

GENI

Global Environment for Network Innovations

System Requirements Document

Document ID: GENI-SE-SY-RQ-01.9

January 16, 2009

Prepared by:

The GENI Project Office

BBN Technologies

10 Moulton Street

Cambridge, MA 02138 USA

NSF Cooperative Agreement CNS-714770

DRAFT - FOR DISCUSSION ONLY

TABLE OF CONTENTS

| | | |
|-------|--|----|
| 1 | DOCUMENT SCOPE..... | 5 |
| 1.1 | PURPOSE OF THIS DOCUMENT..... | 5 |
| 1.2 | CONTEXT FOR THIS DOCUMENT..... | 5 |
| 1.3 | RELATED DOCUMENTS..... | 5 |
| 1.4 | NATIONAL SCIENCE FOUNDATION (NSF) DOCUMENTS..... | 5 |
| 1.5 | GENI DOCUMENTS..... | 5 |
| 1.6 | STANDARDS DOCUMENTS..... | 6 |
| 1.7 | OTHER DOCUMENTS..... | 6 |
| 1.8 | DOCUMENT REVISION HISTORY..... | 6 |
| 1.9 | ACRONYMS..... | 6 |
| 2 | GENI OVERVIEW..... | 8 |
| 2.1 | GENI SUBSYSTEMS..... | 9 |
| 3 | INTRODUCTION TO THIS DOCUMENT..... | 11 |
| 3.1 | GENI SYSTEM REQUIREMENTS..... | 11 |
| 3.2 | ON REQUIREMENTS..... | 11 |
| 4 | SOURCES OF GENI SYSTEM REQUIREMENTS..... | 13 |
| 4.1 | RESEARCH AND EDUCATION..... | 13 |
| 4.1-1 | Network Science and Engineering Research..... | 13 |
| 4.1-2 | Education..... | 13 |
| 4.2 | PROJECT STRATEGY..... | 13 |
| 4.2-1 | Leverage existing technology, infrastructure, and organizations..... | 13 |
| 4.2-2 | Reliable, affordable operation..... | 13 |
| 5 | ARCHITECTURAL REQUIREMENTS..... | 14 |
| 5.1-1 | Programmability..... | 14 |
| 5.1-2 | Sliceability..... | 14 |
| 5.1-3 | Federation..... | 14 |
| 5.1-4 | Aggregate federation..... | 14 |
| 5.1-5 | Clearinghouse federation..... | 14 |
| 5.1-6 | Slice-based experimentation..... | 15 |
| 5.1-7 | Technology heterogeneity..... | 15 |
| 5.1-8 | Future technology insertion..... | 15 |
| 5.1-9 | Multiple Clearinghouses..... | 15 |
| 6 | EXPERIMENT SUPPORT REQUIREMENTS..... | 16 |
| 6.1 | RANGE OF RESEARCHERS AND EXPERIMENTS..... | 16 |
| 6.1-1 | Range of researcher expertise..... | 16 |
| 6.1-2 | Range of experiment lifetimes..... | 16 |
| 6.2 | EASE OF USE..... | 16 |
| 6.2-1 | Resource discovery & scheduling..... | 16 |

| | | |
|--------|--|----|
| 6.2-2 | Software sharing & re-use..... | 16 |
| 6.2-3 | Software development tools..... | 16 |
| 6.3 | SLICE MANAGEMENT..... | 17 |
| 6.3-1 | Slice management tools..... | 17 |
| 6.3-2 | Slice data containment..... | 17 |
| 6.3-3 | Slice resource isolation..... | 17 |
| 6.3-4 | Growing and shrinking slices..... | 17 |
| 6.3-5 | Slice composition..... | 17 |
| 6.4 | EXPERIMENTAL REALISM & CONTROL..... | 18 |
| 6.4-1 | Repeatability..... | 18 |
| 6.4-2 | Realistic Environments..... | 18 |
| 6.4-3 | Intentional failure and/or degradation..... | 18 |
| 6.4-4 | Artifacts..... | 18 |
| 6.4-5 | Resource allocation feedback..... | 18 |
| 6.4-6 | Virtualizing management interfaces..... | 19 |
| 7 | INSTRUMENTATION AND MEASUREMENT REQUIREMENTS..... | 20 |
| 7.1-1 | Component measurements..... | 20 |
| 7.1-2 | Privacy of measured data..... | 20 |
| 7.1-3 | Link measurements..... | 20 |
| 7.1-4 | Operational data..... | 20 |
| 7.1-5 | Measurement transmission and storage..... | 20 |
| 7.1-6 | Measured data access..... | 20 |
| 7.1-7 | Real-time access to measurements..... | 20 |
| 7.1-8 | Component locations..... | 21 |
| 7.1-9 | Time services..... | 21 |
| 8 | USER OPT-IN REQUIREMENTS..... | 22 |
| 8.1-1 | Low barrier to entry for “opt in”..... | 22 |
| 8.1-2 | Opt-in user software installation..... | 22 |
| 8.1-3 | Internet opt-in..... | 22 |
| 8.1-4 | Native GENI opt-in..... | 22 |
| 8.1-5 | Anonymous Opt-in..... | 22 |
| 8.1-6 | Opt-in shutdown..... | 22 |
| 9 | SYSTEM SIZING REQUIREMENTS..... | 23 |
| 9.1-1 | Numbers of researchers..... | 23 |
| 9.1-2 | Numbers of concurrent experiments..... | 23 |
| 9.1-3 | Infrastructure scale..... | 23 |
| 10 | OPERATIONS AND SECURITY REQUIREMENTS..... | 24 |
| 10.1 | POLICY SUPPORT..... | 24 |
| 10.1-1 | Clearinghouse-wide Policies..... | 24 |
| 10.1-2 | Owner’s Policies..... | 24 |
| 10.2 | RELIABLE, PREDICTABLE, AND TRANSPARENT OPERATIONS..... | 24 |
| 10.2-1 | Reliable operations..... | 24 |

