Quick Items
    a. Additional topics / ?
    b. Michael: How do we do automated smoke testing on new releases?
        i. Current testing process
            1. Jenkins based limited tests - checks status page
            2. Seek testers for major changes (e.g., versions of Java, Tomcat); seek testers via Slack, packaging, etc.
                a. Maybe at least one person from the community of testers has to say “tested! It’s ok” to proceed?
                b. IDP possibility: https://github.com/Unicon/avus-testing-framework
        ii. Recent Grouper Build Issue
            1. Grouper automated builds run against a full compose file for automated checking;
            2. Recent glitch slipped past the automated testing.
            3. Need to continue to update automated tests; how much more can we look for beyond return codes;
            4. Possibility for a test -- confirm grouper patch level and fail build before publication?
            5. **[AI]** JIM - Requests to the Grouper Team
                a. Non-zero return code if a patch fails from grouper installer.
                b. Non-zero return code from gsh whenever the process that it kicked off fails.
                c. Ask Grouper team about junit or other testing procedures so they might be incorporated into docker builds to validate functionality where possible before pushing containers.
                d. Ask grouper team about updates to logging software to help with harmonizing docker level time issues for logging app and OS.
        iii. TIER Tagging (in general)
            1. Do we need to be more specific about what is “production” or is the existing documentation good enough in this regard?
            2. It’s not “latest”, at least for Grouper
                a. Branch (in Internet2 github) and Tag (in TIER Dockerhub) naming convention based on patch level of components
                b. Should make change management of container easier and to avoid drift between dev and prod instead of pulling from :LATEST
                c. E.g. 2.3.0-a104-u42-w12-p16, 2.3.0-a103-u42-w12-p16
                    i. 2.3.0=Base version of Grouper
                    ii. A=API patch version

iii. U=UI patch version
iv. W=WebServices patch version
v. TPDP=PSPNG patch version

3. We'll consider maintaining "production" TAG names on the TPD web site.

c.

2. Agenda
    a. Review of old Action Items
        i. Bert - OpenLDAP vs. 389 performance testing
            1. Bert is on vacation; [AI] - Jim to reconnect with Bert next week. Michael can help Bert re: demonstrating the bug
        ii. Keith - Evolveum's perspective on OpenLDAP vs. 389
            1. Evolveum prefers OpenLdap
            2. They document (see June 18 Action Items - 2) several issues with 389 server.
            3. FYI - 389 comes with eduPerson schema built-in. Not the other edu objectclasses.
        iii. Paul - Shibboleth changes - Java and Tomcat
            1. Default conversion to Zulu Java - done
            2. Move off of Tomcat 8.5 - done (switched to Tomcat 9)
            3. **Seeking testers**
                a. used in production now (one location)
                b. some course users were getting out of memory issues
        iv. Scott K - ID Match
            1. Scott will be creating the ID Match container once the code is ready
            2. Update? The code is not ready. :-)
            3. [AI] Scott will check on target dates.
                a. Benn Oshrin expects a pre-release candidate for ID Match the first week of August, 2018. Scott can then prepare a first draft of Docker container for end of August, 2018.
    b. Container Time Zone settings
        i. Current containers default to UTC
            1. [AI] add to TIER container specification
        ii. Is this the correct behavior for logging
        iii. Should we document how to change OS to local time
            1. **[AI]** Optimally, yes (documentation or automation)
    c. TIER Operating System Container Specification Discussion
        i. Alpine vs. Centos vs. Debian vs. ?
        ii. TIER currently required Centos 7
        iii. Some optimizations possible (Alpine and size, Debian and other existing component builds).
            1. ChrisP: footprint vs convenience for the builders IMO

iv. See also June 18, Section 5.a
v. Potential campus security group audit issues
vi. **[AI] Pick up here** on the next call. Should we ask the membership, membership security contacts, other (survey) about this. How easy is it to explain at least what we see are the real issues.

d. Topics below were not discussed on the call
i. Discussion topic: can all/most TIER components be configurable enough such that they used without build-level customization
ii. TIER Rabbit MQ Container Specification
iii. Initial discussion: Shibboleth UI Packaging

June 25, 2018
This week's call is cancelled. Jim is trying to connect with various individuals on Action Items.

June 18, 2018
Attendees (please add yourself):

- Jim Jokl - Virginia
- Michael Gettes - UF
- Scott Koranda - SCG
- Ethan Kromhout - UNC
- John Gasper - Unicon
- James Babb - UW Madison
- Paul Caskey - Internet2
- Keith Hazelton - Internet2/UW-Madison
- Blair Christensen - University of Chicago

Call Agenda and Notes

1. Agenda Bash
    a. Openldap vs. 389
        i. OpenLDAP groups with ~35k to 40k users would take a long time -- approximately six seconds to complete the update. In a practical sense this made normal operations difficult and large group updates were untenable. The 35K was at PSU - which could change depending on the amount of data and caching and other performance factors. A lightly loaded LDAP might see different pain points. There was contact with the openldap devs to address the problem and we (PSU) were told this would not be fixed - related to indexing. This problem has been reported various times over the last 10-15 years from researching the problem. (/mrg)
        ii. USC (Russ) had tested groups ~150k users with the 389 devs to bring forward what Sun had done years ago to fix this branch of the code (AOL

vs. Sun when Netscape split). Large groups are subsecond mods in 389 now.(/mrg)

2. Any quick topics
    a. Fixing time inside the container: add to Dockerfile for ET do:
        i. ENV TZ=America/New_York
        ii. RUN ln -snf /usr/share/zoneinfo/$TZ /etc/localtime && echo $TZ > /etc/timezone
        iii. Logging - everything is Zulu time.
        iv. [discuss next time] Question: should we document this, leave it alone, make it configurable, etc?
    b. Alpine vs. Centos vs. Debian
        i. How would we sell multiple base OS versions to security team?
        ii. How much does the size really matter in a TIER context?
        iii. Alpine 5 MB vs. Centos 7 ~200 MB vs. Ubuntu ~223

Action Items

1. [AI] Paul -Changes to Shibboleth container
    a. Java and Zulu
    b. Moving back to Tomcat 8 or forward to Tomcat 9; known bugs in our current version of Tomcat
    c. A test version is in production
        i. Tomcat 9
        ii. Zulu for Java
        iii. Google doc: https://docs.google.com/document/d/17-0O3Tvty9PONL6wu4PiC6ZWramdyntXmOsq1UpD2tE/edit
        iv. Shibboleth container is now available without a build
        v. Oracle pieces are still commented out and available for use - users would need to do the same build as in the past.
        vi. *Paul seeks **people to test**: both the new container on Zulu Java and Tomcat 9 -*
2. [AI] Keith will ask the midPoint people re: large groups
    a. Keith confirmed evolveum's perspective that they prefer openldap
    b. May need to restart conversation with them depending on the results of performance testing.

https://wiki.evolveum.com/display/midPoint/389+Directory+Server

Drawbacks

Attribute nsUniqueId

The 389ds has a very convenient attribute nsUniqueId that is an attractive choice for account primary identifier. And this mostly works. But it does NOT work for changelog-based live synchronization. Delete deltas in the changelog do NOT have the nsUniqueIdattribute. As the original entry is already deleted at that time then it is not possible for a connector to translate the DN of the deleted entry to a nsUniqueId and the delete delta will not work.

Workaround: change primary account identifier to dn.

Bad Schema

The 389ds is NOT a fully LDAPv3-compliant directory server. It is using non-numeric OIDs, under some circumstances it uses illegal attribute names (such as unhashed#user#password), it is using attributes that are not declared in the schema (firstchangenumber, lastchangenumber), etc. MidPoint 3.2 is bundled with LDAP connector that relies on LDAPv3 compliance of the schema and will fail is 389ds is configured in non-LDAPv3-compliant way. The LDAP connector bundled with midPoint 3.3 was improved to be a more tolerant LDAP client and it will work.

3. [AI] Bert will test 389 vs OpenLdap and see if 389 performs better with large groups.
    a. [AI] Jim to ping Bert on testing
    b.
4. [AI] Keith
    a. The chosen messaging protocol for TIER is AMQP
    b. TIER needs to produce specification for Inbound messaging connector to midPoint
    c. Investigate whether RESTful connector development should be based on Evolveum's Scripted REST connector or on the newer Superclass (abstract) REST model.
    d. These two issues are not time critical now as they are out of scope for the current evolveum SoW for a TIER container.

5. New Topics
    a. Centos 7 base requirement in TIER container standard
        i. Some efficiencies are gained by implementing fixes in just one main OS when there are issues.
        ii. Some efficiencies could be gained by container builders who already maintain in other Linux environments

iii. Still need to chat and decide - many users will have Red Hat licenses making Centos perhaps a tad more attractive.
iv. Other versions of Linux may have newer tool libraries that the component owners need - e.g., php

b. Containers for
   i. RabbitMQ - will need to start on a TIER container soon.
   ii. IDmatch - Scott K will do this build when code base is ready
c. midpoint reference implementation review https://spaces.internet2.edu/display/TPWG/TIER+midPoint+-+Docker+Reference+Implementation
d. Discussion topic for next week: can all TIER components be configurable enough such that they used without build-level customization?
e. Potential campus security group audit issues

June 11, 2018
No meeting scheduled - request that everyone work on action items.

June 4, 2018
Attendees (please add yourself):

- Jim Jokl - Virginia
- Ethan Kromhout - UNC Chapel Hill
- Kevin Ruderman - Boston University
- Scott Cantor - tOSU
- Scott Koranda - SCG
- Paul Caskey - Internet2
- John Gasper - Unicon
- Colin Thompson - UC Merced
- Chris Hubing - Internet2
- David Bickel - Indiana
- Keith Hazelton - UW-Madison
- Blair Christensen, University of Chicago
- Bert Bee-Lindgren, Georgia Tech
- John Bryson, Georgia Tech
- James Notoma, Georgia Tech

Call Agenda and Notes

1. Agenda Bash
   a. Any quick topics
      i. …
      ii. ...
   b. See item #8 below for anything that will consume significant time

2. JAVA discussion

a. Do we continue to require Oracle Java for all TIER components
b. Recent discussion re: allowing OpenJDK as an alternative if it is fully “supported” by the component development group.
c. Discussion
    i. JAVA is changing faster than ever; perspective of the Shibboleth project is that Oracle JAVA is the future. Impression is that JAVA from the OS distos is getting worse as opposed to better.
    ii. AZUL is a possibility. This is a curated version OpenJDK https://www.azul.com/products/zulu-and-zulu-enterprise/ GaTech has been running this in production. Internet2 has just started via the our current Grouper distro.
    iii. Possibilities for builds with AZUL with easy ability to use Oracle JAVA include simply mounting JAVA in the container.
d. Decision
    i. We will switch to Zulu as our default Java
    ii. All containers must include mechanisms and instructions for the use of Oracle JAVA.
    iii. This will also enable us to ship complete containers.

3. Supervisord
    a. We decided that Supervisord would be required in any TIER container that supports more than a single process. Now that there is support in Centos 7 for systemd in a container, do we want to keep this decision.
    b. Logging issues with Supervisord - resolution?
        i. No response ever came back from the Supervisord development team about accepting an RFE about logging from Supervisord itself.
        ii. The TIER logging format can be supported using a "trick" introduced by John Gasper.
    c. Should we limit the internal rate of change within our containers?
    d. Discussion
        i. We have a work around (“ugly, ugly hack”) for logging with Supervisord.
    e. Decision
        i. Retain requirement for Supervisord now, revisit if people run into issues.
        ii. To be clear: the proposal is to move to systemd from supervisord.

4. TIER Reference Implementations & LDAP Server Container
    a. Continue with openldap or migrate to 389 Directory Server
    b. Issue: support for large groups
    c. Discussion
        i. GaTech was seeing issues with large groups with 389 and Grouper - they stopped using really large groups a few years ago
        ii.
    d. Decision

i. We need more data
ii. [AI] Keith will ask the midPoint people re: large groups
iii. [AI] Bert will test 389 vs OpenLdap and see if 389 performs better with large groups.

5. Grouper Status
    a. Are we ready for production
        i. TIER Beacon?
            1. ChrisH is placing this into the application itself
            2. We will pick this up later when that work is done.
        ii. Services ready: Daemon, UI, and WS
        iii. Needs more work: SCIM Server
        iv. What about gsh:
            1. Via Web UI?
            2. Web service to listen for gsh files?
            3. The grouper team will discuss and make a decision
    b. See also JAVA discussion above
        i. We are already using Zulu (done)
    c. See also LDAP discussion above
        i. Pending investigation of ldap servers; if the midPoint people have some secret sauce for large groups, may result in changes to grouper.

6. midPoint Priorities and Update
    a. Shibboleth to protect User Self Service
    b. RabbitMQ / AMQP integration
        i. [AI] Keith -- 1) TIER needs to provide specifications on what is needed for a messaging connector, inbound channel, etc.; 2) The scripted REST Connector was easier to use than the new SuperClass REST connector; Why is the scripting one deprecated?
            1. https://wiki.evolveum.com/pages/viewpage.action?pageId=23167702
    c. Separate container for user self service implementation
    d. Container should support reference implementation

7. COmanage Status
    a. TIER-spec container (e.g., centos-7 based)
        i. Complete. No reported issues.
    b. TIER beacon, logging, etc
        i. Complete. No reported issues.

8. Insert your items here
    a. If OpenJDK/Oracle Java is open to reconsideration, can we also reconsider requirement of CentOS-7 as the base OS? It requires extra work (hence cost)

from the COmanage project since other consumers accept Debian and the official PHP Docker image is Debian based.

b. …

9. Hold for now
    a. Next steps: April 30 1.c (RabbitMQ, IDmatch, midpoint reference implementation, etc.
    b. Revisit the tomcat vs. jetty discussion - Shib has at least one open issue with Tomcat 8.5 (https://issues.shibboleth.net/jira/browse/IDP-1028) - which we just switched to in our container. We'll look at tomcat 9 asap.

May 28, Memorial Day
No call scheduled.

May 21, 2018 & May, 14, 2018
No call scheduled.

May 7, 2018
No meeting scheduled - Global Summit

April 30, 2018
See agenda below - *I expect a short call today*. If you own the packaging of a component, please try to attend today's call.
Attendees (please add yourself):

- Jim Jokl - Virginia
- Bill Kaufman - Internet2 (may need to drop at 4:45ish)
- John Gasper - Unicon
- Sara Jeanes - Internet2
- Chris Hubing - Internet2
- Paul Caskey - Internet2
- Steve Zoppi - Internet2
- Keith Hazelton - UW-Madison

Call Agenda and Notes

1. Wednesday Packaging topic session at Global Summit Trust & Identity Showcase
    a. We have approximately 25 minutes to plan for a (hopefully) interactive session
    b. Some topics to cover (what else is needed)
        i. Docker and Docker Swarm (while trying to not limit orchestration frameworks) --- probably a slide or two; spend some extra time here ; ask devops vs. IAM people in attendance; interaction of IAM people and

Devops people (devops people hand off at Docker layer to IAM folks)? ** conversation

ii. TIER Container Specifications --- one or maybe two slides
iii. Logging && TIER Beacon --- or two slides
iv. Component Reference Implementation(s)
v. Components
    1. Shibboleth
        a. Status
        b. Planned changes
        c. Still needed to reach TIER container compliance
        d. Download link
        e. Documentation link
        f. Ready for production?
        g. Items to be highlighted and discussed
        h. How you can help
            i. We need deployers

    2. Grouper
        a. Status
        b. Planned near-term changes
        c. Still needed to reach TIER container compliance
        d. Download link
            i. https://hub.docker.com/r/tier/grouper/
        e. Documentation link
            i. https://github.internet2.edu/docker/grouper
        f. Test/eval environment link https://github.internet2.edu/docker/grouper/tree/master/test-compose
        g. Ready for production: no
        h. Items to be highlighted and discussed
        i. How you can help
            i. We need testers

    3. COmanage
        a. Status
        b. Planned changes
        c. Still needed to reach TIER container compliance
        d. Download link
        e. Documentation link
        f. Ready for production: no
        g. Items to be highlighted and discussed
        h. How you can help
            i. We need testers

i.

c. Next steps
    i. midPoint
    ii. RabbitMQ - ubuntu version exists ; we may make an TIER version
    iii. Id Match
    iv. IAM Reference Implementation - e.g., the TIER “solution”
d. Discussion
    i. What else can we tell people that we need
        1. Volunteers - contribute to any aspects of the project
    ii. Add contact points for volunteers to a final slide

2. Before Global Summit [AI] Jim will update wiki’s to be re-cast as Reference Implementations instead of Large Scale Deployments
    a. Paul - re-cast the way that we discuss requirements around Swarm
    b. ChrisHu - current RabbitMQ container seems to be working fine. Keith - really no reason to go further from what is there. Jim - does it meet the basic specs of TIER container ie. like using CentOS (Chris checking) -
    c. Keith - have TIER container for MariaDB / LDAP? Jim - Yes - CentOS
    d.
3. …
4. Useful links
    a. How to build small containers with Alpine and the Docker build pattern
    b. Docker EE 2.0 Announcement (with support for Kubernetes in addition to Swarm)
    c. Cornell Cloud Forum Call for Proposals - Due June 1st
        i. Similar to last year’s Forum, we will have session presentations and panels of varying lengths as well as 5 minute Lightning Round sessions.

April 23, 2018
No in-person call today but we do have work for all component container owners.

The main agenda topic for today’s call was going to be preparation for Global Summit. We will instead complete this by email. We have time on the Wednesday of the event for an update/discussion on packaging status.

If you are responsible for the containerization of one of the TIER components please:

1. Reply to Jim Jokl’s email on what you think people need to know about your container,, including:
    a. Current status
    b. Planned changes
    c. Remaining modifications to meet the TIER standards
    d. Download URL and Documentation link
    e. etc.

    f. Most importantly, please include about any concerns you have with the present status, etc. and what needs to be done before Global Summit.
2. I'll work to pull things together into a coherent deck.

The time is supposed to be interactive so hopefully we'll get some good feedback.

Everyone: please list any topics that you want to make sure that we cover here.

1. ..
2. ..
3. ..

April 2, 2018

We will hold what I expect will be a *short* call to check in on action items, agenda, and status items below.

Attendees (please add yourself):

1. Jim Jokl - Virginia
2. Bill Kaufman - Internet2
3. Chris Hubing - Internet2
4. Ethan Kromhout - UNC Chapel Hill
5. Paul Caskey - Internet2
6. John Gasper - Unicon
7. Keith Hazelton - Wisconsin

Call Agenda and Notes

1. Quick items, additional agenda topics
    a. Logging changes
        i. https://spaces.internet2.edu/x/m4ZyBw
        ii. **[AI]** JimJ to update logging spec into TIER container spec
        iii. Logging update
            1. Specify that we need to remove spaces in the Environment and User Supplied Tokens
            2. Our delimiter between records will be a semicolon
            3. No delimiters are allowed in any of our four tokens but may exist in the fifth field (e.g., the native log data).
            4. Examples (from a run of ENV="test ing" (intentional space) and USERTOKEN="Build; 1.2.3"):
                a. *supervisord;console;testing;Build:1.2.3;2018-04-02 18:27:30,778 CRIT Set uid to user 0*
                b. *tomcat;catalina.out;testing;Build:1.2.3;2018-04-02 18:27:32,915 [main] INFO*

org.apache.coyote.http11.Http11NioProtocol- Initializing ProtocolHandler ["https-jsse-nio-443"]

c. *shib-idp;idp-process.log;testing;Build:1.2.3;2018-04-02 18:27:39,348 - INFO [net.shibboleth.idp.log.LogbackLoggingService:240] - Shibboleth IdP Version 3.3.2*

b. Insert your item(s) here
c. ...

2. midPoint Docker Container
    a. Work on the midPoint Docker Container SoW
    b. Evolveum ok with TIER Container Guide
    c. Database: they use MariaDB
        i. TIER will provide container with whatever DB
        ii. Need to have a conversation about setting up containers with correct roles
    d. Need something similar to the Shib IdP container documentation for an example of showing the proper steps to bring things up
        i. Training material:
            1. https://spaces.internet2.edu/display/ShibInstallFest/TIER+Shibboleth+IdP+Training
            2. https://spaces.internet2.edu/display/ShibInstallFest/TIER+Docker+IdP ***
        ii. https://github.com/IdentityPython/SATOSA/tree/master/doc ***
    e. https://spaces.internet2.edu/display/TPWG/TIER+midPoint+Docker+Deployment+-+Large+Production?src=contextnavpagetreemode
    f. Would be nice to have an initial data set for Users and some LDAP entry information so when you start up there is some example data. We may have something like this that folks could optionally load up as a separate step.
        i. JohnG: Grouper has a test-compose (directory under Git that has folders for put many things together) that can be pre-populated and would show them how to set this load data up.
    g. **[AI]** Bill will set up a Slack channel for discussion of the next steps

March 26, 2018

We will hold what I expect will be a *short* call to check in on action items, agenda, and status items below.

Attendees (please add yourself):

1. Jim Jokl - Virginia

2. Scott Koranda - SCG
3. Bill Kaufman - Internet2
4. John Gasper - Unicon
5. Paul Caskey - Internet2
6. David Bickel - Indiana
7. Chris Hubing - Internet2
8. Keith Hazelton - UW-Madison

Call Agenda and Notes

1. Quick items, additional agenda topics
    a. midPoint container
    b. Insert your item(s) here
        i. RabbitMQ container?
            1. https://github.com/docker-library/rabbitmq/blob/94de3d0908514407dfa02a61aa37642a1884de45/3.7/alpine/Dockerfile
            2. https://www.cloudamqp.com/
            3. https://github.com/CentOS/CentOS-Dockerfiles/tree/master/rabbitmq/centos7
            4. Pending some discovery or something that Keith may know, we expect that the “right” answer is for TIER to use #3 above
            5. (Keith) I’ll give the Centos one a trial run. Has anyone already done this?

..

2. Action Items and Updates
    a. Solicit Grouper Testers - Jim J
        i. Done - a few people have promised to test
        ii. Discussion of any Results of testing
            1. **[AI]** Jim J to re-ping the people who agreed to test the grouper build -- include some specific questions for the testers.
            2. ..
        iii. Internet2 Internal Use of Grouper Build
            1. Only positive feedback to date
            2.
    b. Solicit component owners feedback on logging - Jim J
        i. Done
        ii. Discussion of any issues uncovered
        iii. Logfile Softlink (-sfT) to /dev/stdout is a standard container trick - need to do the linking at run time instead of build time.
        iv. Workaround: https://github.internet2.edu/docker/grouper_noVM/issues/10
        v. **[AI]** Paul/John? - does log4j ConsoleAppender buffer so much for us that it will be problematic?
    c. Supervisord logging format change request - Scott K
        i. Request has for code change been made (format of logfile)

        ii. We hope to hear back soon.
    d. Tomcat Logging Format
        i. **[AI]** John G - Determine if we have the same issue with tomcat files - e.g., catalina.out naming
    e. Docker Container Specification - https://spaces.internet2.edu/x/m4ZyBw
        i. Component owner discussion re: issues
            1. ..
            2. ..
        ii. Component owner discussion re: timing of changes
            1. ..
            2. ..
        iii. **[AI]** Jim J to add a new 7.b.iii to allow processing based on naming convention in startup scripts - done
3. New [AI] Jim J: Update and sync the Component Reference Implementations for COmanage, Grouper, Shibboleth, and midPoint to the new TIER container specification.
4. Other Items
    1. Possible Evolveum contract work on containerizing midPoint in conformance with the Docker Container Specification
        a. **[AI]** Keith H will work on a SoW for an Evolveum quote.

March 19, 2018

No packaging call is scheduled for Monday March 19. Please work on and update the action items assigned to you.

Action Items and Updates

1. Solicit Grouper Testers - Jim J
    a. Done - a few people have promised to test
2. Solicit component owners feedback on logging - Jim J
    a. Incomplete
3. Supervisord logging format change request - Scott K
4. Migrate the Docker Container Specifications to Confluence - Jim J
    a. Draft complete: https://spaces.internet2.edu/x/m4ZyBw
5. Component owners please review 4.a for usability and level of effort. The container specification replaces a few standalone AIs for meetings over the past two months.

March 12, 2018

Attendees (please add yourself):

1. Jim Jokl - Virginia
2. Bill Kaufman - Internet2
3. Scott Koranda - SCG
4. John Gasper - Unicon

Call Agenda and Notes

1. Any quick items/topics
    a. First draft of the Centos COmanage container is expected mid-month.
2. Action Items from the last call
    a. Jim - Solicit Grouper Testers
        i. The standalone TIER Grouper image is used to run each of the Grouper roles.
        To test with the Grouper image, use the latest patched build (tier/grouper-multi-purpose:2.3.0-a97-u41-w11-p16), one will probably want to customize the images building local images. The link below has general information on using the image, and a sample of how one might build a (test) environment can be found by looking in the test-compose directory of the project (update each of the Dockerfiles your copy of the test-compose directories to use the above mentioned image).
        ii. https://github.internet2.edu/docker/grouper_noVM/tree/multi-stage-build
    b. Component Owners - verify logging will work
        i. **[AI]** Jim to email each component owner/builder to verify new logging config
3. TIER Container Common Standards
    a. Compatibility/ease of use with SWARM while not breaking other options.
    b. Base Image - Centos 7 - one of:
        i. Standard Centos 7 image we have been using with the addition of supervisord when needed
        ii. Centos 7 image from Dockerhub that includes what is needed to use systemd as init (instead of supervisord)
            1. https://hub.docker.com/r/centos/systemd/
    c. Secret Processing
        i. Assume secrets are mounted in /run/secrets (to support compose in swarm)
        ii. Secret Availability - in-container startup script behavior
            1. Accept the secret in the environment, e.g., COMPONENT_DATABASE_PASSWORD=foobar
            2. If the filename version of the name exists, prefer it: COMPONENT_DATABASE_PASSWORD_FILE=/var/run/secrets/some_file
            3. If both exist, prefer the FILE option a
        iii. Logging
            1. All containers log to stdout
            2. Goal - easily parsable logs for:
                a. Component Name
                b. Native logfile name
                c. Environment (e.g., Prod, Dev, test)

    d. A user supplied token via the environment
3. We will have deployers use the `--log-opt tag`
    a. https://docs.docker.com/config/containers/logging/log_tags/
    b. This solution solves items 3.iii.2.a, 3.iii.2.c, and 3.iii.2.d
4. To solve 3.iii.2.b, we need an inventory of components where we are unable to change logfile format. Components with a single logfile per container should be OK and not need remediation.
    a. Supervisord
        i. Issue: can not change format of logfile to prepend “supervisord.log”.
        ii. Looked at potential for external process to transform log format before writing to stdout but prefer not to use this mechanism due to added complexity
        iii. Scott did some digging
            1. No good news - a source code change is needed
            2. The code is python but they do not use python logging
            3. **[AI]** Scott will ask about possibility for a feature update
    b. Mariadb - should be OK, single logfile per container
    c. OpenLDAP - should be OK, single logfile per container
    d. COmanage (yet--this could become a requirement for upstream)
    e. Shibboleth idp
        i. Shibboleth itself OK via log4j
        ii. Catalina.out tomcat issues
    f. Grouper
        i. Core grouper logging will be ok
        ii. Same issues with Supervisord and tomcat
        iii.

iv.

v.

March 5, 2018

No call today, *instead* please look at the minutes from last week’s call (Feb 26) on TIER container requirements and logging.

We made significant changes to logging requirements from the past discussion and started on the definition requirements for all tier containers. These are in Sections 1.a and 4.

See the Action Items below and please review in general if these changes will work well for your products.

