



GENI Meta-Operations Center Annual Report

10/1/2008-9/30/2009

Jon-Paul Herron - Principal Investigator

Luke Fowler - Co-Principal Investigator

Summary

- *Operating within Budget*
- *Completed all Year 1 Milestones*
- *Worked collecting data from initial set of GENI projects*
- *Functional prototype GMOC environment implemented and available*
- *Published several GENI documents, including Concept of Operations, Emergency Stop, GMOC data exchange formats*

The GMOC project had a successful first year, making good progress on the early definition of operations for GENI, and helping to build the initial systems and processes for federated operational data sharing. These are crucial early building blocks for a facility that must be federated and flexible enough to attract many types of infrastructures, but must also behave coherently and effectively as an integrated facility. These can be used as the bootstraps for the increased services, prototypes, and further definition we expect for Spiral 2.

Much of our work for Spiral 1 was spent working to build a framework for data sharing between GMOC and GENI infrastructure providers. This was handled initially by approaching the control framework projects heading each control framework cluster to try to build a system to share operational data about all the projects within each cluster. Initial discussion was focused on feedback for a draft Common Operational Dataset document describing the most important types of data projects should share with GMOC. A Common Operational Data Format document was also created to document how best to format the data we receive, although this provoked less feedback and interest from the projects.



The results from this effort were fruitful, but inconsistent. At this early stage, the capabilities of the control frameworks and the projects themselves are highly variable. Some projects were able to immediately provide us some data, while others had little or no way of sharing this data at this point. It also became clear that the control framework projects might not necessarily be the right aggregator of operational data for projects within each cluster. Rather, the right architecture might require direct data sharing with individual infrastructure provider projects.

By the end of Spiral 1, GMOC was collecting significant operational data about the PlanetLab infrastructure, though not the other projects within the PlanetLab cluster. We were also gathering some topological data from ProtoGENI, with an expectation of gathering more data in Spiral 2. For ORCA, they hoped to have some data in Spiral 2 as well.

In addition to the project of coordinating operational data sharing and formatting, GMOC developed a prototype GENI-wide view. This view, available from the GMOC website, serves as an early prototype into how to visualize GENI as a facility. There will continue to be significant work to develop this into a suite of maps and other tools to provide health information to all GENI stakeholders. This prototype will change significantly as we begin to receive additional types of data. We've found that even a small amount of data becomes difficult to visualize in a constructive way, and this will become exponentially more difficult as more projects begin to share more information.

Lastly, GMOC has produced 2 documents to help define operations for GENI overall. First, we've produced a document describing a strategy for Emergency Stop. This will be the document used to drive our Spiral 2 developments in supporting Emergency Stop functionality. This document is now in draft format and will be altered after soliciting feedback from other GENI project. Second, we've produced a draft GENI Concept of Operations document to help begin a strategy for Operations of GENI. This document will also likely be altered based on feedback, and will be used to set the direction of GENI operations as a whole.

Major Accomplishments

Milestones Achieved

Q1:

Progress toward Dataset & Format Definition: The first set of GMOC milestones is on track for on-time completion at the 6-month mark of Spiral 1. This set of milestones includes a complete initial GENI Operational Dataset and a complete initial Operational Data Format. Our work to define the basic GENI operational dataset is in progress, with three main elements. It includes discussion of a strawman list of types of data to collect, a discussion of some basic questions

about the data collecting and sharing capabilities of various GENI projects, and research by GMOC into the existing data already included in various projects' data collections.

Q2:

Initial GMOC Data Translator Implementation: While we do not expect in Spiral 1 to have a single unified implementation to translate operational data into our internal format, we have worked with PlanetLab to collect operational data in their existing format (mainly from their Comon application) and translated it into our defined internal data format. In addition, we have started collecting a small set of data from the ProtoGENI project

Initial GMOC Data Exchanger Implementation: As with the Data Translator, we do not expect in Spiral 1 to have a unified exchange system used by all projects. However, we've worked with PlanetLab and ProtoGENI to collect data, using a translation to convert it from the form the projects keep it in, into the Exchange format defined in our Operational Data Exchange Format document, and then using that to collect the data. In the future, the translator can be developed to translate additional data from these and other projects, or projects can send data directly in our defined Data exchange format.

Q3:

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.

Q4:

Draft GENI Concept of Operations The GENI Concept of Operations Document proposes an operational model for GENI infrastructure for problem reception, triage, handling, resolution, and reporting.

Demonstrated initial GMOC Framework implementation at GEC5: GMOC participated in the demo session at GENI GEC5, demoing a prototype version of GMOC functionality with PlanetLab and ProtoGENI data.