| | | |
|--------|--|----|
| 10.2-2 | Predictable operations | 24 |
| 10.2-3 | Visible operational status | 25 |
| 10.2-4 | Help Desk | 25 |
| 10.2-5 | Federated event escalation | 25 |
| 10.2-6 | Federated operations data exchange..... | 25 |
| 10.2-7 | Record keeping duration | 26 |
| 10.3 | SECURE OPERATIONS..... | 26 |
| 10.3-1 | System-wide policy and guidelines for operational security | 26 |
| 10.3-2 | Control framework security | 26 |
| 10.3-3 | Identities..... | 26 |
| 10.3-4 | Identity records..... | 26 |
| 10.3-5 | Authorization..... | 26 |
| 10.3-6 | Accountability | 27 |
| 10.4 | DETECTION AND REMEDIATION OF BAD EXPERIMENTS | 27 |
| 10.4-1 | Rapid detection and decisions | 27 |
| 10.4-2 | Swift, effective neutralization | 27 |
| 10.4-3 | Restricted mode..... | 27 |
| 10.4-4 | Abuse reporting | 27 |
| 10.4-5 | Forensics..... | 28 |

1 Document Scope

1.1 Purpose of this Document

This document specifies GENI system requirements. This document specifies requirements for the system as a whole: that is, all requirements listed in this document pertain to the overall system. They form the basis of further derived requirements that then flow down to the various subsystems, which are in turned captured in Requirements Documents for those subsystems.

1.2 Context for this Document

Figure 1-1 below shows the context for this document within GENI’s overall document tree.

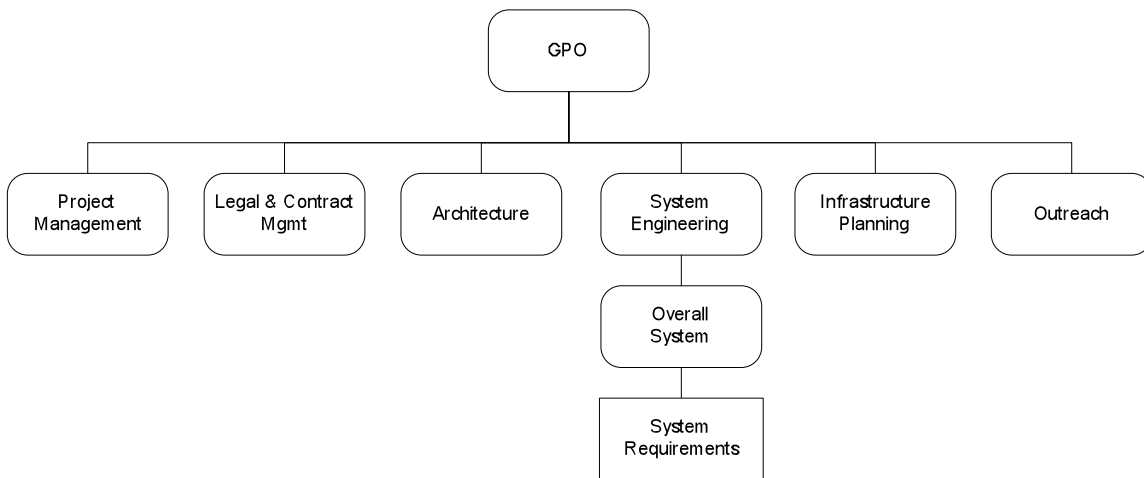


Figure 1-1. Location of this document within the GENI Document Tree.

1.3 Related Documents

The following documents of exact date listed are related to this document, and provide background information, requirements, etc., that are important for this document.

