Minutes of the 9-02-2011 concall

OSIdM4HE-prov concall, 9/2/2011, 16:00 - 17:00 ET

- **Attending:** Rob Carter, Keith Hazelton, Lucas Rockwell,Tom Zeller
- **Minutes of the prior call:** Attendees will review and forward Rob any changes. Meantime, notes are posted to the team workspace on spaces.at.internet2.edu.
- **Agenda bash:** Nothing of significance
- **Action items:**
  - Lucas' diagram added to the wiki - text to be added later
  - Rob added a higher-level interaction diagram to the wiki (slated for brief discussion later in the agenda)
  - Keith's AI to model provisioning using the enterprise integration book's patterns will carry forward
- **New Business**
  - **Paccman Discussion**
    - Rob noted that there had been some discussion of the OSIdM4HE effort during the biweekly MACE-paccman call this week. The discussion was largely a rehash of the framing conversation from Chicago, with most of the interest from the group being focused on potential timelines and what the expected deliverables from the effort might be – whether the effort is focusing on making recommendations to higher ed, producing a "cookbook" for constructing an OS IdM solution on your campus, actually producing an OS IdM "stack", or acquiring funding for a separate project to do that. The gestalt was that the goal is somewhere between the latter two, which seemed to be relatively well-received by pacc-people.
  - **New Diagrams in the Wiki**
    - Rob pointed out the high-level diagram at the top of the rough architecture diagrams table, noting that it's mostly a stake in the ground toward trying to characterize where the different parts of an IdM stack "fit" in the overall ecosystem, and what responsibilities logically fall to which "blocks" in the diagram. He noted that this characterizes only one possible division of labor between the units.
    - Tom noted that provisioning is frequently presented (as is the case in the diagrams and discussions so far within the group) as a sort of "black box", without much detail around exactly what goes on within the provisioning process itself. He suggested that he would like to take an AI to produce a "next level down" analysis of provisioning, and then diagram that more detailed view.
    - Lucas agreed and noted that the top two diagrams in the wiki page for architectural diagrams constitute something of a 50,000- and 20,000-foot view of the space, and that Tom's suggested AI could provide something like a 5,000-foot view of just the provisioning process itself.
    - Lucas also noted his continued belief that with the right kind of loosely-coupled design and the right pieces (some of which, he noted, could just as well be off-the shelf open source or even commercial products) it should be possible for provisioning to be logically independent of the workings of the registries from which data and/or "changes" are consumed. Some standard or set of standard interfaces would have to be provided at one or both ends of the conversation, but the details of the implementation of the two ends could be radically different and completely independent of one another.
    - Tom mentioned that he had recently had a discussion with Chad L. about provisioning data changes and updates, and came to the conclusion with Chad that it may be too early in the lifecycle of the available solutions to pick a specific protocol to use for managing change provisioning. SPML, SCIM, and other protocols are still not entirely set in their approach to change provisioning. They both have models that presuppose changes in some form, but the details can be sketchy and don't always map well onto one another. For this group, he suggested, a good outcome might be to produce a high-level API for representing changes, since changes seem like critical atomics for provisioning systems to manipulate. Something describing what the salient features of a "change" are and what its content is like could be very helpful.
    - Rob agreed, suggesting that there might be different levels of representation, as well, and that a single provisioning facility might want or need to consume changes in one representation but present them to other systems in a different representation. He proposed that change representations might run a gamut from very simple notifications (eg., "something about entry X changed") to very detailed change structure (eg., "at time T value V of attribute A on resource R was changed by agent Ag from value V1 to value V2").
    - Tom noted that SPML stipulates a change representation as part of its definition, and that SCIM has a model based on what it terms "patches". He indicated that the "patch" model is somewhat opaque to him, and seems rather odd but may bear more investigation.
    - Rob agreed that the patch model sounds strange.
    - Tom pointed out that timing may be advantageous for this group to provide some feedback to the SCIM group about its change model.
    - **At this point, Keith joined the call**
    - Keith asked if Tom might have ideas about how best to provide the SCIM working group that feedback?
    - Tom suggested that if this group were to construct a picture of what it thinks a reasonable change model might look like and then compare it to the SCIM "patch" model, he could take suggestions about the differences back to the SCIM working group and propose changes to their model.
    - Rob asked if there might be an AI in that....
    - Tom proposed that the group might want to maintain a list of topics of interest and items we'd like to produce, and place this on that list, since it's probably not something we'd consider a deliverable for the initial check-in with the committee of the whole in mid-September.
