



GENI Meta-Operations Center Quarterly Report

1/1/2009-3/31/2008

Jon-Paul Herron – Principal Investigator

Luke Fowler – Co-Principal Investigator

Summary

- *Operating within Budget*
- *Operational Data actively being collected from PlanetLab*
- *Dataset Discussions underway with all clusters*
- *Operational Dataset Document shared with all clusters and OMIS, mature enough for sharing work to being*
- *GMOC Internal Data Format Document created*
- *GMOC Data Exchange Format Document created*

Major Accomplishments

Milestones Achieved

Actively sharing operational data with PlanetLab: Based on our conversation with the PlanetLab cluster participants, there was consensus that we should begin by gathering existing data from PlanetLab Central. GMOC visited Princeton to discuss this in detail, using our operational dataset document as a guide. The result is that we are now actively gathering data from PlanetLab, storing it in the format described by our Internal Operational Data Format document. This data is still incomplete, in that we are unable to gather some topological data, and in that it is only data covering the PlanetLab project itself, not data about other projects within the PlanetLab cluster. But, it represents both a good initial model for quickly beginning to collect data, and a good start for gathering data from the PlanetLab cluster as other projects in the cluster begin to integrate more with the PlanetLab control framework.

GMOC Operational Dataset Document: The Operational Dataset document has been shared with all GENI control frameworks, and we've updated this document based on input. This document is now mature enough that we've begun to use it to work with projects to begin gathering operational data.

GMOC Operational Data Format Document: Initial versions of two documents have been created and shared with the GENI community, related to this. First, we've created a document that describes the internal data format schema we've developed for storing operational data we collect. Second, we've created a document that describes the data exchange methodology we hope to employ as GENI evolves to gather data from GENI projects, and to share data about GENI with other federated networks such as internationals. These documents are open to input from the community, but we believe these will require less collaboration at this point than our operational dataset document. This is because initially the formats used to exchange data with most projects will be highly customized for each project.

Deliverables Made

Updated Operational Dataset Document: Several new versions of the evolving GMOC Operational Dataset Document were created, incorporating the input from other GENI projects and clusters.

GMOC Internal Operational Data Format Description Document: A new document describing the format GMOC will use to store operational data was created.

GMOC Operational Data Exchange Draft Document: A document describing our approach to data exchange with other projects has been created.

Description of Work Performed During Last Quarter

Activities and Findings

Our activities during Q2 were spent working in depth with control framework clusters and other stakeholders to do two things: gather input on our strawman operational dataset document, and to understand in technical detail what each cluster could provide and how.



This resulted in conversations with staff from every cluster, along with several projects within clusters, and with Steve Schwab to discuss how security and operational data sharing might begin to interface and interact.

Further details about findings and activities can be found in Milestone/Deliverables description, and in the Collaborations sections of this document.

Project Participants

During this time, key participants in GMOC included:

Jon-Paul Herron, PI
Luke Fowler, Co-PI
Chris Small, Senior GMOC Engineer
Camilo Viecco
Dan Doyle, Undergraduate

Publications & Documents

Herron, JP. (2009, March 24-25). *The GENI Meta-Operations Center*. Presented at National Lambda Rail All-Hands Meeting, CallT2, San Diego, CA.

This presentation focused on challenges of operations in a federated/distributed world, and how operations for GENI relate to the future of operations for production networks.

GENI Documents:

Herron, JP. *The GMOC Operational Dataset*.

This is a document produced to understand the initial set of data GMOC hopes to gather from GENI projects.

Viecco, C. *The GMOC Internal Data Format*.

This document lays out the schema and format GMOC will use to store all operational data format.

Viecco, C. *GMOC Data Exchange*.



This document describes a methodology moving forward to describe how projects and GMOC might share data (other than utilization)

Outreach Activities

Undergraduate Development Assistance: Dan Doyle, an undergraduate for the IU Computer Science department continued to assist with research activities into GENI projects and existing data sharing options. A second undergraduate position is available but has not yet been filled.

GMOC attended the IU Informatics Career Fair on Feb 17 to advertise available positions for undergraduate students and discuss the GENI/GMOC projects with IU undergraduates as part of our outreach efforts

Collaborations

Similar to Q1, GMOC worked on gathering input and developing discussions with many GENI stakeholders during this quarter.

PlanetLab: After discussion about GMOC with PlanetLab cluster participants, GMOC staff visited Princeton to discuss overall operational data sharing strategy and detailed capabilities for collaboration. This resulted in further collaboration between GMOC and PlanetLab staff to enable active data sharing.

ProtoGENI: GMOC discussed the strawman GMOC Operational Dataset document with the ProtoGENI cluster participants. There were some comments from this discussion on the document, and further discussions between GMOC and ProtoGENI developers on potential capabilities are underway. A visit with ProtoGENI may be useful, as the PlanetLab visit was.

ORCA: GMOC spoke directly with Jeff Chase to cover the strawman GMOC Operational Dataset document, and discuss future collaboration. We made plans for immediate sharing from the ORCA monitoring software and topology of some projects. Additionally, we talked about some future collaboration on the issue of topology description, and exchange.

ORBIT: ORBIT has provided GMOC some preliminary data output. Further discussions will be needed to gather improved data.

Mid-Atlantic Crossroads/DRAGON: GMOC is working with MAX to understand how to gather operational data from projects. We are working with MAX to understand what data we can/should gather via MAX's integration with the



PlanetLab control framework, and what data we can/should gather directly from MAX

NLR: GMOC has discussed operational dataset sharing with NLR, and now has full capability to gather full initial operational dataset from the NLR GENI elements, once it is in use.

Internet2: GMOC has discussed operational dataset sharing with I2, using the GRNOC's existing access to Internet2 data. This data is all accessible for GENI use, so we are prepared to gather full initial operational data from Internet2's GENI elements, once this is in use.

PerfSONAR: GMOC has opened an ongoing dialogue with the PerfSONAR project, which has already made significant progress in the areas of sharing data among production networks. Specifically, PerfSONAR would provide a potentially useful default method for projects to share operational data with GMOC, or for GMOC to share its data with other federated networks. It contains a topology description

Planned Activities for Q3

Ramp-up of Cluster sharing: During Q3 we plan to accelerate our discussions with GENI projects to begin gathering data from many projects by GEC5. In most cases, this will follow the basic successful model used with the PlanetLab project, using detailed one on one interaction to understand methods and capabilities.

Initial GENI-wide view of operations: For GEC5, we plan to develop a demonstration of the GMOC operational data view, using some basic visualization and data gathered from a broad spectrum of GENI projects. This will be the initial proof of concept.

Initial Investigation of visualization for GENI operations: Part of the effort for initial GMOC operational data view will require investigation into the best methods to visualize the complex, virtualized and federated GENI. A single simple "weathermap" won't work. This provides us an opportunity to re-examine methods of visualization for networks and research.

Finalization of First version of GMOC Documents: We'll also work to bring the current versions of the documents describing the operational dataset and the formats for exchange and storage of operational data, using this current version to build our data-view, and continue to work with projects to gather more and



better operational data. New versions will no doubt be needed, but we will need to finalize this version enough to use them as a guide for our work.