Initial implementation of Emergency Stop and notifications: The GMOC Emergency Stop document proposes a set of conditions and procedures for early Emergency Stop activity in GENI.

Complete initial implementation of prototype GMOC: The GMOC portal demoed at GEC5 serves this initial prototype GMOC function. Continued enhancement and refinement of this software will occur throughout the next year.

Deliverables Made

Q1:

High-level, Strawman dataset: GMOC has developed an initial high-level list of types of data to begin collecting from other GENI projects. This list is by no means finalized or complete; in fact, our intent with this list is to generate and focus our discussion with the other GENI projects. As this high-level list is improved through discussion, it will also begin to include a greater level of detail. This list was presented during our GMOC presentation at GEC3.

Operational Dataset Question List: Additionally, GMOC has formulated an initial set of questions to discuss with other GENI projects, in order to form a fair amount of consensus on what operational data is most sensible to collect.

Q2:

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.

Q3:

GMOC Data Exchanger: The code developed for Data Exchange is stored in an internal GMOC version control system. This code is available on request, and we are working on a self-service system for access to this code.

GMOC Data Translator: The code developed for Data Translation is stored in an internal GMOC version control system. This code is available on request, and we are working on a self-service system for access to this code.

GMOC Data Repository: The GMOC Data Repository system is kept on an internal GMOC server. This data can be made available to GENI as required.

Q4:

Continued improvement of GMOC software: GMOC has worked through Q4 to revise and enhance the GMOC software to increase speed and add to the feature set. The portal used for the GMOC demo at GEC5 has also been moved to production infrastructure.

GENI Emergency Stop Document: The GMOC Emergency Stop document proposes a set of conditions and procedures for early Emergency Stop activity in GENI.

GENI Concept of Operations Document: The GENI Concept of Operations Document proposes an operational model for GENI infrastructure for problem reception, triage, handling, resolution, and reporting.



Summary of Work Performed During Last Year

Activities and Findings

Q1:

Hiring: The GENI Meta-Operations Center posted, advertised, interviewed, and hired Camilo Viecco into the position of GMOC primary developer. Camilo starts on January 8th, and will be tasked with continuing the research into existing models for GMOC data formatting and protocols for sharing, as well as technical discussions about the operational dataset with other GENI projects. Camilo comes to GMOC from the Advanced Network Management Lab at Indiana University, and has extensive expertise in both development and networking.

GMOC also hired one of the two hourly undergraduate assistants to help with the project. The second position is available, but has not yet been filled.
Equipment: The budgeted equipment need to test and develop GMOC software has been specified and is in the procurement process.

Q2:

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.

Q3:

Our activities during Q3 were spent chiefly in 3 areas:

Increased Collection from PlanetLab:

Collection of operational data from PlanetLab began on May 18, 2009. This data includes data from all the PlanetLab (proper) nodes and slices. Only a subset of



the required data as specified by the data requirements is being collected at the moment.

Development/Discussion with ProtoGENI:

Initial work of integration with ProtoGENI stated. We created a ProtoGENI Rspec to GMOC data exchange translator. This allowed GMOC to store ProtoGENI's static topology. Initial work on sharing of operational data has been done but the type and detail of this data is still pending.

Development effort on Prototype GENI operational portal view:

In anticipation of GEC5 demo, we developed an initial prototype of a GENI Operational View system. We see 4 main sets of users for this system:

- 1.) Those interested in viewing GENI data as a whole, such as interested outside parties, GPO, etc.
- 2.) Those interested in viewing GENI data about a particular cluster
- 3.) Those interested in viewing GENI data about a particular aggregate or component
- 4.) The Research user who will be interested in viewing GENI data about their particular slice.

This early prototype focuses primarily on visualizing and graphing GENI clusters, and graphing GENI slices. Because the bulk of our data comes from PlanetLab, this is the data we've focused on showing. As this prototype evolves, we will focus on more data from other projects, topology, and especially on providing the views into slices, including better visualization, for the researcher users of GENI.

Q4:

During this quarter, we made several improvements to the GMOC prototype. These improvements included:

- Work to increase performance of the prototype application
- Work to increase stability
- Made modifications to allow compatibility with Internet Explorer 8 and Opera browsers
- General code clean-up / standardization

During the quarter, we also made contacts with GpENI to host a GpENI node at GMOC, attempted further data exchange with ProtoGENI, and assisted with GENI installs in the Internet2 Network. These activities are detailed in later 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, Senior GMOC Programmer
Dan Doyle, GMOC Undergraduate
Nick Baker, GMOC Undergraduate
Jeff Catania, GMOC Undergraduate

Publications & Documents

Q1:

Herron, JP. (2008, October 20-22). *Developing the GENI Meta-Operations Center*. Presented at Chinese American Networking Symposium, Indiana University, Indianapolis, IN.