February 26, 2018
Attendees (please add yourself):

1. Jim Jokl - Virginia
2. Scott Koranda - SCG
3. Christopher Hoskin - University of Oxford
4. Justin Robinson - Indiana University
5. David Bickel - Indiana University
6. Carey Black - tOSU

Call Agenda and Notes

1. Any quick items/topics
    a. Scott's 2/19 questions on logging
        i. Goal - easy to parse logs for:
            1. Component Name
            2. Native logfile name
            3. Environment (e.g., Prod, Dev, test)
            4. A user supplied token via the environment
        ii. Should we just ask deployers to use `--log-opt tag` ?
            1. https://docs.docker.com/config/containers/logging/log_tags/
            2. This solution solves items 1.a.i 1, 3, and 4
        iii. To solve 1.a.i.2, we need an inventory of components where we can not change logfile format. Components with a single logfile per container should be OK and not need remediation.
            1. Supervisord - maybe fix via syslog? (But does that require another running process?)
            2. Mariadb
            3. OpenLDAP
            4. COmanage (yet--this could become a requirement for upstream)
            5. Shibboleth idp
                a. Shibboleth itself OK via log4j
                b. Catalina.out tomcat issues
            6. Grouper
                a. Core grouper logging will be ok
                b. Same issues with Supervisord and tomcat
                c.
        iv. [**AI**] Component Owners, please verify that this mechanism will work for your software.

1. Owners now need to prepend logs with a single item - the component name.

b. …
c. ...

2. The TIER Grouper Build
    a. Tester Feedback
        i. http://bit.ly/1XNvSmC
        ii. [**AI**] Jim needs to contact some likely suspects directly
        iii.
    b. Other grouper topics (from previous call notes - e.g., Feb 12)
        i. 6.c.ii.2 - (yellow) in progress
        ii. Sanity checking - SCIM
        iii. Container start-up health check work in progress

3. Action Items from February 12
    a. [**AI**] Jim to update status on action items within the Feb 12 call notes.

4. TIER Container General Specifications
   Compatibility/ease of use with SWARM while not breaking other options.
    a. Base Image - Centos 7 - one of:
        i. Standard Centos 7 image we have been using with the addition of supervisord when needed
        ii. Centos 7 image from Dockerhub that includes what is needed to use systemd as init (instead of supervisord)
            1. https://hub.docker.com/r/centos/systemd/
    b. Secret Processing
        i. Assume secrets are mounted in /run/secrets
        ii. Recommended solution example
            1. Accept the secret in the environment, e.g., COMANAGE_REGISTRY_DATABASE_PASSWORD=foobar
            2. If the filename version of the name exists, prefer it: COMANAGE_REGISTRY_DATABASE_PASSWORD_FILE=/var/run/secrets/some_file
    c. Logging
        i. Containers log via stdout,
        ii. See Section 1.a above
    d. Supervisord
        i. Potential issue: can not change format of logfile to prepend “supervisord.log”. Current format:
            2018-02-26 22:01:11,950 INFO spawned: 'shibd' with pid 7
        ii. We would like to make supervisord the default for multi-component containers (this would also simplify 4.a above)

iii. Potential for external process to transform log format before writing to stdout
iv. Potential feature request to the supervisord owners
v. [AI] Scott and Jim will do some digging for a possible solution.

e. …
f. …
g.

February 19, 2018
No TIER Packaging call today. Please review the minutes from our February 12 call and work on Action Items for next week.

February 12, 2018
Attendees (please add yourself):

1. Jim Jokl - Virginia
2. Sara Jeanes, Internet2
3. Scott Koranda, SCG
4. Bill Kaufman - Internet2
5. John Gasper, Unicon
6. Paul Caskey - Internet2
7.

Call Agenda and Notes

1. Any quick items/topics
    a. Jim needs to leave no later than 4:45 pm eastern
    b. ShibIU - see #4 below

2. The TIER Grouper Build
    a. Tester Feedback
        i. http://bit.ly/1XNvSmC
    b. Other grouper topics (see previous call notes below)
        i. 6.c.ii.2 - (yellow) in progress
        ii. Sanity checking - SCIM
        iii. Container start-up health check work in progress

3. TIER Campus Success Meeting Topics
    a. Component Logging Discussion
        i. Common format, conventions
        ii. [**AI**] All logs written to stdout; believe ok but need to test atomic writes

1. Scott verified OK [2018-02-26] on a production system with multiple log files writing to stdout within a single container and no log corruption has been seen.
2. $ docker service logs --tail 300 -f comanage-registry_comanage-registry

   comanage-registry_comanage-registry.1.o2k7b4uw86mf@mwa-registry | 10.255.0.2 - - [26/Feb/2018:20:06:47 +0000] "GET / HTTP/1.1" 302 3950 "-" "-"

iii. ComponentName, LogfileName, Env (mode: Prod, Dev, Test), UserDefinedEnvironmentVar
iv. Any dot in a filename is replaced by a '-'
v. Verify that swarm prepends container ids to each log line
vi. Future: JSON formatted logs
vii. **[AI]** Component owners - please try to estimate how long these changes will take.
   1. 2018-02-26 -- on hold pending completion of discussion on logging

b. Review confluence service definitions
   i. **[AI]** Componen owners, please update this google doc with log names
      1. 2018-02-26 -- no update
   ii. **[AI]** Jim to update confluence sites for new logging plan
      1. 2018-02-26 - done at a generic level
      2. Replaced by TIER container standards section
   iii. https://spaces.internet2.edu/display/TPWG/TIER+COmanage+Docker+Deployment+-+Large+Production
      1. Httpd (5 files)
         a. httpd.access_log, httpd.error_log, httpd.ssl_access_log, httpd.ssl_error_log, httpd.ssl_request_log
      2. Shibboleth SP (2 files)
         a. shibd.log, native.log
      3. Supervisord (1 file)
      4. COmanage (2 files)
         a. error.log, debug.log
      5. MariaDB (? files)
      6. OpenLDAP (1 file)
      7. SATOSA (1 file)
   iv. https://spaces.internet2.edu/display/TPWG/TIER+Grouper+Docker+Deployment+-+Large+Production
      1. Httpd (5 files)
         a. httpd.access_log, httpd.error_log, httpd.ssl_access_log, httpd.ssl_error_log, httpd.ssl_request_log

2. Tomcat (5 files)
    a. tomcat.catalina-out, tomcat.access_request, tomcat.localhost
3. Grouper - component name will be GrouperLoader, GrouperUI, etc.
    a. E.g., GrouperLoader.grouper_error, grouper_debug, grouper_bench, grouper_event
4. Mariadb
5. Supervisord
6. Shibd.log native.log

v. https://spaces.internet2.edu/display/TPWG/TIER+Shibboleth+IdP+Docker+Deployment+-+Large+Production
    1.
vi. https://spaces.internet2.edu/display/TPWG/TIER+midPoint+Docker+Deployment+-+Large+Production
    1.
vii.

c. Other topics
    i. …
    ii. …

4. Other topics
    a. ShibUI
    b. Review Phase 2 Milestones https://drive.google.com/open?id=1vg7NUt3ybhm_jWQmY1U-QurGoLDpzB5w
    c. **[AI]** Please comment as soon as you can by no later than CoB Wednesday.

February 5, 2018

We will **not** hold a TIER Packaging call on 02/05/2018 -- too many of our regular attendees will be away from the office and unable to attend.

January 29, 2018

Attendees (please add yourself):

1. Jim Jokl - Virginia
2. Carey Black - tOSU
3. Scott Koranda - SCG
4. Bill Kaufman - Internet2
5. Keith Hazelton - UW-Madison
6. Chris Hubing - Internet2
7. John Gasper - Unicon
8. David Bickel - Indiana
9. Ethan Kromhout - UNC Chapel Hill (arrived 4:30)

Call Agenda and Notes

1. Any quick items/topics
    a. Insert items here
2. The TIER Grouper build: John Gaspar
    a. Design discussion, what is ready, request for testers
    b. A build is in progress, should be available soon on Dockerhub
        i. https://jenkins.testbed.tier.internet2.edu/job/docker/job/grouper_noVM/job/multi-stage-build/4/console
        ii. https://hub.docker.com/r/tier/grouper-multi-purpose/tags/
    c. Remaining tasks
        i. Naming
            1. What should the final image name (grouper_novm now) - desired name will be be “grouper” unless it turns out to be a major deal.
            2. How do we tag releases: “latest”, plus by patch level, e.g., x-y-z for the three patch levels.
            3. We expect to do weekly updates
        ii. Be able to update to a specific patch level
            1. ***Chris Hyzer*** *has/will provide this capability. Is this ready now - we think so? CH: yes*
            2. Grouper install properties file will need some changes
                a. `grouperInstaller.autorun.installPatchesUpToACertainPatchLevel`
                    i. REF: commit

```
/**
* if should install up to patch levels, comma separated
* e.g. grouper_v2_3_0_api_patch_9, grouper_v2_3_0_ui_patch_10, grouper_v2_3_0_ws_patch_5
*/
```

        iii. Initial sanity checking on SCIM component is needed Grouper TIER SCIM Server
            1. Feedback: ***Chris Hyzer*** *-- how to do a quick automated test?*
            2. CH: Please open a jira and assign to vivek
        iv. A screencast for YouTube of how to build/deploy
            1. In addition to the documentation
    d. Resolution of the 12/22 email thread
        i. We will ask **Chris Hyzer** for health check functionality that will enable us to wait until all underlying containers (e.g., all sources are up) before we move forward with startup. The most critical item is the database - we should not start without it. User configurable for all subject sources.

ii. CH: if the database is not up, you cant start tomcats… if you need an installer call for this please open a jira and assign to me

iii. Yes, this is also a request we will ask **Chris Hyzer** about

iv.

> Sure, is it ok if that goes in the installer instead? Im just thinking that it already has some of that logic, and GSH is intended to connect to grouper, the installer is more of a bootstrap thing like this... you could make a config file and have a dir and run the installer for each of the two cases?
>
> Btw, gsh -registry -runscript will update if not up to date right? Not sure about non-zero on problems though...
>
> Thanks
> Chris
>
> -----Original Message-----
> From: Scott Koranda [mailto:skoranda@gmail.com] On Behalf Of Scott Koranda
> Sent: Friday, December 22, 2017 9:23 AM
> To: tier-packaging@internet2.edu
> Subject: [tier-packaging] Grouper questions maybe RFE related to Docker packaging
>
> Hi,
>
> (Primarily for Chris Hyzer but feedback from everyone encouraged...)
>
> I have a Dockerfile for Grouper that is similar to the nice Dockerfile
> that John Gasper and Chris Hubing are building. Mine goes a bit further
> because when we deploy it with COmanage we can usually make certain
> assumptions and simplifications.
>
> I want the entrypoint script that starts Tomcat to do two things:
>
> 1) Wait patiently until the container(s) running the database(s) are
> reachable. I use the plural because we have both a relational database
> for Grouper state and an LDAP directory for subject sources (managed by
> COmanage).
>
> 2) Create the Grouper tables/indexes if they are not already present.
>
> For (1) I could have scripts that have nothing to do with Grouper, but I
> like the idea of "Grouper" itself being able to tell me if connections
> are ready using the existing hibernate and subject source configuration.
>
> In particular I think it would be nice if gsh.sh could do so. Something
> like
>
> gsh.sh -ready
>
> with a return value of '0' if all is good and ready and non-zero for any
> problems. Then I would just have the entrypoint script loop over that

command.

I don't see the equivalent of a -ready flag, and right now I think gsh.sh returns 0 no matter what happens "inside" of it.

Am I missing something? If not, will you consider such an enhancement?

For (2) I know about 'gsh.sh -registry -reset' but it doesn't quite do what I want. I want a check to see if the tables exist at all and if they do not to have them created, and again with a return value of '0' if it worked and non-zero if not.

Am I missing something about -registry? If not, will you consider such an enhancement?

I think these enhancements would be useful to Chris and John's effort as well (we might have mentioned this on the call, apologies if I missed it).

Thanks,

Scott K

    e. Request for testers
        i. https://github.internet2.edu/docker/grouper_noVM/tree/multi-stage-build

3. midPoint training environment
    a. direct container vs. VM, etc.
    b. Leaning towards the VM idea; move conversation to training discussion (packaging channel on Slack (tier-packaging))

4. Insert your items here

January 22, 2018
Attendees (please add yourself):

1. Jim Jokl - Virginia
2. Ethan Kromhout - UNC
3. David Bickel - Indiana
4. Carey Black - tOSU
5. Keith Hazelton - UW-Madison

Call Agenda and Notes

1. Any quick items/topics
    a. Anything new to discuss on the Grouper build? See also 12/22 thread on tier-packaging -- review on next call
    b. Insert your item(s) here
    c. ...

2. The focus for today's call is continued discussion on midPoint packaging requirements.
    a. Action Items
        i. Ethan will add information on evolveum mailing list to TIER slack channel [AI]
            1. https://wiki.evolveum.com/display/midPoint/Mailing+Lists
            2. This is done
        ii. Ethan will check on midPoint upgrade process [AI] - See minutes for detailed questions (database/code mismatch, handling upgrades, URL, etc.)
            1. Done - text block on its way

> from my experience, everytime I forgot to run upgrade scripts (read:
> during experimenting with midpoint; did not happen in production),
> midPoint refused to start and there will be an error in idm.log / midpoint.log. So midPoint will not start.
>
> I think it's connected to "validate" option in config.xml:
>
> <hibernateHbm2ddl>validate</hibernateHbm2ddl>
> (for embedded H2 repository this can be set to "update" and it will
> update the db structure).
> To be honest, the error message is a bit cryptic. But, nevertheless,
> midPoint will not start.
>
> And yes, it requires the "validate" option to be set. It is the default
> for all databases other than H2.
> There will be an error on startup if the database model is not compatible (e.g. there are missing columns). However, I would say that
> midPoint will continue to operate if the database model is compatible,
> e.g. tables were extended with additional (non-mandatory)

columns.

2.

b. TIER midPoint Docker Deployment - Large Production
    i. Confluence document is complete

c. midPoint and user self service?
    i. Shibboleth SP as part of build for self service -- Ethan will try to prototype.
    ii. Can we externalize user self service?
        1. Separate entry point for admin and self service
        2. The idea is to protect the admin interface more heavily than the self service interface. Can we run on a different server, different port, etc. Additional protection at the network layer is desired.
    iii. https://wiki.evolveum.com/display/midPoint/Self+Services

3. midPoint and Training
    a. What should we request of evolveum re: the upcoming TIER midPoint training scheduled for the end of February
    b. What do we already know midPoint is providing
        i. https://evolveum.com/training-and-certification/midpoint-deployment-fundamentals-end-of-february-2018/
    c. Their usual training topics
    d. Plan: 2018-01-22
        i. Suggestion for midPoint - [AI] jaj to write up
            1. Use docker version of midPoint in their VM
            2. Still provide their VM with demo files, etc., etc.
            3. Longer-term container design using docker secrets - TIER philosophy
        ii. What do we want highlighted in the training - see 2.b.ii
            1. Add a little Docker on Day 1
            2. More in-depth on building a connector (e.g., ldap connector code)
                a. Need a connector that we could edit for ERP (e.g., Banner)
            3. Can we instead add a full day (day 6) of video on connectors
        iii. Should we volunteer to help people (remotely) to help take some of the support load off of evolveum.
    e. ...

January 15, 2018
Holiday - no call today

January 8, 2018
Attendees (please add yourself):

1. Jim Jokl - Virginia
2. Ethan Kromhout - UNC
3. Bill Kaufman - Internet2
4. Sara Jeanes - Internet2

Call Agenda and Notes

1. Any quick items/topics
    a. Anything new to discuss on the Grouper build? See also 12/22 thread on tier-packaging -- review on next call
    b. No call next week (1/15 - Martin Luther King day)
    c. Insert your item(s) here
    d. ...

2. The focus for today's call is continued discussion on midPoint packaging requirements.
    a. TIER midPoint Docker Deployment - Large Production
    b. Do we have an answer to the AI on the behavior of midPoint on upgrades (i.e., code and database version mismatch)? Not yet
    c. How does evolveum recommend that upgrades are handled?
        i. Add URL from evolveum docs on last upgrade?
        ii. And/or Ethan [AI] to include in his discussion with Evolveum
    d. Do they provide scripting for database schema changes, etc.?
        i. Yes, evolveum provides the scripts for updates
        ii. Ethan [AI] will ask evolveum if midPoint is smart enough not to run if the database and application versions do not match in some important way.
    e. midPoint and user self service?
        i. Ethan - have not really worked in this area
        ii. Shibboleth SP as part of build (put midPoint behind an SP)? Ethan - should not be real hard to put this together. It may be as easy as placing the EPPN in REMOTE_USER. Ethan will try to prototype.
        iii. Can we externalize user self service?
        iv. https://wiki.evolveum.com/display/midPoint/Self+Services
        v.
    f. Ethan will add information on evolveum mailing list to TIER slack channel [AI]
        i. https://wiki.evolveum.com/display/midPoint/Mailing+Lists
    g. …
    h. ...

December 18, 2017

Attendees (please add yourself):

1. Jim Jokl - Virginia
2. Ethan Kromhout - UNC
3. Bill Kaufman - Internet2
4. Chris Hubing - Internet2
5. Sara Jeanes - Internet2
6. Keith Hazelton - UW-Madison
7. David Bickel - Indiana University

Call Agenda and Notes

1. Call Schedule
    a. Today (12/18) will be the last TIER Packaging call for the year.
    b. The next call will be on Monday January 8, 2018
2. Any quick items/topics
    a. Quick review of last Tuesday's Grouper container build meeting with John Gasper. GitHub link is HERE
3. The focus for today's call is midPoint packaging requirements.
    a. TIER midPoint Docker Deployment - Large Production current wiki
    b. See notes and action items from the last midPoint discussion on
    c. As far as we know most or all of the existing Connectors are designed to be synchronous with midPoint. Ethan likely was the first to use RabbitMQ in an async fashion.
        i. Should RabbitMQ be part of the core component pieces? How to add it in. Talk more with Evolveum on how to support this.
            1. Ethan has a "crude" RabbitMQ tracer that works with the demo and may be a nice debug tool for folks working with the code
        ii. Be nice to have 1 docker compose to pull in all the minimal parts
    d. EthanK has done an upgrade which required 2-steps
        i. Clean location that comes as part of the package
            1. Deploy war file and run sql script to make any schema changes
                a. can/should this be automated?
    e. NOTE: midPoint version 3.7, Darwin, came out today New with this release: Stand-alone deployment based on Spring Boot
    f. Have midPoint distribution ready and integrated with the midPoint training tentatively set for late February. This could work if the training is dedicated to Internet2/TIER.

December 12, 2017 - Grouper
Attendees (please add yourself):

8. Jim Jokl - Virginia
9. Chris Hubing - Internet2
10. James Babb - UW Madison

11. Carey Black - tOSU
12. Paul Caskey - Internet2
13. John Gasper - Unicon
14. Scott Koranda - SCG

Grouper multi-purpose container code review:

- https://github.internet2.edu/docker/grouper_noVM/tree/multi-purpose-grouper-image
- Perhaps add env variable that contains version and/or patch level
- Web.xml is set up for shib, perhaps document how to do local auth

December 11, 2017
Attendees (please add yourself):

15. Jim Jokl - Virginia
16. Bill Kaufman - Internet2
17. James Babb - UW-Madison
18. Keith Hazelton - UW-Madison
19. Dean Lane - Rice
20. Chris Hubing - Internet2
21. Paul Caskey - Internet2

Call Agenda and Notes

4. Any quick items/topics
5. The focus for today's call is midPoint packaging requirements.
    a. https://spaces.internet2.edu/display/TPWG/TIER+midPoint+Docker+Deployment+-+Large+Production
    b. COmanage Link for reference
        i. https://spaces.internet2.edu/x/64PdBg
6. Wisc Midpoint proof of concept container source: https://github.internet2.edu/TIER/wisc-midpoint
    a. We're still using the 3rd party midPoint container
    b. Need a dB in the environment for midPoint
        i. We expect
7. HA/LoadBalancing for midPoint
    a. https://wiki.evolveum.com/display/midPoint/High+Availability+and+Load+Balancing

MidPoint upgrades: we need to understand if the code will refuse to run when the code and database versions don't match. [AI] asking about how midpoint handles major upgrades...will it just refuse to run, will it auto-upgrade, or will there be really bad behavior?

December 4, 2017
Attendees (please add yourself):

22. Jim Jokl - Virginia
23. John Gasper - Unicon
24. James Babb - UW Madison
25. Carey Black - tOSU (running late)
26. Scott Koranda - SCG
27. Chris Hyzer - Penn
28. Keith Hazelton - UW-Madison
29. Chris Hubing - Internet2
30. Bill Kaufman - Internet2

Call Agenda and Notes

The focus for today's call is the production Grouper distribution.
https://spaces.internet2.edu/display/TPWG/TIER+Grouper+Docker+Deployment+-+Large+Production
Current PR: https://github.internet2.edu/docker/grouper_noVM/pull/6
TIER Beacon - this functionality is already built into the Grouper loader and is on by default:
https://spaces.internet2.edu/display/TWGH/TIER+Instrumentation+-+The+TIER+Beacon
Likely Load Balancer: https://github.com/containous/traefik

- Docker War Story: Zombie processes showing up when not using some sort of init system (tini or dumb-init). In theory, bash also reaps zombie processes too if you started your command with bash -c '/whatever/you/want/to/run.sh'....but that won't pass signals appropriately.

- Open question for next week - Java. We have said in the past that all TIER applications that use Java will use Oracle Java. This adds legal and scripting complexity to applications that may not need Oracle Java. Question: what versions of Java does Grouper officially support? If other than Oracle Java, should we consider something different for Grouper?

- We are working to schedule a detailed Grouper implementation review call for the week of December 11. If you want to attend, please fill in the Doodle poll https://doodle.com/poll/fbktpuqapzb7i6iv

Nov 6, 2017
We will not have a tier-packaging calls in November. Work is progressing well on Grouper and COmanage. We should be able to meet on these components in early December. We want to get versions of these components out for testing before we start the next phase of our work on the suite.

October 29, 2017
October 23, 2017
No Packaging call on these two days - catch up time from TechEx.

October 9, 2017
October 16, 2017
No call on Monday, Oct 9 or Monday Oct 16. I hope to see everyone at Tech Ex.

October 2, 2017
Today's call will likely be short
Attendees (please add yourself):

31. Jim Jokl - Virginia
32. Scott Koranda - SCG
33. Carey Black - tOhio State
34. Bill Kauffman - Internet2
35. Paul Caskey - Internet2
36. Chris Hubing - Internet2
37. Sara Jeanes - Internet2 (late)
38. Scott Cantor - tOSU
39. Chris Phillips - CANARIE
40. Kevin Ruderman - Boston U
41.

Call Agenda and Notes

1. Call logistics at top of this Google Doc
2. Quick Topics
    a. Agenda bash
    b. Insert your item(s) here
    c. ...
3. Agenda
    a. Shibboleth Deployment Document

    b. https://spaces.internet2.edu/display/TPWG/TIER+Shibboleth+IdP+Docker+Deployment+-+Large+Production
    c. ChrisP: Any performance testing being done?
    d.
4. Insert your item(s) here
    a.

September 22, 2017
We will not hold a tier packaging call today at 4:00 eastern. Work is in progress on the production builds along with the shibboleth production document. I'm expecting to be far enough along next week to meet, so please continue to hold this time open on your calendars.

September 22, 2017
Special call around Grouper Packaging

Attendees (please add yourself):
42. Jim Jokl - Virginia
43. Paul Caskey - Internet2
44. Chris Hubing - Internet2
45. Bill Kaufman - Internet2

GitHub repo Chris Hubing is working with
Requirements - TIER Grouper Docker Deployment - Large Production

TechEx Plan
- Postpone HA

Chris Hyzer - ask: - Chris Hubing will send a response

TIER packaging grouper team,

I would like to do a demo or have you give a demo or show a movie of you giving a demo of the new TIER Grouper Packaging at the Sunday morning TechEX 4 hour Grouper seminar. We have ~30 people so it will be a good venue to show the new packaging…

Some questions:

1. When will the new Grouper TIER packaging be available
2. Can someone from your team stop in to our training sometime between 8-12 and give a short overview/demo of the new packaging?

3. Or can you make a short vid of this we can show?
4. Or can the package and docs be available a week before techex so we can evaluate and make a few slides or a demo?

September 18, 2017
No call is scheduled - work is in progress generating builds.

Attendees (please add yourself):

46. Jim Jokl - Virginia
47. John Gasper - Unicon
48. Chris Hubing - Internet2
49. Bill Kaufman - Internet2

Special call with John Gasper of Unicon around the next rev of TIER Grouper and also the work that Unicon has done on a dockerized Grouper

September 11, 2017
Today's call will likely be short
Attendees (please add yourself):

50. Jim Jokl (must leave at 4:45 eastern) - Virginia
51. Scott Koranda - SCG
52. Chris Hyzer - Penn
53. Bill Kaufman - Internet2
54. Chris Hubing - Internet2
55. Sara Jeanes - Internet2
56. David Bickel - Indiana
57. Paul Caskey - Internet2

Today's call will likely be short

September 4, 2017
Labor Day - No packaging call scheduled.

August 28, 2017
Attendees (please add yourself):

58. Jim Jokl - Virginia
59. Chris Hubing - Internet2
60. Keith Hazelton - UW-Madison
61. Carey Black - tOhio State Univ.

62. Chris Hyzer - Penn
63. Bill Kaufman - Internet2
64. Michael Gettes - UFlorida
65. Kevin Ruderman - Boston University
66. Jon Miner - UW-Madison
67. James Babb - UW-Madison

Call Agenda and Notes

5. Call logistics [at top of this Google Doc]
6. Quick Topics
    a. Agenda bash
    b. Insert your item(s) here
    c. ...
7. Agenda
    a. Quick review of COmanage deployment document
    b. Work together to piece together the Grouper deployment document
    c. COmanage https://spaces.internet2.edu/display/TPWG/TIER+COmanage+Docker+Deployment+-+Large+Production
    d. Grouper (see July 17th and August 7th meeting notes below for more ccntext) https://spaces.internet2.edu/display/TPWG/TIER+Grouper+Docker+Deployment+-+Large+Production
8. Insert your item(s) here
    a.

August 21, 2017

I have heard from a couple of people that they will still be at eclipse events at our regular call time so we will cancel for one more week. I'll be on the bridge today Work has been progressing and I expect to see the drafts of the summary deployment documents completed this week. Links to these documents will be placed here when they are ready.

- COmanage https://spaces.internet2.edu/display/TPWG/TIER+COmanage+Docker+Deployment+-+Large+Production
- Grouper https://spaces.internet2.edu/display/TPWG/TIER+Grouper+Docker+Deployment+-+Large+Production
- Shibboleth

August 14, 2017

We will not hold a TIER Packaging call on 8/14. Our primary topic was to complete the COmanage discussion and schedule conflicts mean that some critical people can't join today's call. We will attempt to complete some work via email before next week's call.

August 7, 2017

Attendees (please add yourself):

68. Jim Jokl - Virginia
69. Scott Koranda - SCG
70. Paul Caskey - Internet2
71. Bill Kaufman - Internet2
72. Chris Hubing - Internet2
73. Tom Zeller - Shib
74. Keith Hazelton - UW-Madison

Call Agenda and Notes

9. Call logistics at top of this Google Doc
10. Quick Topics
    a. Agenda bash
    b. Insert your item(s) here
    c. ...
11. Special Agenda - COmanage Docker Deployment
    a. Typical COmanage Deployment Scenarios
        i. In a federation
        ii. Authentication via Shibboleth SP (Requirement is REMOTE_USER and ability to pass some attributes)
        iii. COmanage provisions into a dedicated (non-campus/non-VO general) LDAP server
            1. Mostly openLDAP (maybe one instance of 389)
        iv. Typically a proxy (SATOSA) or (less often) a SAML AA sits in front of the LDAP and provides services
    b. A COmanage instance is multi-tenant (multi-VO)
        i. Typically one LDAP instance per-tenant
    c. A common deployment uses redundant LDAP/Proxy configurations; COmanage itself is generally not high availability. The recommendation is that you don't use the COmanage REST api for anything that is high availability
    d. Supported Databases: MARIADB and Postgres (agnostic)
    e. Grouper is often deployed with a grouper instance - perhaps 50%
        i. COmanage creates the VOs/COs; COmanage groups are often sufficient
            1. Forward and reverse references (isMemberOf and ou=Groups) are maintained

ii. The other half of the deployments want more sophisticated capabilities and use Grouper.
    1. COmanage provisions users into LDAP - standard ou=People records (filled in with CO-related data)
    2. Grouper reads COmanage People objects/users from that LDAP
    3. COmanage provisions groups into Grouper via Grouper Web Services
    4. Grouper PSP (or soon PSP-NG) maintains groups into LDAP (standard ou=Groups); forward and reverse references are maintained.

f. COmanage Person Identifier Creation
    i. Every CO configure COmanage to automatically generate an identifier - opaque with a simple prefix for that CO. e.g., a prefix followed by a six-digit identifier. The opaque identifier facilitates the use of identity linking. It is possible and common to provision additional identifiers that include some more friendly names when later driven by applications.
    ii. A VO-person schema is in the works
g. Automation and/or Documentation
    i. It would be nice to automate the standard configuration (once running) tasks (this may happen anyway)
    ii. A COmanage deployment guide would also be nice.
h. Typical deployments are complex, with multiple moving parts, solving specific problems
i. The COmanage project maintained container
    i. https://github.com/Internet2/comanage-registry-docker
    ii. Not “a” container, multiple containers
        1. Default simple (learn COmanage) container
        2. COmanage + Shibboleth SP
            a. Apache and Shibd (using supervisord)
        3. COmanage + mod_auth_oidc
            a. Apache with mod_auth_oidc
    iii. The containers are Apache and authentication only; an external databases and LDAP are still needed
        1. Full deployments are done in Docker swarm using Docker secrets
        2. All of this is described in the git repository
        3. The COmanage project also provides a packaged LDAP
j. Needs - Next Steps
    i. Everyone to look at the git repository
        1. Compiles shib from source
    ii. Centos vs. Debian?