1.4 National Science Foundation (NSF) Documents

| Document ID | Document Title and Issue Date |
|-------------|-------------------------------|
| N / A | |

1.5 GENI Documents

An initial requirements document was published by the GENI Planning Group as GDD-07-46. This document revises that one, in some places substantially, based on the subsequent maturation of the conceptual design that has taken place since the earlier document’s publication.

| Document ID | Document Title and Issue Date |
|---------------------|--|
| GDD-06-28 | "GENI Research Plan" version 4.3, January 2, 2007 [obsolete] |
| GDD-07-46 | "GENI System Requirements Document" April 2007 [obsolete] |
| GDD-06-08 | "GENI Design Principles," August 2006 |
| GENI-SE-SY-MP-01.1 | "GENI System Engineering Management Plan," May 2008 |
| GENI-SE-SY-SYO-02.0 | "GENI System Overview," version 2.0, September 29, 2008 |

1.6 Standards Documents

| Document ID | Document Title and Issue Date |
|-------------|-------------------------------|
| N / A | |

1.7 Other Documents

| Document ID | Document Title and Issue Date |
|-------------|-------------------------------|
| N / A | |

1.8 Document Revision History

The following table provides the revision history for this document, summarizing the date at which it was revised, who revised it, and a brief summary of the changes.

| Revision | Date | Revised By | Summary of Changes |
|----------|-----------|------------|---|
| 01.1 | 01-May-08 | A. Falk | Initial draft; a major revision of GDD-07-46 |
| 01.2 | 02-Jun-08 | A. Falk | Responses to Heidi's comments, front matter fixes, start on Discussion sections |
| 01.3 | 3-Sept-08 | A. Falk | Major revision based on GPO, TCG review |
| 01.4 | 18-Sep-08 | C. Elliott | Some reorganization, minor rewording, a few more reqt's |
| 01.5 | 14-Oct-08 | A. Falk | Add'l GPO review comments |
| 01.6 | 15-Oct-08 | A. Falk | Add'l GPO review comments |
| 01.7 | 20-Oct-08 | A. Falk | Fixed footer |
| 01.8 | 13-Jan-09 | A. Falk | Revision based review by T. Faber |
| 01.9 | 16-Jan-09 | A. Falk | Further iterations with T. Faber |

1.9 Acronyms

TBD: To Be Determined

TBR: To Be Reviewed

TBS: To Be Specified

2 GENI Overview

The Global Environment for Network Innovations (GENI) is a novel suite of infrastructure now being designed to support experimental research in network science and engineering.

This new research challenges us to understand networks broadly and at multiple layers of abstraction from the physical substrates through the architecture and protocols to networks of people, organizations, and societies. The intellectual space surrounding this challenge is highly interdisciplinary, ranging from new research in network and distributed system design to the theoretical underpinnings of network science, network policy and economics, societal values, and the dynamic interactions of the physical and social spheres with communications networks. Such research holds great promise for new knowledge about the structure, behavior, and dynamics of our most complex systems – networks of networks – with potentially huge social and economic impact.

As a concurrent activity, community planning for the suite of infrastructure that will support NetSE experiments has been underway since 2005. This suite is termed the Global Environment for Network Innovations (GENI). Although its specific requirements will evolve in response to the evolving NetSE research agenda, the infrastructure's conceptual design is now clear enough to support a first spiral of planning and prototyping. The core concepts for the suite of GENI infrastructure are as follows.

- **Programmability** – researchers may download software into GENI-compatible components to control how those components behave;
- **Virtualization and Other Forms of Resource Sharing** – whenever feasible, nodes implement virtual machines, which allow multiple researchers to simultaneously share the infrastructure; and each experiment runs within its own, isolated slice created end-to-end across the experiment's GENI resources;
- **Federation** – different parts of the GENI suite are owned and/or operated by different organizations, and the NSF portion of the GENI suite forms only a part of the overall 'ecosystem'; and
- **Slice-based Experimentation** – A slice is a collection of resources that have been allocated and configured together for the purpose of experimentation. It can contain resources from diverse locations and administrative controls. Slices provide GENI's notion of an experimenter's collection of resources and are the fundamental entity that researchers use to define their experiments. Researchers will remotely discover, reserve, configure, program, debug, operate, manage, and teardown distributed systems established across parts of the GENI suite by manipulating slices and the resources connected to them.

As envisioned in these community plans, the GENI suite will support a wide range of experimental protocols, and data dissemination techniques running over facilities such as fiber optics with next-generation optical switches, novel high-speed routers, city-wide experimental urban radio networks, high-end computational clusters, and sensor grids. The GENI suite is envisioned to be shared among a large number of individual, simultaneous experiments with extensive instrumentation that makes it easy to collect, analyze, and share real measurements.

2.1 GENI Subsystems

This document treats GENI as an interconnected system of software and infrastructure suites under diverse ownership and management. Requirements allocated to the GENI system are intended to specify the minimal behaviors and characteristics necessary for coherent operations.

Whenever possible, the requirements below should permit diversity in technology and administration. For this reason, many possible systems may be developed from the requirements herein. Our goal is to constrain the development community only where necessary to provide a system of interest to the user community, permitting and encouraging innovation wherever possible.

It is not useful to allocate requirements to GENI purely as a black box, ignorant of any specific structure within. This requirements document has been written in the context of the conceptual architecture described in GENI-SE-SY-SYO-02.0. GENI will require the following major subsystems (see Figure 2):

Components & Aggregate Components: Components are devices that have resources which can be discovered, reserved, and programmed or configured by GENI researchers. Components may be organized into aggregates, which share common operations and administration. Researchers may be required to interface with aggregates in order to access components. Components may use virtualization to provide resource isolation between users.