    - Keith noted that in his work with the Bamboo project, he's become aware of a need for a group provisioning strategy. The Bamboo group requires that users of its collaboration platform be able to use whatever native privilege management systems they already have in place (or embedded in their chosen applications – AlFresco, HubZero, Joomla) but needs a mechanism for synchronizing group information between those possibly disparate group stores and the centralized Bamboo Grouper instance (and vice-versa) so that group information can flow between applications within the platform. He asked if that's an already solved problem or if a well-described solution already exists that Bamboo should consider, or whether it's something that this group should look at as a provisioning problem (essentially, a change event generated by a group modification could poke some RESTful API, etc., etc.). He wondered whether that might be something that SCIM could address?
    - Rob agreed that the case is an interesting one and pointed out that in the 50,000-foot diagram in the wiki, the point of the purple bidirectional arrow between the privilege management block and the provisioning block is to indicate that under some models, groups and privileges could be provisionable objects which could both be provisioned into the privilege management mechanism based on registry information and/or provisioned into other consuming systems based on information gleaned from the privileging facility. He noted that the intent isn't to proscribe either approach, but to indicate that either approach might be considered by some viable provisioning and privileging facilities.

- Keith noted that Chris from CANary has apparently expressed interest in seeing a standardized interface definition for provisioning (with a predisposition toward SCIM as a likely candidate) so perhaps discussing a kind of standard change representation and perhaps a set of standard operations with respect to changes may not be too far off the reservation.
- Tom indicated that the SPML group is trying to come up with something along the lines of a "lite" interface definition similar to what SCIM has, but expressed concern about the idea.
- Lucas asked whether the concern was about the SCIM approach, the SPML approach, or the general idea of a standard?
- Tom explained that he's not entirely sure a standard schema is likely to be all that successful in this space. Schemas are good, he pointed out. for comparing and contrasting different solutions' strategies, but they may not in themselves be solutions to any real problems.
- Lucas pointed out that Keith had posted a reference to an article elucidating the fact that SPML already exists, and SCIM might therefore seem redundant. He pointed out that the "S" in SCIM is short for "Simple" – in the way that HTML is something of a "lightweight" subset of its predecessor (SGML), SCIM might be considered a lightweight subset of SPML. It may be a good start toward a solution, but by itself, out of the box, it's not going to support everything everyone wants from a provisioning interface.
- Tom agreed, pointing out that SCIM (and the problem itself) probably are not that simple.
- Keith pointed out that the complexity of our provisioning needs is largely a result of choices we've long made within higher ed – we do it to ourselves, with the result that we inevitably end up with one-offs and custom-cut solutions, even if they're only designed to cover specific edge cases.
- Tom pointed out that much of the discussion in this area boils down to rehashing marketing-speak.
- Lucas explained that to his way of thinking, something like SCIM that's used across multiple sites, has the advantage (over alternative options involving more custom connectors) that a tool written for one environment stands a chance of being useful in another environment, and that tools written in different environments stand some chance of being able to interoperate together in a third. Perhaps the cloud provisioning mechanisms, even though designed for provisioning between cloud provider and subscribers, are still applicable in our enterprise environments, maybe with proper rethinking of roles. The initial version of SCIM seems to be taking a very simple approach to an API – create a simple identity, update some simple attributes, remove a simple identity. Perhaps ignoring the more complex parts of provisioning isn't an unreasonable way to get started.
- Keith noted that sometimes "simple" is really just a matter of choice – of choosing which parts of a problem to solve first – and doesn't necessarily imply long-term limits on the scope of a solution. He then returned to the earlier question of whether it seems as though provisioning classes of object other than "person" seems reasonable for a provisioning facility in this space, and whether we'd consider it to be in-scope for our OSIdM4HE discussions, and what our opinions on existing approaches (SPML, SCIM) ought to be.
- Lucas reiterated his interest in seeing the group try and advocate one or a few standards, as opposed to advocating a wholesale "customize for every consumer" approach to provisioning, and made the point that the group ought to select something (or a small number of things) and work with vendors and other open source projects to facilitate adoption of the selected standard(s), even up to the point of actually developing plugins for those platforms that seem both significant and unwilling or unable to develop them for themselves.
- Keith agreed, noting that it's a solid point that for the current effort, we're deliberately tasked with considering these issues across platforms and application environments, and that if we can't agree on standards when there's such an obvious win to be had as there is in this situation from the model of pushing standards deep into the core and down at the wire level, while pushing customization as close to the edges of the architecture as possible, there's little hope of our solving anything significant in the space. He expressed hope that we can make a very credible case for the "standard core with customized edges" approach being able to meet virtually any site-specific requirements there might be. The trade-off, he indicated, may be the creation of more unique connectors at the edges of the system (where only small amounts of actual code may be required to support them) in return for a high degree of interoperability and flexibility, and then noted that he'd experienced a "moment of clarity".
- Rob then asked if this might be an opportunity to revisit the question of what our deliverable for mid-September should look like?