Attendees (please add yourself):

75. Chris Hubing - Internet2
76. Scott Koranda - SCG
77. Peter DiCamillo - Brown University
78. Chris Hyzer - Penn
79. Bill Thompson - Lafayette College
80. Carl Waldbieser - Lafayette College
81. Blair Christensen - University of Chicago
82. Sara Jeanes - Internet2
83. Kevin Ruderman - Boston U.

Call Agenda and Notes

12. Call logistics at top of this Google Doc
13. Agenda bash
14. TIER Package Release 17070 - https://spaces.internet2.edu/display/TPD/TIER+Package+Delivery
    a. Includes a standalone Shib IDP container that supports either burned, mounted or hybrid configs
        i. https://spaces.internet2.edu/display/TPD/Shibboleth-IdP+Standalone+Container+Release+17070
15. Special Agenda - Continued Grouper Discussion from last week
    a. See notes from July 17 below
    b. Today's discussion started here

July 17, 2017

Attendees (please add yourself):

84. Jim Jokl - Virginia
85. Scott Koranda - SCG
86. Kevin Ruderman - Boston University
87. Blair Christensen - University of Chicago
88. Chris Hubing - Internet2
89. Sara Jeanes - Internet2
90. James Babb - UW Madison
91. Bill Thompson - Lafayette College
92. Carey Black - tOhio State Univ.
93. Scott Cantor - tOSU

94. Michael Gettes - UFlorida
95. Keith Hazelton - UW-Madison
96. Tom Zeller - Shib
97. Peter DiCamillo - Brown University

Call Agenda and Notes

1. Call logistics at top of this Google Doc
2. Special Agenda - Grouper Deployment
   We have several guests joining us today to focus on the Grouper deployment work. Now that we have the Docker containers and VMs available we need to better understand what we need to do (and not do) to facilitate real deployments. As with the Shibboleth work below, for Grouper:
   a. What is the migration strategy from a common campus deployment to the TIER distribution
   b.
   c. A full backup of grouper includes
      i. Database backup
         1. Standard database backup
      ii. Files in filesystem also need backup
      iii. Configuration management somewhere
   d. Availability & default modules
      i. Some schools run the web services components of Grouper in HA mode; database;
      ii. Default TIER design will include the following components at the following availability:
         1. Database (HA) ←- TIER supplied or Campus delivered
         2. Grouper web services (HA)
         3. Grouper user interface (HA)
         4. Grouper loader (HA)
         5. Grouper message bus interface to AMQP (likely using RabbitMQ - need to decide soon) (soon) (HA)
         6. Grouper PSPNG (HA) (ldap provisioning)
         7. Grouper PSP (classic) (yes, skip this module as per Grouper team)
   e. Grouper Web Services Authentication
      i. Apache basic authentication
      ii. LDAP authentication
      iii. Future: certificate
   f. Add ons (additional module support)
      i.
      ii. https://spaces.internet2.edu/display/Grouper/Provisioning+and+Integration Need a mechanism for sites to be able to add their own modules

g. Essential Integrations (outbound) - what we provide beyond d.ii will be
    i. RabbitMQ part of TIER as MariaDB? PLEASE consider this!
        1.
h. Integrations (inbound)
    i. loader
i. Customizations
    i. GDG folder structure in TIER release - is there now
    ii. Folder and group permissions structure
    iii. Should customizations live in TIER package OR Grouper installer?
        1. TIER could pick some defaults and then users could flag them off/on/etc.
    iv. Sources and search config wizard - rather than a blank sources.xml canvas
        1. Some selections for
            a. LDAP
            b. DB
    v. Check out Grouper Loader in the UI: https://spaces.internet2.edu/display/Grouper/Grouper+loader+on+UI
        1. For configuring new loader settings
    vi. Grouper subject API diagnostics in UI: https://spaces.internet2.edu/display/Grouper/Grouper+subject+API+diagnostics+in+UI
    vii. What else to include in Grouper container:
        1. Shib Auth pre configured for Grouper UI
        2. Service principal provisioning?
    viii. Split components into separate containers:
        1. UI node
        2. WS node
        3. OK with shibd running in with each container and Apache
    ix. COmanage to Grouper Provisioner (ScottK)
        1. Needs WS user (read/write)
        2. Needs its own stem, reduce blast radius such that can only manipulate its own stem
    x. COmanage uses supervisord - shibd and apache logs back to docker console
        1. Supervisord works well for them
    xi. Sticky sessions LB required for multiple Grouper UI
        1. COmanage using this Dockerized HAProxy for LB: https://github.com/vfarcic/docker-flow-proxy
    xii. How to hook into **operational infrastructure**?
        1. Eg. logs, security, etc.

July 10, 2017

The July 10, 2017 call is cancelled. We will move forward with our Grouper topic next week.

June 26, 2017

Attendees (please add yourself):

98. Jim Jokl - Virginia
99. Carey Black - tOhio State Univ.
100. Scott Cantor - tOSU
101. Chris Hubing - Internet2
102. Bill Kaufman - Internet2
103. Keith Hazelton - UW-Madison
104. Kevin Ruderman - Boston University
105. Paul Caskey - Internet2
106. Tom Zeller - Shib
107. Chris Phillips - CANARIE
108.

Call Agenda and Notes

3. Call logistics [at top of this Google Doc]
4. Agenda bash
    a. Insert any additional items here
    b. Is this considered Demo or Prod? (Kevin R BostonU)
5. Quick Topics, Action Items, and Updates (if any)
    a. July 3 timeslot - cancel call but send out drafts for review
    b. Grouper default configuration subgroup (Chris H)
        i. Group should form soon - 2017-06-05
            1. Will ping BillT again this week - ChrisH
            2. Was pinged (june 19) - waiting for response
            3. From BillT: “It would be great if the defaults followed the GDG recommendations. We could start with the folder/group layout, and include some scripts to create various folder/grouper/permission structures.”
            4. https://github.com/UniconLabs/grouper-demo-docker/blob/master/seed-data/tier-bootstrap.gsh
            5. We have completed initial checks and will start on the work next (June 26, 2017).
    c. Shibboleth as a managed (cloud) service offering
        i. Jim to bring to TAC
    d. COmanage Docker Distribution
        i. In (nearly) final testing now
        ii. Expecting official release with COmanage 3.1 (3Q 2017)

iii. See https://github.com/Internet2/comanage-registry-docker

e. GEANT Discussion
   i. They may wish to become more involved with our efforts
   ii. Initial conversation was on Friday June 16
   iii. Potential of earlier call slot or special call
   iv. Mario Reale - An overview of the work to evaluate a potential GÉANT Platform for supporting the provisioning of Campus IdPs.
   v. Jim to contact - schedule an alternative meeting on this topic***

f. Next TIER/InCommon newsletter
   i. The next InCommon/TIER newsletter (2nd week of July) will contain a "TIER Corner" where we'll include updates on the AMI releases and the initial testing of the Shibboleth container release. If we have enough feedback by the release date, we'll make an announcement of the stand-alone container version.

6. Shibboleth Future State Deployment Scenarios
   a. See June 12 notes below
   b. Small deployments: Shibboleth Appliance (a few flavors of VM)
   c. Larger sites
      i. Stateless
      ii. Docker Swarm
      iii. Container build support - private docker registry
   d. Next Steps
      i. Draft architectural and operational documentation
   e. Consent-based release
      i. Discussions to start soon

7. Grouper Present and Future State Scenarios
   a. Same sets of discussions as we have had for Shibboleth on the past few calls
   b. We are too light today, who should I invite to our July 10 cal:
      i. Chris Hyzer -
      ii. James Babb -
      iii. Bill Thompson -
      iv. Chicago - David Langenberg -
      v. Brown -
      vi. Michael Gettes -
      vii. Add others in the next few days -- Jim will start to work on the invitations later in the week.

8. Reminder Items
   a. At some point in the future we will replace Tomcat with Jetty in our Shibboleth IdP Docker image. This work will be done after we catch up with much other pending work and likely will not be started until after TechEx.

b. July 17 -- COmanage present and future state deployments
    i. Line up the right set of participants ahead of time - Jim to work with Scott Koranda
    ii.

June 19, 2017

Attendees (please add yourself):

109. Jim Jokl - Virginia
110. Scott Koranda - SCG
111. Carey Black - tOhio State Univ.
112. Scott Cantor - tOSU
113. Chris Hubing - Internet2
114. Keith Hazelton - UW-Madison
115. Paul Caskey - Internet2
116. Bill Kaufman - Internet2
117. David Bickel - Indiana University

Call Agenda and Notes

9. Call logistics at top of this Google Doc
10. Agenda bash
    a. Insert any additional items here
    b.
11. Quick Topics and Updates (if any)
    a. Grouper default configuration subgroup (Chris H)
        i. Group should form soon - 2017-06-05
            1. Will ping BillT again this week - ChrisH
            2. Was pinged - waiting for response
    b. Jetty replacement for Tomcat as a servlet engine for Shibboleth
        i. Delay discussion for now due to other workload; most likely future switch to Jetty
    c. Shibboleth as a managed (cloud) service offering
        i. Jim to bring to TAC and Component Architects
    d. COmanage Docker Distribution
        i. In (nearly) final testing now
        ii. Expecting official release with COmanage 3.1 (3Q 2017)
        iii. See https://github.com/Internet2/comanage-registry-docker
    e. GEANT Discussion
        i. They may wish to become more involved with our efforts
        ii. Initial conversation was on Friday June 16

iii. Potential of earlier call slot or special call
iv. Mario Reale - An overview of the work to evaluate a potential GÉANT Platform for supporting the provisioning of Campus IdPs.

f.

12. Shibboleth GUI high level design document
    a. Process
        i. Initial document from us
        ii. Review full design document Internet2/vendor
    b. Initial draft is here

13. Shibboleth Future State Discussion Continued
    a. See notes from June 12, Section 4
    b.
    c.

June 12, 2017

Attendees (please add yourself):

118. Jim Jokl - Virginia
119. Bill Kaufman - Internet2
120. Scott Koranda - SCG
121. Scott Cantor - tOSU
122. Kevin Ruderman - Boston University
123. Keith Hazelton - UW-Madison
124. Chris Phillips - CANARIE
125. Chris Hubing - Internet2

Call Agenda and Notes

1. Call logistics at top of this Google Doc
2. Agenda bash
    a. Insert any additional items here
    b. For next week, discussion of Shibboleth changes for GUI
3. Quick Topics and Updates (if any)
    a. Grouper default configuration subgroup (Chris H)
        i. Group should form soon - 2017-06-05
            1. Will ping BillT again this week - ChrisH
    b. Shibboleth stand-alone Docker container
        i. Paul is out today, but link to initial test code: https://github.internet2.edu/docker/shib-idp_noVM/blob/master/Dockerfile
    c. Jetty replacement for Tomcat as a servlet engine for Shibboleth
        i. Delay discussion for now due to other workload; most likely future switch to Jetty

    d. Shibboleth as a managed (cloud) service offering
        i. Jim to bring to TAC and Component Architects
    e. Shibboleth GUI high level design
        i. Jim to draft an initial draft description of work for next call
        ii. Concerns re: technical debt of GUI
4. Shibboleth Future State Deployment Scenarios (main topic)
    a. Our present state discussion is in the June 5 notes
    b. Smaller scale deployments will continue to be met via the appliance model
    c. Larger scale future state deployments - three main scenarios
        i. Deployment entirely local
        ii. Deployments entirely in AWS
        iii. Hybrid deployments (campus and AWS)
    d. Larger scale future state deployments - assumptions for what TIER supports now
        i. Database-free - i.e., no Consent
        ii. No back-channel
        iii. Direct Docker as opposed to VM-based
        iv.
    e. Solution Discussion - larger scale future state - strawman
        i. Deployment entirely local
            1. Docker SWARM
            2. Load balancing discussion (TIER (haproxy), campus, etc.)
        ii. Deployment entirely in AWS
            1. Docker SWARM
            2. Load balancing discussion
        iii. Deployment split between campus and AWS
            1. Docker SWARM
            2. Load balancing
                a. DNS-based e.g., F5 global load balancer
                b. AWS-based proxy (e.g., HA proxy)
    f. What additional documentation is needed
        i. Standard maintenance
        ii. Upgrades
        iii. Operations support
5. If time, initial Grouper discussion

June 5, 2017

Attendees (please add yourself):

126. Jim Jokl - Virginia
127. Paul Caskey - Internet2
128. Bill Kaufman - Internet2
129. Chris Hubing - Internet2

130. Carey Black - tOhio State Univ.
131. Scott Cantor - tOSU
132. Tom Zeller - shib
133. Kevin Ruderman - Boston University

Call Agenda and Notes

6. Call logistics [at top of this Google Doc]
7. Agenda bash
    a. (next week) Jetty to replace Tomcat as a servlet engine for Shibboleth?
    b. (next week) Shibboleth as a service
    c. Insert any additional items here
    d. ...
8. Quick Updates (if any)
    a. AMI Component Releases
        i. https://spaces.internet2.edu/pages/viewpage.action?pageId=110336944
    b. Grouper default configuration subgroup (Chris H)
        i. Group should form soon
    c. Projected TechEx Deliverables
    d. Shibboleth stand-alone Docker container
        i. Likely a couple of weeks away
9. Shibboleth Deployment Scenarios (main topic)
    a. Current State

School	A (Scott K)	B (Keith W)	C Shilen P)	D (Scott C)	E (Scott K)	F (Janemarie D)	G (Jim J)	H
Data Centers	3	2	5	2	1		2	
Load Balancer VMs	F5 - 8	2	10 (2x ldp per data center)	2 GLSB NetScalers	2	2	F5 - 2	
SSL Termination	Load Balancer			IdP	IdP		IdP	
OS	Centos		Centos 7 hosting Docker	CentOS 7			Centos	
Consent	No	Yes	Yes	No	No	No	No	

Consent Planned	No	Yes, ?	Yes, CAR	No		No	Yes	
Database	No	Mysql cluster	Yes	No	No	No	No	
MFA	Complex	Complex	Duo	Duo	No (not now)		Duo	
Additional jars	Yes			Yes			No	
Base AuthN	LDAP password	Kerberos against AD	Custom Complex; Kerberos, DUO, Social, etc.	Kerberos and LDAP	Active Directory	CAS (LDAP passwords)	Heavily customized PubCookie	
Attribute Source			LDAP	Database, LDAP, scripts, Grouper	LDAP	LDAP	LDAP	
Session Storage				Client			Client	
Is Artifact Supported				No				
Graceful Configuration Distribution				Yes			yes	
Environment Management (ansible vs. puppet, etc.)				rsync				
Configuration Management				GitLab				
Servlet				Jetty			Jetty	

Engine								
Added Security Constraints				AV, On-host vuln mgmt, CIS benchmarks				
Requirement	Strong Docker Documentation							
Virtualization		VMware		VMware			VMware	
CAS				No	Yes			
Deployment Goal	All local; Docker, not VM	Pure AWS; Elastic Beanstalk and RDS		No Docker may add AWS for backup	Docker or VM appliance	Open to Docker and AWS	Docker - mixed between local and AWS	

b. Future State
    i. What are likely outcomes?
    ii. What do we need to know from the community?
    iii. How many different builds do we need
c. Some Potential Future States
    i. Entry - local - (Appliance VM)
        1. TIER Appliance
        2. Base authentication config: Kerberos/LDAP/AD
        3. No consent
    ii. Entry - hosted
        1. TIER AWS Solution
        2. Two availability zones
        3. Base authentication config: Kerberos/LDAP/AD
        4. No consent

iii. Standard Local (Docker direct)
    1. Campus hardware vs. TIER supplied load balancing
    2. Docker hosting environment - is Centos 7 ok for the OS
        a. OS patching; operations support; etc.
    3.
iv. Standard Hybrid (Docker direct)
    1. Load balancing
    2.
v. Standard Hosted (AWS - Docker direct
    1. AWS load balancing across multiple availability zones
    2.

d.
e. ...

10. Grouper
    a. Deployment Characteristics (now)
        i. Site Xxx
        ii. Site Yyy
        iii.
11. COmanage
    a. Deployment Characteristics (now)
        i.

May 22, 2017

Attendees (please add yourself):

134. Jim Jokl - Virginia
135. Bill Kaufman - Internet2
136. Marlena Erdos -- Shib doc (possibly)
137. Scott Koranda - SCG
138. Chris Hubing - Internet2
139. Shilen Patel - Duke
140. Janemarie Duh - Lafayette
141. Keith Wessel - Illinois
142. Keith Hazelton - UW-Madison
143. Scott Cantor - tOSU

Call Agenda and Notes

1. Call logistics at top of this Google Doc
2. Agenda bash

a. New items; delete items; order changes
b.
c. ...

3. Quick Updates
    a. Grouper distribution
        i. Looks like Oracle changed something on the Java download so this needs to be looked into - ChrisH reviewing. May be an issue with the Country code as well for non-2-character country codes.
    b. Action item updates
        i. AMI distributions
        Release AMIs for Shibboleth, Grouper, and COmanage components with initial documentation - due: June 1, 2017 - [Chris Hubing and Paul Caskey]
        ii. Grouper default configuration
        We agreed to that the TIER Grouper default distribution should match the Grouper Deployment Guide as much as practical, with backout scripts, depending on the time investment needed. Time estimate work in progress [Chris Hubing]
            ChrisH working to connect with BillT to work on setting up a small group to address this.
    c. Projected TechEx deliverables (here)
    d. ...
4. Shibboleth Survey Preparation (Today's Main Topic)
    a. General
        i. What do existing deployments look like?
        ii. What could/would/should an equivalent deployment look like using TIER components?
        iii. What questions do we need to ask to better understand what we need to deliver?
    b. Shibboleth Deployment Notes
        i. Greenfield (old, not edited)
            1. From: nothing
            2. To: operational service → deploy TIER VM
        ii. Scott K - proxy for a large school
            1. From: 8 VM centos behind F5 load balancer with TLS terminated at the load balancer; in three data centers; no back-end database; not currently running consent and don't plan to run it in the future; complex MFA; extra jars have been injected for the past few years; inject intercept flows on user side; password via LDAP;
            2. To: probably won't want to run our VMs; likely want to keep local; want strong documentation for running Docker;
        iii. Keith W

1. From: vmware; two sites; load balanced; database for consent (mysql cluster); complex MFA configuration; password against AD Kerberos.
2. To: pure AWS solution; some of the potential gotcha's involve SSL on Amazon's load balancer. Elastic beanstalk and load balancer. Amazon RDS for a database.

iv. Shilen P
1. From (after we deploy CAR this summer): 9 VMs running Centos 7 in Docker containers in each data center; 5 data centers. Each data center includes - 2 IdPs, 2 CAR/ICMs, 2 CAR/ARPSI, 2 CAR/COPSU, 3 CAR/DBs. Custom login flow/code that includes username/password (MIT Kerberos), MFA with Duo, and what we call OneLink (which includes another MIT Kerberos realm plus external social providers (Google, Facebook, Yahoo, LinkedIn). Also separate Oracle database used to store MFA cookies. Attribute retrieval from 389 Directory Servers.

v. Scott C
1. From:
    a. mix of VM/physical behind GSLB NetScalers
    b. Jetty operating as the web server, LB switches TCP
    c. MFA via Kerberos/LDAP + Duo, fairly vanilla
    d. Attributes via mix of databases, LDAP, scripts, Grouper WS (future)
    e. Fully stateless / cookie and local storage-based, no database or shared cache
    f. Lot of error handling customizations, custom jar for odds and ends
    g. Config managed in git, pushed out via simple scripts to support updating nodes one by one to avoid downtime during any change
2. To: more automation and config management, not interested in Docker, may move nodes out to Amazon in part but probably only as backups

vi. Scott K - proxy for small school (~1500 students, not TIER investor, but InCommon Participant)
1. From: 2 VM nodes behind load balancer, single data center, TLS not terminated at load balancer. LDAP (AD) for authentication. Attributes from LDAP. No database. No consent. No MFA at this time. Fairly "vanilla" deployment except that they do leverage CAS (instead of SAML) for at least one SP.
2. To: continue to run local, open to Docker for deployment but also open to VM appliance (VMware)

vii. Janemarie

1. From: 2 VMs behind a proxy; AuthN via CAS (passwords in LDAP); attributes in same LDAP; MFA in CAS; no database or need for consent in the future;
2. To: hope to move to docker components at the right time; would likely be open to an AWS-based deployment.

viii. Large Production Site A (old, not edited)

1. From: Three VMs running Shibboleth behind a load balancer, scripting in place to deliver configuration changes, etc., etc.
2. To: three TIER VMs behind the same load balancer, TIER scripting (?) for coordination of updates, etc.

ix. Large Production Site B (old, not edited)

1. From: Three VMs running Shibboleth behind a load balancer, scripting in place to deliver configuration changes, etc., etc.
2. To: direct Docker running in in Amazon's container service, TIER scripting, etc.

5. General Survey Discussion Items on Operations (from last week)
   a. Campus operations survey
      i. Topics from Global Summit
         1. VM Management
            a. How to upgrade, maintain, etc.
         2. Database management
            a. How to do backups
         3. Getting TIER-built containers off of the VM to production (ref?)
         4. Load balance deployments
         5. Other VM formats
         6. When will direct docker containers be available?
         7. Direct docker deployment scenarios
      ii. Topics (brainstorming)
         1.
   b. Container vs. VM vs. build
      i. One of our goals was for TIER to do work once and every campus not needing to redo all of this work.
      ii. Remember that VM's are just another form of a container. Docker is just a lighter-weight container.
   c. …

Projected Packaging Deliverables for TechEx 2017

- The focus for TechEx is on enabling/supporting real deployments
- Independent Docker containers (standalone and via VM builds) and appropriate documentation

- Additional VM formats (AMI, etc.) and appropriate documentation
- Shibboleth Configuration via Metadata, associated GUI
- Potentially Grouper preconfigured to match Deployment Guide
- TIER Beacon - specific yes/no (default yes) question, where possible, on if enabled or not
- TIER Production Component Deployments
    - From/to migration cases: what do we recommend for the various existing scenarios
    - Survey to determine what is needed for TIER production use on campus
    - Capture of production deployments
- Any other items from #9 or #10 here
    - Are there any that should not be part of Packaging?
    - Are there any additional tasks that we can take on for TechEx
    - Potential for Packaging from the “here” document
        - Starting now in Packaging WG (for both components and their operating environments)
        - PaulC - looking at using COmanage in TIER/InCommon Shibboleth Training as an SP integration example
        - MarlenaE proposed the idea of a "Quick Start Install Guide" for IdP V3
            - Marlena: What doc is actually desired by TIER for Shib is currently up in the air.
        - PaulC - New InCommon updated training
            - just starting to gel
            - Will have an installer but not sure exactly what form that will take.

May 15, 2017

Attendees (please add yourself):

144. Jim Jokl - Virginia
145. Carey Black -tOhio State Univ.
146. Scott Cantor - tOSU
147. Scott Koranda - SCG
148. Bill Kaufman - Internet2
149. Chris Hubing - Internet2
150. Regrets: Chris Phillips
151. Paul Caskey - Internet2
152. Sara Jeanes - Internet2
153. Keith Hazelton - UW - Madison (late)

Call Agenda and Notes

6. Call logistics at top of this Google Doc

7. Agenda bash
    a. New items; delete items, order changes
    b. Insert your item(s) here …
    c. ...
8. Quick Updates
    a. Additional Virtual Machine images formats
        i. Documentation and testing
        ii. We will produce some documentation and make AMIs available publically
        iii. Target date with minimal docs (a paragraph) -- June 1, 2017
    b. Shibboleth GUI-based configuration - discussion, status, plan
        i. Adding a .jar file to our distro TIER process would be easy.
        ii. Maven interconnect different/difficult for TIER
        iii. This code pieces of this will be in 3.4. 3.4 will likely be available in January/February 2018
        iv. Email conversation Scott, Jim, Paul, etc., on support pre-3.4
    c. ...
9. Discussion Items (Operations)
    a. Campus operations survey
        i. Topics from Global Summit
            1. VM Management
                a. How to upgrade, maintain, etc.
            2. Database management
                a. How to do backups
            3. Getting TIER-built containers off of the VM to production (ref?)
            4. Load balance deployments
            5. Other VM formats
            6. When will direct docker containers be available?
            7. Direct docker deployment scenarios
        ii. Topics (brainstorming)
            1.
    b. Container vs. VM vs. build
        i. One of our goals was for TIER to do work once and every campus not needing to redo all of this work.
        ii. Remember that VM's are just another form of a container. Docker is just a lighter-weight container.
    c. Recruiting (who else do we need to help in this space)?
        i. Shilen Patel - Duke
        ii. James Babb - UW Madison
        iii. Jon Miner - UW-Madison
        iv. Rich Graves - Carleton
        v. Jim to ask Janemarie for ideas
        vi. Matthew Economou (NIH NIAID)
        vii. Keith Wessel - Illinois

viii. TIER technical contacts list?
ix. ...

d. Implementation Documentation
    i. Migration Scenarios (from campus:to TIER) (some possible examples and ideas for survey questions)
        1. Shibboleth
            a. Greenfield
                i. From: nothing
                ii. To: operational service → deploy TIER VM
            b. Large Production Site A
                i. From: Three VMs running Shibboleth behind a load balancer, scripting in place to deliver configuration changes, etc., etc.
                ii. To: three TIER VMs behind the same load balancer, TIER scripting (?) for coordination of updates, etc.
            c. Large Production Site B
                i. From: Three VMs running Shibboleth behind a load balancer, scripting in place to deliver configuration changes, etc., etc.
                ii. To: direct Docker running in in Amazon's container service, TIER scripting, etc.
            d.
        2. Grouper
            a. Greenfield
            b. Large Site
                i. From:
                ii. To: Tier delivered version to local branded/configured version
        3. COmanage
    ii. Deployment stories
        1. Capture deployment stories as TIER products are moved into production.
e. Grouper default configuration
    i. Based on Grouper Deployment Guide (GDG)?
    ii. Implementation via some grouper shell commands
        1. Ability to run/back-out as needed.
    iii. Yes, tentatively based on how much effort would be needed with for scripting
        1. WHO: ChrisHu will work on the time estimates

f. Are all three of our planned distribution mechanisms all still needed
    i. VMs

ii. Docker containers built/maintained on a TIER VM but exported elsewhere for operations
iii. Standalone Docker containers.