Clearinghouses & Control Framework: Clearinghouses operate registries of all principals, slices, and components and are expected to implement some access control policies. They may include trust anchors, initial points of entry, and policy arbiters for federations of GENI components. The control framework includes mechanisms, tools, and services that permit researchers to reserve and obtain access to component resources.

Measurement Subsystem: This subsystem collects measurements from instrumentation and processes, archives them, and manages access to the resulting data. Data collection instruments are considered components, not part of the measurement subsystem.

Administration & Operations: This subsystem includes the tools, services, infrastructure, and staffing to enable contribution of new resources to GENI, response to reports of misbehavior within GENI, assistance to researchers attempting to use GENI, and publication of high level state of operations of GENI components.

Experimenter Tools & Services: This subsystem includes helper tools and services such as those needed to assist researchers in discovering and reserving resources; designing, composing, debugging, deploying, and growing experiments; managing experimental services; managing research teams and component access; and sharing experimental code and lessons-learned.




Figure 2. Major GENI Subsystems

3 Introduction to this Document

This document specifies the overall GENI system requirements. GENI will be designed to meet the needs of the network science and engineering research communities – its users¹. These communities should be considered the *customers* of the GENI system, and the system needs to perform according to their expectations.

3.1 GENI System Requirements

GENI's system requirements address the necessary functions of the system as well as issues relating to its operation, oversight, and evolution. Implementers are advised that purely meeting the letter of the requirement will likely result in an inadequate design. The requirements in this document are supplemented by a set of design principles in GDD-06-08. Design principles are not requirements but are intended to help inform designers as they weigh implementation approaches.

There are several sources of requirements. Where they come from and how they are captured and managed is discussed in the GENI System Engineering Management Plan [GENI-SE-SY-MP-01.1]. Briefly, requirements will come from the user communities, as expressed through activities of the NSF Network Science and Engineering (NetSE) Council; from the design community, as expressed in GENI working groups and by design and prototyping activities; and from the GENI Project Office.

This document specifies requirements for the system as a whole: that is, all requirements listed in this document pertain to the overall system. They form the basis of further derived requirements that then flow down to the various subsystems, which are in turn captured in Requirements Documents for those subsystems.

The organization of requirements isn't intended to allocate them to a particular subsystem or working group but to group them by subject to make this listing easier to read.

3.2 On Requirements

The definition of a requirement is important and merits consideration. The term *requirements* may be defined as “characteristics that identify the accomplishment levels needed to achieve specific objectives under a given set of conditions.” So, requirements dictate *what* the system needs to do and under what *conditions* the system is expected to do it – or more interestingly under what conditions the system might *not* behave as expected. Requirements become a contract between the GENI designers and users.

Requirements may be defined at many levels, and discipline is required to both avoid confusing the various levels and to avoid confusing the task of requirements definition with the task of designing the system itself. Requirements must communicate the needs of the system's users to the designers in a way that allows for unambiguous assessment as to whether the need has been met.

¹ Note that the GENI user community could range far from traditional computer scientists and include researchers from such fields as economy, policy, and sociology.

Here are some principles of requirements that guide the requirements that follow²:

- Requirements should be necessary. Unnecessary requirements overly constrain system designers and may eliminate cost savings or desired system capabilities.
- Each requirement should cover a single parameter. This will make the requirement simpler to understand and verify.
- Each requirement should stand alone. Requirements must be clear and understandable. Phrasing that is easily misinterpreted should be avoided.
- Requirements should be neither too severe nor too lenient. Over-specified requirements will overly constrain the system designers. Under-specified requirements may result in a system that does not meet the users' needs or is otherwise undesirable.
- Requirements should be non-conflicting. Designers must not be asked to do the impossible.
- Requirements should be verifiable. When writing a requirement it is important to consider whether it will be possible to objectively assess whether it has been achieved.
- Requirements should include a rationale. If a requirement is difficult or expensive to achieve, it is important to understand why it exists and the implications of not meeting it.
- Requirements should be phrased using minimal text and descriptive matter and should not include management or statement of work terms (e.g., "develop a safety plan").

² Many of the definitions and processes described here come from the "*TRW Systems Engineering Process Handbook*", 1995, unpublished.

4 Sources of GENI System Requirements

This chapter identifies the ultimate sources of requirements for the GENI system. They reflect the basic mission of the GENI system, and anchor this document's system requirements within their proper context. All GENI requirements should be ultimately traceable, directly or indirectly, to this chapter.

4.1 Research and Education

4.1-1 Network Science and Engineering Research

The GENI system shall support experimental research in network science and engineering, as identified in the "Network Science and Engineering Research Agenda" (TBS).

4.1-2 Education

The GENI system shall support education, as identified in a document TBD.

4.2 Project Strategy

4.2-1 Leverage existing technology, infrastructure, and organizations

The GENI system shall require minimal modifications to existing hardware devices, infrastructure, and software systems, as well as that which becomes available during the project lifetime, as well as minimal modifications to existing organizations, such as entities that own and operate infrastructure, operations staff, etc.

4.2-2 Reliable, affordable operation

Long-term operation of the GENI system shall be reliable and affordable.

5 Architectural Requirements

5.1-1 Programmability

The GENI system shall allow researchers to program and/or configure experimental equipment to maximum extent possible, except where it could violate safety, isolation, security, or operational guidelines.

