

# Grouper call 7-July-2010

## Grouper Call 7-July-2010

### \*Attending\*

Tom Barton, U. Chicago (Chair)  
Shilen Patel, Duke  
Jim Fox, University of Washington  
Chris Hyzer, U. Penn  
Tom Zeller, U. Memphis  
Rob Hebron, Cardiff University  
Steve Olshansky, Internet2  
Emily Eisbruch, Internet2 (scribe)

### \*Action Items\*

#### \*New Action Items\*

[GrouperWG:AI] (Jim) will do performance testing to help evaluate potential issues with the flattened memberships table.  
[GrouperWG:AI] (Shilen) will write up design objectives of the flattened memberships table.

<https://spaces.at.internet2.edu/display/GrouperWG/Flat+Memberships+Design>

[GrouperWG:AI] (Rob) will ask (Gary) to look at the Fluid project to learn more about accessibility of UIs <https://lists.internet2.edu/sympa/arc/grouper-dev/2010-07/msg00010.html>

#### \*Carry Over Action Item\*

[GrouperWG:AI] (Rob) will look at issues relating to testing the ESB Connector.

#### Discussion

#### Early Experience with Grouper 1.6 Release

Jim reported that at U. Washington, Grouper 1.6 is being tested, though it is not in production yet. The plan is to use Grouper 1.6 to process institutional groups and route them into LDAP. Currently U. Washington uses the Grouper API but does not use the Grouper GUI or the Grouper web services. This could change in the future.

Rob reported that Cardiff is testing Grouper 1.6 against Grouper 1.5. No issues have been found, and Cardiff plans to go live with Grouper 1.6 within a month.

Q: Have any performance changes been noticed with Grouper 1.6?

A: Rob: Cardiff has moved away from using Grouper's hooks to using the ESB connector. So the timing is different, but no performance problems have been observed. Grouper 1.6 performance is at least as good as previous versions.

Chris reported that testing at Penn. has gone well and Penn is moving into production with Grouper 1.6 this coming weekend. They are upgrading from Grouper 1.4.

Shilen reported that Duke will most likely upgrade to Grouper 1.6 sometime after the summer.

U. Memphis and U. Chicago plan to go into production with Grouper 1.6 before the start of the fall 2010 term. U. Chicago has been testing Grouper 1.6 and has not encountered problems.

## Log Items Issue

There are issues with unnecessary log messages from the hibernate wrapper. An example is an error logged of "this constraint was violated," where actually the violation was successfully resolved later. It was agreed to provide the ability to pass an option to the hibernate call to suppress such hibernate error messages. If the caller expects to handle the error it can, but by default, it will log the error.

## Flattened Tables

Jim expressed concern about the use of flattened tables in Grouper 1.6.

Concerns included:

- Flattened tables are not the best way to speed up retrievals and searches
- There can be a lag where updates have been made but are not yet reflected in the flattened table.
- Using flattened tables for something other than a downstream directory can give inconsistent results. (Results are eventually consistent, but not exactly consistent.)
- Possible waste of time and space.
- Jim can support making use of flattened tables as an optional feature. But he does not think the flattened table should be made a resource when working with the UI.

Chris noted that tables previously had a row for every immediate and effective membership. Now there is one row for every flattened (effective) membership.

The flattened table should be more efficient when doing queries. Chris advocates using a flattened table for things that don't need updates right away, with the possible exception of things that need to be high performance. Jim stated that in general at U-W there are two queries that need to be high performance.

1. Getting all the groups someone is a member of. (Shib IdP uses that query to get group memberships.)
  2. Getting all the effective members of a group
- U-W routes those high-performance queries thru rewrites from the front end to the LDAP system.

### Using the Flattened Table with Rules

Chris: Flattened tables will become important when rules are implemented in an upcoming Grouper release.

The flattened table will help with the determination of whether a person is being added to a group effectively or not. If someone is removed from a particular group, the flattened events table will be checked to verify if the membership needs to be removed from additional groups or not.

