Fri 10.00am Freedom Room

Scribing Template --Friday, Oct 5, 2012 at 10am

Freedom E Room

TOPIC: Provisioning and Integration use cases & provisioning connectors (Keith H, Rob Carter, Boyd Wilson)

CONVENER: Keith Hazelton / Boyd Wilson

SCRIBE:Rob Carter

# of ATTENDEES: 31

MAIN ISSUES DISCUSSED

- Keith: Rob is heading up CIFER Provisioning and Integration and I'm helping out... General sense of provisioning (getting account information into systems) and getting data integrated... Rob is going to make a pitch for his web site.

    - Rob : pitch

- Don Faulkner:Arkansa

- Foreign students who've completed their study but are still in country. From the Homeland Security perspective, Arkansas owns them, but isn't allowed to own them. Need more buckets to stash them in...

- Warren Curry: Florida

- We've looked at Grouper, but don't have it quite yet... Problem is interesting... Starts with 100,000 people who apply fertilizer or pesticides to land, and need to be licensed. Licensing requires testing and courses taught by the Extension Service. These folks don't start with a UF identity, and then need to get one, get into Sakai, get into the right course in Sakai, *pay fees*. All needs to be done in realtime (while the user waits) -- they need to start their educational engagement immediately. Spans both the financial system and the student / LMS system. Really difficult integration case. Hardest part of the integration was the testing facility -- Sakai testing was separate from identity testing facilit, etc... Trying to create a mechanism for distance ed folks to actually do end-to-end testing was really difficult.

- Tom Barton: Chicago

- Finally getting rid of SSN use in some spaces... Rather than converting all core systems at the source, for various reasons, the plan is to erase or replace the SSNs on the way from the incumbent system to a new system. That boils down to being able to integrate one or more appropriate tools (in our case, our IDM system, where we can crosswalk across systems) to project that mapping capability.   Sometimes there's a batch feed we can intercept and remove SSN data from... others we have to do other things. There are limited points of interaction where we can intercept and transform. Point is -- it's important to have the right set of integration "handles" -- points and interfaces.

- Keith H: Wisconsin

- IN particular - if you're using Penn State's CPR or OpenRegistry, do you have integration stories there? What are some of the targets that information in the CPR needs to get to?

- Jimmy Vucculo : Penn State

- Directories, Kerberos servers, Consumers of data... We made it flexible enough with internal messaging that anything can listen for and process a message. We currently do directory, file service, and calendaring integration, but other things can simply listen for events and process them out of the message stream.

- Keith H

- Scotty Logan had a good story about box.net integration... I'll try to repeat it... The problems we're having, I start to believe, is that their business model from day 1 was personal sharing of data...personal interaction with the cloud. That reminds me of Microsoft in the early 90's -- it was your "personal computer" -- why on earth would you *authenticate* to your *personal* computer? I think we're starting the same journey with cloud vendors -- moving from personal interaction to institutional interaction and from individual use to collaboration with their services. The problem is getting groups to the functionally correct things for folks who want to share things at Stanford (in Scotty's case). The canonical solution is that every object has an owner... If the owner leaves, so does the content, so create faux identities with nullified passwords to own content that needs to persist. The point is that so much of the problem isn't solvable with good glue -- it's about understanding the cloud provider's model and figuring out how to represent your vision for collaboration in the vendor's model. If we can share those stories and solutions, we can at least may help each other...

- Peter : Brown

- Just put Workday into production to replace a home grown HR system. We discovered that wWorkday's organizational principal is supervisors, not departments. Getting group information to be 100% correct is difficult. If a departemnt hires a student and they're not on FA, they show up in the department, but if they're on FA, thier supervisor is in the FA office...

- Rob: SAP story from Duke is similar

- Keith: Perhaps the issue is actually mapping between otherwise non-commensurate logical data models.

- Warren: We built a computation of "primary department" as it were for the IDM -- provisioning gets done by one that's chosen from the multiple departemnts -- we built a specific tool for it.

- Boyd Wilson: We take the state of the data from PeopleSoft, Banner, wherever, and maintain both the original state data and the transformed data -- keeping the original state from the source systems becomes critical in certain mapping cases.

- Warren: Maintaining specific detailed data behind the scenes in the registry is important to deal with aggregation cases for, eg., faculty with multiple appointments. "Primacy" is based on FTEs, other factors -- where it gets interesting is where there's no clear "winner"....