10. Discussion Items (other)
    a. Insert your items here

May 8, 2017

Attendees (please add yourself):

154. Chris Hubing - Internet2
155. Marlena Erdos - Consent Architecture + Shib Documentation
156. Scott Koranda - SCG
157. Janemarie Duh - Lafayette
158. James Babb - UW Madison
159. Carey Black -tOhio State Univ
160. Scott Cantor - tOSU
161. Chris Phillips - CANARIE
162. Paul Caskey - Internet2
163. Sara Jeanes - Internet2
164. Jon Miner - UW-Madison

Call Agenda and Notes

1. Call logistics at top of this Google Doc
2. Agenda bash
    a. New items; delete items, order changes
    b. Insert your item(s) here …
    c. ...
3. Quick Updates
    a. Completion of security work postponed for a few calls
    b. …
4. Current action Items (many from Global Summit)
    a. Integration of Shibboleth code to support configuration via GUI integration
        i. Back-ported code avail
        ii. Will be in 3.4, but 3.4 will contain additional checks/features
        iii. Still need a GUI
        iv. Quickstart exists (based on IdP-Installer: https://github.com/canariecaf/idp-installer-buildtools)
        v. Existing material on Service Provider ‘curriculum’

1. https://collaboration.canarie.ca/elgg/file/view/4779/bcnet2017-accelerating-service-delivery-with-federated-identity-management-workshop (½ day workshop pptx source available.)
2. 30 min presentation deck still awaiting location for hosting at the moment (May 8th)
3. Added after the call AARC work: https://aarc-project.eu/a-hitchhikers-guides-to-the-aai-galaxy/
    a. Chris P: I've been informally coordinating with AARC on content and finding the relevant set for our community (read: early adopters and institutions doing services as opposed to 'platforms'). There is a growing body of work that AARC has that may be referenceable 'as is' but read and determine if they have the right content for the right audience.

vi.

b. Standalone Docker containers
    i. Built in the a TIER VM
    ii. Available via
        docker hub (tier/shibboleth_idp) Last pushed: 14 days ago
        docker hub (tier/tier/shibboleth_sp) Last pushed: 6 months ago
    iii. When the TIER packaging team is ready, the COmanage project Docker material is available at https://github.com/Internet2/comanage-registry-docker
        1. There is an example using Docker stacks and secrets for mod-auth-openidc and MariaDB at https://github.com/Internet2/comanage-registry-docker/blob/master/recipes/production-mod-auth-openidc-mariadb/README.md
c. Additional Virtual Machine environment(s)
    i. What additional VM environment(s) should we support
        1. Responsibility of packaging group/currency?
    ii. Build and testing implications
        1. SSLLabs?
            a. Added after the call by Chris and related to security/testing implementation: Does SELinux get set to 'enforced'? Is it used at all? Why or why not? And if not relevant to packaging, strike the questions. Thx. C
        2. GEANT Greenhouse? https://www.geant.org/Innovation/SIG_TF/Pages/SIG-Greenhouse.aspx
        3.
d. Migration Scenarios
    i. How should/does a campus migrate a production

1. BTW, does ‘account claim’ in service providers get covered in this? This is a migration from ‘not using federated identity’ to ‘using federated identity’ -- a very common bootstrap conversation
    a. Answered: On the list.

April 10, 2017

Attendees (please add yourself):

165. Jim Jokl - Virginia
166. Chris Hubing - Internet2
167. Bill Kaufman - Internet2
168. Carey Black - tOhio State Univ.
169. Scott Cantor - tOSU
170. Paul Caskey - Internet2
171. Keith Hazelton - UWisc
172. Tom Zeller - Shib

Call Agenda and Notes

1. Call logistics at top of this Google Doc
2. Agenda bash
    e.
5. Quick updates
    a. Security topic action items
6. Version check in and testing
    a. Versions: Grouper (still on Tomcat 6) and COmanage (still final beta)
        i. Grouper testing - one more internal test (today) for 4b fixes.
        ii. COmanage - hold off a day or two until we get off of the beta version and the 4c fix.
        iii. Shibboleth - needs some documentation
    b. Grouper SP
        i. The Shibboleth SP in the Grouper build is non-functional. We’ll need to document this. As configured, shibd is not finding the proper libraries. Likely issue is pathing, fails on attempts to download metadata, etc. We will/should get feedback that we should not be running shibd as root. Patch ld library path for release: /opt/shibboleth/lib64 (Patch pushed out to github, needs to be tested)
    c. COmanage is likely to have the same Shibboleth SP pathing issue on its SP.
        i. (Patch pushed out to github, needs to be tested)
7. Continuation of Architects Call discussion
    a. Component lifecycle, updates, etc. (AI - grouper and COmanage discussion)

i. Shibboleth
    1. Shibboleth configuration tree backup/restore scripting to new VM needs operational documentation.
    2. We still need to deal with Tomcat and haproxy configs
ii. Grouper
    1. Questions as to where configuration is stored.
iii. COmanage
    1. Likely mostly documentation (most in database)

b. Messaging and needed documentation (AI - grouper and COmanage discussion)
    i. Shibboleth: document of pathway for larger existing schools to migrate existing implementations.
    ii. Grouper: questions about how much configuration is outside of the database -- context of new version upgrades and transition to TIER-provided components
    iii. COmanage: migration effort should be relatively direct -- context of new version upgrades and transition to TIER-provided components.

April 3, 2017

Attendees (please add yourself):

173. Jim Jokl - Virginia
174. Chris Hubing - Internet2
175. Scott Koranda - SCG
176. Scott Cantor - tOSU

Call Agenda and Notes

8. Call logistics at top of this Google Doc
9. Agenda bash
    a.
10. Quick updates
    a. Quick status update on components for Global Summit
        i. target 10th for next versions for this team.
    b. Shibboleth work update - configuration by metadata tags
        i. Scott has some test code ready
        ii. Can map a variety of types from metadata into properties
        iii. Lots of additional spring config for the property driven approach
        iv. Surprises - extended the code to handle some cases. There is currently no clean way to handle conditional enabling of SAML1. The other items all look ok.
        v. Remaining work likely over the next couple of weeks -- should be in good shape.
11. Security group coordination

a. See March 27 notes - we will collect today's notes in the March 27 space
b. Complete Shibboleth discussion
c. COmanage
d. Grouper
e. Common topics

12. Insert your item(s) here

March 27, 2017

Attendees (please add yourself):

177. Jim Jokl - Virginia
178. Scott Koranda - SCG
179. Paul Caskey - Internet2
180. Marlena Erdos --
181. Scott Cantor - tOSU
182. Chris Hubing - Internet2
183. Bill Kaufman - Internet2
184.

Call Agenda and Notes

Call logistics [at top of this Google Doc]

1. Agenda bash
2. Quick Updates
    a. Check in on anticipated Global Summit deliverable status
    b.
3. Security Group Coordination (*includes March 27 and **April 3** notes*)
    a. The security group has requested that we focus on:
    b. Procedures for how software versions are validated/tested, and how often they are updated.
        i. Components themselves
            1. Where does security start?
            2. How are the 3rd party dependencies protected?
                a. The Shibboleth project focuses on ensuring a secure build process as opposed to having a mechanism to verify the provenance of all of the sub-components.
                b. COmanage: https://spaces.internet2.edu/display/COmanage/Version+Dependencies
                    i. All packages other than PHP are pulled into the COmanage source repository directly.

            ii. Manual process to verify component updates as needed (checksums, etc.).
        c. Grouper: we have a secure build process and check the size of dependencies on startup
    3. Source code control,
        a. Shibboleth: in general, is not based on multiple sign-offs per submission/commit.; write access is limited to a small number of known and authorized developers
        b. COmanage: in general, is not based on multiple sign-offs per submission/commit; write access is limited to a small number of known and authorized developers
        c. Internet2 github
    4. Final product release tests
        a. Shibboleth: no formal process for final product release tests. Automated unit and integrated tests are part of the process.
            i. Resources don't presently exist for more extensive or formalized testing.
            ii. Strong unit testing is in place, integration testing is weak.
        b. COmanage: no formal process for final product release tests, managed by Ben. Testing is done by the team on multiple release candidates.
            i. Project struggles (resource limitations) for automated unit testing.
            ii. Some integration testing is in place.
        c. Grouper: procedure to do automated tests and manual tests before releases
ii. Packaging
    1. Process for new/updated versions of a release
        a. add / or edit new configuration files
        b. Use a combination of Jenkins and Packer to build VMs
            i. Scripting for this process is maintained in Github
            ii. A small number known (to us) are able to maintain the integration pipeline.
        c. Testing
            i. Trivial automated testing is in place to catch some build errors
            ii. Main testing is a manual process.
        d. Updates
            i. Working towards weekly updates (which catch OS updates, etc.).
    2. Security Documentation

        a. TIER components generate security keys. These are burned into containers in many cases. Adequate documentation is needed to ensure operators understand what to do.
        b. Shibboleth Sealer Key Discussion
            i. Synchronization of sealer key refresh
            ii. Mounting externally,
        c. COmanage:
            i. Two salt files are generated at installation and must be carried forward across instances - part of the cake framework. These need to be protected as if they are security keys. We need to work this into the documentation. Database and SMTP server credentials are also in the containers. LDAP, Grouper, GitHub, etc. authentication information is in the database.

c. Logging Procedures
    i. Questions
        1. Logging for audit of build processes?
        2. Default configuration for logging within the various components?
        3. How logs are mounted/aggregated?
    ii. COmanage
        1. Not a large amount of logging - written to filesystem and exported from container.
        2. Future version of COmanage will be configurable to write log data to stdout where it can be captured by normal Docker methods
        3. Standard apache / php logs also exist.
        4. Sensitive data is not written to logs
d. Testing Procedures
    i. Some possible tools
        1. (interest in trialing the use of tenable.io for vulnerability scanning? Chris Hubing is interested…)
        2. Might also want to look at the SWAMP from UW-Madison https://www.mir-swamp.org/ (SWAMP is the Software Assurance Marketplace) for static source code analysis (particularly for Java code)
    ii. Web Application Scanning
        1. Shibboleth: Difficult due to nature of application
        2. COmanage: should be workable but not being done now for lack of resources
        3. Grouper:
    iii. Generic vulnerability scanning against the VM and the Containers

1. We should be able to handle external facing pieces without “too” much trouble as part of the build process.

iv.

e. How to provide auditable proof that what was intended to be in the release was all that changed in the release (down to file level).
   i. Discussion on what is really needed.
   ii. We do verify signatures / checksums for files we download.
   iii. Do we want file change data.
   iv. Can Security send some example attack scenarios that would help us understand how to meet this request? What is meant by the word “intended” to be in the release?
f. Procedures for preparing for and updating end-of-life components
   i. Goal of weekly builds will address general security patches.
   ii. New component releases will be integrated within X weeks
   iii. Docker Container support will be for no longer that the main period of support for the main Components.
      1. We need to have some discussions on support models, provisioning of updated containers, etc.
      2. The general expectation will need to be that people stay current with Docker Container releases.
      3. Lifecycle Management
g. [AI - Jim] Send follow-up questions to Security Group (done - April 3)
   i. Procedures for how software versions are validated/tested, and how often they are updated
      1. In our security discussion context, where does the process start? Is the focus on the Docker/VM packaging of the components or are we including the components themselves?
      2. We assume the interest here focuses on functional testing. Is this correct?
   ii. Logging Procedures

      We are unclear what is being requested here and can think of three possibilities. Can you help us with what to focus on?
      1. Logging associated with the automated build process for auditing.
      2. The default configuration for logging within the various components.
      3. How we mount or otherwise make component log data available at the VM layer or externally.
   iii. How to provide auditable proof that what was intended to be in the release was all that changed in the release (down to file level).
      1. We struggle with the word “intended” in this context. The revision control system maintains changes to files between releases for items such as default configuration, scripting, etc. Is this what is needed?

2. We verify signatures on blobs that we download (e.g., Linux) but do not, for example, have a way to verify that a Linux distribution is free from compromise.
3. If you could send some of the potential attack scenarios, we'll be better able to understand what is needed here.

iv. Procedures for preparing for and updating end-of-life components

1. This area is still a work in progress. Support for any particular Docker container is likely to be shorter than the TIER Component that it runs. We need TIER distribution users to stay current on container builds.

v.

4. Other / New Topics
    a. Insert your item(s) here.

March 20, 2017

We will **not** have a call on Monday, March 20.

Likely topics for our Monday, March 27 include security coordination, new build testing and setting next phase service expectations, mailing list response coordination, and other items that

March 13, 2017

Attendees (please add yourself):

185. Jim Jokl - Virginia
186. Marlena Erdos -- Consent project
187. Keith Hazelton - UW-Madison
188. Bill Kaufman - Internet2
189. Scott Cantor - tOSU
190. James Babb - UW Madison
191. Ryan Larscheidt - UW Madison
192. Jon Miner - UW Madison
193. Carey Black - tOhio State Univ.
194. Paul Caskey - Internet2
195. Scott Koranda - SCG
196. Tom Zeller - shib
197. Chris Hubing - Internet2

Agenda and Notes

Call logistics at top of this Google Doc

5. Agenda bash

6. Remaining discussion from last week's call
7. Quick Updates
    a. Shibboleth via entity attribute configuration status
        i. Finalizing the details on reimbursing the Shibboleth Consortium for the work
        ii. Separate entity attribute for each controlled setting
        iii. Work will include a naming convention for the attributes
        iv. We will need some code change support to do everything we have identified. Most of what we need is available without code changes.
        v. Still looks like about a month of effort.
    b. Shibboleth initial configuration TAC discussion (see 2/27)
    c. ...
8. Deliverable goals for Global Summit (week of 4/24)
    a. COmanage
        i. Expected Global Summit Component Version: 2.0.0
            1. See https://spaces.internet2.edu/display/COmanage/COmanage+Product+Roadmap for release notes
        ii. Docker version now - 1.0.5
        iii. Remaining work brief summary:
            1. Need to upgrade to release 2.0.0
            2. Available for us to look now and this would be a good time for Packaging to start Docker work.
            3. New code but no fundamental infrastructure changes
        iv. COmanage Docker images (beta, not part of release)
            1. https://github.com/Internet2/comanage-registry-docker
    b. Grouper
        i. Expected Global Summit Component Version: 2.3.0
        ii. Remaining work brief summary
            1. Current Docker version: 2.3.0
            2. We need to focus on having Grouper patches in place for Global Summit
    c. Shibboleth IdP
        i. Expected Global Summit Component Version: 3.3.1
        ii. Build/Run VM Distribution
            1. Remaining work brief summary
                ~~a. Current version 3.2.1~~
                ~~b. 3.3.0 is almost ready to go~~
                c. ~~We expect to see~~ 3.3.1 was built
                d. Provision to mount a config instead of burning into containers.
        iii. Independent Container VM Distribution
            1. Remaining work brief summary

    a. Run environment - docker registry
    b. Automation
    c. Provisions to mount a config instead of burning into containers.

9. Other topics
    a. TIER Demo support from teams Overview Document in Progress
        i. COmanage against the Internet2 production version
        ii. Consent demo (from Duke)
        iii. Provision/de-provisioning demo (likely)
        iv. Instrumentation
        v. Right now we are not planning demos of the TIER Docker packages directly.
    b. Insert your item(s) here

March 6, 2017

Attendees (please add yourself):

198. Chris Hubing - Internet2
199. Mike Zawacki - Internet2
200. Sara Jeanes - Internet2
201. James Babb - UW Madison
202. Ryan Larscheidt - UW Madison
203. Keith Hazelton - UW - Madison
204. Carey Black - tOhio State Univ.
205. Scott Cantor - tOSU
206. Jim Van Fleet - Levvel
207. Tom Zeller - Shib
208. Bill Kaufman - Internet2
209. Paul Caskey - Internet2

Agenda and Notes

10. Agenda bash
11. UW Madison Shibboleth IdP containerization - Ryan L / James Babb
    a. In production
    b. Version 3.3
        i. Debian/Jessie
    c. Changed from CentOS to Debian when started pushing to DockerHub
    d. Self-host Oracle JDK, they've accepted the license

e. Building IDP from source, rather than .deb to stay current
f. Single process per container
g. Project in git
    i. 5 branches
    ii. Share attribute-resolver and templates, properties
    iii. 4 environment branches, dev/test, qa, prod, base (used to build off the others) ?
    iv. Multi-branch jenkins
    v. Auto-builds after commit
    vi. Looking to gitlab docker registry storage
        1. Currently distributing container images via private artifactory
        2. Private organization in Dockerhub requested by security team for security scanning in containers
            a. Images are private in Dockerhub
            b. N.b. many security warnings on Dockerhub security scan are spurious (e.g. wrong OS/arch)
    vii. Desiring of creating a Tomcat/Java/Shibboleth image from which their UW Madison images derive
    viii. Assessed Kubernetes as complex so chose Rancher for running container
        1. Rancher image quite large
    ix. Looking to externalize filter, so don't have to do full rebuild of container
    x. Resolver to stay inside image
        1. Targeting for storage in version control
    xi. Separate properties file for secrets - kept on host
        1. Sealer.jks, _____?
        2. 7 mountpoints for config files
            a. Md, idp logs, tomcat logs, and secret file(s)
        3. 4 hosts in production
            a. T / Th over two weeks
            b. Secrets synced on host filesystems
                i. Acceptable based on rate of changes
            c. Aware/evaluating Vault
    xii. Syslog goes off host
        1. Plans to implement ELK stack

12. UW Madison Grouper Docker implementation - James Babb

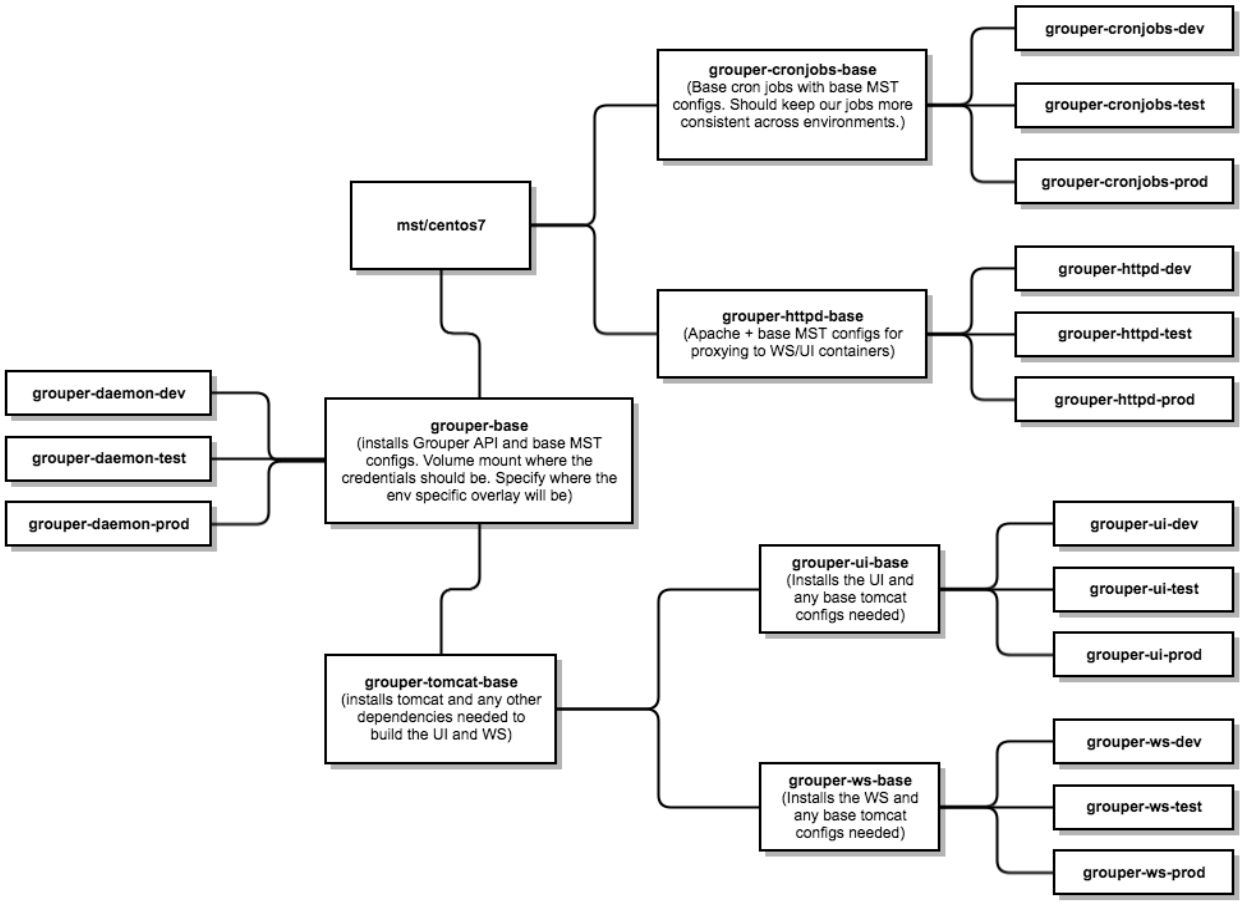

a. Install grouper API via Grouper installer
b. Oracle Java also comes in via local store
c. UI and WS are split out to ease memory management
d. Passwords are not stored in the image
e. Using Git, Jenkins build, Rancher deploys to host
    i. Config stored in git
f. 2 hosts running all containers
g. Rebuild base centos every month (which as a result, rebuildings everything downstream so grouper is also patched then but we can manually rebuild an image if we want to patch sooner)
h. Targeting summer for evaluating in production
i. Reverse proxy runs apache and shibd
    i. Supervisord to run httpd and shibd in same container
    ii. S6 seems to run better on Debian
j. Session store management to support rolling restarts is desired
k. A defined Grouper UI updates folder would be helpful
    i. Would love to be able to re-use tier as-is with only Graphical changes
l.

Meeting Notes and Agenda

February 27, 2017

Attendees (please add yourself):

210. Jim Jokl - Virginia
211. Chris Hubing - Internet2
212. Mike Zawacki - Internet2
213. Paul Caskey - Internet2
214. James Babb - UW Madison
215. Ryan Larscheidt - UW Madison
216. Scott Cantor - tOSU
217. Chris Phillips - CANARIE
218. Tom Zeller - Shib
219. Carey Black - tOhio State Univ.
220. Steve Zoppi - Internet2
221. Janemarie Duh - Lafayette

Agenda and Notes

13. Agenda bash
    a. Additions / changes
    b. #5 below
14. General updates
    a. Quick update on controlling Shibboleth via entity attributes
    b. Quick update re: the work on making native containers available
    c.
15. Shibboleth initial configuration

    During the period of time before we have better initial configuration

    a. Reminders on survey results and earlier discussions
        i. Survey results: https://spaces.internet2.edu/x/CwuVBQ
            1. Simple Spreadsheet - rows: 373 - 515
        ii. Earlier WG discussion and decisions
    b. Current state
        i. Confirmed - Yes -- Load and use InCommon metadata
        ii. Confirmed - Yes -- (and assume an eduPerson based directory) Include support for a default set of attribute definitions (LDAP - name, email; eduPerson -EPPN, Affiliation, primaryAffiliation, ?) We note that we may still need to do something special for AD. We will ask if LDAP or Active Directory and make the appropriate changes.

iii. Confirmed but with warnings -- Row 377 in spreadsheet - Yes - Release EPPN, name, email, affiliation, to all InCommon SPs? (TIER to provide documentation for sites to opt-out if needed)
iv. Confirmed but with warnings -- Yes -- Release EPPN, names, email, affiliation, to SPs with the Research and Scholarship R&S entity category (includes eduGAIN)? (TIER to provide documentation for sites to opt-out if needed along with discussion on why this is generally the “right thing to do” - we also need to ensure that InCommon helps with the education in this area (we believe we are helping InCommon’s agenda)). Warning about assumption of non-reassigned EPPNs.
v. Confirmed - Yes -- Respect a FERPA opt-out attribute to restrict attribute release for some users. (Add some type of configuration to report this issue to the end user).
vi. Confirmed - Yes -- Avoid spurious errors in the logs from external scanners via a properly configured robots.txt
vii. Confirmed - Yes -- Support Enhanced Client or Proxy (ECP) by default ? (potentially make available if configured with a compatible authentication source)
viii. Confirmed -- No -- in general and on Duo now (wait for more implementation maturity before deciding)-- Support multi-factor authentication by default? (what - would the be legal issues if we selected Duo or should we only do TIER-MFA (U2F, PKI, etc.)
ix. Confirmed but only in (future) cases where the configuration is mounted instead of burned into a container -- Yes (with exception of attribute-resolver) -- Automatically reload config files when they are changed (relying-party.xml, attribute-filter.xml, attribute-resolver.xml)?
x. Confirmed - No - Support CAS by default (document HA issues)?
xi. Confirmed - No - not relevant now - grant submitted for funding support and maintenance- Support OpenID Connect by default (when available)?
xii. Confirmed - Yes -NOT support SAML 1 by default?
xiii. Confirmed - Yes- NOT support SAML Attribute Queries?
xiv. Confirmed - No - Update itself automatically (document how a site can do this)?
xv. Confirmed - No - Update itself automatically - operating system security updates only (document how a site can do this)?
xvi. Confirmed - No - Prompt users to consent to attribute release?
xvii. Confirmed - Yes - Add a simple consent type configuration to enable FERPA opt-out override (either per-service or potentially globally) when no attributes would have been released for the user..

c.

16. Next call(s)
    a. March 6
        i. Grouper container architecture call

    b. March 13
        i. Security group coordination
17. Reminder: Please enroll for the TIER Working Group Members and Developers F2F Thursday April 27, 2017.
    a. https://docs.google.com/spreadsheets/d/1IQ9KSKpp8r8s0GVeqfDvFKK-H5rThvBgM4cN0Jv6_yE/edit#gid=0

February 20, 2017

Attendees (please add yourself):

222. Jim Jokl - Virginia
223. James Babb - UW Madison
224. Steve Carmody, Brown
225. Paul Caskey, Internet2
226. Scott Cantor, tOSU
227. Jon Miner - UW-Madison
228. Sara Jeanes, Internet2
229. Tom Zeller - Shib

Agenda and Notes

1. Agenda bash and general updates
    a. Shibboleth initial configuration - expect to proceed via CANARIE tooling
    b. AI [Paul - Jim] - next week's call will focus on the Shib IdP default config prior to integration of the CANARIE code.
    c. …
2. Shibboleth Operational Configuration
    a. Entity attribute tooling discussion
    b. https://wiki.shibboleth.net/confluence/x/VQC_AQ
    c. https://testbed.tier.internet2.edu/cgi-bin/shibcfg/shib-process.py?function=list
    d. [AI] Scott will send Jim an example script - no need to take control of the metadata config file this way.
    e. https://issues.shibboleth.net/jira/browse/OSJ-198
    f. ...
3. Next areas for Packaging
    a. Instrumentation
        i. Update on what is in progress now
            1. Every container sends a beacon once per day
            2. Data sent (4 items): Product (Shib, Grouper, COmanage), version, tier release id (likely date), maintainer of container; Syslog, so the source IP is also known.
            3. It is possible to opt-out; will be documented how; we'd prefer that people release the data
        ii. Next steps discussion

1. We are working to add the Jenkins build number
2. AI - all -- are there additional Packaging related items that we should capture in the near term?

b. Testing
   i. New builds, AMIs
   ii. Plan: work towards making AMIs available for WG testing
c. Security
   i. Coordination with security team
   ii. As part of testing VMs
   iii. AI - Jim to invite Security folks to a near future call
   iv. …
d. Direct Docker support without VMs
e. Insert your item(s) here

4. Other topics
5.

February 13, 2017

Today's call is cancelled but we will try to complete our main agenda item via email.
Please remember to complete the verification that the items listed on Scott's wiki page adequately cover the (vast) majority of your relying party configurations. The URL for the Wiki page is: https://wiki.shibboleth.net/confluence/x/VQC_AQ
Please add any items that you found either to the wiki page or here:

1. …
2. …
3. ...

