

University of Kentucky Quarterly GENI Report

Jim Griffioen and Zongming Fei

Department of Computer Science
University of Kentucky

June 2009

1 Major accomplishments

The following highlights our accomplishments during Q2 2009

1.1 Milestones achieved

During the past quarter we made continued progress toward our milestones. In summary we:

- continued to improve the aggregate reference implementation running at Kentucky.
- continue to work with Utah and members of the security team regarding security issues in the design of our aggregate and measurement system.
- have begun to design an instrumentation/measurement model that can be integrated with the ProtoGENI architecture.

1.2 Deliverables made

None yet.

2 Description of work performed during last quarter

The following provides a more detailed description of our work and results from the last quarter.

2.1 Activities and findings

This quarter we began building on the (evolving) ProtoGENI aggregate system that we installed in the previous quarter. In particular, we spent a good deal of our time getting familiar with the new system, its architecture, and the new software that drives it. As part of this effort, we have begun to develop test scripts that make use of the ProtoGENI interface (API) to reserve resources and create experiments. Not only has this helped us better understand the API, RSPECs, and the overall system architecture, but it has help us uncover some issues with the ProtoGENI design that were going to make it difficult for us to build an instrumentation system. In particular, we found that the current RSPEC approach did not include the information we needed to be able to monitor links in the topology, and, more generally, the information that

is assigned at the time that resources are allocated and redeemed. We raised these issues with the Utah team and had various discussions about how to handle these issues. The solution that the Utah group has decided to use is to redesign the RSPECs to include a “manifest” that provides the additional information we need.

Given a better understanding of the ProtoGENI architecture, we have been developing an instrumentation architecture that will incorporate the Edulab’s measurement capabilities into the ProtoGENI system. Our new approach is different from our original design in that we intend to deploy experiment-specific measurement infrastructure in the form of *measurement controller* nodes that monitor an experiment’s behavior across the various aggregates that comprise the experiment. Not only are experiment measurement results kept private to the experiment, but it distributes the monitoring load across multiple nodes so the monitoring system can scale up as the measurement workload increases. An initial version of our instrumentation design was presented by Jim Griffioen at the GENI Measurement Workshop at the University of Wisconsin. Slides from the talk can be found at <http://groups.geni.net/geni/wiki/GENIMEasWS>. Another important issue that we are trying to resolve is how visible the measurement system should be to users. Our current thinking is that the measurement infrastructure should be as invisible as possible, only making the results of measurement available to users and giving them the ability to control which measurement results they want to see.

In the process of becoming familiar with the ProtoGENI system and developing an instrumentation system for it, we ran into issues related to authenticating users and services between nodes; in particular, between our measurement system resources and the resources being measured. We have had multiple conversations with the Utah team and have discussed these issues with other members of Cluster C—including members of the security team—working with them to come up with security solutions (and a general security framework) that will be useful to all members of the cluster. As of yet, the security/access model has not been finalized, and we continue to discuss various alternatives.

In addition to planning measurement additions to our aggregate, the ProtoGENI team has been making continuous improvements to the code base which we have been applying regularly and then testing to make sure our scripts continue to run with the newly updated system. We have also done some limited testing across aggregates to verify that it works as expected and to become familiar with allocating resources across aggregates.

2.2 Project participants

The following individuals have helped with the project in one way or another during the last quarter:

- Jim Griffioen - Project PI
- Zongming Fei - Project Co-PI
- Hussamuddin Nasir - The project’s primary technician and programmer
- Lowell Pike - Network administrator
- Woody Marvel - Assists in Emulab administration

2.3 Publications

James Griffioen and Zongming Fei, Automatic Creation of Experiment-specific Measurement Infrastructure, In the proceedings of the *First Workshop on Performance Evaluation of Next-Generation Networks (Neteval '09)*, Boston MA, April 2009

2.4 Outreach activities

We participated in the recent GENI Measurement conference and also gave a talk about our work at the Neteval '09 conference.

2.5 Collaborations

Most of our collaborations continue to be with the Utah ProtoGENI team, and we continue to be actively involved in the bi-weekly meetings of the ProtoGENI cluster (Cluster C).

2.6 Other Contributions

None yet.