- Boyd: Clemson

We have connectors... for various target systems... Some Java, some other... Blackboard, Banner, box.net, Exchange,

What I'm really interested in -- how do you connect to the connectors?   Who's interested (lots of folks :-)

Drawing of architecture... (cf. attached image)

- Connectors are designed to be event driven -- largely CRUD events

- Directory/Registry/Vault (Clemson term) + downstream systems... Windows service, written in C, that takes input events and calls the appropriate scripts -- PowerShell on Windows, Perl/Python/Shell on Linux systems... Dispatcher/driver then calls the appropriate scripts... Scripts are separated from the dispatcher... Works across platforms with different scripting languages... But how to implement the map from the vault to the dispatcher?

- SPML? SCIM? Something else? What would people like?

- Deltas only -- not full data... There's an interface for sending a snapshot of full events... When changes occur in the registry, deltas are propagated to the dispatcher...

- Tom B: For each target instance there is a script?

- For each target there is a driver and one or more scripts -- driver runs as a service/daemon...

- Brown: ActiveMQueue message bus is our solution to that interface.

- Boyd: How do you deal with downed systems?

- Pete: MQueue accumulates messages -- the target has to retrieve from the queue

- Chris Phillips: Does password change transit the queue?

- Boyd: There's actually a pull script option that can allow a script to pull data or push data back to the vault

- Chris: Password changes are hard to resolve -- people typically neuter the end points (ctrl-alt-insert is disabled).

- Boyd: We've disabled endpoint password changes, -- makes it easier to force policies more complex than the endpoints support, but...

- Chris: Homegrown?

- Boyd: We're actually using NetIQ as the vault -- from Novell -- it marshalls and pushes events...

- Rob: How do synthetic transactions (that don't map one-to-one from registry to target) get handled?

- Boyd: We have policy that can combine or delay delivery of events ...

- Tom B: The way I used to look at this... there is explicitly an FSM that needs to be operative somewhere (at that point or upstream of it) that will take into account ... much automated provisioning can be accounted for by assigning them as outcomes of transitions from one state to another in an account... Business logic based on transitions between states rather than on transactions. As well as where should changes to the provisioned footprint be made... Seems a good model for discussions with non-programmers... Proposal: Provisioning interface from registry to target needs to be stateful.

- Keith: Much of the art is in defining an event -- we don't want events at the level of field transitions (usually)

- Tom: Need an ontological or conceptual framework for those transitions. Everyone who does this is at least familiar with the problem

- Boyd: What's interesting about that is ... can we solve it in pieces... Much of what we're talking about has to happen upstream. Needs to be up in the registry end of the transaction rather than the target connectors. It's amazing though how much back and forth can be required there.

- Tom: Doesn't though apply that you have to operate your registry in that way... Does CPR support that sort of state model?

- Jimmy: Not really... It's highly configurable...

- Keith: State is an IDM thing -- a downstream system has a state -- a transition is a change of state...

- Warren: States can change in ways that make existing provisioning results nonsensical -- users who have no attribute information in an AA after they lose their credentials, for example

- Chris P: There are some commercial pieces here... There's a lot of desire for open source tools... Other groups have talked about BPML... Is there a Net+ partner to seek out for that?

- Warren: We use BizTalk for that...

- Keith: You're claming that if we had a BPML or BPMN engine we could design the FSMs Tom is talking about there?

- Chris: I'm thinking that inside Grouper there's got to be some form of that... What's the happiness level about how dynamic or static that is?

- TomB: I have wanted that from the beginning in Grouper, but we've never developed it beyond having operational end-times, really -- the stubs are there to build it in, but it's just not happened. In my perfect world, we'd all have these things, and Grouper would consume them, but while it's ready to consume, it's not happened...

- Chris P: One thing that we put a lot of effort into is creating those FSMs... Something more standard could help with that...

- Tom: Going back to Grouper ... A narrow option is to put it in the endpoint... more broadly, it could be a good repository for the provisioning process -- a way of representing some of the states and transitions (BPML).... Could distribute management of the rules at that point... What you want to get to is a central infrastructure that has some business rules, but generally, if they're important and meaningful, you want to be able to delegate...

- Boyd: Target connectors *can* query back for registry state.

- Keith: Down to 3 minutes...

.

-

ACTIVITIES GOING FORWARD / NEXT STEPS

- Add use cases to the CIFER wiki P&I use case pages

https://spaces.at.internet2.edu/display/cifer/Provisioning+and+Integration+Use+Cases+-+%28User%2CHorror%29+Stories

- Rob will nag folks who noted cases today about getting their content into the wiki

- Provisioning team in CIFER and/or Mace-PACCMAN to get Boyd in to do a version of this talk and hammer on it some more

- Clemson is willing to share code -- not necessarily ready yet, though, given absence of a standard interface from registry....Work with Boyd on that...

If slides are used in the session, please ask presenters to convert their slides to PDF and email them to acamp-info@incommon.org

Thank you!