February 6, 2017

Attendees (please add yourself):

230. Jim Jokl - Virginia
231. Scott Koranda - SCG
232. Bill Kaufman - Internet2
233. Chris Hubing - Internet2
234. Chris Phillips - CANARIE
235. Tom Zeller - Shib
236. Janemarie Duh - Lafayette
237. Gabor Eszes - Old Dominion
238. Carey Black - tOhio State Univ.
239. Keith Hazelton - UW - Madison

240. Scott Cantor - tOSU

Agenda and Notes

6. Agenda bash and general updates
    a. Chris and Scott have been chatting re: the CANARIE installer.
        i. Can some pieces of the Canarie IdP installer be pulled into the Shibboleth Consortium?
        ii. Shib consortium needs multi-platform capability whereas the installer is Linux focused now. This can be potentially addressed by preconfiguration prior to download.
        iii. Conversation is fluid and is in progress.
    b. Coordination between working groups before Global Summit
        i. There may be one more COmanage release prior to the meeting
        ii. Shibboleth - expected to be what we are packaging now
        iii. Grouper - expected to be what we are packaging now
        iv. Basic instrumentation will be added to each of the three components.
        v. Email wkaufman @ i2 if you want to be added to the Slack channels
7. Shibboleth Configuration Management
    a. The request from last week was to collect use cases for IdP configuration as per our discussions over the past few weeks re: control via entity attributes https://wiki.shibboleth.net/confluence/x/VQC_AQ
    b. Discussion - additional needs
        i. Enabling consent - need to highlight it in “interceptor-related” or a separate subtopic.
        ii. Should we focus on supporting campus groupings (e.g., enabling bundling like R&S at the campus level)? Or, should we focus on metadata markup managing single attributes?
            1. We don’t break future campus grouping possibilities by not working in this space now.
            2. Our focus now will be to support markup controlling at the element level.
    c. Task for everyone: Please review your campus Shibboleth configuration against https://wiki.shibboleth.net/confluence/x/VQC_AQ and help ensure that the vast majority of your relying party configuration needs could be met by control of the listed configuration elements via metadata markup. Help us identify any missing elements.
8. Reminders / Tasks
    a. Please test the TIER VMs
    b. VirtualBox command line usage issues
    c. Canarie IdP Installer licensing; Canarie knows that they can move to an Internet2 compatible license (Apache2). Current profile on things that the installer deploys is here: https://bit.ly/idpInstaller3-SoftwareProfile

d. Task for Jim: send out reminder of 2.c mid-week.

January 30, 2017

Attendees (please add yourself):

241. Jim Jokl - Virginia
242. Mike Zawacki - Internet2
243. Chris Hubing - Internet2
244. Bill Kaufman - Internet2
245. Sara Jeanes - Internet2
246. Scott Koranda - SCG
247. Scott Cantor - tOSU
248. Paul Caskey - Internet2
249. Chris Phillips - CANARIE
250. Carey Black - tOhio State Univ.
251. Tom Zeller - Shib

Agenda and Notes

9. Agenda bash and general updates
    a. Chris expects to be able to get to a compatible license for the installer.
    b. Brainstorming configuration "cases" in the IdP (compiling in https://wiki.shibboleth.net/confluence/x/VQC_AQ)
        i. Office365, → Needs some UUID/provisioning exercise against MSFT
        ii. Google Apps, →
        iii. Common MFA flows, → out of box, type A: just for this service list, Type B: for everything BUT this set etc..
        iv. Consent (with exceptions for services, or ONLY consent on this service)
        v. Essential attribute set
        vi. Arbitrary sets of attributes for types of services
        vii. Box specific set of attributes; etc.
        viii. Sharepoint specific set of attributes→ SHA1, unsigned assertions
        ix. Something that needs particular Formats for NameID (e.g. slack.com)
        x. Disabling encryption even when a key is present (i.e. the metadata's wrong or you just want to bypass it)<-- I would avoid that, as a policy that you MUST do it(validate). Doesn't that dilute the federation?? :) (Tell that to InCommon, they require keys even for SPs that don't support encryption.)
        xi. Multiple instances for an HA configuration (Do I have the ability to 'hot deploy/rolling deployment??' ← hard.)
        xii.
10. Shibboleth IdP automated configuration

a. Focused discussion re: the minimal set of shibboleth configuration files TIER's tooling will need to manage.
b. Coordination with Shibboleth project

11. Reminders / Tasks
    a. VirtualBox and command-line (non-console) mode support
    b. Approximate File Download Data
        i. Total downloads: 191
        ii. Unique IP address count by filename
            1. TIER-Grouper-R2-V1.ova: 4
            2. TIER-ShibIdP-R2-V5.ova: 8
            3. tier-grouper-r2-v2.ova: 1
            4. TIER-COmanage-R2-V2.ova: 28
            5. TIER-ShibIdP-R2-V6.ova: 37
            6. TIER-Grouper-R2-V2.ova: 31

January 9, 2017

Attendees (please add yourself):

252. Jim Jokl - Virginia
253. Mike Zawacki - Internet2
254. Sara Jeanes - Internet2
255. Chris Hubing - Internet2
256. Scott Koranda - SCG
257. Chris Phillips - CANARIE
258. Bill Kaufman - Internet2
259. Carey Black - Ohio State Univ.
260. Scott Cantor - tOSU
261. Tom Zeller - Shib IdP
262. Keith Hazelton - UW-Madison
263. Paul Caskey - Internet2

Agenda and Notes

1. Agenda bash and general updates
    a. Jim to look at component download logs before the next call
    b. Issues to be investigated
        i. Virtualbox in command line mode for use with the TIER components (import successful, would not start -- Shibboleth IdP VM)
    c. We may be able to find some collaborators -- See Chris' Dec 19 email
    d.
2. Shibboleth Ease-of-Use
    a. See notes (below) from December 12. 2016

b. What do we need in the Campus Metadata Management Tool requirements to support Shibboleth configuration entity metadata markup?
    i. There is a Jagger instance available for testing in the testbed - https://jagger.testbed.tier.internet2.edu/
        1. Send email to chubing@internet2.edu with your ePPn to get access
    ii. Generate metadata configuration for Shibboleth (and possibly other software: simplesamlphp?). The metadata config file format is generally pretty static between versions. This is new work.
    iii. Add in nameid support
    iv. Ability to import/process existing metadata files, now mandatory.
    v. Ability to add per-entity data to records.
    vi. Should/will per-entity metadata influence what we are doing?
        1. We would need to develop some form of proxy to do this work.
    vii. Should the tool work by groups with (a) a small number of entity attributes controlling Shibboleth behavior or (b) is there an entity attribute for each Shib function (e.g., one entity attribute per released attribute)? Or (c) both?
c. What do we need done with Shibboleth itself to support this mechanism
    i. Could Shib consortium work on configs based on entity tags? - yes, likely, if consensus on meaning of tag.
    ii.
3. Slack channel #tier-packaging now has notifications from jenkins when new builds are completed
4. Insert your item(s) here

December 19, 2016

Attendees (please add yourself):

264. Jim Jokl - Virginia
265. Chris Hubing - Internet2
266. Scott Cantor - tOSU
267. Paul Caskey - Internet2
268. Sara Jeanes - Internet2
269. Keith Hazelton - Wisconsin
270. Bill Kaufman - Internet2

If you have time please do some additional testing on the containers on the TIER Testbed.

Shib IdP ease of use: see 2 below
Direction is to use the CANARIE installer as a good clean start

December 12, 2016

The Monday, December 12 meeting is cancelled. Please see the email list for requested work.

Agenda and Notes

1. Reminder: test the Docker component releases
2. Shibboleth Configuration/Management discussion
   Now that we have the container releases on track, we have been back to the discussions on the task of easing the initial configuration and operational workload of running a Shibboleth IdP on the past few calls. Possibilities have ranged from developing tools to manage Shibboleth configuration files, bootstrap configuration tools, and other similar ideas. Over the past couple of calls and some mailing list discussion, we appear to be converging towards:
   a. Leveraging and working with CANARIE on their existing Shibboleth Installer for performing the initial configuration work prior to the container build step.
   b. Linking the Shibboleth ease-of-management work with our earlier campus metadata management tool effort to provide a path for automating many Shibboleth IdP functions. The metadata management tool would be used to mark up metadata to suit campus needs for many common functions, e.g., attribute release. This path would need some more detailed investigation as it is not being used in production anywhere. The primary task that we wouldn't automate via this path is adding a new attribute. Need to be able to preserve changes in the upgrade path.
3. TIER Release Components: Initiating builds - anything that still needs to be automated
   a. [AI] JimJ - update notes for each component to change CentOS password immediately upon installation ESPECIALLY if deployed on a public network
   b. We should script this to force this going forward
4. Reminder: No calls on Monday December 26 and Monday January 2. Happy Holidays
5.

December 5, 2016

Attendees (please add yourself):

271. Jim Jokl - Virginia
272. Chris Phillips - CANARIE
273. Paul Caskey - Internet2
274. Keith Hazelton - UW-Madison
275. Bill Kaufman - Internet2
276. Scott Cantor - tOSU
277. Chris Hubing - Internet2

278. Mike Zawacki - Internet2
279. Tom Zeller - Shib
280. Chris Hyzer - Penn

Agenda and Notes

1. Action Items
    a. Use cases for Shib ie. Top 4-5 common shib IDP integrations for next time (e.g. Office365, GSuite) See Section 2.c from our November 28 notes. [AI]
2. Today's call will focus on the CANARIE Shibboleth IdP installer and its use in TIER.
    a. Summary of today's call for discussion next time. A possible course of action
        - Metadata management tool scope additions for marking entity attributes
            1. Potentially very little to no work on the shibboleth config itself
            2. Tag entities with attribute release, mfa, other characteristics
            3. Need ability to import existing campus metadata
            4. IdP operators work in a very different, but likely easier way.
        - CANARIE Installer mods one-time configuration of the TIER Docker config tree.
    b. https://canarie.zoom.us/j/331462513 (screen share for this -- muted laptop and landline audio please)
        - IdP installer - then hand over the keys with pre-built environment
        - There is an interview process to understand what the user wants to do
        - Pre-flight check is done before running the installer to make sure everything is in place
        - Can paste an existing pre-populated config file and import it so fields are then visible in GUI. Fields that are mandatory are color-coded and also auto-populate available fields.
        - Has FedSSO Features area to enhance FedSSO operation.
        - Uses jetty instead of tomcat
        - Requires minimum attributes. Set for eduroam and FedSSO
        - Currently supports IdP v3.2.1
        - ScottC: seems like some of this is what should be provided as default with Shib
        - Jim: possibility of moving forward with modifying the CANARIE installer/config tool to support TIER. Scott: should be doable if scoped properly. See 2.a. Above for Jim's details.

November 28, 2016

Attendees (please add yourself):

281. Jim Jokl - Virginia
282. Bill Kaufman - Internet2 (phone only today)
283. Mike Zawacki - Internet2
284. Chris Hubing - Internet2
285. Chris Phillips - CANARIE
286. Paul Caskey - Internet2
287. Janemarie Duh - Lafayette
288. Keith Hazelton - UW-Madison
289. Tom Zeller - Shib
290. Scott Cantor - tOSU

Agenda and Notes

1. Component Testing
    a. Any new information on Component Testing? With the holidays, there may not be a lot of change.
    b. https://testbed.tier.internet2.edu/secure/download_vm/
    c. Releases of the Shibboleth IdP and COmanage are on the site
    d. Grouper is getting closer

2. Back to the discussion on Shibboleth IdP Configuration Management
    a. Files that are candidates to potentially, automatically generated:
        i. Idp.properties, relying-party.xml, saml-nameid.xml, attribute-resolver.xml, attribute-filter.xml?
        ii. Chris Phillips - CANARIE - uses a URL for his attribute-resolver
    b. Ability to drop metadata files into a directory and have it be automatically ingested is a possibility
    c. Top 4-5 common shib IDP integrations for next time (e.g. Office365, GSuite) [AI]
        i.

3. Meeting outcome summary

November 21, 2016

No call this week - please review action items. Have a great Thanksgiving Holiday week.

November 14, 2016

Attendees (please add yourself):

291. Jim Jokl - Virginia

292. Chris Phillips
293. Sara Jeanes - Internet2
294. Mike Zawacki - Internet2
295. Scott Koranda - SCG
296. Chris Hubing - Internet2
297. Bill Kaufman - Internet2
298. Tom Zeller - Shib
299. Matthew X. Economou - NIH/NIAID (Contractor)
300. Scott Cantor - tOSU

Agenda and Notes

1. Component Testing
    a. Update from callers the status of their packaged Shibboleth IdP testing https://testbed.tier.internet2.edu/secure/download_vm/
        i. Scott noted that this release is now behind the officially supported version upstream (3.2.1 vs. 3.3.0) and since TIER doesn't have a mechanism for incorporating code patches to address security issues, the download site should note this discrepancy for the time being.
    b. Quick status update on the initial Grouper build work
    c. Quick status update on the initial COmanage build work
        i. Directory layout not complete, jimj says he has a fix
        ii. SMTP needs come up as COmanage needs to use email
            1. assumes username/password needed for relay host - might not always be the case, won't send mail if these values are absent (therefore invite functionality doesn't work)
2. Shibboleth Campus Metadata Tool
    a. Request for contributions status
    b. Other possibilities: USC, Stanford, Duke, CMU?, Jagger
3. Shibboleth configuration management
    a. Using Salt as a mechanism for managing the Shibboleth IdP configuration tree
        i. https://saltstack.com/community/
        ii. https://testbed.tier.internet2.edu/ShibbolethSalt.pdf
    b. Refreshed on the CANARIE IdP Installer
        i. http://bit.ly/idpinstaller

4. 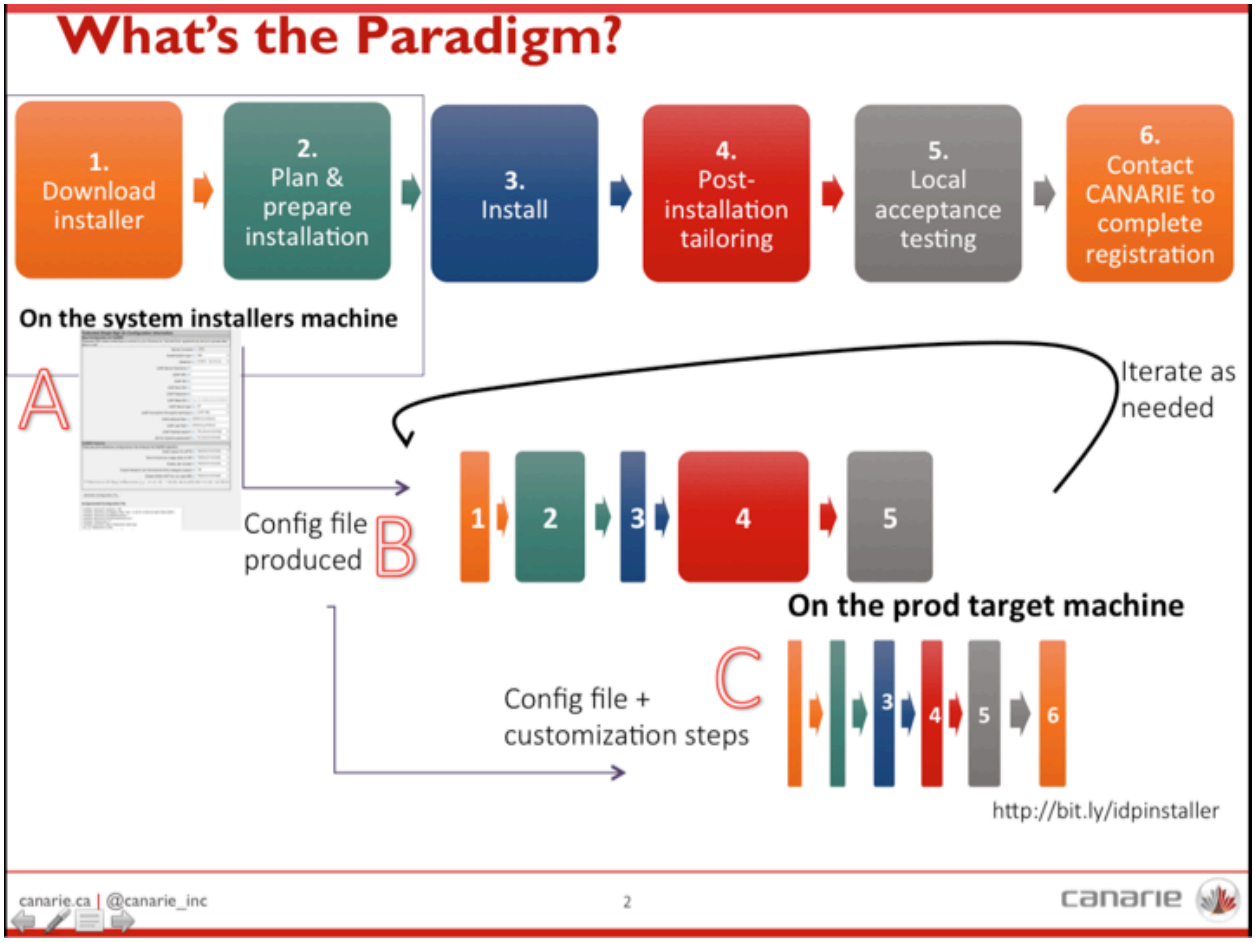
What's the Paradigm?
1. Download installer
2. Plan & prepare installation
3. Install
4. Post-installation tailoring
5. Local acceptance testing
6. Contact CANARIE to complete registration
On the system installers machine
A
Config file produced
B
Iterate as needed
1
2
3
4
5
On the prod target machine
Config file + customization steps
C
3
4
5
6
http://bit.ly/idpinstaller
canarie.ca | @canarie_inc
2
canarie

   i.

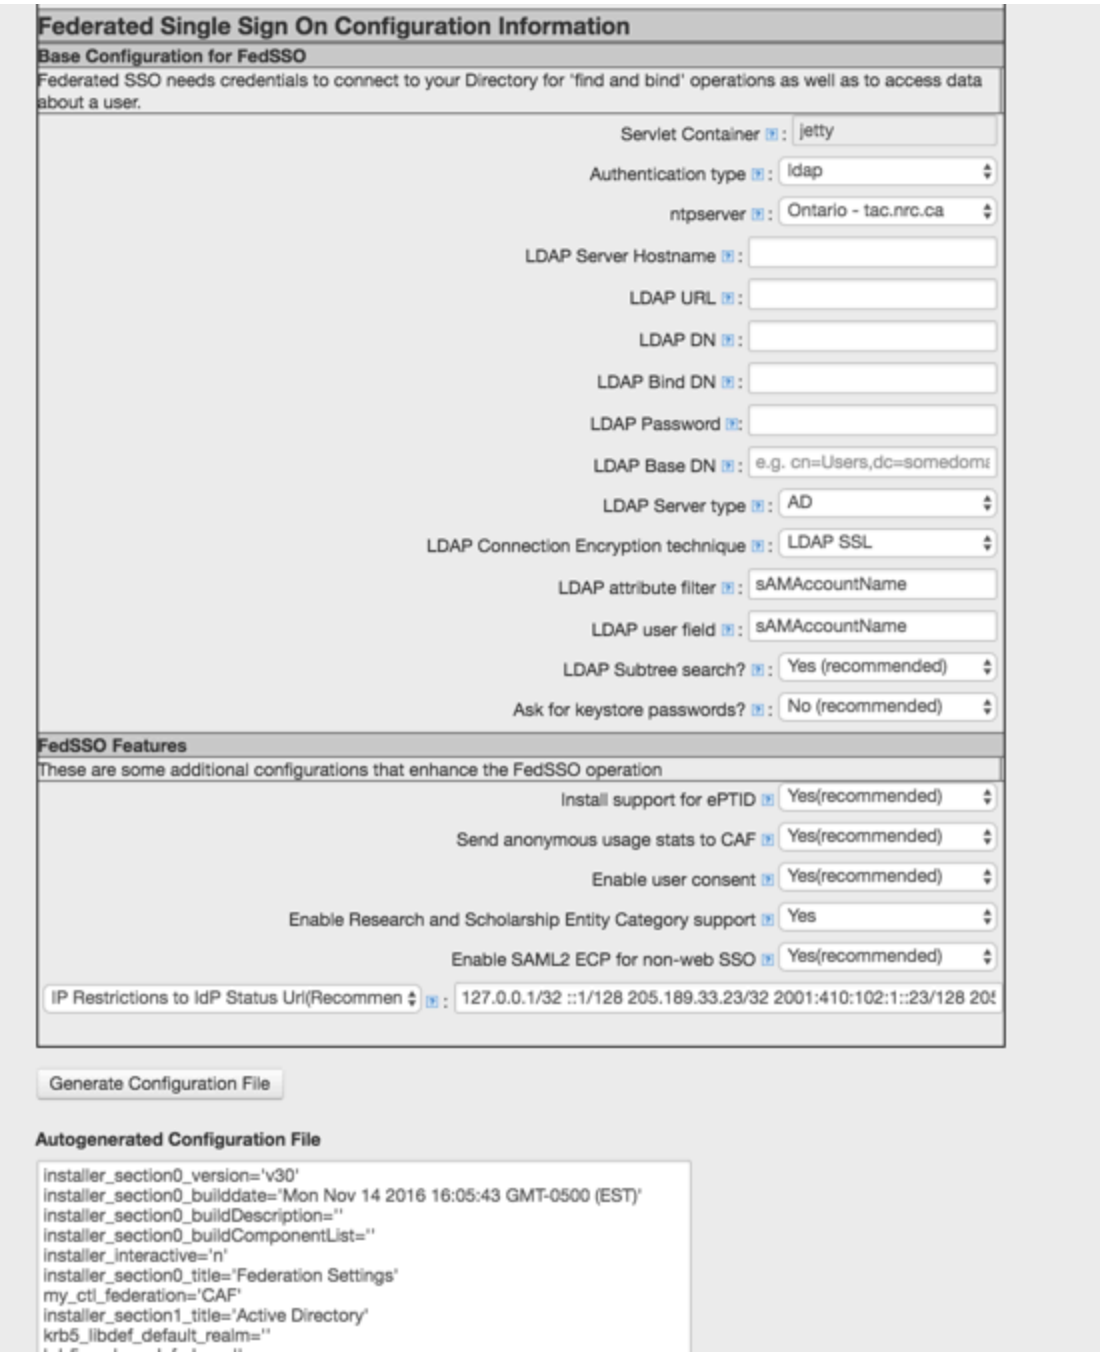

ii.
iii. See also our notes from February 8, 2016
iv. Test drive everything: https://github.com/canariecaf/idp-installer-buildtools
    1. Sample config for idp-installer: https://github.com/canariecaf/idp-installer-buildtools/blob/master/idp/config.template

b. Next steps
    i. …
    ii. requirements

5. Insert your item(s) here

a. …
b. ...

November 7, 2016

Jim Jokl is travelling today and will be unable to make the call. We will meet next week at our regular time.

October 31, 2016

Attendees (please add yourself):

- Gabor Eszes (Old Dominion)
- Jim Jokl - Virginia
- Mike Zawacki - Internet2
- Scott Koranda - SCG
- Chris Hubing - Internet2
- Paul Caskey - Internet2
- Keith Hazelton - UW-Madison
- Chris Hyzer - Penn

Agenda and Notes

We will try to end the call early - some people will need to leave to make it home for Halloween before dark.

1. Component Testing -- Status of testing for
    a. Shibboleth IdP
        i. Has anyone tested the mid-October release - a little
        ii. A couple more people volunteered to do some testing.
    b. COmanage
    c. Grouper
2. Remaining tasks for core packaging for these components
3. Expanded requirements for Campus Metadata Management tool
    a. Process
    b. URL with updated content: https://spaces.internet2.edu/display/TPWG/Campus+Shibboleth+Metadata+Management+Tool
    c. https://spaces.internet2.edu/display/InCFederation/Contacts+in+Metadata
    d. https://spaces.internet2.edu/display/InCFederation/Metadata+Administration
4. Insert your items here
    a. …
    b. ...
5. Two weeks from now: start on Shibboleth IdP configuration management work

Meeting Notes

October 24, 2016

Attendees (please add yourself):

- Mike Zawacki - Internet2
- Bill Kaufman - Internet2
- Scott Koranda - SCG
- Chris Hubing - Internet2
- Sara Jeanes - Internet2
- Tom Zeller - Shib
- Scott Cantor - tOSU
- Keith Hazelton - UW-Madison
- Chris Hyzer - Penn
- Paul Caskey - Internet2

Agenda

6. Component Testing
    a. Jim: We now have what I suspect will be pretty close to our production version of the Docker/VM implementation of the Shibboleth IdP. See https://testbed.tier.internet2.edu/secure/download_vm/ for download and for a link to the release notes. I need to make some significant enhancements to the documentation, but what exists now should be good enough for this group. Note also the setup scripting is relatively clean if you select the "connect to the TIER Testbed option" and pretty rough otherwise. Some additional scripting to ease that process a bit will be coming in the near future.
    b. Chris Hyzer has testing in progress. [Waiting for Levvel to respond] We are looking for more feedback on the Grouper packaging, e.g., problems found, what else is needed for a production-ready version? https://docs.google.com/document/d/1GUUyZIH5TWW2SkzDbFtApJrGrJlRlJKRpSDw_1qoDr0/edit
    c. Scott Koranda has testing in progress. We are looking for more feedback on the COmanage packaging, e.g., problems found, what else is needed for a production-ready version? - Stuck with VB OVA and needs help from Levvel. Nothing seems to be listening on 80/443. AMI starts up but Shib is not Federated and so cannot login. Scott put out a strawman in email to fix this design issue. Ability to inject configuration files before docker is run.

https://docs.google.com/document/d/1SCL1MHJw2H5KpYFHcnFFRdN72U96fy6-PEyqWckEcgg/edit

Testing email lists:

Tier-pack-coman@internet2.edu
Tier-pack-grouper@internet2.edu

The group agreed that Scott K should move forward in working with Levvel to get a Federated Shib integrated in for COmanage to authenticate.

Next in the queue for the WG will be to pick up our discussions on Shibboleth configuration management.

October 17, 2016

We will not hold a packaging call today. The Shibboleth IdP work is targeted for completion this week and COmanage testing will be ready for discussion on our call next week. Things are moving forward nicely.

Those of you with some Grouper experience, please consider using some of your new free time to give that build a try.

October 10, 2016

Note: we will not hold an in-person call this week.

Agenda

7. Component Testing
    a. We are looking for more feedback on the Grouper packaging, e.g., problems found, what else is needed for a production-ready version? https://docs.google.com/document/d/1GUUyZIH5TWW2SkzDbFtApJrGrJlRlJKRpSDw_1qoDr0/edit
    b. We are looking for more feedback on the COmanage packaging, e.g., problems found, what else is needed for a production-ready version? https://docs.google.com/document/d/1SCL1MHJw2H5KpYFHcnFFRdN72U96fy6-PEyqWckEcgg/edit
    c. Shibboleth-IdP - we expect to have a final version ready for testing in approximately two weeks. This next version will have the operational pieces in place to swap production containers behind the load balancer.
8. Expect to see more detailed documentation on requirements for the campus metadata management tool later this week.
9. Next in the queue for the WG will be to pick up our discussions on Shibboleth configuration management.

October 3, 2016

Attendees (please add yourself):

1. Jim Jokl (Virginia)
2. Bill Kaufman (Internet2)
3. Scott Cantor (tOSU)
4. Keith Hazelton (UW-Madison)

5. Tom Zeller (Shib IdP)
6. Janemarie Duh (Lafayette)
7. Mike Zawacki (Internet2)

Agenda

1. Quick update on core component packaging status (Levvel Status Report - Sept 30)
2. The Monday, October 3 call will focus on the campus metadata management tool. The original functionality that we specified at the URL below. Attached to the Confluence page below is the design document for implementing the requested functionality (this is also what was emailed out earlier today)
   https://spaces.internet2.edu/display/TPWG/Campus+Shibboleth+Metadata+Management+Tool
   Our goal for today's call is add any missing detail to the specifications (e.g., APIs) and review the attached design document.
3. Insert your items here

Minutes

[AI] Need to develop a few sentence requirement regarding attribute-based authorization for authenticating with the tool to make changes. Scott - should not have authentication baked in. They could deliver a default solution as long as it can easily be turned off.

[AI] Need to identify how the tool will integrate into the TIER package/environment - this is probably a TIER-responsibility

Provide Swagger documented API information to vendor

September 26, 2016

No meeting held due to overlap with TechEx 2016

September 19, 2016

Everyone,

We will not hold a TIER packaging call today. We had hoped to have a semi-final Shibboleth IdP release to discuss today but we are still waiting on a couple of pieces.