Discussion: Programmability is a very important goal for GENI. This requirement is intended to maximize the programmability of components.

5.1-2 Sliceability

The GENI system will permit concurrent controlled sharing of resources - sliceability - where possible.

Discussion: Sliceability is a powerful strategy for permitting multiple researchers to share physical resources in ways appropriate to their experiments. An obvious mechanism for sliceability is virtualization of components that can support it. Other components may need to be sliced by partitioning resources such as memory or spectrum allocation.

5.1-3 Federation

The GENI system shall support federation, the ability to create slices that incorporate multiple independently administered resources, which resources may be owned and operated by a variety of administrative entities including universities, non-profit organizations, for-profit organizations, and governmental entities both within the United States and elsewhere.

Discussion: A key strategy for GENI's growth and ability to include new technologies is to make it easy to for facilities to connect to GENI. Federation should enable facility owners to make component resources available to GENI users without giving up ownership or complete control.

5.1-4 Aggregate federation

The GENI system shall permit components and aggregates owned or administered by one organization to make use of clearinghouses administered by another.

Discussion: This requirement is intended to permit contribution of resources without requiring the contributor to run their own clearinghouse. Because of the clearinghouse's role as trust anchor, its operator may have policies that an aggregate needs to meet before it will be allowed to join. Note that federation is not well understood at this time and this requirement may change as the concept matures.

5.1-5 Clearinghouse federation

The GENI system shall permit creation and operation of slices that span multiple clearinghouses.

Discussion: This will allow NSF portions of the GENI ecosystem to interoperate with non-US portions and/or commercial portions. For example, this may requirement will require supporting the

exchange of authorization information to enable GENI users access to federated components and users associated with federated components access to GENI. Note that federation is not well understood at this time and this requirement may change as the concept matures.

5.1-6 Slice-based experimentation

The GENI system shall support slice-based experimentation, in which a researcher experiment runs within an “end to end” slice that provides containment and some degree of isolation for the experiment.

5.1-7 Technology heterogeneity

The GENI system shall support experiments employing a wide class of computation, storage, and networking technologies, spanning the spectrum of wired and wireless technologies available today.

Discussion: Technology diversity is a goal for GENI. As is the explicit inclusion of wireless technologies. The GENI design should be technology agnostic but have sufficient diversity that research proposals can be stressed using a variety of technologies.

5.1-8 Future technology insertion

The GENI system shall include explicitly defined procedures and system interfaces to facilitate incorporation of additional technologies, including those that do not exist today.

Discussion: This will allow GENI to be useful to a broad range of researchers, remain useful beyond the project’s duration, support GENI’s role as a low-friction vehicle for deployment of new technologies by both academic researchers and industrial partners, and foster close collaboration between “device researchers” and “systems researchers.” Again, the GENI design should be technology agnostic.

5.1-9 Multiple Clearinghouses

The GENI system shall permit multiple Clearinghouses. Aggregates and researchers may be registered in more than one clearinghouse at the same time.

Discussion: This is an important capability that will allow concurrent diversity of resource allocation strategies, usage and registration policies, etc. (It isn’t clear yet whether slices should be registered in more than one clearinghouse at a time.)

6 Experiment Support Requirements

This chapter provides requirements as regards support of research experiments, such as the types of researchers that must be supported, tools and services provided, etc.

6.1 Range of researchers and experiments

6.1-1 Range of researcher expertise

The GENI system shall support researchers from a wide range of expertise, i.e., ranging from novice students to “power users.”

Discussion: This requirement will likely motivate a wide range of tools, archived sample code and experiments, educational materials, and diagnostics for new users. It would also motivate low-level programmatic interfaces for ‘power users’ to gain access to GENI’s more sophisticated capabilities.

6.1-2 Range of experiment lifetimes

The GENI system shall support slices from a wide range of lifetimes, i.e., ranging from 1 minute in duration to slices that run continuously for weeks to years.

Discussion: The ability to support long-lived experiments is a key goal for the GENI system as it is necessary to attract real users to experimental services. As experimental services mature become more stable, they may be considered less as experiments and as infrastructure itself. Criteria for how this occurs (and the implications) will need to be explored.

6.2 Ease of use

It is important that the ‘bar’ for GENI use be as low as possible. Power users may desire low level control and interfaces to GENI components but many users will benefit from tools that will make it easier to establish and manage slices. Open interfaces and public code repositories will enable the user/developer community to build and extend tools.

6.2-1 Resource discovery & scheduling

The GENI system shall provide tools for resource discovery and scheduling.

6.2-2 Software sharing & re-use

The GENI system shall provide tools and interfaces to facilitate sharing and re-use of software developed for research experiments.

6.2-3 Software development tools

The GENI system shall provide tools for configuring, managing, monitoring, and debugging components running experimental software and/or configurations.

Discussion: Note that in most cases GENI components will not be co-located with the programmer making remote debugging tools essential.

6.3 Slice Management

6.3-1 Slice management tools

The GENI system shall provide tools for configuring, managing, monitoring, and debugging end-to-end slices.