For at least some of the rules, what fires them is a flattened change, not an immediate or effective change. (Example: For some rules, we don't care if someone is removed from a group, we want to know if they are removed from a group AND they have no other effective memberships anymore)

Chris acknowledged that there would be a small delay of that information propagating down using flattened table.

Q: TomB: are some of these questions related to rules resolvable with a standard API call instead of with a flattened table?

A: Shilen: There can be timing problems with an API call. For example, if a member has been removed from a group in two ways at about the same time, then depending on how timing works out, the system might think the person is a member of the other group in response to each query and so the query would return an incorrect answer. There is a short time where you'd get the wrong answer to question about whether or not the person is an effective member. The flat table "smooths out time" and eliminates this problem.

By using the flattened table approach, when an immediate membership add or delete happens, a record is entered to the change log. The change log daemon looks at the change within a short interval, then looks at the flattened table to see, based on what's already in there, if the change needs to be added there in the flattened table or not.

It was noted that the update to the flattened table could happen quicker if the process was asynchronous, but if there's a huge queue there is no guaranteed time for the update to be available.

Q: How do things work with LDAP change logs?

A: LDAP just processes sequentially, though there is some variation depending on the version.

Jim noted the U-W has groups with 100,000 or 200,000 members. If those groups are included those as members of other groups, the flat table could be so big as to hurt performance, with the database spending a lot of time updating and searching the table.

It would be beneficial to do some performance testing. Jim noted that after finishing some other projects, he hopes to have time to test the flat table performance issues.

[GrouperWG:AI] (Jim) will do performance testing to help evaluate potential issues with the flattened memberships table.

TomB stated that a written summary of the design objectives of the flattened table would be helpful.

[GrouperWG:AI] (Shilen) will write up design objectives of the flattened memberships table. <https://spaces.at.internet2.edu/display/GrouperWG/Flat+Memberships+Design>

### Error Reporting for Web Services

Chris stated that using web services, there are sometimes errors/exceptions that are not really errors, just an unresolvable (not found) member, for example. It was agreed that a good solution is to put a different return code on such conditions so they do not show up as exceptions.

### Advance CAMP Debrief and Takeaways

TomB, Chris, TomZ and Shilen attended Advance CAMP in Raleigh, June 23-25, 2010.

TomB mentioned that some attendees at Advance CAMP were surprised by the unconference format. The format was interesting and probably could be improved, but it produced substantial work items.

Action Items emerging from Advance CAMP are being tracked at <https://spaces.at.internet2.edu/display/ACAMPActionItems/Home>

The Advance CAMP program committee will keep in contact with the action item leads to monitor how things are going.

TomZ is moving ahead as the lead on the action item:

"Determine how federated provisioning should work and participate in SPML standards work to support it "

TomZ noted that there is some interest in embedding some SPML inside SAML. Spec work is happening now.

Oracle's purchase of Sun and increasing costs on existing IdM systems has led to much interest on campuses of looking for provisioning solutions and open source options.

Chris is the lead on the action item:

"Determine how federated group management will work in COmanage"

This action item will involve Grouper handling of federation issues, and TomB noted that this action item should be reflected in the Grouper roadmap

TomB is the lead on the action item: "Define a Standard API for Groups"  
(Eric Westfall, Chris H and Ray Davis are group members)

There was some interest at ACAMP in a permissions API that has been talked about in on MACE-paccman calls. There was also interest in attribute and permissions frameworks.

### **Accessibility of Grouper UIs**

A recent inquiry was received on the Grouper-Users list of whether Grouper is ADA Compliant

<https://lists.internet2.edu/sympa/arc/grouper-users/2010-07/msg00000.html>

<http://www.ada.gov/stdspdf.htm>

[GrouperWG:AI] (Rob) will ask (Gary) to look at the Fluid project to learn more about accessibility of UIs

<http://fluidproject.org/>

<https://lists.internet2.edu/sympa/arc/grouper-dev/2010-07/msg00010.html>

Next Call: Wed. July 21 at noon ET.