Getting Ready for Production with Grouper

Here are steps for after Grouper is installed and prior to going live in the production environment. For a more extensive overview, see also the Planning Guide.

- Review the Grouper Deployment Guide.
- Go to the UI and create folders, groups, attributes, permissions, etc.
- Further configure Grouper via the API.
- Configure a Loader job
- Set up Notifications (change log)
- Plan load balancing. Note: the UI needs sticky load balancing, and generally people do not cluster sessions. For the WS you do not need sticky load balancing, though if you had it you could have better performance with caching and prevent caching errors.
- Set up the Grouper client. You might want to make a zip with some default server names for an environment in your institutions (one for test and one for prod?), and put that on a web server at your institution and link to it from your institution's wiki about Grouper
- Set up Web Services. Note you need to configure how you want authentication to work. E.g. at Penn we do HTTP basic auth which does a kerberos bind based on the kerberos service principal sent in. This module is included in Grouper. Penn also has a DB table for kerberos principals and a subject source to read them. We have a proprietary UI to manage these.
- Set up the Provisioning Framework in order to provision groups, memberships and stems/folders.

Q: Where can I see an example of using build scripts to set up various environments after using the Grouper Installer?

A: Here is an example from Penn on Managing Grouper in Multiple Environments (note that this example uses ant, not Maven)

See Also

Grouper Planning

Grouper Training Videos

Grouper Training slides (including group naming best practices)

Glossary