New Topics: Campuses and Operations

Heidi Picher Dempsey, GPO
Denver, Colorado July 26, 2011
www.geni.net

Sponsored by the National Science Foundation

Campus suggestions for GENI new work

1. **Resource Access and "Opt-in"** Generally speaking campuses want to be able to approve and log access to their resources, and having the right architecture and tools in place to do this is important, and currently not adequate in prototypes. For OpenFlow in particular, it would be good to have a **tool that helps admins** understand what's already accepted and how new requests would impact current state, because this is complex and each request builds on other previously approved requests. In particular, admins can get many requests per day for months and it isn't feasible to remember it all. (Flowvisors already have thousands of rules in relatively small campus deployments**. Auto-approval** (for all or for high-level rules) or **delegated approval** (to another organization) is one way of solving the problem; better tools for each campus admin is another. Note that if opt-in is too hard, the "approval" step no longer means anything, so there are security, availability,and possibly experiment implications.

Campus new work (cont)

2. **Resource Isolation.** The prototypes in mesoscale currently don't really have adequate isolation to protect experiments from each other in all components. For example, a single campus FlowVisor with many rules slows down, which impacts all experiments.
3. **Experiment Topology.** Outages, configuration problems, and lack of isolation can affect network topology (not new of course). Should we inform experimenters of this, and if so how do we trace objects like switch interfaces to human experimenters? Should we just pass on all outage info to mailing lists, web pages etc. that experimenters can subscribe to/check or not? Do we need software tools help?
4. **Emergency Stop could use an automatic programmed interface to GENI components.** Also possibly support for the idea that Emergency Stop capability should be something you can delegate to another organization (e.g. an operator group).

Campus new work (cont)

5. **Relating GENI URNs to people** is hard for operations and is already important for things like Emergency Stop. There are workarounds now, but they won't scale.
6. **Relating all GENI "objects" is too hard** in our small deployments. Relating slices, slivers, resources, VLANs, flows, nodes and interfaces is hard. We can't "look up" the relationships anywhere and currently use workarounds like naming conventions.
7. **Security and data privacy for sharing monitoring** data requires design support for access controls, proxies etc. that don't currently exist, but we are monitoring already.