

GMOC: GENI Concept of Operations

Executive Summary

The GENI facility will require a model of operations moving forward that is both **responsive** enough to the needs of GENI users and stakeholders (researchers using the facility, users opting-in to GENI, and GENI operators) and **flexible** enough to evolve with the federated, reconfigurable nature of the facility.

We expect that GENI research users will require a facility that is highly available, with a clear and simple process for getting support, which provides a robust set of highly detailed data about its use and operations, and which communicates maintenance and experiment affecting issues in a responsive and transparent way.

Many options for how GENI might provide operational services have been discussed. The options discussed so far fit one of three main approaches. First, GENI may choose to have little or no central GENI operational focus. Individual aggregates and control frameworks would each be responsible for operations of their own pieces of the GENI facility. Second, GENI could choose to have a centrally-managed operational center, which would be responsible to a large extent with monitoring and operating the facility. This could be in concert with operators of individual components, but the central GENI operations center would have an active role in operating GENI components. Third, GENI could choose to leave operations of GENI components and systems to individual GENI aggregates and control frameworks, but provide an integrated suite of services to coordinate operations, and provide the unified interface for GENI operations to the GENI users and stakeholders.

This third approach, the “Meta-Operational” model, represents the best way to leverage the existing expertise and in-depth understanding of the GENI components, while also providing the appropriate integrated operational interface for the GENI community. It also would ensure transparency, consistency, and accountability.

Within this general approach, there are many duties can be divided, and this needn’t be spelled out in great detail. Rather, the details should be left to evolve as GENI does. The focus now should be on the immediate needs of GENI as the facility begins to support real experiments in Spiral 2.

At times during the evolution of GENI, in Spiral 2 and later, events will “trigger” certain operational requirements of GENI. In Spiral 2, a highly active community of GENI experimenters is anticipated. This will trigger the requirement for a unified

contact for experimenters to report trouble, as well as the system to connect those GENI experimenters with the correct GENI federate operators who will troubleshoot and repair any issues. Other possible Spiral 2 trigger events include connecting GENI experiments to external networks, inter-cluster experimentation and the deployment of a GENI-wide federated security infrastructure.

This means that the earliest focus for operations, to be implemented in Spiral 2, must focus on:

- coordination of contact information
- basic help desk functionality
- basic reporting on health of the facility and slices
- processes for initial emergency shutdown

Introduction

Key Stakeholders for GENI Operations

There are many stakeholders for GENI operators: The opt-in users of GENI, GENI experimenters (both those running experiments as well as those planning experiments), the site operators, the project operators, the cluster operators, the operators of the common GENI infrastructure, and the GPO. Each of these groups has different requirements and would most likely desire a different perspective into GENI and a different event notification policy.

Anticipated User Requirements

Opt-in users would like to know the status of their GENI and how it affects their use of GENI. Experimenters would like to know the status of their experiments, side effects caused by other experiments and notification of scheduled and unscheduled downtimes for resources that their experiments depend on. Site operators would like to see how their resources are being used and by whom, receive notifications for interconnect downtimes and schedule maintenance for their systems. Project operators would like to observe the complete network. A summary of the user requirements can be seen in table 1:

	Scheduled Maintenance Notification	Unexpected downtime notification	Primary visibility	Shutdown/Reset requests?
Experimenters (planning)	No	No	All nodes/interconnects	No
Opt-in Users	No	Opt-in (for local)	Their connection/nodes/experiment	No
Experimenters (running)	Own	Own	Their experiment (slice)	Slice, Sliver

Site operators	Own	Own and projects	Their nodes and local slices	Local Slice
Project operators	Their sites and interconnects	Project +cluster common and interconnects	Their slices/nodes	Project Slice, node, interconnect
Cluster operators	Their projects, sites and interconnects	Cluster wide resources and interconnects	Their slices/nodes	Cluster slice, node interconnect
GENI operators	All	Any cross-cluster used resource	All slices/nodes	All slices, nodes and interconnects
GPO	No	No	All slices/nodes	Interconnects and unresponsive projects

Table 1. Expected GENI user requirements.

Operational Model

Generally, the options for a GENI operational model fall along a continuum between a Centralized Model and a Totally Distributed Model. Neither of these two is likely the best model, nor even possible in the GENI environment, but understanding each will help us to identify a good potential operational model within this continuum.

Centralized Model

At one end of the continuum, the Centralized Model would require all components on GENI to be operated by a central GENI operations organization. This central GENI organization would directly investigate, monitor, and troubleshoot GENI components system-wide.

