

# University of Kentucky Quarterly GENI Report

Jim Griffioen and Zongming Fei

Department of Computer Science  
University of Kentucky

September 2009

## 1 Major accomplishments

The following highlights our accomplishments during Q3 2009.

### 1.1 Milestones achieved

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

- supported operation of the Kentucky ProtoGENI prototype, enabling basic experimentation through the ProtoGENI clearinghouse.
- designed an instrumentation/measurement system and are about to release a document describing the architecture/design of the instrumentation toolset and planned tools.
- began integrating existing tools into the Kentucky prototype testbed.
- continued discussions with the Utah group regarding the design of the measurement system and methods for authentication and secure access to slice and measurement resources.

### 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

In the previous quarter we started the process of defining an architecture for our measurement infrastructure, and had made some key design decisions. In particular, we decided to use experiment-specific (i.e., slice-specific) measurement nodes, creating a local measurement system within each experiment. This design matches the standard usage model in which users are primarily interested in collecting information about their own experiment. It allows users to keep measurement data private and local, but yet allows data to be made public if desired.

Our task this quarter has been to flesh out the details of our architecture and map it onto the ProtoGENI infrastructure. This has been rather challenging, in part because ProtoGENI itself is under development and does not yet have a set of guidelines or best practices describing the “approved” or “correct” way to hook into, and interoperate with, the ProtoGENI infrastructure. In order to flesh out the details of our design, we have had to build components, connect

them into ProtoGENI via the API, and try them out, often finding problems/limitations of the API that in turn required redesigning our architecture (and software). In that sense, fleshing out a detailed design for our measurement system has been an iterative, and often slow, process. As part of this work, we began using the *manifests* that Utah recently added to their APIs. While the manifest abstraction is a welcome addition, the manifest implementation and API calls need further improvement and have caused us to rethink how and when topology information should be obtained by our measurement nodes.

While this iterative “experiment/redesign” process delayed us from settling in on a workable (ProtoGENI-compatible) architecture/design for our measurement/instrumentation system, we are now at a point where we have a reasonably complete and workable design for the measurement/instrumentation system. We are nearing completion of a document describing the architecture and the types of tools we plan to support. In regards to the user interface and display of data, we decided to take a lazy-evaluation approach. The raw data collected by the system will only be processed into RRD database files. Conversion into a format suitable for viewing will be delayed until it is needed by users. Instead of generating graphs and image files immediately for potential use in the future, we will implement scripts that generate graphs from the RRD data on demand. We intend to use a content management system (such as Drupal) to allow users to customize the information they want to display based on their needs, generating graphs and tables only for the information of interest to the user. This is a significant change from our previous architecture and gives the users much more control over the interface to the measurement system.

One upside of diving into software development even before the design was complete is that we already have a beta version of our measurement nodes up and running, collecting and displaying basic network and operating system information. Although we were not been able to use the existing Edulab code—because it heavily tailored to the Emulab APIs—we were able to quickly rewrite several of the tools to the ProtoGENI API as a proof of concept that data can be collected and displayed using our new measurement node architecture. The prototype does not yet utilize a content management system, but the hooks are there to add one in the future. Authentication issues (described in our previous quarterly report) have not yet been addressed in the ProtoGENI code, and so our current prototype uses hard coded authentication to facilitate communication between nodes. We continue to hold discussions with Utah regarding these authentication issues.

We have been using the Kentucky ProtoGENI prototype for our experimentation on a regular basis this past quarter. All experiments and credentials are managed through the ProtoGENI clearinghouse, and have been operating quite reliably. The Kentucky system is open to other users, but so far has primarily been used for simple tests by other cluster C members.

Recently we have begun exploring other tools that we can incorporate into our system including netflow information and data collected via perfsnar. Integration with these tools is still in the early design stage, but we hope that our architecture will allow these tools to be supported as well.

## 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
- Xiongqi Wu - Research Assistant
- Jeremy Reed - Research Assistant (half time)
- Lowell Pike - Network administrator
- Woody Marvel - Assists in Emulab administration

## 2.3 Publications

None this quarter.

## **2.4 Outreach activities**

None this quarter.

## **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.

## **2.6 Other Contributions**

None yet.