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.

Herron, JP., Fowler, L., Small, C. (2008, October 28-30). *Introduction to the GENI Meta-Operations Center*. Presented at GENI Engineering Conference 3, Palo Alto, CA.

This presentation was a required introduction to the GMOC project concentrating on the deliverables, approach, and expectations this will have on other GENI projects.

Herron, JP., Fowler L. (2008, December 8-12). *The GENI Meta-Operations Center*. Poster presented at IEEE eScience 2008 conference, Indianapolis, IN.

This was a presentation of a poster on GMOC. During this session, we fielded questions from reporters and others interested in GMOC and in GENI.

GMOC Website: A website for the GENI Meta-Operations Center was launched at <http://gmoc.grnoc.iu.edu>. This is also available via the main GRNOC page at <http://www.grnoc.iu.edu>. This site initially contains information about the GMOC project's goals, and approach, including the proposal, presentations, and other introductory information. As the project accelerates, this will also house information about which GENI projects are sharing data with GMOC, and



what kinds of data each is sharing. It will also be the site containing the GENI-wide operational system view data as it's developed.

Q2:

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)

Q3:

Herron, JP. (2009, June 10). *The GENI Meta-Operations Center*. Presented at Terena Networking Conference 2009, University of Malaga, Malaga, Spain.

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. Specifically, this related to projects like Federica in Europe.

GENI Documents:



Viecco, Camilo H. Proposal (2009 June 15) : *Use of URN's as GENI Identifiers*, Published at <http://gmoc.grnoc.iu.edu/gmoc/file-bin/urn-proposal3.pdf> and presented to the Cluster C for integration.

The GENI project current uses GGID for identifiers of GENI resources. This is problematic for several reasons such as the bundling of identification and authentication information and the inability to sub-allocate the UUID id space. This problems are documented and a proposal to solve this problems by using URNs and a grammar for these GENI URN's is presented. This proposed format is currently being adopted by the ProtoGENI cluster.

Q4:

GENI Documents:

Herron, JP. Fowler, L. Viecco, C. Small, C. *Emergency Shutdown*, published at <http://groups.geni.net/geni/attachment/wiki/GENIMetaOps/Emergency%20Shutdown%20System%20Description3.pdf>

The GMOC Emergency Stop document proposes a set of conditions and procedures for early Emergency Stop activity in GENI.

Herron, JP. Fowler, L. Viecco, C. Small, C. *GMOC: GENI Concept of Operations*, published at http://groups.geni.net/geni/attachment/wiki/GENIMetaOps/GENI_Concept_of_Operations-final.pdf

The GENI Concept of Operations Document proposes an operational model for GENI infrastructure for problem reception, triage, handling, resolution, and reporting.

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

Q1:

GMOC has began the process of collaboratively working to define the GENI Common Operational Dataset with the other GENI projects that will be sharing data with GMOC. Initial discussions are with the main control framework cluster projects, as well as infrastructure projects that have existing relationships with GRNOC, such as Internet2, NLR, and MAX GigaPoP. These initial discussions should open the door to further dialogue on technical details and sharing of existing data information with GMOC.

GMOC has formulated the initial questions to discuss with these projects, and requested time to discuss with PIs or their assigned delegates on these topics.

ProtoGENI: GMOC has had a discussion with ProtoGENI, and gathered a good deal of information and data produced by ProtoGENI that will be relevant to GMOC efforts. Capabilities are already in place in EmuLab for Emergency Stop and many of the topology and status data GMOC will need. More detailed discussions and GMOC research into the data ProtoGENI provided us will reveal how best to interact. But initially, it appears that we will be able to bring up operational data sharing with ProtoGENI quickly.

Mid-Atlantic Crossroads/DRAGON: GMOC has had discussions with MAX GigaPoP, and formulated a plan for further collaboration. This collaboration will involve a combination of direct polling of MAX components by GMOC and gathering of data already collected from this equipment by the DRAGON project. Further discussions will continue, with actually data gathering beginning sometime in January.

NLR: GMOC has also had discussions with NLR, and has come up with an initial plan for data sharing on NLR components. Since these components are already monitored and measured by the Global Research NOC at IU, NLR has agreed to allow GMOC to make use of any data that is already publicly available for GENI use. This should allow extensive data sharing on NLR components, but if netflow (or similar user) data is needed, further discussion would be necessary. Using the direct data collections should allow these components to be monitored much more easily and quickly.