Advantages:

- Single Contact Point for GENI users to report problems
- Maybe allows for faster problem investigation, with no hand-offs

Disadvantages:

- Requires a single entity be capable of troubleshooting every GENI component, which is likely highly difficult and expensive, if not impossible
- Would greatly hinder GENI's ability to add interesting resources quickly and would slow each project's innovation by adding overhead to ready new resources for central operations.
- Doesn't take advantage of the first-hand expertise of each project's staff to understand their own components
- Doesn't leverage the existing operational infrastructure already in place in some projects.

Totally Distributed Model

At the other end of the continuum, operations would all be handled by the individual GENI projects. Each GENI project would be responsible for fielding problem reports and requests from users, monitoring, troubleshooting, and reporting for their components.

Advantages:

- Leverages existing projects
- Allows maximum flexibility for projects
- Might lower risk somewhat by allowing multiple operational methods to be prototyped.

Disadvantages:

- Puts the work of understanding GENI architecture and organization onto the GENI user, who would be required to request help from the right group(s).
- Standards, service levels, operational methods, operational data sharing, and reporting will be inconsistent across projects
- Requires each GENI project to understand the GENI architecture well enough to understand inter-relation of projects.
- Complex problems that span projects will require projects to coordinate on their own on an ad-hoc basis, which could delay problem resolution from frequent hand-offs.
- As GENI begins to interconnect with other networks, those other networks will likely require a single operational point of contact.

Given these advantages and disadvantages, as well as the initial expected requirements of GENI operations, a good initial model would be one that provides the Centralized Models benefits of a unified operational interface to GENI users, as well as a consistent level of service, communication path, and coordination of multiple inter-connected GENI projects with the Distributed Models benefits of allowing flexibility and innovation within each project, leveraging projects' existing operations structures and expertise about their own components, and allowing prototyping of multiple operational models.

One model that would help balance these considerations would be a Meta-Operational Model.

The Meta-Operations Model

In this model, component and aggregate level operational tasks would be distributed. Each project would be responsible for the detailed operations and maintenance of the aggregates and components for that project. This would be done

using whatever methods are appropriate for that particular project, which would ensure flexibility and maximize reliability of each project individually.

In addition, overall coordination, service levels, data collection, and reporting for GENI would be performed by an organization outside of any particular project. This organization would be responsible for ensuring that the operations for each project are operating in a way that best serves users of the GENI facility as a whole. In this way, this organization performs “Meta-Operations” duties.

Specific Tasks for this organization would include:

- Serve as the operational “front-door” for GENI, doing basic monitoring of availability of GENI resources and fielding problem reports for GENI from GENI stakeholders
- Serve as the primary communication path among all related GENI projects and users to minimize confusion and maximize visibility
- Directing problem reports to the appropriate GENI project operational contacts
- Provide a unified ticketing system to track problems GENI-wide
- Provide a unified place for GENI stakeholders to look for operational data in a consistent format
- Gather GENI-wide statistics and reporting on availability, usage, etc for GPO & other interested GENI stakeholders
- Facilitate communication between multiple parties for complex problems that span multiple projects
- Perform some GENI-wide operational duties such as emergency shutdown or isolation.

In essence, this model would allow for projects to work freely with minimal operational overhead on their own projects, with the Meta-operations organization providing the coordination and communications among projects, the consistency in service levels and reporting, and the unified front-door interface for the GENI community for operations.

Key Task Areas

Event Notification

One key area for a GENI operational system is event notification. A meta operations organization would be responsible for several event notification tasks, including:

- Receive scheduled outage notifications from infrastructure providers; interpret those notifications; notify appropriate GENI parties of the events.
- Detect unscheduled outage events and receive reports of unscheduled outages. Create event notification for these outage events, analyze impact,

and notify the appropriate parties (experimenters, infrastructure providers, international federates, etc.)

Problem/Question Report Reception and Tracking

A meta operations organization should serve as an entry point for experimenters or other GENI participants who are having a problem with GENI infrastructure. GENI participants should be able to treat a meta operations organization as a single point of contact (by phone, e-mail, and other mechanisms) to report problems related to their use of GENI infrastructure.

A meta operations organization should also serve as a “home” for tracking of these requests by providing a central trouble ticket system which all parties involved in the issue can document current status, next steps, resolution, etc.