We will email as soon as the Shibboleth release is ready for testing. If you have some time now, please look at the Grouper distribution.

Thanks, Jim

September 12, 2016 4:00 ET

Since no one chimed in as being ready to discuss the Shibboleth or Grouper builds, we will not hold a call today at 4:00.

Please try to use your free hour to start your work on testing one or both of the VMs.

We will try to schedule some time later in the week to collect feedback on the builds.

Jim

No meeting on Monday, Sept. 5

We will not be meeting due to the Labor Day weekend

August 29, 2016 4:00 eastern (Regular Call timeslot)

Attendees (please add yourself):

1. Jim Jokl (Virginia)
2. Mike Zawacki (Internet2)
3. John Gasper (Unicon)
4. Scott Cantor (tOSU)
5. Scott Koranda (SCG)
6. Tom Zeller (Shib IdP)
7. Bill Kaufman (Internet2)
8. Sara Jeanes (Internet2)

Agenda and Minutes

Today's call is a Shibboleth IdP subgroup call.

1. SSL Termination & Load Balancing
    a. Trade-offs on SSL termination location
        i. Any Feature Issues
    b. Quiescing servers
    c. Certificate AuthN
    d. Load Balancing Software Options
        i. …
        ii. ...
2. IdP Secrets (keys and passwords)
    a. Update for group
3. Other IdP Issues
    a. ..
    b. ..

REMINDER: If you have not already done so, please sign up for the TIER Developers and WG Members F2F
Thursday, Sept 29, noon - 3pm
https://docs.google.com/spreadsheets/d/1PxaqjtlMRGe3AHL1hTh6TTFHTiB5-eo_1FxpVZyMgil/edit#gid=0

August 22, 2016 4:00 eastern (Regular Call timeslot)

Attendees (please add yourself):

1. Jim Jokl (Virginia)
2. Mike Zawacki (Internet2)
3. Scott Koranda (SCG)
4. Gabor Eszes (Old Dominion)
5. Bill Kaufman (Internet2)
6. Sara Jeanes (Internet2)
7. Tom Zeller (Shib IdP)
8. Keith Hazelton (UW-Madison)
9. Scott Cantor (tOSU)
10. Chris Hyzer (Penn)

Agenda and Minutes

1. Shibboleth IDP initial testing discussion

    a. Notes and download link
    b. Has anyone completed a full test yet?
        i. no
    c. Notes for levvel
        i. No ipv4 bindings for tomcat
        ii. Configuration errors
        iii. https://wiki.shibboleth.net/confluence/display/IDP30/SecurityAndNetworking
        iv. Scott just updated that page with some details on lifecycle of the various keys
2. Action Items (see below)
3. Additional topics
    a. Insert your item(s) here
4. Fri. discussion with Levvel about how Shib deals with secrets, etc.
    a. Scott C and Tom Z need to chat with Levvel to help them get it

August 15, 2016 4:00 eastern (Regular Call timeslot)

Attendees (please add yourself):

1. Jim Jokl (Virginia)
2. Scott Koranda (SCG)
3. Sara Jeanes (Internet2)
4. Paul Caskey (Internet2)
5. Keith Hazelton (UW-Madison)
6. Bill Kaufman (Internet2)
7. Scott Cantor (tOSU)
8. Tom Zeller (Shib IdP)
9. Chris Hyzer (Grouper)
10. Niva Agmon (Temple U)

Agenda and Minutes

1. Action items from the August 8 call
    a. Scott: IdP on Tomcat and Logging
    b. Scott: Properties file approach to assist with protection of secrets
        i. Done on 8/8 - jim to get to levvel.io
    c. Tom/Jim: Two base configuration trees for IdP
        i. Discussed several possibilities for interim config bootstrap scripts.
    d. TIER: Oracle Java distribution

i. End user will need to agree to T&Cs and then download the jre
ii. We will not be able to distribute the Oracle Java directly as part of the package.
iii. The "StackOverflow" answer (wget/curl examples): http://stackoverflow.com/questions/10268583/downloading-java-jdk-on-linux-via-wget-is-shown-license-page-instead

e. TIER: Docker repo services
i. Still working on this - we know the existing Internet2 contract will not cover this work.
ii. Jim - can you send Steve and I something that describes the overall intent/rationale around this and how it was arrived at? - thx Bill -- will do

2. Status
a. Shibboleth - expecting a first version to work with this week.
b. Grouper startup call https://spaces.internet2.edu/display/Grouper/Grouper+TIER+Packaging
c. COmanage startup call (all good -- Scott K)

3. Other Topics
a. Insert your item(s) here
i. Test Plan for Grouper - Chris H to draft a short plan

https://spaces.internet2.edu/display/Grouper/Grouper+TIER+Packaging

Note, the UI has the SP problem that COmanage has. Needs creds in WS to make a sample cal

ii. Test Plan for COmanage - Scott K to summarize thoughts in email
1. COmanage requires a source of authentication to do anything beyond "smoke testing". The source is a Shibboleth SP. The SP requires an IdP to do anything useful. Do we assume Levvel.io will go through a federation exercise? Or do we only want them to do "smoke testing"?

August 8, 2016 4:00 eastern (Regular Call timeslot)

Attendees (please add yourself):

1. Jim Jokl (Virginia)
2. Scott Koranda (SCG)
3. Bill Kaufman (Internet2)
4. Sara Jeanes (Internet2)
5. Scott Cantor (tOSU)
6. Emily Eisbruch, Internet2
7. Janemarie Duh (Lafayette College)
8. Chris Phillips / CANARIE
9. Chris Hyzer Penn
10. Tom Zeller (Shib IdP)
11. Niva Agmon (Temple U)

Agenda

1. Follow-up on Action Items
    a. July 25
        i. Scott: IdP on Tomcat and logging
            1. Some feedback was received - Scott will try to summarize in the near future.
            2. http://marc.info/?t=146947963600004&r=1&w=2
        ii. TIER Staff: Possibility for Docker repo services for schools
            1. Discussion is in progress
        iii. Scott: properties file approach to assist with protection of private keys
            1. Will pull this together by the next call
    b. July 18
        i. Shibboleth configuration automation --- CANARIE discussion -- On Hold
    c. July 11
        i. JimJ/TomZ: Shibboleth configuration trees for (a) tied into the TIER testbed for automated test and (b) tied into InCommon with the default settings -- see June 30 call notes
            1. J
2. Potential interaction with 1.a.iii and vaultproject.io discussion on the list last week
    a. Jim will follow-up and report back next week.
3. Brief Discussion on DevOps pipeline document
    a. https://spaces.internet2.edu/display/TPWG/Core+Packaging+Build+Documents
    b. Expect updated version to review next week.

4. Other Items
    a. Expected Grouper/COmanage deliverables for TechX? (for the purposes of resource planning)
        i. Contract with Level would call for having those ready by TechX, but that is appearing unlikely as work focuses on Shibboleth IdP. Should know more in a week or two. Grouper phase 2 meeting has happened?
        ii. Scott Koranda requests a design meeting with Levvel.io.
        iii.

July 25, 2016 4:00 eastern (Regular Call timeslot)

Attendees (please add yourself):

1. Jim Jokl (Virginia)
2. Gabor Eszes (Old Dominion)
3. Janemarie Duh (Lafayette)
4. Bill Kaufman - (Internet2)
5. Mike Zawacki - (Internet2)
6. John Gasper (Unicon)
7. Keith Hazelton (UW-Madison)
8. Tom Zeller (Shib IdP)
9. Scott Cantor (tOSU)
10. Scott Koranda (SCG)
11. Paul Caskey (Internet2)
12. Drew Zebrowski (Temple)

Agenda and Notes

Agenda

a. Quick Update on Campus Metadata Management tool
    a. https://spaces.internet2.edu/display/TPWG/Campus+Shibboleth+Metadata+Management+Tool
    b. Plus notes from June 6 below
b. Shibboleth Design Document
    a. https://docs.google.com/document/d/1sCrzQiriuVHS0g8wruHlgmSf9wDWhoMB9x5gRUZQWcc/edit?usp=sharing
    b. Notes will be made in google doc
    c. **[AI] Scott C to put out call to find out who is running idP over Tomcat and how tomcat is configured, especially with respect to logging.**

    d. **[AI] TIER staff to review possible provision of Docker Trusted Registry as a service for campus built container images.**
    e. **[AI] TIER will need to provide levvel.io with, in effect, a list of all private keys and other secrets which need to be protected and not be part of the campus Docker images. **provide additional properties file as part of the config**
c. Future Call -- Devops Pipeline: https://spaces.internet2.edu/display/TPWG/Core+Packaging+Build+Documents
    a. Notes on PDF here ...
    b.

July 18, 2016 4:00 eastern (Regular Call timeslot)

Attendees (please add yourself):

1. Scott Koranda (SCG)
2. Scott Cantor (tOSU)
3. Jim Jokl (Virginia)
4. John Gasper (Unicon)
5. Sara Jeanes (Internet2)
6. Brian Savage (BC)
7. Paul Caskey, Internet2
8. Tom Zeller (Shib IdP)
9. Drew Zebrowski - Temple U
10. Keith Hazelton - UW-Madison

Agenda and Notes

Primary agenda: two time-critical reports

1. Shibboleth: https://docs.google.com/document/d/1v1qHO6dvev8a7zZcptBEl4N2mpdydm4Sas7H8lE5DRM/edit
    a. Notes will be made in google doc
2. Pipeline: https://spaces.internet2.edu/display/TPWG/Core+Packaging+Build+Documents
    a. Notes on PDF here ...
    b.

Other Topics (time available)

1. Chris' email on CANARIE configuration tool - see 7/11/2016 email
    a. From Chris(July 18): Chris & Jim talked and Jim J commented that this is not likely to be reached on this call. Will collaborate with Jim on agenda priority and will attend then -- CP
2.

Note: https://github.com/UniconLabs/dockerized-idp-testbed
Something Gasper is sharing via a larger email to the committee (via Mike Grady): Dockerfiles are much like source code, .java files for example. Once compiled they should be treated as immutable, and one would not store env specific settings in the .java file. Like java classes, Docker images can be "inherited" to extended the image's functionality and make it more specific to the case we are trying to solve. Docker containers are like Java objects. They are instantiations of Docker images (i.e. Java classes). When a Java process is killed the object goes away. Likewise Docker containers are usually treated as ephemeral. The exception is a storage container, which stores persistent data. But neither the IdP nor Grouper really have persistent data as far as their images/containers are concerned. When a container is stopped, it can be restarted or removed. But they should also be immutable; any config change should be treated like changing the value of a String object… a new object is created and the old one garbage collected.

July 11, 2016 4:00 eastern (Regular Call timeslot)

Attendees (please add yourself):

1. Paul Caskey, Internet2
2. Sara Jeanes, Internet2
3. Mike Zawacki - Internet2
4. Scott Koranda - SCG
5. Keith Hazelton - UW-Madison
6. Jim Jokl - Virginia
7. Bill Kaufman - Internet2
8. Chris Phillips - CANARIE
9. Tom Zeller - Shib IdP

Agenda and Notes

1. Update on the work of the Shibboleth subgroup and Levvel.io
    a. (see minutes June 30)
2. Creation of Grouper and COmanage Subteams
    a. We also need subteams and mailing lists for Grouper and COmanage (similar to June 27 Item 2.a)

    b. We scheduled both Grouper and COmamage interviews early in the levvel.io process to provide time to make potential requests of the component teams.
    c. Grouper Volunteers: JimJ,
    d. COmanage Volunteers: JimJ, Scott K,
3. Other Items
    a. Insert your item(s) here
4. Shibboleth Subteam (i.e., 2.a.ii from June 27 below) Work
    a. Discussion: action items from July 6 call
    b. Creation of configuration trees for Levvel.io
        i. Config Tree tied into testbed for initial testing (start with idp.testbed or existing Docker VM or ?). Need this config tree soon. TomZ and JimJ
        ii. Creation of subtree to match our June 6 call
    c. Discussion / changes / approval of Levvel.io notes from our Thursday call
        i. Levvel sent notes to the subteam email list on Friday July 8 at 11:13 am ShibbolethIDPMeetingMinutes7716
    d. Other topics
5. TIER Working Group meeting at TechEx 2016 Thursday, Sept. 29 noon to 3 pm
    a. Please register: http://tinyurl.com/hdgknrv
    b. Lunchtime informal meetings for working groups? Please leave your feedback below:
        i.

July 6, 2016 (Shib IdP Default Configuration Call)

1:30 eastern time

Attendees (please add yourself):

1. Jim Jokl - Virginia
2. Scott Koranda SCG
3. Scott Cantor, tOSU
4. Bill Kaufman, Internet2
5. Sara Jeanes, Internet2
6. Gabor Eszes - Old Dominion
7. Keith Hazelton - UW-Madison
8. Tom Zeller - Shib IdP

Agenda: Shibboleth Default Configuration for Docker Packaging Work

1. What needs to be different from the current default - see Feb 15 notes
    a. Changes needed default config: a, b, c, d, e (define and document), f (guidance for Levvel.io), (i) yes, but subtract attribute-resolver.xml and potentially add metadata sources (Scott to check - (AI)), l, m, p, and q (either the default attribute page or some new page).
    b. Scott will provide some logging config data default configs (AI)
    c. Active Directory - later, when we get the rest of the automation in place.
2. Decisions
    a. No database this time (impact will be client-level consent storage)
    b.

June 30, 2016 (Shib IdP Acceptance Criteria discussion)

Attendees (please add yourself):

1. Jim Jokl - Virginia
2. Mike Zawacki - Internet2
3. Nick Roy - Internet2
4. Bill Kaufman - Internet2
5. Sara Jeanes - Internet2
6. Gabor Eszes - Old Dominion
7. Tom Zeller - Shib IdP
8. Scott Cantor - tOSU
9. Scott Koranda - SCG
10. Steve Carmody, Brown
11. Janemarie Duh, Lafayette
12. Keith Hazelton, UW-Madison (joining half-way through)
13. Paul Caskey, Internet2 (last 15 minutes)

Agenda: Vendor Acceptance Criteria for Shibboleth IdP

1. Background
    a. Our General Core packaging Requirements Refresh https://spaces.internet2.edu/display/TPWG/Core+Packaging+Assumptions+Refresh
    b. Back-end Infrastructure in the TestBed

        i. Existing infrastructure (LDAP, Kerberos, etc.) can be used to assist with testing and test automation.
        ii. Additional tools can be provided

2. What set of tests will we apply to determine that the vendor has successfully completed their work to build Docker and VM images of the Shibboleth IdP?
    a. Starting from the current Shibboleth IdP distribution and leveraging a default configuration provided by TIER, produce functional Docker and VM images via an automated build process..
    b. Shibboleth Configuration Files
        i. Ability to edit existing, and add additional configuration files, and have these changes persist across starts and stops of the container or VM
        ii. The configuration hot-reloading capability of the software is supported. Configuration changes need to take effect without restarting the component where supported by the component.
        iii. Ability of configuration data to persist across updates and upgrades of the OS and TIER components.
        iv. Ability for system operators to view downloaded metadata files and have these files persist across restarts of the container (ideally this data would be read-only outside of the container).
    c. Log access
        i. The Shibboleth IdP log files are/can-be available for processing in real-time outside of the container.
        ii. … same for OS/subcomponent logs (e.g., tomcat/apache/etc.)
    d. OS / Infrastructure Configuration
        i. Configuration for infrastructure components (e.g., tomcat) is clearly documented.
        ii. What does an end-user need to do, for example, to increase the memory allocated to tomcat.
        iii. Servlet engine and other equivalent configuration as provided by TIER. (we need to do this). Hardening SSL configuration. The TOMCAT_BASE configuration has to be externalized equivalent to 2.b.i and 2.b.iii above.
        iv. The container's clock is synchronized to real time.
        v. Verify Oracle java distribution requirements - click-through?
    e. Design Verification
        i. The provided images match the features / configuration described in the design documentation.

f. Pass an automated set of simple set of post-install checks?
    i. ??
    ii. ??
g. Pass a manual set of simple post-install checks?
    i. ??
    ii. ??
h. Session data storage -- [flesh out details]
i. ??

June 27, 2016

Attendees (please add yourself):

Mike Zawacki - Internet2
Keith Hazelton - UW-Madison
Bill Kaufman - Internet2
Scott Koranda - SCG
Joanna Rojas - Duke University
Paul Caskey - Internet2
Jim Jokl - Virginia
Scott Cantor - tOSU
Janemarie Duh - Lafayette College
Sara Jeanes - Internet2
Tom Zeller - Shib IdP
Niva Agmon - Temple U

Agenda

1. Procurement Status Update

2. Early Deliverables
    a. Component mailing list membership
        i. DevOps Pipeline List: KeithH, SaraJ, TomZ, BillK, ScottK
        ii. Shibboleth IdP List: TomZ , BillK, ScottK, Janemarie, ScottC
    b. Acceptance Criteria
        i. Shibboleth IdP - Packaged Image Acceptance Documents (six weeks starting now-ish)
        ii. Based on our packaging assumptions document, https://spaces.internet2.edu/display/TPWG/Core+Packaging+Assumption

s+Refresh, focus in on key points for the Acceptance Criteria for the Shibboleth IdP.

1. Automated test
   Pointed at a backend test infrastructure with attributes, AuthN, what set tests can we automate?
2. Acceptance (what would constitute “acceptance” for working group members)
   a. Access to edit all configuration files
      i. Ability to add new files, eg. a custom authn flow.
      ii.
   b. Log access
   c. Separation of configuration from application and os and upgrades; what do I maintain; what does TIER maintain;
   d. Simple set of post-install checks

iii. We will hold a call at 2:30 eastern Thursday 30 to focus on Acceptance for the Shibboleth IdP
   1. **AI: Mike to set up audio bridge, send invite**

iv. (Week 7) Grouper - Packaged Image Acceptance Documents
v. (Week 7) COmanage - Packaged Image Acceptance Documents

c. AI: Mike to set up audio bridge

3. Insert your item(s) here
   a. TIER WOrking Group meeting at TechEx 2016
      i. Please register: http://tinyurl.com/hdgknrv
      ii. Meeting details sent to mailing list: https://lists.internet2.edu/sympa/arc/tier-packaging/2016-06/msg00033.html
   b.

June 20, 2016

Attendees (please add yourself):

Mike Zawacki - Internet2
Scott Koranda - SCG
Keith Hazelton - UW-Madison
Jim Jokl - Virginia

Bill Kaufman - Internet2
Scott Cantor - tOSU
John Gasper - Unicon (unofficially)
Niva Agmon - Temple U.
Drew Zebrowski - Temple U.
Chris Hyzer - U Penn

Agenda

The agenda for our June 20 meeting will focus on high availability needs for the initial three components.

1. Background:
    a. Shibboleth IdP: https://wiki.shibboleth.net/confluence/display/IDP30/Clustering
    b. COmanage: https://spaces.internet2.edu/display/COmanage/Registry+Installation+-+High+Availability+Considerations
    c. Grouper - https://spaces.internet2.edu/display/Grouper/Grouper+high+availability
    d. (For future discussion) midPoint (Entity Registry and Provisioning) Clustering and High-Availability Setup: https://wiki.evolveum.com/pages/viewpage.action?pageId=11075783

2. Requirements Discussion
    a. What are the expected use HA cases for the components
        i. Shibboleth - common deployment expectation would be high availability (with some level of maintenance window (or not). Immediate recovery - doesn't go down.
        ii. Grouper - common deployment needs: web services component - immediate recovery (used for real time AuthZ decisions); other aspects of grouper: one-day recovery. As Grouper maturity builds on campus, more campuses will rely on Grouper for real-time AuthZ.
        iii. COmanage - typical deployments include a non-HA registry and a HA LDAP.
        iv. We all agree we need HA AuthN/AuthZ sources -- e.g., LDAP, etc.
            1. Keith H: Is there enough LDAP activity/dependency to make HA worthwhile?
                a. Jim: Yes, and it's needed for Shib, so the dependency makes it worth including
    b. Which components can have stateless or partial configuration options
        i. Shibboleth - with no server side storage we lose cross-node Replay, SAML artifact, CAS, (memcached can handle most of what we need). Without a database we lose
            1. Database stored persistent IDs
            2. Server-side consent persistence
        ii. Grouper Web Services -- what is needed for HA

1. Database
2. UI servers - cluster stateful session behind load balancer
3. Web services servers - cluster cleanly behind load balancer - stateless
4. Daemon- Can run multiple instances
5. Note: other Grouper subcomponents are also stateless and run trivially behind a load balancer.

iii. COmanage

1. See 2.a.iii above -- registry itself not typically HA

c. other

3. Common Subcomponent Discussion
    a. LDAP
    b. Database
        i. Two HA possibilities for Grouper - rw backup; ro backup (can stay on-line)
        ii. TIER: evaluate difficulty of HA solution - installation and operations (consult MySQL DBA experts)
        iii.
    c. other
4. Component Discussion
    a. ..
    b. ..
    c. ..
5. Other Topics
    a. Packaging Working Group BoF for TechEx 2016? Deadline to submit **June 30th.**

June 6, 2016

Attendees (please add yourself):
Mike Zawacki - Internet2
Bill Kaufman - Internet2
Jim Jokl - Virginia
Janemarie Duh - Lafayette College
Sara Jeanes - Internet2
Scott Cantor, tOSU
Paul Caskey, Internet2
Keith Hazelton - UW-Madison
Tom Zeller - Shib IdP
Drew Zebrowski - Temple U.

Agenda

1. Release 1 testing feedback
    a. Keith - Ran through Shib IdP all the way. Worked well, no surprises. Didn't try to hook it up to anything else.
    b. Janemarie - Ran into some problems with networking. Was able to import machine but couldn't get IP (both wired and wireless). Local network ended up killing connectivity due to NIC being set to bridge mode. It was seen as a security violation due to too many mac addresses on a single port. Once that was sorted out, got a static address and should be able to install.
    c. Bill - Getting IP but looks like Shib is waiting to hear back from remote end.
    d.
2. Core packaging assumptions refresh -- see link
    a. Log discussion - common formats (yes/no), time, etc. updates by component groups?
    b. Logging to facilitate health checks and instrumentation
    c. Is monitoring something TIER will provide as a 'component' or will it just be something like providing a source of monitoring data (e.g. Nagios)
    d. Change "Monitoring" to "Local Process Management"
3. Campus Shibboleth metadata management tool - see link
    a. Mdui section: mandatory, recommended, optional; (be able to enforce); Mandatory: displayname, description; all others are optional
    b. Enforce https on logo
    c. verify/eliminate special characters;
    d. Add collect Security Contact Data
    e. Add collection of Support Contact Data
    f. Add collection of phone number of the user to the contacts
    g. Fill in the Bindings section
    h. Handle logout
    i. Certificates (as opposed to certificate)
        i. Be able to enforce key sizes
        ii.
4. Frequency of calls and subgroup work
    a. Proposal: Break down into component subgroups, use this timeslot to have focused work meetings.
    b. Next week's meeting to cover High Avaiaibility for for TIER starting with the needs of Shibboleth, Grouper, and COmanage
    c. Defining HA for TIER
    d. Need to draft definition and requirements for HA, present to component architects
    e. IdP: https://wiki.shibboleth.net/confluence/display/IDP30/Clustering
    f. Jim to follow up on COmanage and Grouper

i. COmanage: https://spaces.internet2.edu/display/COmanage/Registry+Installation+-+High+Availability+Considerations
ii.

5. Other items
    a. ?

May 23, 2016

Attendees (please add yourself):
Nick Roy (Internet2)
Mike Zawacki - Internet2
Scott Koranda - SCG
Bill Kaufman - Internet2
Jim Jokl - Virginia
Scott Cantor, tOSU
Keith Hazelton - UW-Msn
Steve Carmody, Brown

Janemarie Duh, Lafayette
Tom Zeller, Shib IdP
Paul Caskey, Internet2
Brian Savage - Boston College
Niva Agmon - Temple U

Agenda:

1. Feedback on the Shibboleth container VM
    a. **Volunteers to fully test VM ahead of next week's call?**
        i. **Janemarie D**
        ii. **Steve Carmody**
        iii. **Keith H**
    b. Paul C - is it possible to put together a version of the VM which exposes the logs to export/review?
        i. Jim: I can make that change
2. Requirements for local campus metadata management.

Local Campus Metadata Management Requirements

One of the items identified early in our discussions on what is needed to make Shibboleth easier to deploy and use effectively on campus was the notion of a campus local metadata management tool. Many schools have created local solutions and many others manage local metadata by hand. Our question here is what would be the requirements for a campus federation local metadata management tool. If such a tool looks useful, does it make sense for TIER to adopt or create one.

Requirements

1. End Users
    a. Enable campus users to authenticate using Shibboleth and request that their SP be added to the campus metadata distribution.
        i. Simple web form to prompt the SP operator for standard/simplified information. Included with the default web form is text that explains the needed elements and where to get the information.
        ii. Error checking on the entered data.
    b. The ability for the identity that created the entry and designates to edit the entry (including the list of designates).
    c. The ability for the identity that created the entry and designates to delete the entry.
    d. Ability to check on the processing status of the request, including the ability for pure self-service - i.e., automatic addition to the metadata.
    e.
2. System Operators
    a. The ability for campus system operators to edit any metadata entry.
    b. The ability for campus system operators to approve/reject requests.
    c. The ability for campus system operators to, perhaps manually, directly edit the metadata for an entity.
    d.
3. General
    a. Handle only the case of the addition of SP metadata.
    b. Send email to End Users whenever the data for an entity that they own is modified.
    c. Be trivially simple to containerize, install, and operate.
    d. Keep track of previous versions of entity metadata
    e. Options to publish metadata (a) on a regular cycle, (b) after testing by a sysadmin, or (c) in semi-real time after a new record is submitted.
    f. Support for regular review and signoff by entity owner
    g. Ability to know that sets of entities are linked (e.g., production and test)
4. Information Requested
    a. Contact information / Dept information
    b. Certificate(s)

c. Virtual host(s) / SAML endpoints
d. MDUI extension info
e. Requested attributes and justification for sensitive attributes
f. Software implementation details (text box or a few standard answers plus “other”)

May 9, 2016

The May 9 call is canceled. Please use the time to work on taking a look at the Shibboleth VM and start on requirements for a potential campus metadata management tool.

Meeting Notes

May 2, 2016

Attendees (please add yourself):
Scott Koranda - SCG
Jim Jokl - Virginia
Kevin Foote - UOregon
Scott Cantor - tOSU
Drew Zebrowski - Temple U.
Keith Hazelton - UW-Madison
Mike Zawacki - Internet2
Niva Agmon - Temple U.
Janemarie Duh - Lafayette
Chris Hyzr - Penn
Gabor Eszes - Old Dominion

Agenda

- We will hold a short call today
- Any discussion on Docker / VM status for COmanage
- Shibboleth IdP Docker VM
    - https://testbed.tier.internet2.edu/

    - https://spaces.internet2.edu/display/TPD/Shibboleth-IdP+Virtual+Machine+Documentation
    - **AI: JIM - Produce set of Docker commands that make it easier to log into the container to review its operation and capabilities (Due 5/5)**
- Next Steps
    - Campus Federation requirements

April 25, 2016

Attendees (please add yourself):
Mike Zawacki - Internet2
Scott Koranda - SCG
Keith Hazelton - UW-Madison
Jim Jokl - Virginia
Kevin Foote - UOregon
Scott Cantor - tOSU
Paul Caskey - Internet2
Drew Zebrowski - Temple U.
Niva Agmon - Temple U.
Brian Savage - Boston College
Tom Zeller - Shib IdP
Janemarie Duh- Lafayette

Agenda

- We will hold a brief call today
- Update on Docker / VM status - close on Shibboleth IdP
    - 95% complete as of this meeting. Anticipated completion this week. Look for update here: https://spaces.internet2.edu/display/TPD
        - This group will review before general announcement.
    - Once IdP piece is done, group will focus in on usability.
    - Goal is to move toward doing the builds ourselves, rather than spending contractor cycles on builds.
    - Questions
        - Q: Any summary of what was brought into IdP?
        - A: Need to get KERBEROS identity into place, etc. Lots of small clean-up tasks rather than large tasks.
        - Q: So does the IdP that runs in the container hit the KCD test bed to look for valid credentials? Is the test bed something TIER is running
        - A: Yes. And yes, TIER is running the test bed.