Discussion: A slice is a collection of resources on different components. Slice tools will therefore need to support coordination of heterogeneous technologies.

6.3-2 Slice data containment

The GENI system shall contain mechanisms for isolation among slices (except those deliberately interconnected) to ensure that system attacks in one portion of the system cannot “escape” and attack other experiments.

Discussion: It will be hard or maybe impossible to guarantee isolation under some conditions, e.g., when slices connect to the Internet. One way to view this requirement is that a component should be able indicate whether isolation is possible. This could allow composition of isolated slivers into an isolated slice. (What else would be required?)

6.3-3 Slice resource isolation

The GENI system shall support controlled isolation between slices so that they do not interfere with each other. These isolation mechanisms must be sufficiently robust to make reproducible experiments possible. Users must be able to assess the level of isolation they are receiving and the interference from other slices.

Discussion: The required amount of isolation will vary for different types of experiments. Rather than create specific isolation requirements, GENI will require that isolation mechanisms are included in shared resources and that users can understand the isolation they are receiving. Then users can make their own assessment as to whether it is sufficient.

6.3-4 Growing and shrinking slices

The GENI system shall provide mechanisms for adding resources to, or subtracting resources from, an existing slice, e.g., to grow or shrink an experiment.

Discussion: To support long-lived experiments a researcher should not be forced to tear down and re-create a slice to make changes to its composition. This would allow portions of the slice to keep running while slivers were added or deleted. Some operational disruption should be expected when a slice grows or shrinks. This requirement applies primarily to the control plane in that a researcher should not lose all a slice’s allocations when resources are added or subtracted.

6.3-5 Slice composition

The GENI system must support controlled interconnection of slices to each other and to the current Internet, allowing researchers to build directly on each other’s work, and to draw on existing Internet users and resources.

Discussion: The experimenters are going to have to make sure that their experiments can talk meaningfully to one another (or can be translated to do so). After all, these may be two non-IP

experiments running under different economic and technical models. Additionally, this may require some sort of ‘robustness’ metric so that a researcher who is making use of an existing slice has some idea of the stability of that slice.

6.4 Experimental Realism & Control

6.4-1 Repeatability

The GENI system shall support predictable and repeatable behavior for experiments running on some portions of the overall infrastructure suite.

Discussion: GENI needs to include sufficient mechanisms and control that repeatable experiments can run. This isn’t to say that all of the GENI needs this but some experiments will want things like controlled background traffic and the ability to have tight control on resources available. A possible derived requirement from this might be that each component or subsystem must determine if it is capable of repeatable behavior and, if so, should make it clear to users how to use it in that fashion.

6.4-2 Realistic Environments

The GENI system shall provide a realistic platform to test systems that range from centralized to distributed on a regional, campus, or end-node basis.

Discussion: A design goal for GENI is to provide a more realistic platform for research than can be found in testbeds and Internet overlays. A diverse set of network types is necessary to achieve this goal.

6.4-3 Intentional failure and/or degradation

The GENI system shall support intentional failure and/or degradation on command of any virtualized components. Failures may be single or en-masse to support simulations of massive infrastructure outages.

6.4-4 Artifacts

All unrealistic behavior in GENI components, such as timing jitter from virtualization or emulation, shall be identified, minimized, and specifically documented.

Discussion: users need to be able to learn of artifacts to understand how their experiment’s performance might be affected.

6.4-2 Resource allocation feedback

The GENI system shall provide feedback about what resources a slice actually receives to enable researchers to evaluate the validity of their results.

Discussion: In resource allocation it may be possible to request a range or unspecified resources. To achieve the goal of supporting repeatable experiments, it will be necessary that a researcher be able to learn the specific amount of resources are allocated to a slice.

6.4-3 Virtualizing management interfaces

To the extent possible, the GENI system shall virtualize all physical management interfaces used by system operators.

Discussion: An important benefit of virtualization is that researchers can get access to management interfaces of devices previously only available to the administrator. This should enable experimental capabilities, e.g., in the area of network management and device configuration.

7 Instrumentation and Measurement Requirements

This chapter provides high-level system requirements for instrumentation and measurement. These areas are currently poorly understood, and the following requirements are likely to evolve substantially as this area become better understood.

7.1-1 Component measurements

The GENI system shall support on-line collection, storage, and distribution of component measurements in support of measurement-based quantitative research.

7.1-2 Privacy of measured data

The GENI system shall support measurement collection mechanisms capable of meeting the privacy needs of experiments and end-users as described in [TBS] *GENI Privacy Policy*.

Discussion: The GENI Privacy Policy document may be helpful in establishing a set of consistent privacy controls structuring the interaction between GENI researchers and end-users of experimental services.

7.1-3 Link measurements

The GENI system shall include infrastructure to allow measurement of optical, wired and wireless links.

7.1-4 Operational data

The GENI system shall provide operational data on components to researchers.

7.1-5 Measurement transmission and storage

The GENI system shall provide networking and storage services for measurements.

7.1-6 Measured data access

The GENI system shall allow researchers to control access to their collected measurements while stored in measurement archives.

Discussion: Access control is required because, in some cases, measurements must be kept confidential until personal information can be revealed or deleted.

