

## GENI Project Progress Report

**Project Title:** Intelligent Data Movement Service on GENI

**Principal Investigators:**

Ezra Kissel <ezkissel@indiana.edu>

Martin Swany <swany@iu.edu>

### 1. Major accomplishments: 5/2014–7/2014

Milestones achieved:

- **S6.c.1** - Deliver and demonstrate prototype DMS software that integrates with IBP, XSP, Phoebus, and UNIS to manage data distribution.
  - DMS development began ahead of schedule during last period (GEC19). New developments include automatic, on-demand provisioning of GENI resources when needed, and automatic tracking and renewal of persistent allocations.
- **S6.c.2** - Create specialized storage and acceleration appliance images to be created on-demand by the DMS based on resource utilization.
  - Images were updated to include improved support for both IG and EG deployments, reflecting the developments in IG/EG image conversion. Storage depot image update includes more robust configuration script for public/private interface discovery and service registration.
- **S6.c.3** - Deliver configuration and management scripts for prototype services.
  - Tools and scripts are included within appliance images. The DMS service and supporting configuration scripts are available upon request from the IDMS development git repository.
- **S6.c.4** - Deliver AuthN/AuthZ design document.
  - AA considerations were shelved as future work since our focus shifted to delivering a working IDMS prototype that could be used by experimenters with current tools.
- **S6.c.5** - Demonstrate I&M tools in collecting and measuring prototype services running in GENI slices.
  - Web-based IDMS monitoring was updated with publish/subscribe model to scale as more clients view status page. Components from GEMINI I&M are used to provide storage and network monitoring statistics.

Deliverables made:

- Updated appliance images for storage depot and Phoebus Gateway published to project Wiki.
- Documentation describing custom image creation and conversion for InstaGENI and ExoGENI.
- Long-lived IDMS slice(s) created to support external experimenter use. Tutorial page published on project Wiki and presentation delivered at upcoming experimenter tools session at GEC20.

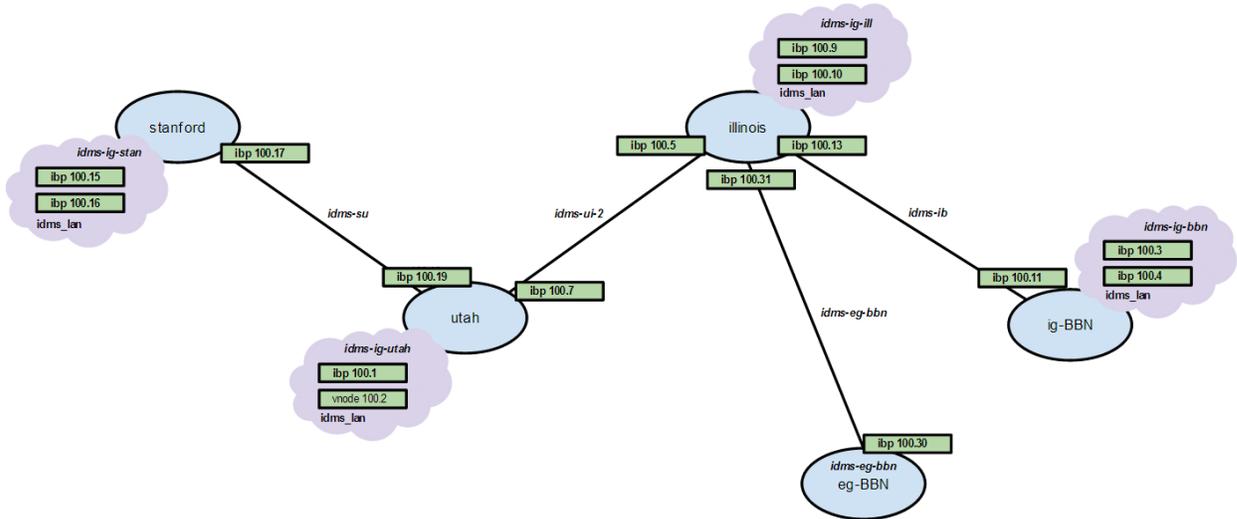
### 2. Description of work performed

Activities and findings:

We focused on the following development areas since the previous period (GEC19):

1. The DMS was extended to support the on-demand creation of storage resources based on utilization monitoring, and to auto-renew persistent IDMS allocations used in the long-lived slice(s). When new GENI resources are added to the IDMS experiment, the manifests are converted and tracked within the UNIS, giving us a complete topological representation of reserved resources at each AM along with associated metadata such as expiration times. The DMS then queries monitoring statistics for each storage resource (e.g., available bytes) and calculates if new resources should be added, or if under-utilized resources can be freed. If any slivers are close to expiration, the DMS renews resources based on the max allowable per AM. The AM interface in DMS uses the GCF python lib directly and support AM API v{1-3} calls.

- Working with the ExoGENI and InstaGENI/ProtoGENI developers, we were able to formalize the steps necessary to take existing ExoGENI VM images and convert them for use on InstaGENI. This work has led to further discussion on streamlining the conversion process via AM API calls; however, the installation of specific client tools for each aggregate type remains an involved and time-consuming requirement. Further details may be found documented on the project Wiki.
- A long-lived IDMS experiment was created to support external experimenter use and to investigate the process of adding on-demand resources into a “base” set of slices. A diagram of this base experiment is shown below.



We made use of the shared vlan feature at ProtoGENI aggregates to enable the attachment of resources from multiple slices. As of GEC20, IDMS shared vlans are available at the Stanford, Utah, Illinois, and BBN PG aggregates. GENI stitching was used to connect each of the shared vlan “islands”, and allowed us to include EG resources at BBN. Ongoing work will make use of the mesoscale OF shared vlans to more extensively connect EG and IG resources. Further details regarding the experiment configuration and opt-in use may be found on the tutorial page: <http://groups.geni.net/geni/wiki/sol4/IDMS/ExpTutorial>

Project participants:

Ezra Kissel [ezkissel@indiana.edu](mailto:ezkissel@indiana.edu) (PI)  
 Martin Swany [swany@iu.edu](mailto:swany@iu.edu) (co-PI)  
 Akshay Dorwat [adorwat@indiana.edu](mailto:adorwat@indiana.edu) (RA)

Publications:

N/A

Internal project documents:

- Available on request from IDMS git repository.

External publications:

- Published on project Wiki: IDMS experimenter tutorial, custom image creation/conversion HOWTO, and updated image table with example RSpec.

Outreach activities:

N/A

Collaborations:

- Continued knowledge exchange with Paul Ruth (Science Shakedown) and ExoGENI team concerning our shared goals across respective sol4 experiments, e.g., image creation, stitched topologies, and storage.
- Continued outreach with geni-help, orca-users, and protogeni-users communities to ensure IDMS operation across leading aggregate types and backbone nodes.

Other Contributions:

N/A