        - Scott - voices concerns with IdP making calls to test instance. Jim: This initial release is intended as a “first look” to encourage adoption, familiarizaiton
        - (Question around COmanage… to be covered below)

- Discuss any COmanage VM testing results. Has anyone else tried this out?
    - Keith: Took software as far as it was built out (e.g. to login page), no issues.
    - Jim: Biggest single thing we can do is get that component into the test bed so that it can be fully tested.
    - Keith: Some familiarity with command line interface needed - something to consider longer term for the long tail.
    - 2 paths forward: First would be federating test bed. Second would be setting up infrastructure to allow an admin to inject her test IdP metadata so that it too could be federated with test COmanage SPs.
        - First is not much work but requires rolling a new COmanage container/VM.
    - Needed: Quick list of minor changes/updates. **AI: Scott Koranda to assemble this before the next call**

April 18, 2016

Attendees (please add yourself):

- Scott Koranda (SCG)
- Nick Roy (Internet2)
- Jim Jokl - Virginia
- Kevin Foote - uoregon
- Janemarie Duh - Lafayette
- Steve Carmody, Brown
- Mike Zawacki - Internet2
- Keith Hazelton - UW-Madison
- Drew Zebrowski - Temple U
- Tom Zeller - Shib IdP
- Chris Phillips - CANARIE
- Paul Caskey - Internet2
- Niva Agmon - Temple U

Agenda

- Update on Docker / VM status
- Testing the COmanage VM release

    - Any of us with time & inclination should try the COmanage VM install per the release notes and share any feedback on our WG list (for now).
- https://spaces.internet2.edu/display/TPD
- Start thinking about next steps; Next couple months:
    - Catching & fixing glitches in R1
    - Figure out what we want the initial setups and configurations to look like
    - Anticipated timeline for next version? A few weeks. Then collect feedback, work on next rev. Feedback to be solicited internally (i.e. devs and WGs, not public), ala Agile method
- Feedback

April 11, 2016

Attendees (please add yourself):
Chris Hyzer (Penn)
Scott Koranda (SCG)
Jim Jokl - Virginia
Mike Zawacki - Internet2

Limited call, recap of Packaging Group's contribution to TIER Release 1

March 28, 2016

Attendees (please add yourself):
Jim Jokl - Virginia
Mike Zawacki - Internet2
Scott Koranda - SCG
Keith Hazelton - UW-Madison
Nick Roy - Internet2
Paul Caskey - Internet2
Kevin Foote - UOregon
Steve Carmody, Brown
Janemarie Duh - Lafayette
Tom Zeller - Shib IdP
Niva Agmon - Temple U

Agenda
We will hold a very brief call today to discuss status and work in progress.

Goal is to have Docker container done before TIER Release One. Allow for review and comment ahead of release.

- Reversed the order of work plan - rather than prioritizing focus on config portions, created containers first
- Running behind at the moment
- Added contractor time to complete Docker/container portion
- Group will continue working on remaining pieces
- Also added VM portion to deliver containers, given lack of comfort with Docker amongst community (though most survey respondents identified Docker as a priority for future deployment)
- Main work of group will pick up again after April 16 release date

Survey results: https://spaces.internet2.edu/x/CwuVBQ

Q: Sense for number of institutions which are prepared to deploy, manage Docker containers right now?
A: A definite minority - less than 10%

Meeting Notes

March 14, 2016

Attendees (please add yourself):
Jim Jokl - Virginia
Brett Bieber - University of Nebraska-Lincoln
Keith Hazelton - UW-Madison
Kevin Foote - uoregon
Gabor Eszes - Old Dominion
Drew Zebrowski - Temple U
Janemarie Duh - Lafayette College
Tom Zeller - Shib IdP
Niva Agmon - Temple U
Steve Carmody, Brown
Brian Savage - Boston College

Agenda

The focus for today's call will be on Docker training and what is/will-be needed by the community. If you have a few minutes before the call, please take a look at: https://training.docker.com/instructor-led-training so we can discuss potential course content.

Other training options: https://training.docker.com/self-paced-training (about 3 hrs of webinars, Intro, fundamentals, operations)

Discussion synopsis:

1. Our goal for this discussion on a potential Docker training event was as a second path around two year hurdle that we saw in the survey data re: when sites thought they'd generally be ready to use container-based solutions in production. The first path is our planned VM environment to enable sites to treat the TIER containers black boxes.
2. The people on our team who are familiar with Docker did not make today's call so we did not make progress on what an appropriate training agenda is.
3. TechEx is likely not the best place for core Docker training. The people who attend the event are generally not the people who have campus operational responsibility.
4. WebEx-type training might be appropriate but is often complicated by staff still doing their day jobs while training is in progress.
5. A special event that couples some container training with TIER specific training might work.
6. There was a general feeling that the TIER-specific black box documentation/training may be all that is really needed, especially for the smaller schools. The larger schools will be working in this space eventually anyway.

3-7-16 Meeting

TODAY'S MEETING HAS BEEN CANCELLED. WE WILL MEET NEXT WEEK AT 3/14 AT THE USUAL TIME

2-29-2016 Meeting Notes

Attendees (please add yourself):
Jim Jokl - Virginia
Scott Koranda - SCG
Nick Roy - Internet2/InCommon
Scotty Logan - Stanford (mostly lurking today)
Keith Hazelton - UW-Madison
Chris Phillips - CANARIE (may have to leave early)
Kevin Foote - uoregon

Scott Cantor - tOSU
Niva Agmon - Temple U
Drew Zebrowski - Temple U
Janemarie Duh - Lafayette
Tom Zeller - Shib IdP
Chris Hyzer - Penn

Agenda

- Grouper Requirements
    - Include the Shibboleth SP in the packaging
    - Include MySQL (or potentially MariaDB) in the Grouper packaging
    - Use tomcat (FYI: we can use real Tomcat with Shibboleth, just not what comes with the OS)
    - High-level architecture diagram
    - Package web service and ui separately (separate Docker containers)?
    - Potentially include
        - scripting for ldap subject source configuration
        - Potentially include script for SP configuration
- (if time) COmanage Requirements https://spaces.internet2.edu/display/COmanage/COmanage+Technical+Manual
    - Reference architecture diagram
    - Need PHP (any modern php will be ok)
    - Need Apache HTTP Server
    - Need relational database, test with MySQL and Postgres, when given choice dev team installs Postgres, no known Oracle deployments, no objection to going with MySQL
    - Need authentication layer, most likely Shibboleth SP
    - Inject details for first admin (givenName, sn, identifier (ePPN usually))
    - Inject SMTP email configuration details?
    - Set up cron job?
    - Separate vs. enterprise LDAP? Survey says existing enterprise ldap.
    - Attribute authority - likely not an April release question
    - SAML IdP/SP proxy - likely not an April release question
    - Quickstart Applications -
- Survey Data: https://spaces.internet2.edu/x/CwuVBQ

2-22-2016 Meeting Notes

Attendees (please add yourself):
Jim Jokl - Virginia
Mike Zawacki - Internet2
Scott Koranda - SCG
Gabor Eszes - Old Dominion
Kevin Foote - uoregon
Scotty Logan - Stanford
Steve Carmody, Brown
Brian Savage - Boston College
Paul Caskey, Internet2
Scott Cantor, tOSU
Tom Zeller, Shib IdP
Niva Agmon, Temple U
Chris Hyzer - Penn

Agenda

- TIER Working Group and Developers Meeting to be held following Global Summit
    - May 19, 2016 , 9am - 12:30pm in Chicago
    - **Important**: Please indicate on this Google Spreadsheet if you will attend
- Complete Shibboleth requirements discussion (we will continue this work in the 2-15-2016 notes section). Results: https://spaces.internet2.edu/x/CwuVBQ
- Grouper Requirements

2-15-2016 Meeting Notes

Attendees (please add yourself)
Gabor Eszes - Old Dominion
Mike Zawacki - Internet2
Scott Koranda - SCG
Jim Jokl - Virginia
Steve Carmody, Brown

Tom Zeller - Shib IdP
Mark McCahill - Duke University
Scott Cantor - OSU
Niva Agmon - Temple U
Paul Caskey - Internet2
Keith Hazelton - UW-Madison
Janemarie Duh - Lafayette

(ACTION ITEM, Scott Cantor - ask marvin - needs server-side support) re: offering default support of CAS

Agenda

(a) Update on tooling call for Shibboleth initial configuration
(b) Finish Shibboleth (at least for our initial work)
    (i) We let our conversation drift from requirements to solutions on the last call. Today we need to remain focused on what we need for Shibboleth itself.
    (ii) Shibboleth IdP default configuration settings
    (iii) Shibboleth pain points to be addressed over time
    (iv) Remember to refresh your memory on the survey results (https://spaces.internet2.edu/x/CwuVBQ)
(c) Start on Grouper Requirements

Shibboleth Default Configuration

(a) Yes -- Load and use InCommon metadata
(b) Yes -- (and assume an eduPerson based directory) Include support for a default set of attribute definitions (LDAP - name, email; eduPerson -EPPN, Affiliation, primaryAffiliation, ?) We note that we may still need to do something special for AD.
(c) Yes - Release EPPN, name, email, affiliation, eduPersonTargetedId to all InCommon SPs? (TIER to provide documentation for sites to opt-out if needed)
(d) Yes -- Release EPPN, name, email, affiliation, eduPersonTargetedId to SPs with the Research and Scholarship R&S entity category (includes eduGAIN)? (TIER to provide documentation for sites to opt-out if needed along with discussion on why this is generally the “right thing to do” - we also need to ensure that InCommon helps with the education in this area (we believe we are helping InCommon’s agenda)).
(e) Yes -- Respect a FERPA opt-out attribute to restrict attribute release for some users. (Add some type of configuration to report this issue to the end user).
(f) Yes -- Avoid spurious errors in the logs from external scanners via a properly configured robots.txt
(g) Yes -- Support Enhanced Client or Proxy (ECP) by default ? (potentially make available if configured with a compatible authentication source)

(h) No -- in general and on Duo now (wait for more implementation maturity before deciding)-- Support multi-factor authentication by default? (what - would the be legal issues if we selected Duo or should we only do TIER-MFA (U2F, PKI, etc.)
(i) Yes (with exception of attribute-resolver) -- Automatically reload config files when they are changed (relying-party.xml, attribute-filter.xml, attribute-resolver.xml)?
(j) No - Support CAS by default (document HA issues)?
(k) No - not relevant now - grant submitted for funding support and maintenance- Support OpenID Connect by default (when available)?
(l) Yes -NOT support SAML 1 by default?
(m) Yes- NOT support SAML Attribute Queries?
(n) No - Update itself automatically (document a site can do this)?
(o) No - Update itself automatically - security updates only (document how a site can do this)?
(p) No - Prompt users to consent to attribute release?
(q) Yes - Add a simple consent type configuration to enable FERPA opt-out over-ride (either per-service or potentially globally) when no attributes would have been released for the user..

Authentication

(1)

What we need for Shibboleth

(1) Data on what configuration changes sites make (away from our defaults)
(2) Web interface to wrap command line tools that are already available (likely low hanging fruit)
(3) Log analysis - starting perhaps with some common log format work
(4)

2-8-2016 Meeting Notes

Attendees (please add yourself)
Jim Jokl - University of Virginia
Mark McCahill - Duke University
Scotty Logan - Stanford
Gabor Eszes - Old Dominion
Janemarie Duh - Lafayette College
Brett Bieber - University of Nebraska-Lincoln
Tom Zeller - Shib IdP
Chris Phillips - CANARIE
Scott Cantor - tOSU
Paul Caskey - Internet2
Drew Zebrowski - Temple U
Niva Agmon Temple U

Agenda

(a) Review of where we ended last week's call
(b) Discuss "Proposed Solution" section at https://spaces.internet2.edu/x/6QqVBQ
(c) Discuss Shibboleth section of the survey at https://spaces.internet2.edu/x/CwuVBQ in the context of our needs for default configuration settings, configuration, ease of use, etc., etc.

ToDo

1. call about IdP Installer and survey results
    a. Jim J
    b. Chris P
    c. Scotty L

Packaging

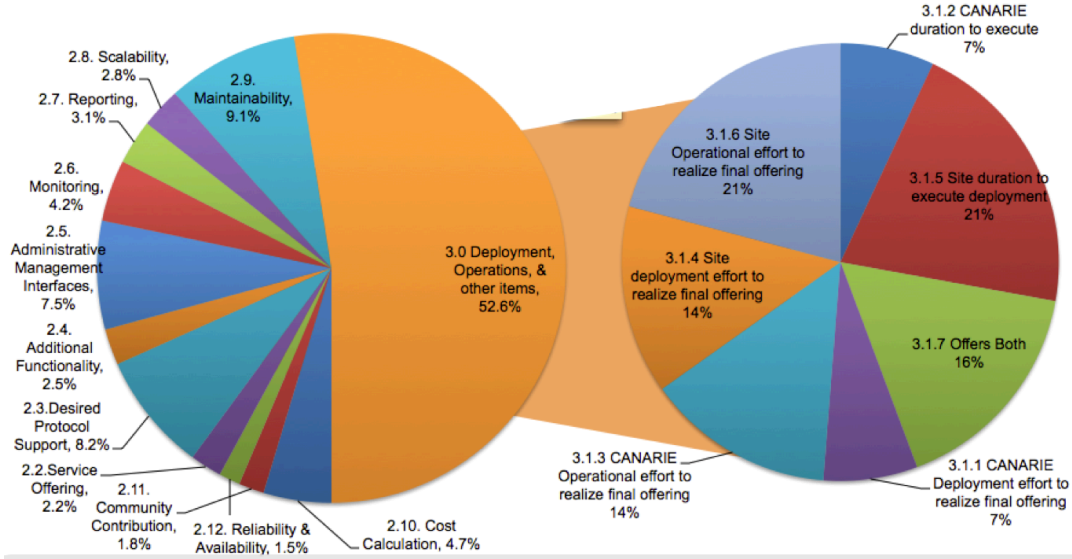

Specifying the Solution

- Weighted scorecard approach
- Over 15 solutions and combinations evaluated
- Facets of operation are as important as what it does

Packaging Survey Results (core packaging)

1) Physical Servers - ~50% now to 21% two years out
2) Virtual Servers: ~40% now to 25% two years out
3) Virtual appliance – steady over two years (now, 1 year, 2 years) 33% to 28% to 38%
4) Docker: ~9% to ~50% two years out
5) Other also went up: hosted, cloud/hosted, Cloud/SAAS, AWS, Cloud/SaaS/IaaS, physical purpose-built appliance, SaaS, SaaS, IaaS/PaaS, Not comfortable with IO performance in our current virtual environment, Not Sure,
6) ~60% – 40% on cloud vs. local; but later 81% state they prefer to be cloud agnostic
7) What do you prefer for packaging: 28% wanted existing packaging; 21 (Docker) + 13 (VM/Appliance) + 15 (Cloud IaaS) + 11 (Cloud SaaS) => 60% for solutions covered by the container-based packaging proposal).

What are the local configs anyway (e.g., Shib IdP but all components will have equivalent needs)?

1. datastore url, credentials
2. tls cert

3. additional relying parties
4. custom attribute definitions
5. custom attribute filters
6. cosmetic / look & feel

What are the operational tasks anyway?

1. adding new relying parties
2. new attribute definitions
3. new attribute filters
4. rotate certs
5. rotate credentials
6. notice a pattern that these lists are very similar? how do we make the operational load easier?

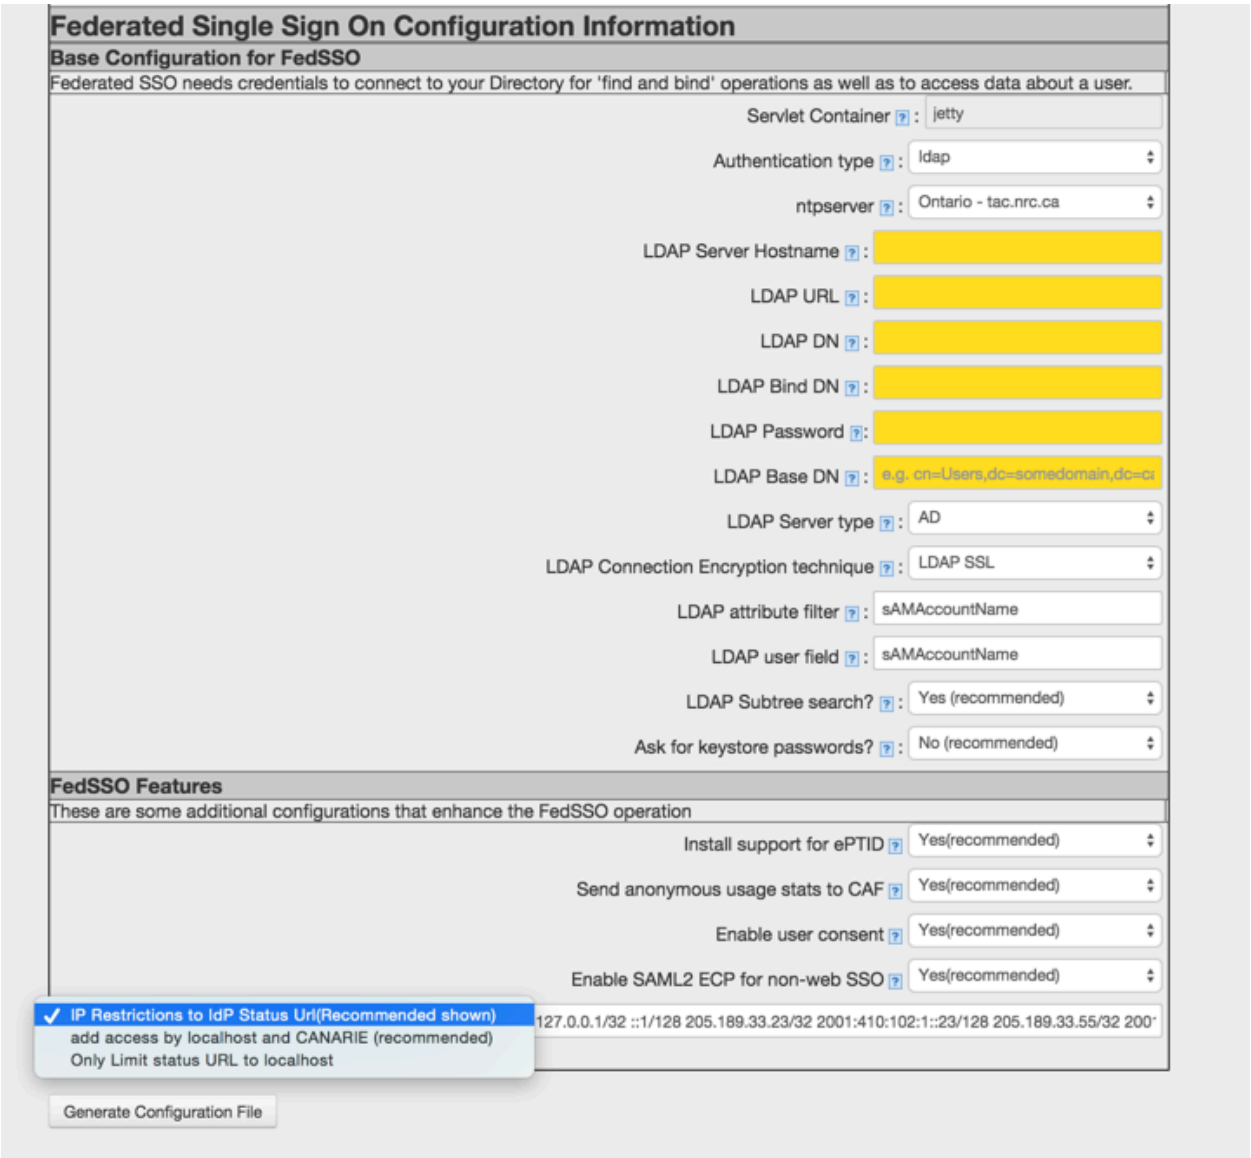

The CANARIE automated installer process http://bit.ly/idpinstaller

This link is to the landing page we maintain locally for the installer. Current 'v3 of the installer' that would exhibit a stable install is the RC3 one (see bit.ly link for more).
The link resolves to: https://collaboration.canarie.ca/elgg/groups/profile/847/idp-installer

The github repository is at https://github.com/canariecaf/idp-installer-CAF/tree/3.0.0-CAF-RC3

2-1-2016 Meeting Notes

Attendees (please add yourself)

Jim Jokl - University of Virginia
Keith Hazelton - UW-Madison
Scott Koranda - SCG
Mike Zawacki - Internet2
Kevin Foote - uoregon
Nick Roy - Internet2
Gabor Eszes - Old Dominion
Chris Phillips - CANARIE
Paul Caskey - Internet2
Scotty Logan - Stanford
Tom Zeller - Shib IdP
Brian Savage - Boston College
Scott Cantor - tOSU / Shib
Janemarie Duh - Lafayette
Mark McCahill - Duke

Agenda

1. Quick review of 1-25-2016 meeting notes
2. Packaging Proposal Discussion - see: https://spaces.internet2.edu/x/6QqVBQ
3. Shibboleth default configuration and ease of use - see survey: https://spaces.internet2.edu/x/CwuVBQ
4.

1-25-16 Meeting Notes

Attendees (*please add yourself*):

Nick Roy - InCommon/Internet2
Mike Zawacki - Internet2
Scotty Logan - Stanford
Scott Koranda - SCG
Brian Savage - Boston College
Gabor Eszes - Old Dominion
Janemarie Duh - Lafayette College
James Jokl - University of Virginia

Brett Bieber - Nebraska
Keith Hazelton - UW-Madison
Tom Zeller - Shib
Drew Zebrowski - Temple U.
Scott Cantor - tOSU / Shib
Kevin Foote - uoregon
Niva Agmon Temple U.
Paul Caskey - Internet2

AI FOR ALL - Review these notes, add, tweak, update as needed

Review of survey results

Most state that they would be comfortable with containerization 1-3 years out, but not currently.
Preference of local IdM solution could be troubling (56% said yes)
Interested in learning about support for hybrid model.
Note - “Greatest impediment to new tech/solutions” stated lack of local resources to pursue it

- Could indicate a need to insulate deployers from need for tech expertise

Also need to consider how to provide enough support to ease deployment but portable enough to handle changes down the road

- Component upgrade question from Jim: How often does format for basic config files change?
    - Depends on the tech. Worth noting that such upgrades would likely be uncommon.

Proposal from Scotty: use Packer, provide VMs with containers. Other ideas?

- Lots of data on what people want for defaults, but what about providing needed site customization?
- Chris: Would you package image for all possible environments? How to handle provisions for plugins, etc? Will need to address these questions.

Scotty: Would be good to have a view on Splunk usage, too. Maybe offer way to have installers automagically dump logs into local Splunk instance?

Packer discussion (including VM component)
Jim: Scotty, what are you using now?

- Grouper component, Shib IdP, deployable via AWS

Jim: Build environment wasn’t part of initial April deliverables. What do we have on hand now?

Build RPMs - would need to be picked up from community
Docker image of Shib IdP. But how would deployer handle local config issues/conflicts?

Jim - for next call: Look over results, come with suggestions on moving to next phase

---

12-21-15 Meeting Notes

Attendees (*please add yourself*):
Scott Koranda - SCG
Jim Jokl - Virginia
Nick Roy - InCommon/Internet2
Paul Caskey - Internet2
Kevin Foote - uoregon
Chris Hyzer - penn

Agenda: review the draft survey instrument: http://www.questionpro.com/t/AK1buZTO63

Dial-in numbers:
+1-734-615-7474 (English I2, Please use if you do not pay for Long Distance)
+1-866-411-0013 (English I2, toll free US/Canada Only)
Access Code: 0125971

12-17-15 Survey work meeting

Note different audio bridge info:

+1-734-615-7474 (Please use if you do not pay for Long Distance),
+1-866-411-0013 (toll free US/Canada Only)
Access codes: 0107375#

Attendees *(please add yourself):*
Nick Roy - InCommon/Internet2
Mike Zawacki - InCommon/Internet2
Jim Jokl - Virginia
Tom Zeller - shib
Janemarie Duh - Lafayette
Misagh Moayyed - Unicon
Chris Hyzer - Penn
Gabor Eszes - Old Dominion
Paul Caskey - Internet2

Finalization of Survey. Please start with Grouper section

12-14-15 Meeting Notes

Attendees *(please add yourself):*
Nick Roy - InCommon/Internet2
Scott Koranda - SCG
Tom Zeller - shib
Mike Zawacki - Internet2
Brett Bieber - University of Nebraska-Lincoln
Drew Zebrowski - Temple U.
Jim Jokl - Virginia
Janemarie Duh - Lafayette College
Kevin Foote - uoregon
Brian Savage - Boston College
Paul Caskey, Internet2
Chris Phillips
Niva Agmon Temple U.
Keith Hazelton - UW-Madison
Steve Carmody, Brown
Gabor Eszes - Old Dominion

Agenda:

Today's meeting will focus on a review of the survey: Survey Prep Work

---

12-7-15 Meeting Notes

Attendees *(please add yourself):*
Mike Zawacki - Internet2

Janemarie Duh - Lafayette College
Tom Zeller - Shibboleth IdP
Drew Zebrowski - Temple U.
Jim Jokl - Virginia
Scott Koranda - SCG
Brett Bieber - Nebraska
Keith Hazelton - UW-Madison
Brian Savage - Boston College
Chris Phillips - CANARIE / IdP-Installer
Keivn Foote - uoregon
Niva Agmon Temple University
Gabor Eszes - Old Dominion Univ

Can the teleconference url be added to this doc?

- *Meeting is audio only, with bridge info included at the top of this doc.*

(Review of demographic and core surveys. Updates were made during meeting)

AI: Review survey segments, change or suggest as needed

11-30-15 Meeting Notes (Cancelled)

Today's meeting has been cancelled. We will convene again per usual at Monday, 12-7

11-23-15 Meeting Notes

Attendees *(please add yourself):*
Mike Zawacki - Internet2
Scott Koranda - SCG
Keith Hazelton - UW-Madison

Drew Zebrowski - Temple U.
Scotty Logan - Stanford
Kevin Foote - UOregon
Chris Phillips - CANARIE
Jim Jokl - Virginia
Janemarie Duh - Lafayette
Brian Savage - Boston College
Steve Carmody, Brown
Niva Agmon, Temple U
Chris Hyzer, Penn
Paul Caskey, Internet2

- Review of Survey questions (as filled in below)
- Ensure you have access to the Shib survey: https://docs.google.com/document/d/1-noQWM0ur0A-pIXVZBJT8Gchfe3rZmROMYTLG956ag8
    - Most/all work done for this call was tracked in this document
        - **AI - ?: Rework needed for 3(b)vi - "Translating log errors into meaningful feedback. (very hard and would build an expectation that we'd be unlikely to be able to deal with)." (DONE - replaced by question 3(b)vii during call)**
- **AI - ALL: Need to add respondent roles/demographics are captured at top of each survey**
- **AI - ALL: Those wishing to similarly review/revise other subgroup surveys. Please add your name(s) to the subgroup heading and begin work over the coming week. Revisions to be discussed at next meeting.**

Janemarie Duh - What tool will we use for surveys? Jim: UVA has a tool for surveying. Will do an initial draft of Shib survey in that and we can finalize decision later. **AI - JIM: Will try to have that available before next call**

11-16-15 Meeting Notes

Attendees *(please add yourself):*

Mike Zawacki - Internet2
Mark McCahill - Duke University
Derek Owens - University of Notre Dame
Brian Savage - Boston College
Janemarie Duh - Lafayette College
Chris Phillips - CANARIE
Jim Jokl - Virginia
Kevin Foote - UOregon
Niva Agmon - Temple University
Paul Caskey - Internet2
Drew Zebrowski - Temple U
Steve Carmody - Brown University
Keith Hazelton - UW-Madison
Tom Zeller - shib
Scotty Logan - Stanford

Work for our November 16/23 calls

Survey Prep Work
Our focus before the next call is to start the work of developing the set of questions that each of our areas (Shibboleth, Grouper, COmanage, and Core Packaging) need answered before we can move on and start to work clustering and solutions. The next step after the questions are ready is to select out groupings of target audiences and then move forward with survey distribution. Please use the sections below to collect questions. A few people have volunteered to start the process for a some of the areas. The service areas (Shibboleth, Grouper, and COmanage) should focus on issues directly related to their product (e.g., default configurations, tool needs, external dependencies (e.g., with Shibboleth, AuthN/Group data), etc.).