7.1-7 Real-time access to measurements

The GENI system shall provide real-time access to measured data.

Discussion: The intent of this requirement is to ensure that measured data is available for immediate use. The measurement handling system shouldn't unreasonably delay measurements. A potential application would be to permit development of (experimental) management systems which use measured data as an input. Note that we really don't understand the costs of this requirement and

whether it needs to apply everywhere. This requirement is likely to evolve as the conceptual design for handling measurements matures.

7.1-8 Component locations

The GENI system shall make information about the physical location of all components available to researchers.

7.1-9 Time services

The GENI system shall support a common timeframe for all measurements, synchronized to within TBD microseconds across the system.

This requirement is very poorly understood at present, and is likely to evolve considerably.

8 User Opt-In Requirements

8.1-1 Low barrier to entry for “opt in”

The GENI system shall provide low barrier to entry for opt-in users (end-users of experiments running in slices).

Discussion: This requirement could be improved (how low is “low”?). The goal is that GENI should not make it difficult for end-users to opt-into GENI-hosted services.

8.1-2 Opt-in user software installation

The GENI system shall include tools that enable researchers to install new software in popular operating systems on end-user machines (computers, cell phones, etc.), with user consent.

Discussion: End-users may choose to run experimental code on a system, allowing it to join in a GENI slice.

8.1-3 Internet opt-in

The GENI system shall permit opt-in users to connect via the Internet.

Discussion: Internet-connected end-users may connect to an experiment which may or may not use IP within GENI.

8.1-4 Native GENI opt-in

The GENI system shall permit opt-in users to directly connect where possible.

Discussion: GENI shouldn't preclude opt-in users from connecting computers or other systems directly to GENI-enabled infrastructure.

8.1-5 Anonymous Opt-in

The GENI system shall permit anonymous opt-in user participation in experiments.

8.1-6 Opt-in shutdown

The GENI system shall include mechanisms by which opt-in users can be switched out of experiments, for example if the experiment crashes.

Discussion: Opt-in user participation and resource contributions should be considered valuable and mechanisms should be developed to avoid causing harm to their systems or degrading their network service if they join a GENI slice.

9 System Sizing Requirements

9.1-1 Numbers of researchers

The GENI system shall support at least 100,000 [TBR] authorized researchers and 10,000 organizations at any given time.

Discussion: Concurrent experiments allow multiple researchers to share the facility. This is needed to facilitate long-running experiments. The number must be large enough so that the capacity for experiments on GENI is not a limiting resource.

9.1-2 Numbers of concurrent experiments

The GENI system shall support at least 1,000 [TBR] continuous, concurrent experiments.

Discussion: Concurrent experiments allow multiple researchers to share the facility. This is needed to facilitate long-running experiments. The number must be large enough so that the capacity for experiments on GENI is not a limiting resource.

9.1-3 Infrastructure scale

The GENI system shall support at least 10,000,000 [TBR] components on at least 1,000 [TBR] federated infrastructure suites.

Discussion: This requirement is intended to ensure GENI can accommodate sufficient numbers and diversity of systems to permit large-scale distributed systems experiments.

Discussion: Real-time access is required because a) slices may use measurements to change their behavior or b) experiments may be so long lived that it would be inconvenient to wait until the experiment is concluded to provide access to the data.

10 Operations and Security Requirements

This chapter provides high-level system requirements for operational aspects of the federated infrastructure, including issues such as mechanisms to implement policy, federated operations, system security, record-keeping, etc.

10.1 Policy Support

10.1-1 Clearinghouse-wide Policies

The GENI system shall provide mechanisms to implement clearinghouse-wide resource allocation policies.

Discussion: This will allow funding agencies or other component contributors to put overall constraints on how their components will be used. An example of this would be limiting the amount of resources a single graduate student can acquire. There is no requirement that a Clearinghouse should have only a single mechanism and, in fact, support for multiple mechanisms is desirable as there will likely be shifts in mechanism over time as technology improves or needs change.

10.1-2 Owner's Policies

The GENI system shall allow owners of individual aggregates and/or components to implement and enforce their own policies for resource management of these individual aggregates and/or components.

Discussion: This policy might be implemented at the component manager, in a broker, or in the clearinghouse.

10.2 Reliable, Predictable, and Transparent Operations

10.2-1 Reliable operations

The GENI system shall have high-enough reliability and availability so that it supports network science and engineering experiments (????).

10.2-1 Predictable operations

The GENI system shall provide researchers with good mechanisms for understanding the reliability of the infrastructure and software on which they intend to run an experiment, before they start running the experiment.

Discussion: Experimenters should have an idea of the robustness of the infrastructure they are using. Some experimenters will have a high tolerance for failure, others less so. This requirements is structured to permit a wide range of reliability in GENI components but still keep researchers informed about what they are using. This requirement might be generalized to 'robustness classes' where parameters such as mean-time-to-repair or online support hours are included.