Problem Triage

A meta operations organization would be responsible for initial problem triage and for engaging the appropriate GENI resources needed to fix the problem. This includes categorizing the problem, determining severity and priority, and determine parties to be involved in addressing the problem. Using a standard set of procedures, using the severity/priority as an input, a meta operations organization would then engage the appropriate parties to address the problem, and perform time-based problem escalation based on procedures agreed upon by the GENI community.

Federate Coordination

It is expected that many experimenters using the GENI infrastructure will be stitching together slivers from multiple, separately managed pieces of GENI infrastructure in order to form a coherent slice. A meta operations organization would function as a coordination center, connecting experimenters with appropriate infrastructure operators to answer questions, resolve problems, or to connect multiple federates together to resolve an issue with an experimenter whose experiment spans the federates’ infrastructure.

Operational data sharing

A meta operations organization can provide a unified view of operational data across the GENI infrastructure. Such a view is useful for multiple audiences and purposes.

A GENI experimenter who has a slice that spans multiple aggregates will benefit from a single view of their slice which depicts data such as operational status, network link utilization, etc.

The GPO, NSF, and other interested parties will have a need to determine how GENI is performing as a whole, how it is being utilized at a high level, and how its' use has changed over time.

A GENI infrastructure provider may have a need to view operational data pertaining to 'adjacent' infrastructure.

All of these use cases rely on a common set of GENI operational data. Different views of this data should be created in order to facilitate the different audiences for GENI operational data.

Operational Reporting

A meta operations organization would have a unique view into the operational state of the GENI infrastructure as a whole. One important function of such an organization would be to gather and collate this data. A meta operations organization would produce tools to allow interested parties to view data about the availability, reliability, and performance of the GENI infrastructure.

Operational Triggers

As developments are made in the development of GENI there will be new requirements of the operations of the GENI infrastructure.

Active Experimentation

The use of active experiments on GENI will require a number of changes in the way GENI Operations will be handled. Live end-to-end experimentation is a publicized GENI goal before the end of Spiral 2.

Experimenters will need a method to report issues with the GENI infrastructure and software and have the report forwarded to the correct party. This in turn will require GENI Operations to be able to have contact information or a mechanism to forward the request. It will also require collection of who can validly submit requests to GENI Operations. The point or points of contact for GENI operational Requests will need to be publicized and well known to the GENI community.

An active experiment creates the need for notifications when there is a change, maintenance or outage in the underlying substrate that may affect an experiment. In order to notify each experimenter GENI Operations would need a method to contact all experimenters affected.

Inter-Cluster Experimentation

The creation of experiments that span multiple clusters present another challenge to the operation of GENI. A problem may span multiple projects or clusters may need the contact and cooperation of a number of groups to troubleshoot, determine the location of a problem and correct it. Cross-Cluster experimentation may be assisted by verification of the correct operation of the cluster and its components by a meta operations organization.

External Interconnections

Experiments that interact with external networks are a critical issue. External networks have an expectation that any interconnections are reliable and will not cause any disruption in the operation of their networks. This includes the use of the GENI connection to spread malware and other malicious traffic.

A meta operations organization would be responsible to receive reports of issues emanating from GENI that are affecting external networks. The reports will be passed to the appropriate contacts. The meta operations organization may have the ability, as GENI further progresses, to disable the specific connection to isolate the problem.

Security

The deployment of a GENI-wide security infrastructure in Spiral 2 is uncertain. However as a security infrastructure is deployed, an operations function would need to be closely integrated. Operational data collected would need to be protected using the GENI security infrastructure. Also authentication and authorization of request from GENI users would also need to be integrated with the security infrastructure.

A meta operations organization would also be responsible of the determination if the GENI components are operating according to the security policy determined by the GPO and the GENI community. Notification of violations of the security policy would be distributed to the appropriate parties.

Operational Requirements in Spiral 2

The most important task on spiral 2 is on governance. Projects must agree on: a common classification on incidents/requests; a common set of policies on what operations will they support and the timeframes for action when faced with different incidents/requests by category; agreement on what mechanisms for authentication and authorization are acceptable; determination of what data is

public and private; and a commitment for a common interface for interoperation at the control plane level.

In the technical side, the tasks would be to: augment the reporting/tracking systems to use the agreed authentication and classification mechanisms; augment the global visibility of GENI to reduce false/incorrect requests; augment the global visibility of GENI to better interoperate and facilitate action when requests are received and to begin work on automation of requests for those of lowest impact and priority.