- Keith suggested as an option our putting on the table in mid-September (with the committee of the whole) a scope for provisioning that would encompass implementing fundamental Cr/U/D operations between systems with independent stores for persons, privileges, groups, and roles, with an initial emphasis on persons and groups and later expansion to include roles and privileges. He further suggested that we might indicate a preference for using some sort of RESTful or REST-like interface for the services that provide support for these operations, and we might leverage the thinking that's gone into SCIM, but aim to address real-world requirements (rather than theoretical models) by actually implementing SCIM or SCIM-derived interfaces into specific consumer systems (eg., Sakai). He suggested possibly taking a real-world requirement from Bamboo or some actual UCSF task that needs implementing and walk it through a story that addresses it using such a SCIM-ish solution. He added that in such a discussion, we could conceivably talk about various classes of problems, too – bidirectional synchronization of information about person or group objects between participating systems; pushing authoritative data from a central hub to one or more consumers, etc. Such a collection of models might then, he pointed out, create for each consuming system and obvious list of sub-projects that would need to be implemented – the list of interfaces for performing each of the activities in the description.
- Rob raised the question of whether such a deliverable would likely meet the needs of those in the larger group whose goal is to construct a case for funding the OSIdM4HE effort – what it might be like, essentially, to be a Bob Morgan hearing that report from the provisioning subgroup. He asked how such a deliverable (something outlining the interfaces that need to be developed by each provider and each consumer in the provisioning "dance") might be substantively different from what FIFER seems to have been attempting to do?
- Keith explained that while FIFER isn't completely done, it doesn't seem to have committed itself to the whole program, as it were – to issuing a simplified interface definition that multiple systems can implement in their own contexts and then interoperate through. That seems to be the FIFER goal, but lacking some real-world urgency in the FIFER discussions, it wasn't entirely clear whether that would ever actually be realized. On the other hand, he pointed out, perhaps OSIdM4HE could be the catalyst to drive that discussion in the FIFER group to some completion.
- Rob asked if there was, then, perhaps an AI in evaluating the gap between the interface definition we might feel is needed and what FIFER has or is working toward?
- Keith asked who the "right" folks in the FIFER group might be?
- Rob noted that his only real interaction with FIFER had also been via Benn Oshrin, providing him some some background and use case ideas, and that he wasn't really sure who else might be a good contact for discussing FIFER in more detail.
- Keith pointed out that his view of FIFER is from the perspective of a lay person, but that he does work fairly frequently with Benn Oshrin on coManage efforts tied to the Bamboo project. He said he'd be willing to discuss the question of a gap analysis with Benn when next they meet to discuss coManage issues, and perhaps invite him to an upcoming call, noting that with only two weeks or so to go until the next whole-group discussion, such a meeting would need to happen relatively quickly.
- Rob noted that the discussion could probably be fairly brief, if scheduling Benn's time becomes difficult...

        - Keith pointed out that much of what he's seen from FIFER has amounted to the FIFER group going through Chris Hyzer's Grouper API model and picking the parts that they like/feel a strong need for, and agreed to take an AI to put together a description of his (by now, fading) moment of clarity for the group and also present it to Benn for a FIFER perspective. He asked whether folks would agree that such a discussion with Benn might be fruitful at this point?
        - Tom asked for clarification of Keith's proposal, and Keith rephrased it as putting an outline of what he'd suggested we might present in mid-September before Benn and asking for his take on it from the perspective of what he's seen via the FIFER effort.
        - Lucas asked whether FIFER has any position with respect to provisioning?
        - Keith indicated that FIFER has been primarily focused on APIs rather than specific actions – on models rather than implementations – and that that may be the key to answering the question of how this might differ from FIFER – the proposed deliverable for this effort (a clear direction for how to plug things together and build the "plumbing" necessary to perform provisioning) may be similar but complementary to FIFER's work.
        - Before time ran out, Lucas requested that someone consider adding the "documentation" theme to the OSIdM4HE wiki space (or at least the provisioning portion of it) to provide streamlined navigation through multiple layers of child pages in the wiki.
        - Keith agreed to talk to the folks who set up the wiki space about making the change, and also to accept an AI for initiating contact with Benn.
    - **Action Items**
        - Roll forward AI for Lucas to consider text to add describing his diagram in the architecture diagrams table
        - Roll forward AI for Keith to map provisioning onto the enterprise integration patterns
        - New AI for Tom to propose a more detailed breakdown of the provisioning space (perhaps with a diagram at the "5,000 foot" level, beyond the detail in the first two diagrams in the architecture diagrams table in th ewiki
        - New AI for Keith to discuss getting the "documentation" theme added to the wiki space
        - New AI for Keith to talk to Benn Oshrin about his moment of clarity and its possible interaction with FIFER
        - New AI for Keith to document his moment of clarity for the group
    - **Adjourned**