As a potential requirement, GENI subsystems and aggregates shall declare and conform to one of the availability classes below:

Class A: 99% (i.e., unavailable 88 hours/year)

Class B: 95% (i.e., unavailable 18 days/year)

Class C: 90% (i.e., unavailable 36 days/year)

Class D: none

10.2-2 Visible operational status

The GENI system shall make sufficient data available that researchers and maintainers will be able to evaluate the availability and operational status of the system. Discussion: GENI components should export sufficient data to support development of services that represent the ‘health’ of GENI. This will be useful in informing users about the status and availability of potential components and may also support development of new management tools. In principle an operator should be able to gather information on a component-level granularity, but as the system gets larger, tracking every component will become infeasible. It is key that aggregation strategies do not hide important availability details and that operators can drill down when necessary. It is also important that the amount of availability and operational information remain manageable.

10.2-3 Help Desk

The GENI system shall provide help-desk services to assist researchers in using GENI and diagnosing problems.

Discussion: A goal for GENI is permit and encourage users who may not be experts, such as students. The role of the Help Desk is to provide assistance in diagnosing problems, for example, determining whether difficulties in a slice are due to buggy code or failures in the substrate. Note that this need not be a centralized facility and could be provided online by a combination of operators, developers, and other users.

10.2-4 Federated event escalation

The GENI system shall provide operations and management support for event management and escalation, including security events, within GENI and with those organizations that interconnect with GENI.

Discussion: GENI will need policies, procedures, relationships, and mechanisms to support communication of the details relating to network events, such as unexpected outages or security-related incidents, to affected and interested parties.

10.2-5 Federated operations data exchange

The GENI system shall support operational and management data exchange according to *[TBS] GENI O&M Policy* between GENI and operators/owners of federated components, aggregates, and networks.

Discussion: The GPO should establish one or more common data exchange formats and policies to support operations and Help Desk support.

10.2-6 Record keeping duration

The GENI system retain usage logs for 60 (TBR) days.

Discussion: Logs should be retained for sufficient period to allow attribution of bad behavior to a slice but not so long as to become a burden on component providers.

10.3 Secure Operations

10.3-1 System-wide policy and guidelines for operational security

The GENI system shall provide clear policy and guidelines for operational security, which shall be adopted and enforced by all federated entities within the overall system.

Discussion: Developing security policies should be done carefully lest a concern for one kind of research (say, new applications used by opt-in Internet users) inhibit other kinds of research (say, research in live malware). GENI should permit a wide range of research and not institute policies that take a lowest-common-denominator approach. One possibility is that system-wide security policy might be more of a meta-policy discussing things like the need for any aggregates to have a stated security policy that covered some specific topics and included certain contact information.

10.3-1 Control framework security

The GENI system's control framework shall be built and operated to best security practices, with a goal that it will not be compromised by an "outsider threat" during its design lifetime.

10.3-2 Identities

All GENI components, institutions, researchers, and slices shall be assigned a unique identity.

Discussion: Every GENI principal should have a different identity to support accountability. This does not require that a principal can have only one identity. (This is hard to enforce.)

10.3-3 Identity records

All parties assigning GENI identities shall record, verify and maintain the real world identity and contact information of the entity or party responsible for the entity's actions within the GENI system.

Discussion: This requirement is driven by the need for accountability of user behavior.

10.3-4 Authorization

The GENI system shall require authorization before allocation of resources according to *TBS User Authorization Policy*.

Discussion: In general, authorization is required to acquire GENI resources. However, this isn't intended to excluded anonymous use in limited cases. A separate document will be developed to discuss authorization policies.

10.3-5 Accountability

The GENI system shall permit network activity to be traced back to the responsible slice (& researcher).

Discussion: The primary motivation for this requirement is to identify sources of bad traffic.

10.4 Detection and remediation of bad experiments

10.4-1 Rapid detection and decisions

The GENI system shall provide mechanisms for rapidly detecting, analyzing, and neutralizing experiments that are currently causing harm to other entities, within the GENI system or external.

10.4-2 Swift, effective neutralization

The GENI system shall be able to neutralize a bad experiment within 1 second of the decision to do so, no matter how large or virulent the experiment.

Discussion: This requirement may impose constraints on the system that will inhibit participation in GENI. Meeting the requirement would require some additional infrastructure and design as well as influencing how local resources allocated to GENI are positioned in networks (i.e., behind "circuit breakers"). It could make attracting component donors harder with possible effects on user opt-in. Relaxing the timing may reduce the constraints on experimenters and donors without placing the Internet in much more jeopardy.

10.4-3 Restricted mode

Should the GENI system enter a period where activities of some components cannot be adequately monitored or controlled, it shall automatically restrict those activities by other means to a point where safety can be assured (e.g., by shutting down a slice or bringing GENI as a whole into a safe state).

10.4-4 Abuse reporting

The GENI system shall provide a point of contact for reporting abusive behavior.

Discussion: Many GENI-related operators will be able to use the GENI clearinghouse and operations to identify the source of bad traffic. But to permit external entities, such as commercial ISPs, to cause GENI to kill slices that are generating bad traffic, a central reporting point is required.

10.4-5 Forensics

The GENI system shall provide mechanisms for performing audits of system activity for after-the-fact analysis of system failures or abuses.

Discussion: The primary motivation for this requirement is to identify causes and sources of bad traffic.