Document was annotated and tweaked during this call

General demographics re: who is filling in the form and their role (role from a constrained set of responses)

Chris: What is the goal with this list?

Jim: Determine the set of core questions to determine what will be in the final packing (or at least stimulate conversation of same)
Question: Question 10 seems to be aimed at those who haven't yet deployed service. Should we include question aimed at those who have already deployed but want to optimize?
Scotty: Perhaps use skip logic on survey to craft those sorts of questions more to suit respondent's use/needs?
Jim: Questions were meant to be structured that way
Jeanmaire: Shib quiz should capture that
(Question #10 was restructured during call)

Jim: Need to address default attribute release - that will cover questions around 3rd party IAMs like 365/AD, Google, etc.

Scotty: Could it make more sense to include packages that would correctly configure AD and other solutions to work with/in lieu of Shib?
Jim: Could be. If the right tool for some of those sites isn't Shib we should address that.

Chris: Presentation from TechEx15 had good gap analysis between AD and Shib. Could serve as resource. https://refeds.org/meetings/30th-meeting-oct-2015 -- specific presentation: https://refeds.org/wp-content/uploads/2015/10/refeds-chris-nick-fedtech-techex15-asPresented.pptx
Chris: What is TIER/InCommon/I2 willing to support?
Jim: This group is meant to recommend packaging of solutions, not just installers but considering additional installers/tools/upgrades needed. Should try to get to a small, manageable number that TIER can support over time and consider small/medium schools with limited IT expertise.
Chris: Off the cuff it will be very difficult to fully swap ADFS in for Shibboleth and get the same functionality as Shibboleth delivers without killing the support team to support both configurations (see refeds presentation above for supporting this perspective) . AD could be supported but ADFS would require a large, specialized team to deploy and support since there are things that Shibboleth can just Do that ADFS cannot.

Per Jim: From Shib document (see above) this would be a good question for core packaging survey: *For your identity services, would you prefer (a) a virtual appliance, or (b) a managed cloud-based service?*

11-09-15 Meeting Notes

Attendees (***Please add yourself***):

ChrisH, Nick Roy, Jim Jokl (jaj@virginia.edu), Keith Hazelton, (keith.hazelton@wisc.edu), Brett Bieber (bieber@unl.edu), Derek Owens (dowens@nd.edu), Kevin F (kpfoote@uoregon.edu)., Scott Koranda, Mark McCahill (mccahill@duke.edu), Brian Savage (brian.savage@bc.edu), Mike Zawacki, Janemarie Duh (duhj@lafayette.edu), Tom Zeller,Drew Zebrowski(drew@temple.edu), Chris Phillips, Scott Cantor, Paul C, Scotty Logan, Steve Carmody.

HINT: You can turn your name into a mailto: link

1) Review Charter
https://spaces.internet2.edu/display/TPWG/TIER+Packaging+Working+Group

- Question from Jim: Does the Charter look reasonable? Timelines realistic, other questions?
- Brett: Seems like COmanage isn't as tightly integrated into charter language. Consistency in tech referenced would be helpful.
- Jim: **AI - UPDATE LANGUAGE TO REFLECT THAT (done - 20151109)**
- Mark: In addition to packaging some manner of testing would be good. Ideally something operators can run on their own instance.
- Paul: Config check is good, but should also test on different levels (e.g. appropriate attribute release, etc).
- Scott: Surprised that "packaging" included admin interface development.
- Jim: That's a result of merging a few different groups.
- Scott: Concerned that the scope of that work is potentially huge, questions of technical feasibility.
- Jim: Refering to "ease of use" portions?
- Scott: Yes
- Jim: We may need to do some outreach, seek out additional SMEs
- ChrisH: Could we use existing admin interfaces?
- Scott: Possibly, yes, if we're careful of scope.
- Jim: We want to make sure we get packaging work done in parallel with other tasks. Going to be challenging.
- Chris: Much of the work on packing thus far has been less interested in ease of use/GUIs. Feel that adding that to the first set of deliverables may not be realistic.
- Jim: TIER more concerned with mapping out parts/deliverables of a greater whole. It could be that the GUI pieces will need to be integrated into future phases/releases.
- Janemarie: Will IAM maturity levels figure into release schedule/strategy?
- Jim: Focus of deliverable #1 was to think ahead on who to engage with on downstream steps.
- Chris: Potential missing is callout for what to do about maint releases after initial install. That would be a bigger issue than packing question.
- Jim: That's part of the "upgrade" language.

- Chris: Should we declare a scope of what's in/out?
- Scotty: We've been working on containerization of data and configs. Possible approach here? Will Grouper take similar approach?
- ChrisH: Not planned at this time.
- Scott: Additional point re: upgrades - need to make those as painless as possible or people won't upgrade.
- Paul: Would it be feasible for first release to allow for config backups?
- Jim: Should move on - this call is intended more as organizational/operational overview. **AI: Jim to update charter to reflect concerns voiced on the above topics (was covered in the problem statement - added some language to last Mission bullet - 20151109)**
- Chris: What platform is this intended for? Is that mentioned in the the deliverables?
- Jim: Assumption is we'll support whatever platforms are compatible with the underlying tech.
- Scott: Would be easier to work at the protocol level, especially considering stated timeline.
- Jim: Gets to TIER mission of modularity of deployment - don't have to run full suite.
- Chris: Consider bootstrapping of organizations into tech that's new to them. Should this be added to deliverables? Also out of scope but should we consider migration policy/recommendations?
- Jim: Good thing to consider, probably best parking lotted

2) Groups and subgroups

1) Shibboleth (Janemarie, Steve Carmody, kevin foote)
  a) Goal - trivial to configure, deploy, operate, and maintain.
  b) What are existing pain points in this space? Improvement suggestions?
  c) What additional tools are needed, if any, for configuration, operation, upgrade?
  d) What needs to be packaged, defaults - authn iface, etc.
    - IdP First
    - SP Second

2) Grouper (Brett Bieber)
  a) Goal - trivial to configure, deploy, operate, and maintain.
  b) What are existing pain points in this space? Improvement suggestions?
  c) What needs to be packaged, defaults
  b) What additional tools are needed, if any, for configuration, operation, upgrade?

3) COmanage (Scott Koranda skoranda@sphericalcowgroup.com, Steve Carmody)
  a) Goal - trivial to configure, deploy, operate, and maintain.
  b) What are existing pain points in this space? Improvement suggestions?
  c) What needs to be packaged, defaults
  d) What additional tools are needed, if any, for configuration, operation, upgrade?

4) Core packaging (Scotty, Brett Bieber, Brian Savage)
  a) Goal - trivial to configure, deploy, operate, and maintain.
  b) What is needed from the component groups?
  c) What can be done before (b) is ready?

- Jim: What is a good group/sub-group strategy, especially with regard to interviewing schools/adoptees?
- Scott K: COmanage is substantially different from Grouper and Shib IdP. For example, it has not been deployed by campuses, only research and professional organizations. Still, a COmanage subgroup makes sense. I will represent COmanage team. May not need to consume as much bandwidth as other groups since it's not as widely deployed.
- ScottC: For Shibboleth consider that IdPs and SPs are very different animals. SP part of a larger environment/infrastructure whereas IdPs are more self contained
- Jim: Good point. We'll most likely focus first on IdP
- Chris: Question around how complete a given installer should be. Does every option need to fit inside a single VM?
- Jim: Nope.
- Scotty: Should we interview first, determine needs, and then build around that?
- Jim: Question is whether it's efficient to discover that as a whole group? Or should we parallelize that effort?
- Steve Carmody: Should we consult InCommon as to org types that they foresee using TIER?
- Jim: What I was trying to get it was breaking into 4 groups to achieve deliverable #1/user research. Question is should we split into groups now or after that initial research?
- Steve Carmody: Initially we have to make a guess at what's needed by these orgs.
- Janemarie: Alternate IDP WG generated list of potential institutions that could be used to better understand who might use the products of this team
- Jim: We have to get to the point of understanding the needs for each of the components of the working group. Concerned that trying to develop totality of of understanding with the entire group might get unwieldy, hence suggestion for sub-groups.

3) Work schedule, meeting times

4) Next steps.

Jim: Can we get volunteers to come up with survey questions for each sub-group?
AI for all: Please add your name next to the relevant group in this document. Add your email address as well.

Survey Prep Work

Demographic Information

1. Responder: Name, email address (2x), etc.
2. Institution Name
3. Faculty & Staff count
4. Undergraduate Student Count
5. Graduate Student Count
6. Carnegie classification (with lookup link http://carnegieclassifications.iu.edu)
7. Are or do you want to provide alumni with services that depend on your core IdM solution (e.g., email for life, crm, transcript, etc., etc). If so, how many person objects would this add to the counts above?
8. Are or do you want to provide "friends of the university" with services that depend on your core IdM solution (e.g., email, crm, etc., etc). If so, how many person records would this add to the counts above?
9. Approximately how many staff support the technical aspects of your Identity and Access Management program?
10. From this list, please indicate the greatest impediments to adoption of new software or services:
    a. executive buy-in
    b. readiness of existing identity data
    c. lack of local technical expertise/staff
    d. incompatibility with our environment
    e. lack of support
    f. other…
11. Were the people responsible for both the technical and policy aspects of this survey consulted before answers were provided? (Yes/No)
12. Are there issues on your campus relative to TIER that are not directly related to packaging Shibboleth, COmanage, or Grouper? If so, please explain (text box).

Core Packaging

1. Does your institution have a central directory (LDAP/X.500)? If so, what software do you run? Check all that apply:
    a. Microsoft Active Directory Domain Services
    b. Microsoft Active Directory LDAP Directory Services (AD-LDS)
    c. OpenLDAP
    d. RedHat Directory Server/Fedora 389
    e. ForgeRock OpenDJ
    f. Oracle Directory Server EE
    g. Oracle Unified Directory

    h. Apache Directory Server
    i. IBM Tivoli Directory Server
    j. Novell eDirectory
    k. Other (fill in the blank)
2. Does your institution provide central authentication for users using a single sign-on (SSO) service? If so, please describe that service (add check boxes and “other”)
    a. CAS
    b. Shibboleth Identity Provider
    c. CoSign
    d. Pubcookie
    e. SimpleSAMLphp
    f. Microsoft Active Directory Federation Services
    g. CA SiteMinder
    h. Oracle Enterprise Single Sign-On
    i. IBM Tivoli Security Access Manager for Enterprise Single Sign-On
    j. Other (fill in the blank)

3. What deployment model does your institution prefer for running its IdM infrastructure? Check all that apply: (what do you want today and what will you want in a year or two?)
    a. Physical servers
    b. Virtual Machines
    c. Virtual Appliance
    d. Docker Containers
    e. Other (fill in the blank)

4. Which operating systems are supported by your school for production use in server environments? Check all that apply: (what do you have today and what will you have in a year or two?)
    a. Windows
    b. Linux (do you have a preference about the Linux distro?)
    c. other (*BSD?)
5. Level of comfort building/managing Java servlet environments 1 - 10 where 1 is little confidence and 10 is very confident (e.g., we run this type of environment in production)?
6. I want my IdM solution in the cloud so that all of my services will continue to run when my campus has issues (Yes/No)?
7. I want my IdM solution local. I have policy or operational needs that require a local solution (Yes/No)?
8. We are cloud/local agnostic and need to be able to place components either locally and/or in the cloud (yes/no)
9. If the following components were available from TIER as a packaged solution, which order would you like to deploy/adopt them in? (please rank)
   Shib IdP, COmanage, Grouper, Shib SP

10. What automated user and group provisioning tools do you rely upon, if any? [list of tools, plus text box (I would recommend to just leave this as a text box to allow free flow reply.)]
11. Would you be willing and able to allow access to your central authentication and directory services from a cloud service? {yes/no}
12. For your TIER implementation, do you have a strong preference for:
    a. local installation using existing packaging solution
    b. local installation from pre-configured containers (Docker, etc)
    c. a preconfigured VM/appliance that you run and maintain locally (“local” virtual appliance)
    d. a virtual appliance managed remotely
    e. a cloud-based virtual infrastructure (IaaS)
    f. a managed cloud-based service (SaaS)?
    g. Other (fill in the blank)
13. Would you like to see official support for running behind a load balancer (x-forwarded-for headers; include in logs / container config? including configuration documentation for a common load balancer (e.g., F5). {yes/no}
14. Would you like to see official support for offloading SSL to a load balancer? {yes/no}
    a. For which components?
        i. Shibboleth IdP
        ii. Grouper
        iii. CoManage
        iv. Shibboleth SP
15. Do you expect to operate the TIER apps in a high availability (HA) environment that includes more than one active node?
    a. I don’t care about HA
    b. I want HA for authentication only
    c. I want HA for authentication/authorization
    d. I want HA for authentication, authorization, and read-only data sources
    e. I want HA for authentication, authorization, and live data sources

Shibboleth

https://docs.google.com/a/lafayette.edu/document/d/1-noQWM0ur0A-pIXVZBJT8Gchfe3rZmROMYTLG956ag8

Shibboleth is the defacto SAML-based federation software standard for research and higher education institutions in the InCommon Federation. But while it is robust and reliable, configuring and managing it takes skills and resources that may not be readily available to many institutions. One of the goals of the TIER Packaging Working Group is to make the Shibboleth Identity Provider (IdP) software a pleasure to configure, deploy, operate, and maintain for a campus environment with standard minimal requirements. We are working on identifying the tasks that are required to manage it, what default capabilities might be part of a packaged Shibboleth IdP, and what features would make it a good citizen of the InCommon Federation. To help TIER, please fill out the survey below and and share your experiences with installing, configuring, and managing a Shibboleth IdP at your institution. If your institution decided against running a local Shibboleth IdP, we would like to know what led to that decision.

1. Institutional Profile (**move to demographics**)
    a. Carnegie classification (go to http://carnegieclassifications.iu.edu/ for a list of schools and click on “InstitutionLookup”)
    b. How many undergraduates does your institution have (checkboxes for several ranges)
    c. Total size of community (students, faculty, staff; i.e. how many user objects in LDAP)
2. Institution IT Profile (**move to demographics**)
    a. Primary server platform (checkboxes - Linux, Windows; do we care at this stage about Linux variants?)
    b. Central Directory service
        i. LDAP (eg OpenLDAP, Fedora 389 Directory Server, other) (I don’t think we care about the implementation)
        ii. Active Directory
        iii. If you do not have a central directory (LDAP/AD), what is your ERP solution (e.g., Banner, Peoplesoft, etc.)?
    c. Does your IT staff install and manage Java servlets (Tomcat/Jetty, example apps)
    d. VMWare image
        i. Amazon EC-Does your site use virtual environments or containerized packaging (check all that apply)
        ii. 2 service (AWS)
        iii. Docker
        iv. Other (Puppet?, Vagrant; not suitable for PROD deploy?; Google Cloud?)

e. Do you have an LDAP user attribute indicating that the person has opt'ed-out under FERPA?
f. Would your site be willing and able to allow access to your central authentication and directory services from a cloud service?

3. InCommon and Federation
    a. Is your campus a member of InCommon ?
    b. If yes
        i. Does your campus use the Certificate Service (yes/no)
        ii. Is your institution:
            1. Not running an IdP
            2. Currently managing a Shibboleth IdP (yes/no)
            3. Using an outsourced Shibboleth IdP solution
            4. Using an IdP implemented from some other source (e.g., ADFS, SimpleSAMLphp).
                a. If yes, which do you use and why? (I don't think that there are that many. Do we really need to know what they are using?)
    c. If you are not using any IdP, are there reasons your campus has not deployed an IdP?
        i. InCommon Value Proposition does not appear sufficient
        ii. The difficulty of deploying and managing an IDP is perceived as too costly
            1. No local experience with servlet containers/Java servlets
            2. Current Shibboleth documentation does not meet local needs
            3. Have you investigated any of the implementation partners, or outsourcing options ?
        iii. Cost of joining InCommon
        iv. An appropriate support program for institutions that do not have the required technical skills seems to be lacking
        v. Support in navigating the InCommon processes required for operating an IdP within the federation appears to be lacking
        vi. Other (briefly describe)

Shibboleth IdP

1. Would your site prefer: (indicate your order of preference) (**<-- move to core**)

a. The current Shibboleth packaging
b. A virtual container containing a fully operational Shibboleth IDP. The Container would include: an operating System, servlet container, the Shibboleth IDP, a DB, and GUI tools to manage the configuration.
c. An installer capable of installing and configuring everything needed to produce a functional IDP (servlet container, IDP, SQL DB, and GUI tools to manage the configuration); the site would be responsible for the underlying OS.
d. Use of a fully functional IDP operating as a cloud-based service (hosted model).
    i. How much would your site VALUE this service (checkboxes)

2. Would your site be concerned if this IDP were pre-configured to: **(KEEP)**
    a. Load and use InCommon metadata,
    b. Include support for a default set of attribute definitions (LDAP - name, email; eduPerson - EPPN, Affiliation, primaryAffiliation, ?)
    c. Release EPPN, name, email, affiliation, eduPersonTargetedId to SPs with the Research and Scholarship R&S entity category? (removed “only to IC member SPs”) (documentation for sites to opt-out if needed ?)
    d. Respect a FERPA opt-out attribute to restrict attribute release for some users.
    e. Prompt users to consent to attribute release ?
        i. Relying parties :
            1. all ?
            2. some ?
                a. which ?
        ii. Attributes :
            1. all ?
            2. some ? (i.e. per-attribute consent)
                a. which ?
    f. Avoid spurious errors in the logs from external scanners via a properly configured robots.txt
    g. Support Enhanced Client or Proxy (ECP) by default ?
    h. Support multi-factor authentication by default ?
        i. DUO (Support for multiple authentication contexts/Multi Factor Authentication) ?
    i. Automatically reload config files when they are changed (relying-party.xml, attribute-filter.xml, attribute-resolver.xml)
    j. Support consent and logout using :
        i. Client storage ?

            1. HTML5 Local Storage ?
            2. cookie storage only ? (the default)
        ii. Server storage ?
            1. database ?
            2. memcache ?
    k. Support CAS by default ?
    l. Support OpenID Connect by default ? (not available as of 3.2)
    m. NOT support SAML 1?
    n. NOT support SAML Attribute Queries?
    o. Update itself automatically ?
        i. Security updates only ?
3. This IDP was accompanied by a configuration GUI that allowed:**(KEEP)**
    a. Selectable AuthN integrations (choice of AD, LDAP (over TLS) with username/password, or LDAP with SASL/GSSAPI), etc)
    b. Enable mapping and configuration of eduPersonUniqueId.
    c. Specify whether the anonymous relying party is completely untrusted, or somewhat trusted.
4. Prioritize the following primary functionality:**(KEEP)**
    a. A Configuration Management GUI that that would address 80% of the use cases. (For the other 20%, configuration files would have to be hand-edited.)
    b. A Management GUI to manage local federation metadata.
    c. Non-XML configuration files that could be compiled into the IDP's XML files.
    d. A Management GUI that could obtain a list of user attributes in the local LDAP server, and generate IDP config elements for those attributes.
    e. Configuration support for operating behind a load balancer (x-forwarded-for; include in logs / container config? including configuration documentation for a common load balancer (e.g., F-5))
    f. (remove this -- too general) Tooling in general - what are there day-to-day operational challenges that could be made easier with automation or additional tools.
5. Prioritize the following advanced functionality:**(KEEP)**
    a. Log analysis reports for service management (e.g., top 10 SPs, etc.)
    b. Enable a configuration to notify IdP admins of errors occurring on the idp to increase operational awareness. (replace the above)
    c. Operate a Shibboleth IDP in a high availability (HA) mode that includes more than one active IDP node

    d. x509 certificate management (self signed for federation, and then commercial cert installation) would be helpful. It is actually a large chunk of cross platform challenges due to underlying openssl/NSS layer
    e. Attribute Release Support for “common” relying parties such as Internet2 services, Net+ providers, and common cloud providers, (this may need IC to maintain a “registry”)
    f. Provide a mechanism to share the DataSealer secret key between IdP nodes.

Grouper

1. What systems do you hope to have integrate with Grouper? This potentially includes feed group membership information into Grouper or provisioning these systems from Grouper. (check any that apply)
    a. Active Directory
    b. Banner
    c. Canvas
    d. Google Apps
    e. LDAP
    f. Moodle
    g. Office 365
    h. PeopleSoft
    i. RDBMS
    j. Sakai
    k. Other (text box)
2. What kind of subject source would you like to connect to?
    a. LDAP
    b. SQL
    c. Other {text box}
3. Have you installed Grouper before? (if yes, rate 1-5)
    a. (if yes installed) Describe your installation experience, how did you install Grouper? {text box}
    b. had you watched the Grouper training videos? {yes/no}
    c. (if yes installed) what did you like about the installation procedure? {text box}
    d. (if yes installed) what could be improved? {text box}
4. Have you ever patched Grouper? (rate 1-5 plus text boxes)
    a. (if yes patched) what did you like about the patching process? {text box}
    b. (if yes patched) what could be improved in the patching process? {text box}
5. Have you ever upgraded Grouper? (rate 1-5 plus text boxes) describe the process you used
    a. (if yes upgraded) what did you like about the upgrade process? {text box}

    b. (if yes upgraded) what could be improved in the upgrade process? {text box}
6. If you run Grouper at your institution, are there runtime tasks (maintenance, troubleshooting, configuring, etc) that are tedious or time consuming or difficult that you would like improved?
7. Do you want more features in the Grouper installer? (rate installer 1-5)
    a. If so, please describe {text box}
8. Any other suggestions for installing, configuring, upgrading, patching, or operating Grouper? {text box}
9. Which environments do you use or would you expect to use? prod? test? dev? performance? training? other?
10. The Grouper UI requires authentication. Would you be concerned if the TIER packaged version of Grouper was packaged with the Shibboleth Service Provider?

COmanage

todo: structure like grouper questions
maybe a short blurb to explain what comanage is, why you'd want to run it

(Scott Koranda,)

Background information for packaging workgroup consideration

- COmanage documentation
- COmanage technical manual
- COmanage installation details (current)
- PHP application
- COmanage is multi-tenant and designed to support multiple Collaborative Organizations (COs). Sometimes COs are "large" such as a large international astrophysics project that may run its own COmanage instance (i.e., LIGO) and sometimes they are "small" such as a group of a couple of researchers.
- Minimal configuration is done on the command line. At the initial time of deployment to onboard the first user (known as the initial platform administrator) one types into a command line the login identifier (usually ePPN), given name, and family name. After that all configuration is done through the web application itself.
- The primary packaging-related issue for deployers that we have seen until now is that most deployments need or want a number of related tools or services to fully leverage the collaboration platform. They struggle most with deploying and configuring the related tools as opposed to COmanage itself. The tools/services include:
    - A SAML SP of some type is needed since COmanage only consumes federated identity. We most often deploy the Shibboleth Native SP (Shib SP) for Apache HTTP Server (Apache). COmanage itself is agnostic since it simply looks at $REMOTE_USER and other Apache environment variables as configured so

theoretically any SSO/SAML tool could be used but so far the leading choice is the Shib SP.
    - COmanage can provision to an LDAP directory and most deployers want this for easy integration with certains kinds of applications. The usual choice is OpenLDAP but again COmanage is agnostic about the specific flavor of LDAP directory server. Sometimes the directory server runs on the same host and other times it does not.
    - Most deployers find they want to deploy a SAML attribute authority (AA) so that SAML SPs can consume ePPN from and IdP and then use it to query the AA to obtain CO-specific attributes about the user including group memberships. Since COmanage is multi-tenant this really means (usually) multiple AAs. Until now we have used the Shibboleth IdP configured as an AA but since it is not inherently multi-tenant there is interest in leveraging a different tool, probably based on the pySAML2 codebase and developed by SUNET/SWAMID (Roland Hedberg).
    - Many deployers, especially larger deployments, find they also want a SAML-to-SAML proxy or gateway. A common use case is to present a large number of IdPs as a single IdP to a commercial SP that can only consume SAML metadata for a single IdP. The proxy is configured to receive the ePPN from the home organization IdP, then query COmanage for CO-specific attributes about the user include group memberships and then assert them to the SP. At this time we are using primarily the SaToSa codebase, also developed by SUNET/SWAMID.

Draft survey questions

COmanage Registry manages federated identities for collaborative organizations (COs). A CO can be as simple as a few researchers working together or as complex as a large international science project with thousands of researchers across many countries. COmanage Registry helps enroll and onboard CO members using their federated identities, manage groups, identifiers, attributes, SSH keys, and provision person and group data to applications.

1. Is your institution presently running COmanage?
    a. Never heard of it.
    b. Do not understand what COmanage could do for my organization.
    c. Have not identified inter-organizational collaborative use cases.
    d. Have not identified need for research virtual organizations at my organization.
    e. Lack of time or expertise with COmanage.
    f. Other.
2. COmanage requires a service provider to consume federated identity asserted by an identity provider or login server. Would you be concerned if the TIER packaged version of COmanage was packaged with the Shibboleth Service Provider?

3. COmanage is most often deployed with other components. These components can be packaged with COmanage to result in a completely self-contained deployment or existing campus components could be used. The COmanage team recommends the use of separate instances of the components unless there are clear reasons to do otherwise. For each component, would you use the provided one to maintain a clean separation between COmanage and your existing systems, or use your existing deployment of the given component instead?
    a. Shibboleth Service Provider (for consuming federated identity)
    b. LDAP (for exposing person and group data to other applications)
    c. Grouper (for finer grained group management that requires set math)
    d. SAML Attribute Authority (for exposing person and group data via SAML2)
    e. SAML IdP/SP proxy (for assisting with attribute release by IdPs)
4. Upon deploying COmanage, what additional applications would you consider integrating with?
    a. Wikis (eg. Confluence, Dokuwiki, Foswiki, Moin)
    b. Mailing list servers (eg. Sympa, Mailman3)
    c. Calendaring and event invitations (eg. Bedework)
    d. Conferencing (eg. BigBlueButton)
    e. Google Apps for Education
    f. SSH servers
    g. other

Redundant now, probably delete:

5. COmanage Registry can provision person data and group memberships for a collaborative organization (CO) to an LDAP directory for easy consumption by some applications. A common deployment pattern is to deploy an LDAP directory specifically for use with COmanage, but an existing LDAP directory may also be used. A directory used specifically with COmanage keeps a clean separation between core enterprise business functions and supporting COs but requires extra operational support. Would you be concerned if the TIER packaged version of COmanage defaults to a packaged separate LDAP directory?
6. COmanage Registry can provision memberships in collaborative organizations (COs) and CO groups to Grouper to accommodate more complex group math and integration with service authorization controls. A common deployment pattern is to deploy a Grouper instance specifically for use with COmanage, but an existing enterprise Grouper deployment may also be used. Grouper used specifically with COmanage keeps a clean separation between core enterprise business functions and supporting COs but requires extra operational support. Would your institution be concerned if COmanage comes packaged using a separate Grouper deployment for use only with COmanage?
7. A common deployment pattern for supporting COs using COmanage Registry is to deploy a SAML2 attribute authority so that SAML2 service providers can query for user attributes managed by the CO. This allows applications protected by different Service

Providers to retrieve attribute data for a person based on a shared identifier. The Shibboleth Identity Provider (Shibboleth IdP) may be used as the attribute authority. Would your iIs your institution prefer a separate Shibboleth IdP deployment for use only with COmanage Registry as a SAML2 attribute authority or to leverage an existing enterprise Shibboleth IdP deployment used to provide federated authentication for your institutional users?

Things to pass on to other teams

Shibboleth IdP

Improve error messages; add unique ids to each error shown to users, and include that id in the error log.

Add error codes for each error message; collect stats from community to improve documentation on how to diagnose each error.