Q2:

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

Q3:

Worked with SPP and ProtoGENI groups to deploy hardware to be co-located in the Internet2 network.

Discussed collaboration with Franz Fidler of the Embedded Real-Time measurement project on optical level measurement, including how substrate measurements can be integrated into the GMOC Operational datasets. Showed optical measurements tools developed and data collected by the GlobalNOC.

Discussed topology format options and identifier format with other GENI aggregates. Discussed collaboration with Iliia Baldine for future data sharing with ORCA.

Similar to Q1 & Q2 GMOC worked on gathering input and developing discussions with many GENI stakeholders during this quarter.

PlanetLab: GMOC continued to work with PlanetLab to gather operational data about the Planetlab project itself, working to fix some minor bugs in the PlanetLab software, and working out a reasonable interval for operational data polling

ProtoGENI:

GMOC worked with ProtoGENI to collect Topology data from Emulab-Utah. GMOC also worked to ProtoGENI for the use of a common URN based identifiers for GENI resources. Started discussions for determining what operational data can be shared and the periodicity for polling data.

RSPEC Meeting:

Discussed with the Rspec Meeting attendees on how topology data is represented (NDL, NDL variations and GMOC exchange format) and continues discussion on identifiers for GENI resources.

Measurement Meeting: Discussed with Measurement Workshop attendees how measurement plane data, when available, may be integrated into the GMOC data sets.

Q4:

Continued discussion with projects about data sharing: GMOC continued discussing data sharing activities with various GENI projects. More data sharing activities are expected to occur based on milestones for GMOC and other GENI projects during Y2.

GpENI: Collaboration with GpENI to host GpENI node at GMOC for investigation of integration of GpENI data into GMOC data sets. Hardware expected to be deployed during Q1 of Y2.

GENI infrastructure installs: GMOC assisted in the planning and installation of the GENI infrastructure installations in the Internet2 network. GMOC is investigating possible data sets that can be gleaned from this infrastructure for inclusion in the GMOC data sets.

Planned Activities for Y2

We plan to complete the following activities during Year 2. This list of activities is copied from our Year 2 Statement of Work.

1. Complete Spiral 2 Prototyping Plans. Document plans for gathering and sharing operational data with Spiral 2 control framework projects, and with several exemplar projects that cover different types of aggregate data. Initial plans will include at least 1) going over the Common Operational Dataset document, and developing a plan to gather as much of this data as each project can produce, in the format of their choosing; 2) going over the Operational Data Format and developing a plan to collect any data projects can produce in that format. In addition to the control framework projects, the exemplar projects will cover these technologies (candidate projects listed in parentheses. If candidate projects are unable to complete these tasks, alternate projects will be identified):
 - 1) backbone and regional networks (I2, NLR, MAX) (December 30, 2009)
 - 2) compute resources (CMU)
 - 3) sensors (Kansei)
 - 4) wireless networks (TBD after Sol2 funding finalized)
 - 5) cloud computing (TBD after Sol2 funding finalized)
 - 6) end-user opt-in (Million-node GENI)(December 31)
2. Prototype manual emergency shutdown. Implement initial emergency shutdown process with control framework and exemplar projects. This process will be based on the emergency stop documentation from the



GMOC's spiral 1 work. Tasks for this milestone include:

- 1) Develop database of contact information for emergency shutdown with each project
 - 2) Discuss expectations with each project for emergency shutdown process (response times, feedback loop, etc)
 - 3) Conduct an initial drill to test the process
- (March 30)

3. Prototype slice view of operational data sharing. Develop a per-slice operational view capability. View should include topology of a specified slice, including information about components being used, active operational status of slivers and components, and utilization data for slice traffic. This prototype will be enhanced to use data from real exemplar projects if appropriate data is available. (March 30)
4. Complete revised Concept of Operations document. Revise the Concept of Operations document drafted in Spiral 1 based on lessons learned from exemplar interactions, emergency shutdown, and visualization prototypes. May also include any new drafts for Common Operational Data set document and Operational data format documents. Document will be reviewed and discussed at GEC8. (May 31)
5. Complete Spiral 2 prototype data sharing system. Complete implementation and testing and put prototype data sharing in place for all projects identified in milestone 1 Prototyping Plan. Data collected will include either the full set of required data in the Common Operational Dataset document or a subset of critical data, depending on each project's capabilities, and the discussions in milestone 1. (September 30)
6. Prototype multi-cluster GENI-wide visualization. Generate a prototype GENI map that visualizes multiple clusters at a time, in order to show integration within and between clusters, and the activity of slices across clusters. (September 30)